Title
Essays on Network Games with Incomplete Information, with Applications in Finance

Author
Leister, Christian Matthew

Publication Date
2015

Peer reviewed|Thesis/dissertation

Essays on Network Games with Incomplete Information, with Applications in Finance

by

Christian Matthew Leister

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Shachar Kariv, Chair
Professor Haluk I. Ergin
Professor Nicolae B. Gârleanu

Spring 2015

Essays on Network Games with Incomplete Information, with Applications in Finance

Abstract

Essays on Network Games with Incomplete Information, with Applications in Finance

by

Christian Matthew Leister

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Shachar Kariv, Chair

This dissertations includes three (3) chapters, each adding to the growing network games literature that incorporates incomplete information. Financial over-the-counter markets give motivating applications. (1) *Trading Networks and Equilibrium Intermediation* studies the efficiency of trade in networks. A network of intermediaries facilitates 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. 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. (2) *Interdealer Trade: Risk, Liquidity, and the role of Market Inventory* further studies traders facing private shocks, placed in a dynamic setting. Trades between ex ante symmetric, inventory carrying intermediaries ("dealers") are motivated by divergent liquidity needs of the counter parties. Market prices and asset flows are pinned by dealers' indifference between providing intermediation services and retaining liquidity to be utilized in subsequent interdealer markets. More active interdealer markets simultaneously increase the value to intermediation and the option-value to providing these services. Under infrequent shocks, interdealer trade boosts the availability of liquidity in the broader market. This boost decays with market inventory, which serves as a constraint on interdealer activity. Through this market mechanism, prices vary inversely with both search frictions between dealers and on their total current holdings. (3) *Information Acquisition and Response in Peer-effects Networks* endogenizes the quality of information that market participants carry in a general peer effects model. When pairwise peer effects are symmetric, asymmetries in acquired information are inefficiently low relative to the utilitarian benchmark. And with information privately acquired, all players face strictly positive gains to overstating their informativeness as to strategically influence the beliefs and behaviors of neighbors. If strategic substitutes in actions are present and significant, low centrality players move against their signals in anticipation of their neighbors' actions. A blueprint for optimal policy design is developed. Applications to market efficiency in financial crises and two-sided markets are discussed.

Acknowledgments

The following chapters have greatly benefited from a number of people. While each chapter includes its own set of acknowledgments on each of their respective title pages, here I'd like to express particular thanks to five extremely important influences. First, my advisor Shachar Kariv has been the quintessence of a supportive mentor over the past six years. He has been a persistent inspiration and I was truly blessed to stumble onto his undergraduate course as my first formal introduction to game theory seven years ago. Since, he has espoused a growing passion for the field and its potential to better understand the diversity of pertinent multi-agent relationships that shape conflict in societies. I am truly proud to call this field my profession, and even more Shachar my mentor. Nicolae Gârleanu has also been an extremely active advisor to my work. Nicolae's keen understanding of underlining, driving mechanisms has pushed the development of my work from early states into their later development. My deep interest in the potential of network and game theory to contribute to the financial discussion is a product of his engagement. Haluk Ergin has also been an invaluable influence. Haluk has pushed the rigor of my modeling as much as anyone. My work has prospered significantly from his tutelage. My coauthor Maciej Kotowski has also contributed to all of my work. Our long skypes and back-and-forth edits have been extremely educational and fulfilling. Finally, Sanjeev Goyal has been a great influence over the past year. Sanjeev's guidance and encouragement during my visits to Cambridge have been extraordinary, and I can not sufficiently express my appreciation for his engagement in my work.

To my parents

Contents

1 General Introduction

The following three chapters study network games of different forms, each under some version of incomplete information. Throughout all of these works, theory is persistently applied toward better understanding the dynamics of over-the-counter exchange, both financial and in more general settings (though with emphasis on the first). The ways in which the network exists as a friction in the economy vary across the chapters. The first two chapters, focusing on intermediation, have the network capturing the set of feasible trades in the market. The final chapter abstracts at a higher level, allowing the network to capture the set of "peer effects", or preference dependencies in the economy. The interaction of incomplete information with the constraints of the given network's structure, and how these frictions interplay to influence equilibrium behavior and market efficiency, establish an encompassing theme maintained throughout the dissertation.

Financial intermediaries, or synonymously inventory carrying "dealers" and "market makers", have been extensively studied by the financial economics literature. Dealers' equilibrium bid-ask spread establishes one important unit of analysis, quantifying market inefficiencies including dealers' captured rents while establishing an important component of round-trip transaction costs faced by investors. A classic market microstructure literature starting with Garman (1976) [7] and shortly thereafter with Amihud and Mendelson (1980) [1] focuses on markets in which all trades are processed through a monopolistic market maker that optimally influences bid-ask spread.[1] Glosten and Milgrom (1985) [8] later show how asymmetric information amongst investors implies equilibrium bid-ask spread under the presence of a risk neutral, zero-profit market maker. Departing from the assumption of a single market maker facilitating all transactions, Grossman and Miller (1988) [9] reinvented the literature by allowing for investors to trade directly, while showing that discounting, random investor arrival, and the immediate accessibility of a dealer's liquidity can rationalize dealers' bid-ask spread.[2]

While these seminal works made crucial contributions to our understanding of financial intermediation, the financial networks literature takes an important departure from these models. In many over-the-counter financial markets, including corporate and municipal bonds, "dark pools", and markets for rare assets, the market's architecture can imply a number of trades prior to an asset finding an efficient home.[3] Heterogeneity in the structure of the market, in the form of an incomplete set of persistent trading relationships, can be captured using a network. In such settings, market participants that are directly connected in the network are free to trade with each other, while those distant in the network are excluded from transacting with each other. Gale and Kariv (2007,2009) [5, 6] provide early work

[1] Amihud and Mendelson (1986) [2] extended their setting to a competitive dealer market with cross-sectional heterogeneity in asset trading costs, which drives market equilibrium bid-ask spread.

[2] Thus, dealers create value by their presence in the market. For contributions in similar spirit incorporating search, see Rubinstein and Wolinsky (1987) [18] and Duffie, Gârleanu and Pedersen (2005) [4] for a more contemporary contribution.

[3] For municipal bonds, for example, Li and Schürhoff (2012) [16] see roughly 20% of all transactions involving two or more dealers, with extensive transaction chains involving up to seven dealers.

establishing the efficiency of trade in these general networked markets. Given the absence of trading costs, moderate discounting and sufficient access to dealer liquidity, they establish efficiency in both the final allocation of assets and in their pricing, with intermediaries demanding zero bid-ask spread in equilibrium. Gale and Kariv (2007) [5] end by suggesting future work studying the interplay of other market frictions with the market's structural incompleteness. With incomplete and asymmetric information a fundamental friction in real financial exchange, the following chapters advance this agenda, incorporating incomplete (and in the fourth chapter, endogenous) information into network markets.

As the first chapter, "***Trading Networks and Equilibrium Intermediation***" [13] with Maciej Kotowski, is the first paper known to the authors that explores the role of private trading costs on the efficiency of trade in networks. The paper assesses the resilience of equilibrium intermediary networks in facilitating trade in the presence of these dealer-level shocks. The paper separately explores two dimensions of network resilience in the presence of cost shocks. The first dimension addresses network stability, defining an unstable network by the presence of a partnership of adjacent traders that optimally behave as a single trader.[4] The second dimension addresses the extent of competitive intermediation by allowing entry into the market. Traders entering a particular row of the network both compete with others in the row and provide valuable intermediation services to upstream and downstream traders. A generic under entry is established[5], with thin and unbalanced network structures forming in equilibrium.[6] The paper finishes by addressing how these dimensions interplay, showing that stability is most obtainable in thin equilibrium network structures when multiple equilibria persist.

"***Interdealer Trade: Risk, Liquidity, and the role of Market Inventory***" [14] further studies the role of complementarities between dealers facing private random shocks. In a dynamic setting, ex ante identical dealers intermediate an upstream supply and downstream demand. High valuing consumers arrive via a Poisson process, while dealers hold inventories in order to time their arrival. In the event that a dealer is unable to sell to a high valuing consumer prior to realizing a liquidity shock, an interdealer market allows her to transfer her asset to a dealer carrying available liquidity. In each period's market equilibrium, dealers choose either to intermediate asset flows or retain their liquidity to be used in future dealer markets. Interdealer trading frictions are captured by a process that generates the market's network structure. In this setting, a basic complementarity between dealers taking on these two distinct market roles is established. With dealers observing market flows and aggregate inventories, a novel link between market inventory and the asset's price is established working solely through the dynamics of interdealer trade. The paper shows that the need for interdealer trade heightens as all dealers face more frequent shocks. Without interdealer trade, dealers' private risks would be fully priced in upstream markets.

[4] The paper shows that formation of partnerships only reduces aggregate welfare –in this setting, given by the likelihood that the network successfully transfers the asset to a consumer.

[5] This under entry in intermediation markets contrasts with the classic over entry in competitive production markets, as illustrated by Mankiw and Whinston (1986) [17].

[6] Further, multiple equilibrium structures may arise, with significantly thin networks being dominated by a maximal –yet still inefficiently thin– equilibrium.

Thus, active interdealer trade boosts the asset's price in the broader market, with market inventory the essential state variable used toward assessing both the aggregate need and accessibility of dealer-provided liquidity.

Finally, "***Information Acquisition and Response in Peer-effects Networks***" [15] takes a higher-level approach in capturing network effects, and endogenizes the quality of information that market participants carry. A set of "players" simultaneously act in a second stage under a network of signed, weighted, and directed peer effects. Prior to this stage, each player invests in payoff relevant information. With correlation in players' payoffs introduced, signals take on a dual role: they inform both of the state of the world and of what other players observe. In this setting, the essential network property driving the direction of equilibrium inefficiencies is the extent of symmetry in pairwise peer effects. In highly symmetric economies, such as industrial organizations or social networks where pairs either both compete or both coordinate with each other, information investments are "*bunched*": the equilibrium gap between the most informed and least informed players is inefficiently small. In economies in which anti-symmetric relationships are prominent, such as in financial crises as discussed below, the directions of inefficiencies in information acquisition reverse. The paper further shows that with information privately acquired, players face strategic use to acquiring information not internalized in equilibrium.[7] When pairwise peer effects are symmetric, these strategic forces take on a clear positive direction: players throughout the network face the incentive to exaggerate their informativeness. The size of the strategic value to acquiring information is driven by each player's connectedness in the network: by their "influence" on the network.[8]

The paper finishes with an application to financial markets in liquidity crises. When traders are unconstrained in their asset positions, strategic substitutes in competitive market equilibrium implies strategic substitutes in information acquisition. Traders therefore over acquire information. In crises, some subset of distressed traders face severe funding constraints. In the spirit of the liquidity spirals studied in Brunnermeier and Pedersen (2009) [3], these liquidity constrained traders are assumed to exhibit upward sloping demands for assets, as they retain their inventories when short-term prices are high. A paradigm shift is realized as the proportion of constrained traders grows large. Traders throughout the market now under acquire information. Constrained traders impose positive externalities on each other as they aim to coordinate on assets in high states of the world. Unconstrained traders fail to internalize the value that their information investments provide to constrained traders. The application then applies the paper's policy analysis: identification of constrained traders (e.g. via stress tests) coupled with certification of their information investments influences these traders to strategical acquire additional information. This encourages the constrained side of the market to collectively acquire more information, increasing aggregate welfare.

The dissertation is organized respecting the above order, both because this gives the chronology of the development of the papers and because this provides a logical order in their

[7] Put differently, if players could publicly invest in information, they would adjust their acquisition choices as to directly influence others' beliefs regarding how informed they are.

[8] Under anti-symmetric relationships these implications again reverse: all players face negative strategic value to information, with the incentive to *understate* informativeness.

theoretical contributions. The first two [13], [14] extend the trade-on-networks literature to incorporate incomplete information in the form of trading costs and demand uncertainty, the third [15] studies inefficiencies in heterogeneous settings when information is rationally acquired. The three works logically build on each other, each informing the next with relevant insights. The chapters together offer a dialogue, informing of the role of incomplete and asymmetric information in over-the-counter exchange, while more broadly adding to a rich and growing literature on information in network games.

References

[1] Amihud, Yakov; Mendelson, Haim. "Dealership market: Market-making with inventory". *Journal of Financial Economics*, 8(1):31–53, March 1980.

[2] Amihud, Yakov; Mendelson, Haim. "Asset pricing and the bid-ask spread". *Journal of Financial Economics*, 17(2):223–49, December 1986.

[3] Brunnermeier, Markus; Pedersen, Lasse Heje. "Market Liquidity and Funding Liquidity". *The Review of Financial Studies*, 22(6):2201–38, June 2009.

[4] Duffie, Darrell; Grleanu, Nicolae; Pedersen, Lasse Heje. "Over-the-Counter Markets". *Econometrica*, 73(6):1815–47, November 2005.

[5] Gale, Douglas; Kariv, Shachar. "Financial Networks". *American Economic Review: Papers and Proceedings*, 97(2):99–103, 2007.

[6] Gale, Douglas; Kariv, Shachar. "Trading in Networks: A Normal Form Game Experiment". *American Economic Journal: Microeconomics*, 1(2):114–132, 2009.

[7] Garman, Mark. "Market Microstructure". *Journal of Financial Economics*, 3:257–275, 1976.

[8] Glosten, Lawrence; Milgrom, Paul. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders". *Journal of Financial Economics*, 14(1):71–100, March 1985.

[9] Grossman, Sanford; Miller, Merton. "Liquidity and Market Structure". *Journal of Finance*, 43(3):617–33, July 1988.

[10] Grossman, Sanford; Stiglitz, Jospeh. "On the Impossibility of Informationally Efficient Markets". *American Economic Review*, 70(3):393–408, June 1980.

[11] Hellwig, Christian; Veldkamp, Laura. "Knowing What Others Know: Coordination Motives in Information Acquisition". *The Review of Economic Studies*, 76:223–251, 2009.

[12] Ho, Thomas; Stoll, Hans. "The Dynamics of Dealer Markets under Competition". *Journal of Finance*, 38(4):1053–74, September 1983.

[13] Kotowski, Maciej; Leister, C. Matthew. "Trading Networks and Equilibrium Intermediation". *mimeo*, 2014.

[14] Leister, C. Matthew. "Interdealer Trade: Dealer Risks, Liquidity, and the role of Market Inventory". *mimeo*, 2014 A.

[15] Leister, C. Matthew. "Information Acquisition and Response in Peer-effects Networks". *mimeo*, 2014 B.

[16] Li, Dan; Schürhoff, Norman. "Market Microstructure". *Journal of Financial Economics*, 3:257–275, 2012.

[17] Mankiw, Gregory; Whinston, Michael. "Free Entry and Social Inefficiency". *RAND Journal of Economics*, 17(1):48–58, 1986.

[18] Rubinstein, Ariel; Wolinsky, Asher. "Middlemen". *Quarterly Journal of Economics*, 102(3):581–93, August 1987.

Trading Networks and Equilibrium Intermediation

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

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

[*]John F. Kennedy School of Government, Harvard University. E-mail: maciej_kotowski@hks.harvard.edu
[†]University of California, Berkeley. E-mail: leister@berkeley.edu
[1]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.

[Figure 1]

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 2), 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" [21, 22], are easily

[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.

discernible in the equilibrium networks of our model.

We develop our argument progressively. Section 2 introduces our model. Throughout, we take the presence of intermediaries as given and we focus on the interactions among them.[2] In section 3 we examine price-formation and exchange taking the network structure as given. Section 4 considers network stability and we propose our model of network formation in section 5. Section 6 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 7. 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.

2 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)[12]. 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 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

[2] We do not model the underlying reason a particular market features intermediaries, though such arrangements are undoubtably common. Spulber(1996)[32] 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)[12]. Choi et al. (2014)[4] examine a related class of networks, which they call "multipartite networks."

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.

[Figure 2]

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. 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]

[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 $\mathbf{n} = (n_r, \mathbf{n}_{-r})$.

[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.

(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.

2.1 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)[19] or Patil (2011)[29], 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 [30], bargaining [7, 23, 9, 6, 31], or posted prices [4].[9] Manea (2014)[24] 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.

[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 (see Myerson (1982)[26]) with resale (see Zheng (2002) [35]) is a possible generalization.

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 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 (see Kremer (1993)[20])—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)[15] 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)[12], Manea (2014)[24], Wright and Wong (2014)[34], or Choi et al.(2014)[4]—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.

[10] Nagurney and Qiang (2009)[27] employ networks with similar structures to ours to describe production inside a firm.

3 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.

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)[12] and Gale et al. (2015)[13] 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.

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

[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.

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.)

4 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 5 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 $\bar{m} = \max\{r : m_r \geq 1\}$ and $\underline{m} = \min\{r : m_r \geq 1\}$ 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.

[Figure 3]

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 trader in the partnership $\mathbf{m}$ 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.

[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.

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

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.

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 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,

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

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 [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

[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]		+

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 anti-trust 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.

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

4.1 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)[17] propose a stability notion whereby a fixed set of agents can form and drop links. Ostrovsky (2008)[28] models supply-chain networks and proposes a generalization of "stability," in the sense of Gale Shapley (1962)[10], to that class of problems. Manea (2014)[24] 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 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.

[16] This observation accommodates Ostrovsky's (2008)[28] 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)[28].

4.2 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.

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.

[17] Echoing Theorem 4, $\chi_{\mathbf{m}}(\mathbf{n})$ can be shown to equal a sum of expected profits.

5 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 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)[14] 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.

5.1 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

[18]See Jackson (2008)[16] for a survey of network formation.

[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.

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 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.

[Figure 4]

[Figure 5]

5.2 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 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. (1997)[21, 22].[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. (1997)[22], 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

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

[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.

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 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.

5.3 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

$1 \le r \le R$ 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 $\hat{\mathbf{n}} \ge \mathbf{n}^*$.[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 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

[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

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.

[Figure 6]

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)[1] who show that small sectoral 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

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

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.

6 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 $R = 2$.

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.

[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.

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.

[Figure 7]

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.

7 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

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

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)[8].[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)[3], Kranton and Minehart (2001)[19], or more recently Condorelli and Galeotti (2012a)[5], 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 by GaleKariv (2007,2009)[11, 12] 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)[30] 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

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

[28] Jackson (2008)[16] provides a comprehensive survey of the literature on economic networks.

[29] Wright and Wong (2014)[34] also examine chains of intermediation, though in a search-theoretic context.

organizations and, like here, has been noted to imply cross-cutting implications for efficiency Mankiw and Whinston (1986)[25]. Whereas Mankiw and Whinston (1986)[25] 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)[4] 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)[18] and Arrow (1975)[2] 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.

References

[1] Acemoglu, Daron; Carvalho, Vasco; Ozdaglar, Asuman; Tahbaz-Salehi, Alireza. "The Network Origins of Aggregate Fluctuations". *Econometrica*, 80(5):1977–2016, 2012.

[2] Arrow, Kenneth. "Vertical Integration and Communication". *The Bell Journal of Economics*, 6(1):173–183, 1975.

[3] Bala, Venkatesh; Goyal, Sanjeev. "A Noncooperative Model of Network Formation". *Econometrica*, 68(5):1181–1229, 2000.

[4] Choi, Syngjoo; Galeotti, Andrea; Goyal, Sanjeev. "Trading in networks: Theory and Experiments". *Working Paper 2014/08, Cambridge-INET Institute*, 2014.

[5] Condorelli, Daniel; Galeotti, Andrea. "Bilateral trading in networks". *mimeo*, 2012a.

[6] Condorelli, Daniel; Galeotti, Andrea. "Endogenous trading networks". *mimeo*, 2012b.

[7] Corominas-Bosch, Margarida. "Bargaining in a Network of Buyers and Sellers". *Journal of Economic Theory,*, 115(1):35–77, 2004.

[8] Economides, Nicholas. "The Economics of Networks". *International Journal of Industrial Organization*, 14(6):673–699, 1996.

[9] Elliott, Matt. "Inefficiencies in Networked Markets". *mimeo*, 2013.

[10] Gale, David; Shapley, Lloyd. "College Admissions and the Stability of Marriage". *The American Mathematical Monthly*, 69(1):9–15, 1962.

[11] Gale, Douglas; Kariv, Shachar. "Financial Networks". *American Economic Review: Papers and Proceedings*, 97(2):99–103, 2007.

[12] Gale, Douglas; Kariv, Shachar. "Trading in Networks: A Normal Form Game Experiment". *American Economic Journal: Microeconomics*, 1(2):114–132, 2009.

[13] Gale, Douglas; Kariv, Shachar; Kotowski, Maciej; Leister, C. Matthew. *work in progress*, 2015.

[14] Gary-Bobo, Robert Jean. "On the Existence of Equilibrium Points in a Class of Asymmetric Market Entry Games". *Games and Economic Behavior*, 2(3):239–246, 1990.

[15] Gofman, Michael. "A Network-based Analysis of Over-the-counter Markets". *mimeo*, 2011.

[16] Jackson, Matt. "Social and Economic Networks". *Princeton University Press, Princeton, NJ*, 2008.

[17] Jackson, Matt; Wolinsky, Asher. "A Strategic Model of Social and Economic Networks". *Journal of Economic Theory*, 71(1):44–74, 1996.

[18] Kranton, Rachel; Minehart, Deborah. "Networks versus Vertical Integration". *Journal of Economics*, 31(3):570–601, 2000.

[19] Kranton, Rachel; Minehart, Deborah. "A theory of Buyer-seller Networks". *American Economic Review*, 91(3):485–508, 2001.

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

[21] Lee, Hau; Padmanabhan, V.; Whang, Seungjin. "The bullwhip effect in supply chains". *MIT Sloan Management Review*, 38(3), 1997a.

[22] Lee, Hau; Padmanabhan, V.; Whang, Seungjin. "Information Distortion in a Supply Chain: The Bullwhip Effect". *Management Science,*, 43(4):546–558, 1997b.

[23] Manea, Mihai. "Bargaining in Stationary Networks". *American Economic Review*, 101(5):2042–80, 2011.

[24] Manea, Mihai. "Intermediation in networks". *mimeo*, 2014.

[25] Mankiw, Gregory; Whinston, Michael. "Free Entry and Social Inefficiency". *RAND Journal of Economics*, 17(1):48–58, 1986.

[26] Myerson, Roger. "Optimal Auction Design". *Mathematics of Operations Research*, 6(1):58–73, 1981.

[27] Nagurney, Anna; Qiang, Qiang. "Fragile Networks: Identifying Vulnerabilities and Synergies in an Uncertain World". *John Wiley Sons, Hoboken, NJ*, 2009.

[28] Ostrovsky, Michael. "Stability in Supply Chain Networks". *American Economic Review*, 98(3):897–923, 2008.

[29] Patil, Hemant. "Buyer-seller Networks with Demand Shocks and Intermediation". *Review of Economic Design*, 15(2):121–145, 2011.

[30] Rubinstein, Ariel; Wolinsky, Asher. "Middlemen". *Quarterly Journal of Economics*, 102(3):581–594, 1987.

[31] Siedlarek, Jan-Peter. "Intermediation in networks". *mimeo*, 2012.

[32] Spulber, Daniel. "Market Microstructure and Intermediation". *Journal of Economic Perspectives*, 10(3):135–152, 1996.

[33] Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders". *Journal of Finance*, 16(1):8–37, 1961.

[34] Wright, Randall; Wong, Yuet-Yee. "Buyers, Sellers, and Middlement: Variations on Searchtheoretic Themes". *International Economic Review*, 52(2):375–397, 2014.

[35] Zheng, Charles. "Optimal Auction with Resale". *Econometrica*, 70(6):2197–2224, 2002.

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 8.[30]

[Figure 8]

(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 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.

[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.

[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.

(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, $\tilde{\nu}_r = \tilde{\delta}_{r-1}\tilde{b}_{r-1}$. $\square$

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.

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. $\square$

Proof. 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)[33] 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. 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. 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. 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. Theorem 5 (1) From (2), $\pi_r(\mathbf{n}^*) \leq p(1-p)^{n_r-1}$; therefore,

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

(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 $\pi_r(\mathbf{n}^0) - \kappa \geq 0$ for all r. Define $\mathscr{Q}_r(\mathbf{n}) = \{\tilde{n}_r \in \mathbb{N} \colon \pi_r(\tilde{n}_r, \mathbf{n}_{-r}) - \kappa \geq 0, n_r \leq \tilde{n}_r \leq \bar{n}\}$ 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 $\mathscr{Q}_R(\mathbf{n}) \neq \emptyset$. 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. 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. Corollary 1 By Theorem 6, the set of equilibria are a directed set. This set is finite. The conclusion follows. □

Proof. 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

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

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. 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,

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

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 $\bar{r} = \max\{r \colon n_r^* = n_1^*\}$. Since $\mathbf{n}^*$ is an equilibrium,

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

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, $\mathbf{n}^* \leq \hat{\mathbf{n}}$. □

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 $\underline{m} = \min\{r \colon m_r \geq 1\}$ and $\bar{m} = \max\{r \colon m_r \geq 1\}$. For example, consider Figure 9, 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 9]

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 $r \notin \{\underline{m}, \ldots, \bar{m}-1\}$ 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 $r \in \{\underline{m}, \ldots, \bar{m}-1\}$ 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 $r \in \{\underline{m}, \ldots, \bar{m}-1\}$ 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, $\tilde{\nu}_{\mathbf{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. $\square$

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, $\tilde{\nu}_r = \prod_{k=\underline{m}}^{r} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$. $\square$

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*

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

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 $\tilde{b}_{\underline{m}-1} \le v$, it follows that $\tilde{\nu}_{\mathbf{m}} \le v$. 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 $\tilde{\nu}_{\mathbf{m}} \ge \tilde{b}_{r-1}$.
   - 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

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

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

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

   where, given $\sigma^*_{-\mathbf{m}}$ and the assumed evolution of 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)[33] 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.

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)[21].

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:

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

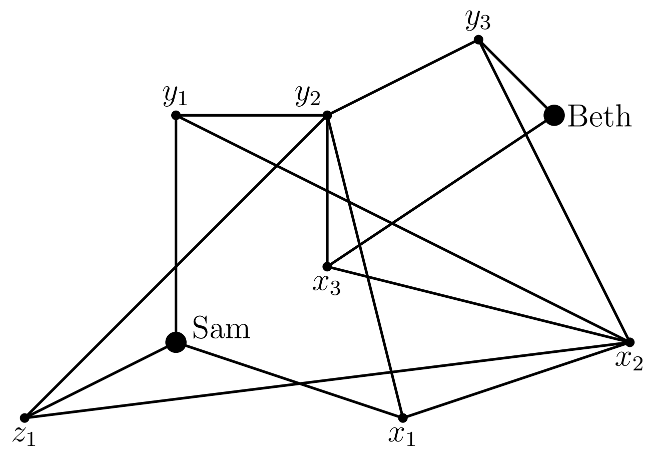

Figure 1: A trading network.

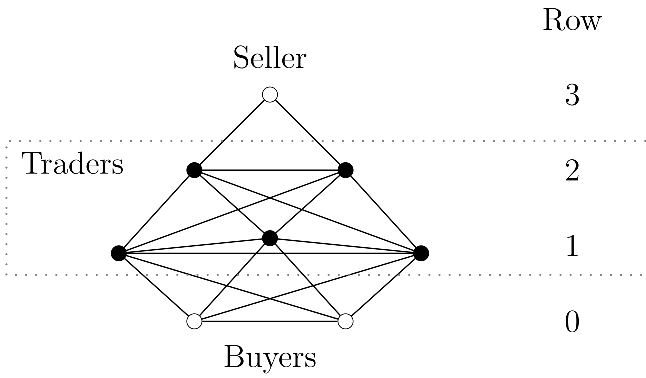

Figure 2: A trading network with configuration $\mathbf{n} = (3, 2)$.

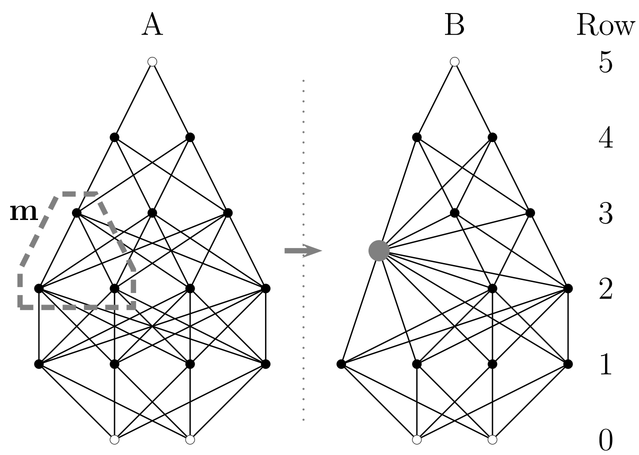

Figure 3: Formation of the partnership $\mathbf{m} = (0, 2, 1, 0)$. (Within-row links are omitted for clarity.)

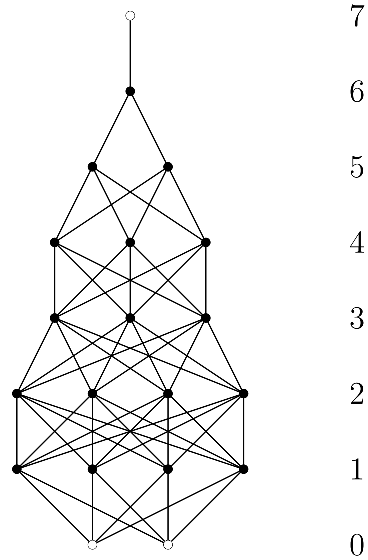

Figure 4: The equilibrium $\mathbf{n}$ in Example 3. (Within-row links are omitted for clarity.)

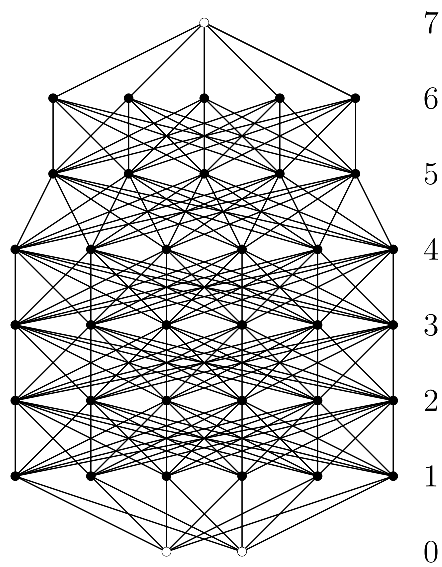

Figure 5: The equilibrium $\mathbf{n}'$ in Example 3. (Within-row links are omitted for clarity.)

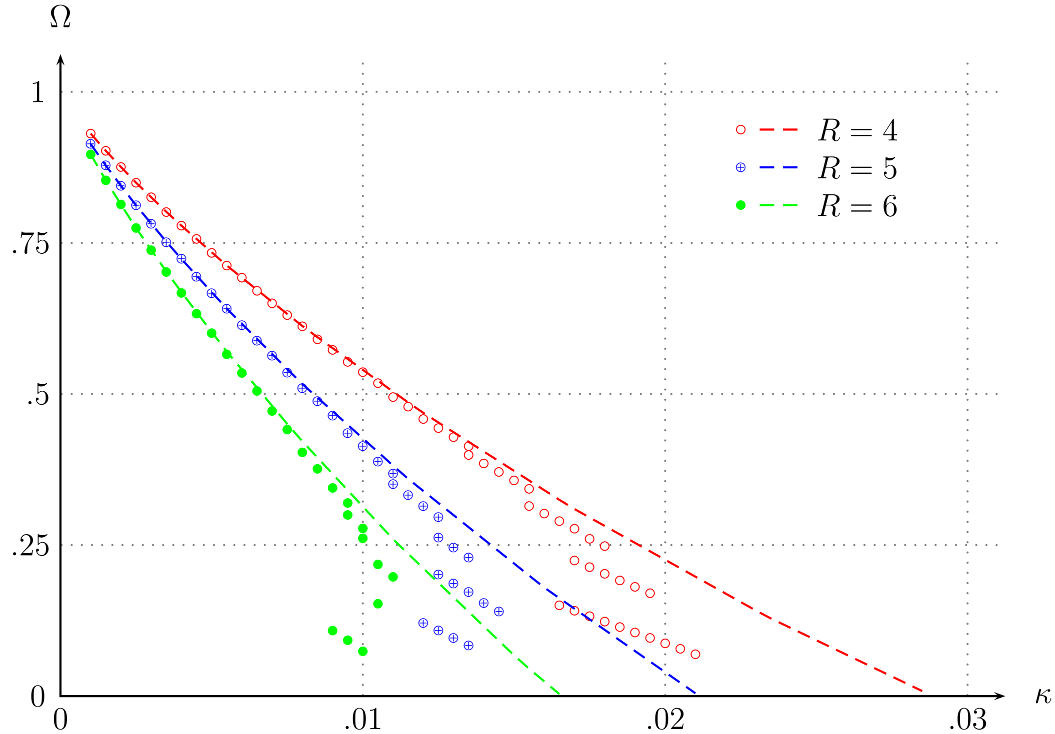

Figure 6: Welfare and entry costs in Example 5.

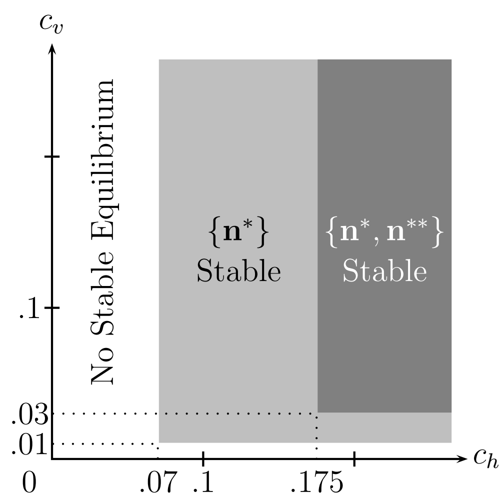

Figure 7: Stable equilibria as a function of (c_h, c_v) in Example 6.

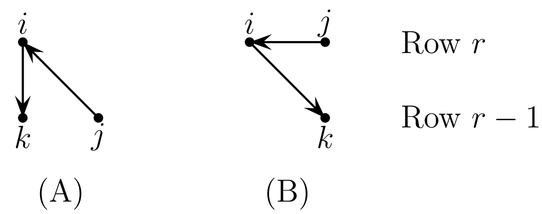

Figure 8: Potentially profitable transactions.

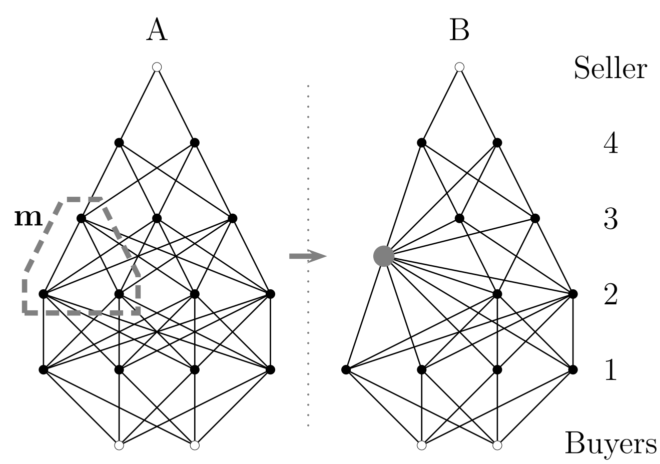

Figure 9: The formation of the partnership $\mathbf{m} = (0, 2, 1, 0)$. (Within-row links are not illustrated for clarity.)

Interdealer Trade: Risk, Liquidity, and the role of Market Inventory

C. Matthew Leister*

Abstract

This paper studies interdealer trade as an underlining market mechanism linking market inventories to asset prices and liquidity. Trades between ex ante symmetric intermediaries ("dealers") are motivated by divergent liquidity needs of the transacting parties. These trades provide a hedge against inter-temporal inventory risks. Market prices and asset flows are pinned by dealers' indifference between providing intermediation services and retaining liquidity to be utilized in subsequent interdealer markets. Thus, more active interdealer markets simultaneously increase the value to intermediation and the option-value to providing these services. Given moderate private risks, interdealer trade boosts the availability of liquidity in the broader market. This boost decays with market inventory, which serves as a constraint on interdealer activity. Through this market mechanism, prices vary inversely with both search frictions between dealers and on their total current holdings.

Keywords: Intermediation, Dealer Market, Interdealer Trade, Market Liquidity, Market Inventory, Trading Networks, Assignment Game.

JEL: C71; C78; D21; G12.

*University of California, Berkeley. E-mail: leister@berkeley.edu

[1]I'd like to thank Ivan Balbuzanov, Haluk Ergin, Nicolae Gârleanu, Shachar Kariv, and Maciej Kotowski, as well as Rachel Kranton for a motivating discussion.

9 Introduction

Trades between financial intermediaries ("interdealer trade") comprise a quarter or more of total volume in both centralized markets[1] and over-the-counter markets[2]. Yet, the theoretical motivation for these transactions has only been partially studied.[3] The ultimate impact of interdealer activity on asset prices in a dynamic market setting remains an open question. This paper contributes with a dynamic framework of interdealer trade, which is derived as a value-creating market equilibrium response to dealers' private risks.

The efficiency in the market's allocation of assets and liquidity depend on the abilities of financial intermediaries to seamlessly transfer inventory flows from suppliers to the highest valuing investors. If these dealers realize their own liquidity needs through this process –as this paper captures– then the need for interdealer trade naturally arises. And with dealers actively trading assets amongst each other, augmenting both supply and demand from investors, the prices that market participants face intimately depend on the collective trading behaviors of dealers.

Toward capturing this dynamic, the following market equilibrium model incorporates interdealer trades, playing an essential role in the intermediation process. These transactions are motivated as rational responses to the asymmetric liquidity needs of the dealers. Search frictions limiting the extent that dealers can find and transact with each other further shape expectations. Under this broad setting, interdealer markets are cast as funding sources boosting the availability of liquidity in the broader market.

There can be a number of sources to asymmetries across dealers' liquidity needs. For example, portfolio hedging alone can derive heterogeneity, with dealers responding to market conditions asymmetrically as a function of their private investment strategies. Dealers may also face constraints past on from the funding needs of their managing firms. In the presence of these risks, carried inventories can face compulsory liquidations. And when investors are unable to provide dealers with liquidity in a timely manner, other dealers can 'fill the void' by purchasing assets at higher prices. Thus, time varying heterogeneity in dealers' needs for liquidity introduces a common use for an active interdealer market, insulating dealers from a convolution of idiosyncratic risks and uncertainty in demand.

The model design is intended to capture gains to interdealer trade at a general level, while abstracting away from other market intricacies. Toward the latter, dealers are cast as risk-neutral profit maximizers, who physically connect an 'upstream' supply market to a 'downstream' demand market. Earnings are derived from trading profits.[4] Discounting (i.e. costs of capital) introduces immediacy, with expectations over future trade pinning continuation values. As dealers time the arrival of high valuing demand, equilibrium inventories and prices in both the upstream and interdealer markets are obtained. Each period, dealers can trade between each other in an 'interdealer market' when the arrival of downstream demand falls short of the liquidity collectively required by the dealers. Interdealer markets are modeled with a search process followed by an assignment game. This environment is flexible enough to capture the various constraints that dealers may face as they search for each

[1] e.g. the London Stock Exchange; see Reiss and Werner (1998) [24]

[2] e.g. the U.S. corporate bond market; see Herndershott and Madhavan (2013) [17]

[3] See Section 9.1 for a discussion of important contributions

[4] This abstracts away from collected cash flows (coupons, dividends).

other, and broad enough to allow for multiple interdealer prices as they realize heterogeneous bargaining positions in the market.

After constructing the market environment, the paper characterizes the market's steady state. First, the familiar notion of pairwise stability is employed, which constrains the set of interdealer transactions and prices. Using classic results from the two-sided matching literature, a generic monotonic dependence of expected interdealer gains on market inventories is established. With a cooperative game theoretic flavor, this approach shifts the focus toward the transfer of value and away from the many complications embodied in the equilibrium strategies driving interdealer activity. A basic relationship between intermediation and the provision of liquidity by dealers manifests itself through interdealer trade. These two market roles are seen to function as mutually exclusive strategic compliments. Precisely, the equilibrium extent of total market inventory inflates the value captured by dealers providing liquidity. Conversely, as the provision of liquidity deteriorates (e.g. as market inventory increases) so too do the incentives to retain and carry inventories.

The paper then formally establishes links from market inventories to the asset's upstream price. The expected provision of liquidity from dealers increases with the frequency of private shocks. This is because dealers place greater likelihood in the need for interdealer trade, thus increasing the option value to retaining liquidity. Further, the sensitivity of the asset's upstream price to market inventory increases with the frequency of private shocks. This is precisely because greater market inventory implies an expected degradation in interdealer markets effectively insulating risk. This dependence between prices and market inventories originate from the dynamics of interdealer trade.

The paper concludes by offering various extensions, and assesses the robustness of the model's results upon loosening its assumptions. Extensions in both the inventory space and the model's binary form of liquidity risk preserve the effects summarized above. Finally, the discussion offers potential avenues for future work, as well as potential empirical implications.

The following sections are organized as follows. Below, related models are reviewed. Sections two and three present the setup, formal model and develop the paper's treatment of interdealer trade. Section four characterizes equilibrium dealer behavior. Section five gives discussion and concludes. A list of model variables is provided after the conclusion. Formal proofs to results and details for the numerical solutions are collected in two Appendixes.

9.1 Related Literature

As the following model captures, the market's response to dealers' private liquidity risks is to pass on these costs to investors in the form of bid-ask spread. A deep literature has extensively modeled many of the costs faced by constrained monopolist intermediaries (e.g. Glosten and Milgrom (1985) [12] and Amihud and Mendelson (1980) [2]). This paper, on the other hand, extends a complimentary literature analyzing dealer costs in a broader market setting. In particular, the model captures a dependence of trading costs on frictions that constrain interdealer activity.

More related theoretical contributions include Ho and Stoll (1983) [18] who explain interdealer trade as portfolio rebalancing and inventory risk sharing.[5] While their analysis

[5]As such, risk aversion plays an essential role in deriving value to interdealer transactions. In this paper's

does not derive market equilibrium inventories, the following setting's more stylized setup allows for an endogenous derivation of inventory flows. A comparable risk sharing story is derived, but with risk neutral dealers trading in response to time varying liquidity needs. By deriving inventories in market equilibrium, the causal link between prices and inventory works in both directions.

The models of Grossman and Miller (1988) [13] and Duffie, Garleanu, and Peterson (2005) [8] motivate intermediation as market-makers supply immediacy, absorbing asset flows from customers arriving to the market with asynchronous timing. The following, rather, abstracts away from mechanisms motivating the need for intermediation, leaving dealers to physically connect upstream and downstream markets that would otherwise be left disconnected. In this setting, dealers facing private risks effectively provide immediacy to each other.

While Duffie, Garleanu, and Peterson (2005) [8] take important steps toward deriving steady state inventory flows into and out of the dealer market, their model remains silent on the role of interdealer trade, allowing dealers to transfer assets between each other without friction. The authors show that steady state trading costs vanish as the dealer market becomes increasingly competitive. Without friction in the interdealer market, their unique interdealer price is pinned by the short side of the market. The following model, on the other hand, focuses in on the dealer market by introducing heterogeneity in the form of private liquidity risk in presence of interdealer search frictions. Thus, the inventory effects that are captured here motivate future work quantifying the efficiency implications in markets where investors can directly search for each other.

This paper also relates to the literature on trading networks. The work of Kranton and Minehart (2001) [22] introduces matching into a broader buyer-seller setting. In their private values setup, the authors establish a basic efficiency result that extends classical assignment game results called on below (Section 11.2). The model of Kotowksi and Leister (2014) [21] is also closely related. The authors establish a generic under entry as a result of complementarities between vertically oriented intermediaries. Though the following analysis does not study entry, it does establish a parallel complementarity between dealers carrying assets and those providing liquidity. Thus, complementarity between intermediaries taking on different but supporting roles is seen to extend outside of a fixed-network setting.

Finally, this paper owes much credit to the work of Rubinstein and Wolinksy (1987) [26] who study steady state measures of buyers, sellers, and middlemen in a market with a continuum of agents. Particularly relevant, Rubinstein and Wolinksy include the option value to vacant intermediaries into continuation values of asset holders, and vice versa. The following model establishes this codependency as an essential ingredient to deriving the market's equilibrium provision of liquidity. With this step, interdealer activity comes as a market equilibrium response to dealer level risks.

setting, on the other hand, dealers' time preferences take on this role.

10 The Model

N risk neutral dealers physically connect an 'upstream' supply market to a 'downstream' demand market. Direct trades between the upstream and downstream markets are excluded. Upstream supply can be interpreted to include net-of direct trades between 'end-users'. Dealers' payoffs are derived solely from current and discounted future trading profits. The analysis will maintain that dealers be competitive in the upstream market. These assumptions abstract away from cash flows and dealer markups, and will reduce dealer payoffs in a tractable way.

Time is discrete, with t taking values in the set of integers: $t \in \mathbb{Z}$. $r \in (0, 1)$ gives a constant per-period cost of inventory (i.e. cost of capital). For notational simplicity, denote the constant discount factor $\delta := \frac{1}{1+r}$. In each period, some subset of the dealers may purchase assets from the upstream supply. Through subsequent periods, dealers can potentially trade assets amongst each other in an interdealer market as their idiosyncratic liquidity needs are realized. Such trades are made with the ultimate goal of selling the assets downstream to buyers willing to purchase them at price v. The arrival of this demand is independent across periods, occurring with some probability p for each asset holder. The value of p may depend on the total number of assets held by dealers at any given time (more on this below). Asset holders can always sell downstream at price $\underline{v} < \delta v$. Thus, asset holders face enough risk over future demand as to motivate the need to hold inventories in anticipation of the subsequent arrival of high valuing demand.

Each dealer may be hit with a one period capacity shock with probability q, in which case the dealer can neither bid for nor carry an asset as inventory and must sell off any asset held. Absent of a capacity shock, dealers remain constrained to hold at most one asset at a time, with $a = 1$ ($a \in \{0, 1\}$) denoting the state of holding an asset. $N_1^t \leq N$ denotes the current market inventory at time t, equal to the total number of dealers holding assets. This leaves $N_0^t := N - N_1^t$ dealers without assets. $V_{a,N_1^t}^t$ denotes the equilibrium value to being in state a at time t given market inventory N_1^t, prior to any resolution of uncertainty in that period (i.e. capacity shocks, the arrival of high demand downstream, or the resolution of uncertainty regarding the interdealer market; see below). Again, dealers holding an asset are termed *asset holders* while those not in possession of an asset are *vacants*.

[Figure]

The model time line is depicted in Figure 1. At time t dealers holding assets derive some value above remaining vacant, denoted $G^t(N_1^t) := V_{1,N_1^t}^t - V_{0,N_1^t-1}^t$. $V_{0,N_1^t-1}^t$ captures the option value to retaining liquidity for future use. Then, asset holders observe demand along with the arrival or non-arrival of a liquidity shock. In the event of a liquidity shock, asset holders unable to sell for price v are forced to move to the interdealer market to attempt to collect gains above liquidation value $\underline{v}$. That is, dealers potentially gain from entering the interdealer market if they either do not have an asset but harbor the capacity to absorb an asset (i.e. did not receive a liquidity shock themselves), or do have an asset but can neither sell it downstream for price v nor hold the asset due to a liquidity shock. If shocked asset holder are unable to sell in the interdealer market, they then liquidate and collect $\underline{v}$. Once all asset holders have either sold their assets downstream, retained them as inventory, or discarded them in the interdealer market, the market moves to period $t+1$. Intermediaries that acquired an asset in the interdealer market in period t or held on to a previously obtained asset remain asset holders in period $t+1$.

For completeness, we can summarize the dealers' contingent action sets with the following table:

	event	action set
asset holders	period t	$\emptyset$
vacants		$\{B, L\}$
asset holders	demand arrival	$\{S^v, S^{\underline{v}}, \mathcal{S}^D, H\}$
	demand shock only	$\{S^{\underline{v}}, \mathcal{S}^D, H\}$
	demand and liquidity shock	$\{S^{\underline{v}}, \mathcal{S}^D\}$
vacants	no liquidity shock	$\{\mathcal{B}^D, L\}$
	liquidity shock	$\emptyset$

B denotes the action of buying an asset from an upstream supplier, L of retaining liquidity, H of holding an asset between periods as carried inventory, and S^v [$S^{\underline{v}}$] of selling downstream for v [$\underline{\mathrm{v}}$]. $\mathcal{S}^D$ and $\mathcal{B}^D$ abbreviate the action sets to asset holders and vacants in the interdealer market, respectively.

Though these formal non-cooperative primitives formalize the set of feasible histories, the model will proceed to analyze the cooperative implications of the setup. As such, the particular actions of dealers will remain suppressed, as they come ancillary in interest to the values created through intermediation and transferred through interdealer trade.

10.1 Market equilibrium

In what follows, the condition that upstream price equals the value gap is derived in equilibrium. That is, with dealers demanding assets in the upstream market, the model endogenously determines the number of asset in circulation, requiring only that the upstream market clears in each period. Thus, N_1^t will capture the degree of intermediation supplied by the dealer market at any period t. As such, N_1^t constitutes the basic state variable of the model.

At each time t, N_1^t is determined leaving $N_0^t := N - N_1^t$ vacant dealers available to bid for these assets in the interdealer market. N_1^t will assumed to be common knowledge in period t. Each asset holder is able to sell in period t for price v with probability $p > 0$. I only assume that p is weakly decreasing as a function of market inventory N_1^t. That is, asset holders may crowd each other out as market inventory increases. Decay in p will factor into values and asset inventory in equilibrium, but will not drive the model's key results. With $p(N_1^t)$ taken exogenously, the analysis instead focuses on the role of the interdealer market. Market inventory will play a crucial role pinning the value that asset holders derive from the interdealer market as a hedge against liquidity risk. As more dealers carry assets and market inventory grows, p will drop along with this hedge value, together decreasing the value that dealers derive from intermediating assets. Accordingly, period t inventories N_1^t adjust. As opposed to requiring stationarity in market inventory, the model solves for stationary beliefs over the process determining market inventories. That is, N_1^t will adjust from period to period, while the process determining future market inventories will be a function of t only through current inventory N_1^t.

Though all results will be obtained under this basic setup, I offer two examples extending demand to allow for multiple assets. First, one could specify the following correlation over the asset holders' demand realizations by requiring the entire pool of asset holders to face the downstream demand together. An integer number $\widetilde{D}$ of high demanders arrive in period t, denoting $F(\cdot; N_1^t)$ and $f(\cdot; N_1^t)$ the cumulative and density functions of $\widetilde{D}$.[6] We parameterize these functions by N_1^t as the demand process may depend on the total number of assets held by investors, which will equal total issuance minus N_1^t.[7] Then, the prior probability of each asset holder selling her asset at price v will be given by

$$p\left(N_1^t\right) = 1 - F\left(N_1^t - 1; N_1^t\right) + \sum_{\widetilde{D}=0}^{N_1^t - 1} \frac{\widetilde{D}}{N_1^t} f\left(\widetilde{D}; N_1^t\right). \tag{1}$$

This gives the probability that all N_1^t assets are purchased $(1 - F(N_1^t - 1; N_1^t))$ plus the likelihood that her asset is purchased in cases where not all assets are purchased. Naturally, this probability is decreasing in N_1^t. I'll refer to this setup as the *co-demand* setup.

Second, when asset holders face private demand realizations, market inventory will no longer apply downward pressure on p: p remains constant as N_1^t varies. I establish this *segmented demand* setup as an alternative benchmark. Formally, in this case $\widetilde{D}$ follows a binomial distribution conditioning on total market inventory N_1^t:

$$F\left(\widetilde{D}; N_1^t\right) = \binom{\widetilde{D}}{N_1^t} p^{\widetilde{D}} \left(1 - p\right)^{N_1^t - \widetilde{D}}. \tag{2}$$

However, that each asset holder faces a constant (and equal) likelihood p of selling downstream at price v is all that will be needed for the application of this setup.

These two setups can be viewed to give two extremes. If downward demand uncertainty exhibits intermediate correlation with market inventory, one could expect that market behavior falls somewhere between these two cases. As I will show below, our basic results only requiring p be weakly decreasing with N_1^t. Hence, the particular setup of downstream demand plays an ancillary role in what follows. What remains crucial is that dealers face demand uncertainty (at times) motivating the need to carry inventories. A formal treatment of beliefs over future carried inventories– which will intricately depend on expectations over interdealer trade– is left to the following section. For now, let $\hat{N}_1^{t-1}$ denote the total number of assets that are carried as inventory into period t.

In period t, if an asset holder is unable to sell downstream then her net present value becomes the greater of $\underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$ and $\delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]$, from either liquidating the asset in the current period for value $\underline{v}$ or carrying the asset over to the subsequent period as

[6] I use tildes throughout to denote random variables that are realized between periods t and $t+1$.

[7] For N_1^t small relative to total issuance, this dependence should (approximately) drop out.

inventory. This assumes that any future upstream purchases are made at price $G^{t+1}(N_1^{t+1})$. Again, this step is derived below in upstream market equilibrium. The Bellman equations for this setup are:

$$V_{0,N_1^t}^t = (1-q)\,\mathcal{V}_{0,N_1^t}^{ns} + q\delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right], \tag{3a}$$

$$V_{1,N_1^t}^t = \left(\begin{array}{c} (1-p)\left(\begin{array}{c} (1-q)\max\left\{\underline{v}+\delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right], \delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]\right\} \\ +q\mathcal{V}_{1,N_1^t}^{s} \end{array}\right) \\ +p\left(v+\delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]\right) \end{array}\right). \tag{3b}$$

where $\mathcal{V}_{0,N_1^t}^{ns}$ denotes the expected value to a vacant intermediary entering the interdealer market (upon realizing no liquidity shock), and $\mathcal{V}_1^s$ the expected value to an asset holder entering the interdealer market (upon realizing both liquidity and demand shocks). $\mathcal{V}_{0,N_1^t}^{ns}$ and $\mathcal{V}_{1,N_1^t}^{s}$ depend on the respective likelihoods of buying and selling in the interdealer market, as well as each intermediary's reserve utility conditioning on not buying and selling. As such, $\mathcal{V}_{0,N_1^t}^{ns}$ and $\mathcal{V}_{1,N_1^t}^{s}$ are both functions of $E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$, and $E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]$ (more later).

When $\delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right] < \underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$, asset holders optimally liquidate over carrying inventories, in which case $\mathcal{V}_{0,N_1}^{ns} = \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$ and $\mathcal{V}_{1,N_1}^{s} = \underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$. Without gains from trade in the interdealer market this leaves the respective agents with values:

$$V_{0,N_1^t}^t = \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right], \tag{4a}$$

$$V_{1,N_1^t}^t = pv + (1-p)\,\underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]. \tag{4b}$$

Vacants in the market are left with their continuation payoffs. I will refer to this behavior as *asset funneling*. Most of our analysis will be focused on the case with $\delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right] \geq \underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$. Though, the complete equilibrium characterization will incorporate asset funneling (Proposition 4, below).

At the beginning of each period t, vacants purchase assets from the upstream supply. Upstream suppliers will be able to extract only the value gap $G^t(N_1^t)$ from any demanding intermediary. N_1^t will vary from period to period depending on the arrival of demand and realized liquidity needs. That is, in each period the number of assets sold downstream will likely not coincide with the ensuing volume demanded from the upstream supply. Variability over $\widetilde{D}$ as well as variability in the interdealer market will drive variability in N_1^t.

To pin the process for N_1^t and close the model, upstream asset supply is modeled with an arbitrary upward sloping inverse supply function $\Phi(N_1^t; \hat{N}_1^{t-1})$, which yields the price of assets $N_1^t - \hat{N}_1^{t-1}$ flowing into the dealer market. We condition on $\hat{N}_1^{t-1}$, as the effect of an increase in the current carried market inventory $\hat{N}_1^{t-1}$ will impose the following counteracting effects on prices. First, for each N_1^t, asset flow $N_1^t - \hat{N}_1^{t-1}$ drops with $\hat{N}_1^{t-1}$: the period t market price of assets purchased upstream decreases as fewer demanding vacants inhabit the market. On the other hand, If $\hat{N}_1^{t-1}$ and N_1^t are large relative to the total issuance of the assets then an increase in the level of carried assets may impose upward pressure on prices by depleting the stock of assets that investors hold.[8] In such cases, Φ can potentially decrease with $\hat{N}_1^{t-1}$ for each N_1^t. By limiting the size of the dealer market relative to total issuance, Φ will decrease with $\hat{N}_1^{t-1}$.

Assumption 1. $\Phi\left(N_1^t;\hat{N}_1^{t-1}\right)$ *is non-increasing in* $\hat{N}_1^{t-1}$ *for each* N_1^t.

Beyond an inverse dependence of Φ on $\hat{N}_1^{t-1}$, Φ will need to remain stationary for a steady state to obtain. Requiring that Φ be stationary stresses the role of downstream demand risk. With real broker-dealers facing both future supply and demand uncertainty, an extension of the model could allow for stochastic Φ^t with dealers strategically 'buying low'.[9] Here, however, downstream demand uncertainty drives interdealer transactions. THis comes naturally, as the stochastic arrival of demand necessitates inventories in order to anticipate the arrival of high demand. Coupled with liquidity risk, this introduces the potential for gains to interdealer trade.

To obtain the asset's period t market price, the ensuing market inventory N_1^t is pinned by the upstream market clearing condition:

$$\Phi\left(N_1^t;\hat{N}_1^{t-1}\right) \leq G^t(N_1^t), \text{ with} \tag{5a}$$

$$\Phi\left(N_1^t+1;\hat{N}_1^{t-1}\right) > G^t(N_1^t+1). \tag{5b}$$

That is, at the beginning of each period t vacant intermediaries from the previous period demand assets until the value gap $G^t(N_1^t)$ is driven down to marginal costs to suppliers. To obtain this formally, Section 12 establishes that $G^t(N_1^t)$ is decreasing in N_1^t, constituting a generic law of demand. Then, setting upstream inverse supply Φ to the value gap pins the inflow of assets $N_1^t-\hat{N}_1^{t-1}\geq 0$. Thus, the process generatng N_1^t is determined by the number of assets carried over from the previous period ($\hat{N}_1^{t-1}$) as well as the willingness of investors to supply additional assets. In what follows, I use $N_1^t(\hat{N}_1^{t-1})$ to denote the endogenous function yielding N_1^t from $\hat{N}_1^{t-1}$ as implied by (5a)-(5b).

Through the interim, after asset holders have realized the arrival (or non arrival) of high demand, entrants into the interdealer market evaluate the prospects of leaving with or without an asset. The expected values to holding and not holding an asset depend on the beliefs over period $t+1$ market inventory, N_1^{t+1}. That is, N_1^{t+1} influences both the likelihood of selling assets tomorrow as well as expectations over future inventories and values in subsequent periods. As seen below in Section 11, expected continuation values $E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$ and $E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]$ to entering period $t+1$ with and without an asset, respectively, will play crucial roles in pinning gains to trade in the interdealer market.

[8] To better see this, at the extreme of $\hat{N}_1^{t-1}$ equaling total issuance (ie. dealers hold all assets) supply must fall to zero for all prices. On the other hand, in the corporate bond market seasoned bonds will likely be primarily held by investors with a small portion of total issuance carried by dealers. In this case we should expect $\Phi(N_1^t;\hat{N}_1^{t-1})\approx\Phi(N_1^t-\hat{N}_1^{t-1})$ so that the effect of an increase in carried inventories reduces to a horizontal shift in supply.

[9] Allowing Φ^t to vary will not change the upstream market clearing condition (5a)-(5b), but will complicate expected continuation values $E_t\left[V_{0,N_1^{t+1}}^{t+1}\right]$ and $E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]$ (see below).

10.2 Interpretation

The precise structure and timing offered in Figure 1 should be taken only to formalize the discrete time setup of the model. We can better interpret the setting by placing it into a continuous time analogue, as follows. At any time, dealers can choose to hold assets (take on inventory), purchasing them from upstream supply at competitive market prices. Then, having optimally determined inventories, asset holders face demand and liquidity risk over the following horizon. The probability that high demand arrives before a liquidity shock is given by p. With probability $(1-p)$ this event does not occur in which case one of two sub-events are realized: (i) a liquidity shock does not arrive (with probability $(1-p)q$), or (ii) a shock arrives before the arrival (or non-arrival) of demand (with probability $(1-p)(1-q)$). Thus, the measures defined with p and q can be easily interpreted to capture cumulative likelihoods for this partition of outcomes.

The model caries an important interpretation by requiring that the upstream price be set to the value gap $G^t(N_1^t)$ in equilibrium. As will be seen below, endogenizing market inventories via the market clearing condition (5a)-(5b) offers an equilibrium definition to 'market liquidity'. Precisely, the value to serving as a vacant and supplying liquidity to asset holders derives from the prospect of purchasing an asset from another dealer at a bargain (relative to the upstream market price). This option value to vacants bares the incentives to providing future liquidity in the interdealer market. Thus, the model derives the extent of dealer provided liquidity as a market equilibrium property.

The endogenous variable that will determine equilibrium costs of trading will be the value gap $G^t(N_1^t)$. That is, any ensured round trip trade through the dealer market entails selling to a dealer at ask price $G^t(N_1^t)$ and buying at high price v (demanding price $\underline{v}$ does not guarantee an intermediary will sell). As such, $v - G^t(N_1^t)$ proxies for the equilibrium compensation that intermediaries demand for the risks that they face. As holding the asset becomes more risky to dealers, $G^t(N_1^t)$ will decrease and the necessary compensation required for intermediation will increase in equilibrium. Thus, G^t internalizes the degree of uncertainty that intermediaries face as they transfer assets and complete the market. With alternative sources of uncertainty such as informed trading or return and inventory risk abstracted away, the behavior of the value gap $v-G^t(N_1^t)$ captures the response of trading costs to the various dimensions of the model.

11 The Interdealer Market

The values $\mathcal{V}^{ns}_{0,N_1}$ and $\mathcal{V}^{s}_{1,N_1}$ will define the expected value captured through interdealer trade, and play key roles in equilibrium values and prices. Before formally characterizing $\mathcal{V}^{ns}_{0,N_1}$ and $\mathcal{V}^{s}_{1,N_1}$, we can first bound these values by removing the interdealer market in all periods, and prevent interdealer transactions as in the second benchmark above (allowing $q > 0$). Again, $\mathcal{V}^{ns}_{0,N_1}$ falls to $\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right] = 0$ as the use of vacants' liquidity becomes excluded. Correspondingly, $\mathcal{V}^{s}_{1,N_1}$ falls to $\underline{v}$ as shocked asset holders are left with no option but to liquidate at price $\underline{v}$. Thus, 0 and $\underline{v}$ constitute extreme lower bounds to $\mathcal{V}^{ns}_{0,N_1}$ and $\mathcal{V}^{s}_{1,N_1}$, respectively. The distance these values lie above their bounds will depend on how much gains from trade dealers in each respective state can expect to capture in the interdealer market. This will depend on both the likelihood of transacting with a counter-party in the interdealer market and the expected price and corresponding welfare split from any such transaction. Before revisiting this natural decomposition to interdealer gains, the following pursues a more formal treatment of the interdealer market.

A realization of $\widetilde{D}$ leaves $\widetilde{N}_1 := \max\{0, N_1^t - \widetilde{D}\}$ intermediaries holding assets that are susceptible to capacity shocks, in which case shocked asset holders will search for other dealers willing to buy their asset in the interdealer market. If none are found, the asset is sold off for value $\underline{v}$. Denote $\widetilde{N}_1^s \in \{1, \ldots, \widetilde{N}_1\}$ the number of such shocked asset holders. As many as N_0^t dealers will be potentially available– upon not realizing capacity shocks themselves– to pick up these assets in the interdealer market.[10] Correspondingly, denote $\widetilde{N}_0^{ns} \leq N_0^t$ the number of unshocked vacants. Together, the random variables $\widetilde{N}_1^s$ and $\widetilde{N}_0^{ns}$ determine the size of the interdealer market in each period. Under the above market structure, distributions over $\widetilde{N}_1^s$ and $\widetilde{N}_0^{ns}$ will live in a family of binomial distributions. The corresponding expressions for the co-demand and segmented demand setups are provided in the Appendix. In what follows, however, I use $g_1\left(\cdot\,|N_1^t\right)$ to denote the distribution of $\widetilde{N}_1^s$ and $h_0\left(\cdot\,|N_1^t\right)$ for the distribution of $\widetilde{N}_0^{ns}$. I will also use $g_1^S\left(\cdot\,|N_1^t\right)$ to denote beliefs over $\widetilde{N}_1^s$ for asset holders entering the interdealer market (S for seller) and $h_0^B\left(\cdot\,|N_1^t\right)$ for the beliefs over $\widetilde{N}_0^{ns}$ of entering vacants. These beliefs will rationally be formed taking $g_1\left(\cdot\,|N_1^t\right)$ and $h_0\left(\cdot\,|N_1^t\right)$, respectively, and condition on the given agent entering the interdealer market.[11]

All of these distributions will condition on market inventory N_1^t, which will play a key role throughout the model. The following conditions on $g_1\left(\cdot\,|N_1^t\right)$ and $h_0\left(\cdot\,|N_1^t\right)$ will naturally hold for the co-demand and segmented demand setups. I will assume them to generally hold, taking them as regularity conditions for dealers' expectations over the size of the interdealer market.

Assumption 2. *Upon an increase in either N_1^t or q :*

1. $g_1\left(\widetilde{N}_1^s\,|N_1^t\right)$ exhibits a first order stochastic shift up, and,

[10] Again, this assumes that selling intermediaries are unavailable to demand assets in the interdealer market until, at earliest, the following period; see footnote 4.

[11] Defining the updated beliefs is done for the sake of concreteness, though with large N we should expect conditional distributions to approximate the unconditional primitives. Note I do not update beliefs over the other side of the interdealer market, as independence between $\widetilde{N}_1^s$ and $\widetilde{N}_0^{ns}$ upon conditioning on N_1^t is presumed.

2. $h_0\left(\widetilde{N}_0^{ns}\,|N_1^t\right)$ *exhibits a first order stochastic shift down.*

3. For $q=0$*:*

$$g_1\left(0\,\middle|N_1^t\right)=g_1^S\left(1\,\middle|N_1^t\right)=h_0\left(N_0^t\,\middle|N_1^t\right)=h_0^B\left(N_0^t\,\middle|N_1^t\right)=1. \tag{6}$$

Assumption 2.1 implies that expectations over the number of dealers in need of liquidity increase with market inventory and with the likelihood of liquidity shocks (holding all else equal). 2.2, on the other hand, implies that expectations over the available liquidity offered from vacants decreases with N_1^t and q. 2.3 gives rational conditions on beliefs upon removing liquidity risk from the market.[12]

With beliefs over the size of the interdealer market defined, I can now describe the model's structural treatment of interdealer search frictions. Take the interdealer market with $\widetilde{N}_0^{ns}$ buyers and $\widetilde{N}_1^s$ sellers that can potentially meet to trade. With an abuse of notation, I denote the set of buyers and sellers $\widetilde{N}_0^{ns}$ and $\widetilde{N}_1^s$, respectively. An exogenous search technology connects sellers to potential buyers by collecting *buyer sets* $\widetilde{\mathcal{B}}_i\subseteq\widetilde{N}_0^{ns}$ for each seller $i\in\widetilde{N}_1^s$. The set of transactions between sellers and buyers will be limited to pairs (i,j) with $j\in\widetilde{\mathcal{B}}_i$. I also require that sellers and buyers only transact with at most one other agent. The set of realized buyer sets is denoted $\widetilde{\mathcal{B}}$, with $b(\cdot|\widetilde{N}_0^{ns},\widetilde{N}_1^s):2^{\widetilde{N}_0^{ns}\widetilde{N}_1^s}\rightarrow[0,1]$ denoting the stationary density over $\widetilde{\mathcal{B}}$ conditional on the interdealer market size $(\widetilde{N}_0^{ns},\widetilde{N}_1^s)$. For simplicity, I require that this density be anonymous over agents' identities.

Assumption 3. *(*anonymity*) The probability of any* $\widetilde{\mathcal{B}}$ *forming equals the probability of* $\widetilde{\mathcal{B}}'$ *obtained from* $\widetilde{\mathcal{B}}$ *by permuting the sellers' and/or the buyers' indices.*

To interpret this search technology, one can imagine sellers in $\widetilde{N}_1^s$ corralling buyers into the sets $\widetilde{\mathcal{B}}$. Which buyers, as well as the total number of buyers that each seller can gather, is determined as an outcome of the search process, $b(\cdot|\widetilde{N}_0^{ns},\widetilde{N}_1^s)$. Assumption 3 requires that each seller face the same abilities and limitations to finding sellers in the market.

An equivalent interpretation of the realized sets $\widetilde{\mathcal{B}}$ is to take them as a bipartite network connecting buyers and sellers. Each link of the realized network represents a potential trade. The set of neighbors of each seller i is given with $\widetilde{\mathcal{B}}_i$. Importantly, buyers may reside in more than one buyer set. Further, the given network may or may not be 'connected', with separate components describing distinct interdealer markets separated by (short) time intervals or other contemporaneous constraints. With these interpretations in hand, I will refer to realized buyer sets $\widetilde{\mathcal{B}}$ as a "network".

Upon realization of $\mathcal{B}$, the ensuing interdealer trade can be described as follows. If a buyer and seller meet alone the agents may bargain via some well-defined bargaining protocol. If instead a seller is given a group of buyers exclusively, the seller can auction off the asset under some auction protocol. Or, if some buyers are shared by two or more sellers, some well defined trading protocol will determine which (if any) buyers purchase assets and from whom. As further explained in Section 11.2, I will only require that the resulting set of transactions and prices be stable, in the sense that no two connected dealers can deviate in

[12] The likelihood $g_1^S\left(1\,|N_1^t\right)=1$ conditions on the asset holder entering the secondary market, a probability zero event when $q=0$.

a mutually beneficial way. Leaving further discussion for the sequel, I will apply pairwise stability as the most liberal requirement that can be placed on the resulting allocation of assets.

Prior to each transaction, buyers and sellers are assigned reserve utilities $\delta E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]$ and $\delta E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]+\underline{v}$, respectively. In the event of a trade, these agents will collectively realize value:

$$\delta E_t\left[V^{t+1}_{1,N^{t+1}_1}\right]-\delta E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]-\underline{v}\geq 0. \tag{7}$$

As such, each interdealer transaction in the interdealer market following period t is taken as some split of the *value set*[13] given with Figure 2. I will denote this subset of $[0,v]^2$ by Δ.

[Figure 2]

From a technical perspective, gains to interdealer trade are scaled in this way for simplicity and notational ease. Still, extensions of the model could require interim beliefs over market inventories to update on information gathered prior to each trade. For example, traders could condition on some coarsening of the information contained with $\widetilde{\mathcal{B}}$, or on the set of transactions since period t. However, as discussed in Section 10.2, the process $b(\cdot|\widetilde{N}^{ns}_0,\widetilde{N}^s_1)$ and our discrete-time setup is meant as an approximation to a more detailed continuous-time interdealer market. Thus, scaling expectations over interdealer gains by period t expected-discounted continuation values comes as a natural simplification.

With gains to interdealer trade equally scaled, this setting yields a well-defined assignment game in the classic matching literature (see Shapley and Scarf 1973, or Roth and Sotomayor 1990), but where the given structure to the interdealer market constrains the game. Formally, upon normalizing disagreement values to zero, asset holders and vacants in the interdealer market that are not matched under $\widetilde{\mathcal{B}}$ create zero worth when assigned to each other. All other asset holder-vacant pairs create value given by (7). When referring to the assignment game below, I will assume this value to be positive and normalize these gains to 1, which is without loss of generality. This step also lends to the interpretation of transfers as welfare-shares. I will also refer to asset holders as *sellers* and vacants as *buyers*, in the spirit of the matching literature. I return to the assignment game below, but first proceed to establish a unique steady state of the model.

[13]Also referred to as a "bargaining set", though I refrain from this terminology as transactions may be the result of multilateral bargaining or auctions.

11.1 Steady State Existence and Uniqueness

In this subsection I formally restrict exchange in the following crucial ways to yield a unique steady state set of values. However, these assumptions will later be motivated in Section 11.2. The first assumption restricts the set of transactions to be determined by the realized sets $\widetilde{\mathcal{B}}$:

Assumption 4. *For any arrangement of buyer sets resulting from the search process, the set of transactions (and set of disagreements) is independent of* $E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]$ *and* $E_t\left[V^{t+1}_{1,N^{t+1}_1}\right]$.

At first glance, this assumption may seem strong. For all transactions that do not occur due to the relevant buyer and seller link being excluded in $\widetilde{\mathcal{B}}$, this independence comes trivially. For unexecuted transactions, the assumption gives more bite. Importantly, intermediaries of each type (buyers/sellers) are symmetric in their reservation values and potential gains to each transaction. Thus, each value set for all transactions move "in-step" with $E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]$ and $E_t\left[V^{t+1}_{1,N^{t+1}_1}\right]$. One should then find it natural to expect any bargaining advantages, or more generally, trading advantages from one potential transaction over another to be captured solely through the trading network $\widetilde{\mathcal{B}}$.

In this light, Assumption 4 comes a practical assumption on the realized set of transactions. It implies that as the conditional values $E_t\left[V^{t+1}_{0,N^{t+1}_1}\right]$ and $E_t\left[V^{t+1}_{1,N^{t+1}_1}\right]$ change, the period t allocation of gains resulting from interdealer trade do not. Precisely, alternative transactions with other potential buyers should not be prefered to any seller as the agreement set shifts. Similarly, for buyers in multiple buyer sets, alternative transactions with alternative potential sellers should remain weakly not preferred. As the values at stake change continuously for all agents on each side of the market, the set of mutually beneficial transactions remains fixed.

I next impose a technical restriction on the set of transfers resulting from interdealer transactions. Denoting the set of interdealer transactions $\widetilde{X}$ resulting from buyer sets $\widetilde{\mathcal{B}}$, let $\Lambda_{x,\widetilde{\mathcal{B}}} : \mathbb{R}^2 \to \Delta$ map continuation values to the agreement point in Δ, for each $x \in \widetilde{X}$.

Assumption 5. *For any vectors* (y,z) *and* (y',z') *in* $\mathbb{R}^2$*, the function* $\Lambda_{x,\widetilde{\mathcal{B}}}$ *satisfies:*

$$\left\|\Lambda_{x,\widetilde{\mathcal{B}}}(y,z) - \Lambda_{x,\widetilde{\mathcal{B}}}(y',z')\right\|_\infty \leq \|(y,z)-(y',z')\|_\infty . \tag{8}$$ [14]

Assumption 5 maintains that $\Lambda_{x,\widetilde{\mathcal{B}}}$ is non-expansionary. The restrictiveness of this assumption may at first be unclear. However, given the model's natural restrictions to linear value sets, this property is easily obtained under solutions to common social welfare problems. In particular, when set share s is given to the seller and share $1-s$ to the buyer, $\Lambda_{x,\widetilde{\mathcal{B}}}$ will be non-expansionary under any norm, as shown in the Appendix C. Then, it is easy to show that the Nash and Rawlsian welfare problems both yield a constant share split for linear value sets, as we have here. Thus in this light, Assumption 5 can be viewed as a weakest restriction on interdealer transfers needed to establish a unique stationary solution.

[14] $\|\cdot\|_\infty$ gives the sup-norm: $\|f\|_\infty = \sup_X \{|f| : x \in X\}$, for any real valued function f defined on $X \subseteq \mathbb{R}$.

At the beginning of each period, expected values for both states $a \in \{0,1\}$ are derived taking expectations over competition and future demand in upstream, downstream, and interdealer markets. Values out of the interdealer market will also depend on expectations over future networks $\widetilde{\mathcal{B}}$. Denote $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}(\cdot)$ for $\mathcal{V}^{ns}_{0,N_1^t}(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s, \widetilde{\mathcal{B}})$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}(\cdot)$ for $\mathcal{V}^{s}_{1,N_1^t}(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s, \widetilde{\mathcal{B}})$, the corresponding expected values conditional on the size of the interdealer market and on the realized network $\widetilde{\mathcal{B}}$.[15] As described above, stationarity and anonymity in the distribution $b(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$ is maintained throughout. Under only Assumptions 4 and 5 and without explicitly solving for the continuation values $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}$ we can establish existence of a unique steady state solution to (3).

Proposition 1. *There exist unique steady state values V_{0,N_1} and V_{1,N_1} in $[0, \delta v]$ and $[\underline{v}, v]$ (respectively) for each $N_1 \in \{0, \ldots, N\}$.*

PROOF: By contraction mapping; see Appendix C.

Proposition 1 shows that under the above general setting with moderate structure imposed on the interdealer market, we can expect unique steady state values for V_{0,N_1^t} and V_{1,N_1^t}, for each market inventory $N_1^t \in \{0, \ldots, N\}$. Thus, given market inventory N_1^t, Proposition 1 ensures unique (jointly determined) values to both holding an asset in the market and to providing liquidity as a vacant. These values will be time dependent solely through the inventory state N_1^t. Correspondingly, I will hence forth drop the superscripts on values, writing V_{0,N_1^t} and V_{1,N_1^t}.

11.2 Interdealer Trade: stability and monotonicity

This section applies classic results in the assignment game literature to gain the traction needed for the paper's main descriptive results; see Roth and Sotomayor (1992) [25]. Take the assignment game $(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij})$ under the following restrictions from $\widetilde{\mathcal{B}}$: if $j \in \widetilde{\mathcal{B}}_i$ for $i \in \widetilde{N}_1^s$, then $\alpha_{ij} = 1$, with $\alpha_{ij} = 0$ otherwise. Define a stable assignment as (i) a set of transactions (pairings) $\widetilde{X} \subseteq \widetilde{N}_1^s \times \widetilde{N}_0^{ns}$ between buyers and sellers with each agent in at most one pair and with $j \in \widetilde{\mathcal{B}}_i$ for each $(i,j) \in \widetilde{X}$, and (ii) a set of values (x_i, y_j) such that the following holds:

$$1. \qquad x_i, y_j \geq 0 \text{ for each } i \in \widetilde{N}_1^s,\ j \in \widetilde{N}_0^{ns}, \tag{9a}$$

$$2. \qquad x_i + y_j \geq 1 \text{ for each } \ (i,j) \ \text{ with } j \in \widetilde{\mathcal{B}}_i. \tag{9b}$$

1. gives agent rationality, as each agent can demand her disagreement value 0. 2. gives pairwise stability, insuring that no buyer-seller pair unmatched in the assignment can form to create value for themselves. Existence of a stable assignment are easily established using standard linear algebra results.[16].

[15] Formally, the only random variable that these values will be left with to take expectations over will be the particular stable assignment of the ensuing assignment game, when more than one stable assignment exists.

[16] See Theorem 8.6 in Roth and Sotomayor (1992) [25]

[Figure 3]

An example of an assignment problem restricted in the above manner is depicted in Figure 3, with dashed lines indicating that the buyer is in the connected seller's buyer set, and solid lines indicating a transaction. In this realization of the interdealer market sellers constitute the long side of the market with $\widetilde{N}_1^s = 6$ and $\widetilde{N}_0^{ns} = 4$. One example of a stable assignment is also given with the solid lines. A necessary condition for the given shares to be stable is $s \geq s'$, so that i and j' do not form a blocking pair.

Had the network been complete (all buyers connected to all sellers), each buyer would be assigned to a seller and would capture all gains in any stable assignment. However, due to the restrictions that the network $\widetilde{\mathcal{B}}$ imposes, the set of stable assignments can give a seller some non-zero gain (seller i, for example). This is an example of a seller "lucking out" under the realized buyer network: she is connected to a buyer not connected to any alternative seller. However, we can expect that such realizations become fleetingly rare as the size of the pool of buyers becomes small relative to the number of sellers.

Lemma 1. *Under the assignment problem* $\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij}\right)$*, if* $k \in \widetilde{N}_1^s$ *[*$k \in \widetilde{N}_0^{ns}$*] with distinct* $k', k'' \in \widetilde{\mathcal{B}}_k$ *[*$k', k'' \in \left\{i : k \in \widetilde{\mathcal{B}}_i\right\}$*] such that* k' *and* k'' *are not in* $\widetilde{\mathcal{B}}_i$ *for all other* $i \in \widetilde{N}_1^s$ *[*$\widetilde{\mathcal{B}}_{k'}$ *and* $\widetilde{\mathcal{B}}_{k''}$ *are empty], then all stable assignments give* k *value* 1.

PROOF: Assume k realizes value φ less than 1 in some stable assignment. k' and k'' at best obtain value $1-\varphi$. Without loss of generality assume this to be k'. Then, $\{k, k''\}$ form blocking pair in which k obtains value $\varphi + \frac{1-\varphi}{2} > \varphi$ and k'' obtains value $\frac{1-\varphi}{2}$, contradicting stability. □

Lemma 1 gives an intuitive condition on $\widetilde{\mathcal{B}}$ that suffices for all stable assignments to allocate all gains to some trader k. If an agent k on either side of the network is connected to two individuals– each not connected to any other agent– then k captures all welfare in all stable assignments. Accordingly, one should expect that the likelihood of realizing a network in which the conditions of Lemma 1 hold increases as the interdealer market becomes more lopsided. Thus, we take Lemma 1 along with this intuition to motivate the following assumption[17]:

Assumption 6. *For each* N_1^t, $\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}_{0,N_1^t}^{ns}$ *takes the following limit conditions:*

1. *approaches* $\delta E_t\left[V_{0,N_1^{t+1}}\right]$ *(the buyer reservation value) as* $\widetilde{N}_1^s \to N_1^t$ *and* $\widetilde{N}_0^{ns} \to 0$ *(jointly),*
2. *approaches* $\delta E_t\left[V_{1,N_1^{t+1}}\right] - \underline{v}$ *as* $\widetilde{N}_1^s \to 0$ *and* $\widetilde{N}_0^{ns} \to N_0^t$ *(jointly).*

Accordingly, for each N_1^t, $\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}_{1,N_1^t}^{s}$ *takes the following limit conditions:*

[17] This assumption subsumes the assumption on the process $b(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$ that the conditions of Lemma 1 are satisfied with probability close to one as the relative sizes of the interdealer market depart from each other.

1. *approaches* $\delta E_t\left[V_{1,N_1^{t+1}}\right]$ *as* $\widetilde{N}_1^s \to N_1^t$ *and* $\widetilde{N}_0^{ns} \to 0$ *(jointly),*
2. *approaches* $\delta E_t\left[V_{0,N_1^{t+1}}\right] + \underline{v}$ *(the seller reservation value) as* $\widetilde{N}_1^s \to 0$ *and* $\widetilde{N}_0^{ns} \to N_0^t$ *(jointly).*

Formally, these properties will not hold for arbitrary $b(\widetilde{\mathcal{B}}|\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$.[18] None the less, I maintain the above conditions viewing them as practical assumptions on the extremes of the interdealer market.

Towards characterizing the general properties of the market's steady state values and prices, establishing monotonic gains to interdealer trade over the space of $(\widetilde{N}_1^s, \widetilde{N}_0^{ns})$ remains a crucial step. Towards this goal, I call on the following entry results from the assigment game literature. To first prime the reader, classic results establish the set of stable assignments as a complete lattice, including *buyer optimal* and *seller optimal* solutions. Precisely, given assignment problem $(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij})$ we can establish value vectors $(\overline{x}, \underline{y})$ and $(\underline{x}, \overline{y})$ in $[0,1]^{|\widetilde{N}_1^s|} \times [0,1]^{|\widetilde{N}_0^{ns}|}$ that define the seller and buyer optimal payoff profiles, respectively, from the set of all stable assignments. For any other payoff profile (x, y) resulting from some stable assignment of $(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij})$, it must then be that $x_i \geq \overline{x}_i$ $[x_i \leq \overline{x}_i]$ for each seller $i \in \widetilde{N}_1^s$, and $y_j \leq \overline{y}_j$ $[y_j \geq \underline{y}_j]$ for each buyer $j \in \widetilde{N}_0^{ns}$.

With these definitions in hand, take the following *entrant games*:

Definition 1. *Given assignment problem* $\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij}\right)$*, then:*

1. *A* seller entrant game *is given by* $\left(\widetilde{N}_1^s \cup \{i'\}, \widetilde{N}_0^{ns}, \alpha'_{ij}\right)$ *if* $\alpha_{ij} = \alpha'_{ij}$ *when* $i \neq i'$.
2. *A* buyer entrant game *is given by* $\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns} \cup \{j'\}, \alpha'_{ij}\right)$ *if* $\alpha_{ij} = \alpha'_{ij}$ *when* $j \neq j'$.

Simply put, the values created between non-entrants are independent of the values created with the entrant. In the context of network $\widetilde{\mathcal{B}}$, the inclusion of the entrant does not add or remove non-entrant buyers from any non-entrant seller's buyer set. Further, the entrant does not affect continuation values holding N_1^t fixed. With these tools in hand, I provide the following lemma:

Lemma 2. *Given any bipartite network* $\widetilde{\mathcal{B}}$ *connecting each seller* $i \in \widetilde{N}_1^s$ *to potential buyers* $\widetilde{\mathcal{B}}_i \subseteq \widetilde{N}_0^{ns}$*, let* $(\overline{x}, \underline{y})$ *[*$(\underline{x}, \overline{y})$*] be the seller [buyer] optimal stable assignment of the problem* $\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns}, \alpha_{ij}\right)$*. We have:*

1. *for* $(\overline{x}', \underline{y}')$ *[*$(\underline{x}', \overline{y}')$*], the seller [buyer] optimal stable assignment of any seller entrant game* $\left(\widetilde{N}_1^s \cup \{i'\}, \widetilde{N}_0^{ns}, \alpha'_{ij}\right)$*:*
   (a) $\overline{x}'_i \leq \overline{x}_i$ *[*$\underline{x}'_i \leq \underline{x}_i$*] for each* $i \in \widetilde{N}_1^s$*, and*
   (b) $\overline{y}'_j \geq \overline{y}_j$ *[*$\underline{y}'_j \geq \underline{y}_j$*] for each* $j \in \widetilde{N}_0^{ns}$*;*

[18] One can always construct a process $b(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$ that leaves agents on the long side of the market capturing gains with positive probability, for any $\widetilde{N}_0^{ns}, \widetilde{N}_1^s > 0$.

2. *for* $\left(\overline{x}', \underline{y}'\right)$ *[*$(\underline{x}', \overline{y}')$*], the seller [buyer] optimal stable assignment of any buyer entrant game*
$\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns} \cup \{j'\}, \alpha'_{ij}\right)$,

   (a) $\overline{x}'_i \geq \overline{x}_i$ *[*$\underline{x}'_i \geq \underline{x}_i$*] for each* $i \in \widetilde{N}_1^s$*, and*

   (b) $\overline{y}'_j \leq \overline{y}_j$ *[*$\underline{y}'_j \leq \underline{y}_j$*] for each* $j \in \widetilde{N}_0^{ns}$.

PROOF: (Proposition 8.17 in Roth and Sotomayor (1992) [25])

In words, Lemma 2 shows that adding an agent on one side of the market– holding the network over all other agents fixed– will at most benefit agents on the adjacent side of the market and hurt agents on the entrant side.

Finally, before leveraging on Lemma 2 and establishing the desired monotonicity, I define the following notion of independent entry. Taking any $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$, any $\widetilde{\mathcal{B}}$ over $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$, and some seller entrant i', let $\alpha(\widetilde{\mathcal{B}}, i')$ denote the set of networks $\widetilde{\mathcal{B}}'$ on $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s \cup \{i'\})$ that take $\widetilde{\mathcal{B}}$ and add only links from buyers to i':

$$\alpha\left(\widetilde{\mathcal{B}}, i'\right) := \left\{\widetilde{\mathcal{B}}' |\, (i,j) \in \widetilde{\mathcal{B}} \;\Leftrightarrow (i,j) \in \widetilde{\mathcal{B}}', \forall i \in \widetilde{N}_1^s\right\}. \tag{10}$$

Similarly, denote the set of networks $\widetilde{\mathcal{B}}'$ on $(\widetilde{N}_0^{ns} \cup \{j'\}, \widetilde{N}_1^s)$ that take $\widetilde{\mathcal{B}}$ and add only links from sellers to j':

$$\beta\left(\widetilde{\mathcal{B}}, j'\right) := \left\{\widetilde{\mathcal{B}}' |\, (i,j) \in \widetilde{\mathcal{B}} \;\Leftrightarrow (i,j) \in \widetilde{\mathcal{B}}', \forall j \in \widetilde{N}_0^{ns}\right\}. \tag{11}$$

Assumption 7. *(*independent entry*) For any* $\left(\widetilde{N}_0^{ns}, \widetilde{N}_1^s\right)$ *and* $\widetilde{\mathcal{B}}$ *over* $\left(\widetilde{N}_0^{ns}, \widetilde{N}_1^s\right)$*:*

$$\begin{aligned} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) &= \sum_{\widetilde{\mathcal{B}}' \in \alpha\left(\widetilde{\mathcal{B}}, i'\right)} b\left(\widetilde{\mathcal{B}}' \middle| \widetilde{N}_0^{ns} \cup \{i'\}, \widetilde{N}_1^s\right) \\ &= \sum_{\widetilde{\mathcal{B}}' \in \beta\left(\widetilde{\mathcal{B}}, j'\right)} b\left(\widetilde{\mathcal{B}}' \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s \cup \{j'\}\right), \end{aligned} \tag{12}$$

for any seller entrant i' *or buyer entrant* j'.

That is, the marginal distribution of networks $\widetilde{\mathcal{B}}$ over nodes $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$ is unchanged upon entry. Assumption 7 implies that entry does not affect the distribution of links over submarket $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$. The property will maintain the possibility of link correlation, and allow for particular structures to form more frequently than others.

To better understand the flexibility of Assumption 7, one could consider the following stronger notion:

Definition 2. $b\left(\cdot \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right)$ *satisfies* independent link formation *if for any distinct*

$\left(\widetilde{N}_0^{ns},\widetilde{N}_1^{s}\right)$ *and* $\left(\widetilde{N}_0^{ns\prime},\widetilde{N}_1^{s\prime}\right)$ *and any* (i,j) *with* $i\in\widetilde{N}_1^{s}\cap\widetilde{N}_1^{s\prime}$ *and* $j\in\widetilde{N}_0^{ns}\cap\widetilde{N}_0^{ns\prime}$*:*

$$\frac{b\left(\widetilde{\mathcal{B}}\cup(i,j)\left|\widetilde{N}_0^{ns},\widetilde{N}_1^{s}\right.\right)}{b\left(\widetilde{\mathcal{B}}\left|\widetilde{N}_0^{ns},\widetilde{N}_1^{s}\right.\right)}=\frac{b\left(\widetilde{\mathcal{B}}'\cup(i,j)\left|\widetilde{N}_0^{ns\prime},\widetilde{N}_1^{s\prime}\right.\right)}{b\left(\widetilde{\mathcal{B}}'\left|\widetilde{N}_0^{ns\prime},\widetilde{N}_1^{s\prime}\right.\right)},\tag{13}$$

where $j\notin\widetilde{\mathcal{B}}_i\cup\widetilde{\mathcal{B}}_i'$*.*

Independent link formation requires that information about the presence of other buyers or sellers in the interdealer market, or information regarding other links, does not affect the likelihood that a particular link between a seller i and buyer j forms. Now, consider an example in which each seller i links to at most $\lfloor\frac{1}{3}\widetilde{N}_0^{ns}\rfloor$ buyers. Here, independent entry may be satisfied if entrants do not affect the distribution of preexisting links. However, link independence is violated (with or without the entrant), as link inclusions are negatively correlated. Even more, Assumption 2 coupled with Assumption 3 (anonymity) implies the special case of a Poisson random network (excluding links between asset holders and between vacants). Thus, I employ Assumption 7, taking it as a desirably weak condition required to establish the following monotonicity:

Proposition 2. *For each* N_1^t*, under independent entry and assuming either only buyer-optimal or seller-optimal stable assignments from the interdealer market, we have that:*

1. $\sum_{\widetilde{\mathcal{B}}}b\left(\widetilde{\mathcal{B}}\left|\widetilde{N}_0^{ns},\widetilde{N}_1^{s}\right.\right)\widetilde{\mathcal{V}}_{0,N_1^t}^{ns}$ *is increasing in* $\widetilde{N}_1^{s}$ *and decreasing in* $\widetilde{N}_0^{ns}$*, and*
2. $\sum_{\widetilde{\mathcal{B}}}b\left(\widetilde{\mathcal{B}}\left|\widetilde{N}_0^{ns},\widetilde{N}_1^{s}\right.\right)\widetilde{\mathcal{V}}_{1,N_1^t}^{s}$ *is increasing in* $\widetilde{N}_0^{ns}$ *and decreasing in* $\widetilde{N}_1^{s}$*.*

PROOF: First, fixing N_1^t fixes $E_t\left[V_{0,N_1^{t+1}}\right]$ and $E_t\left[V_{1,N_1^{t+1}}\right]$ by Proposition 1, and thus pins the value set Δ. Upon including an additional seller i' to $\widetilde{N}_1^{s}$ (obtaining seller set $i'\cup\widetilde{N}_1^{s}$), the likelihood of subgraph $\mathcal{B}$ (i.e. by excluding all links to i') occurring must equal $b(\mathcal{B}|\widetilde{N}_0^{ns},\widetilde{N}_1^{s})$ by independent entry. Further, given any realized set $\mathcal{B}_{i'}$ and conditioning on supgraph $\mathcal{B}$, the value to each $i\in\widetilde{N}_1^{s}$ can only go down and the value to each $j\in\widetilde{N}_0^{ns}$ can only go up by Lemma 2. Thus, taking the average value over all realizations $\mathcal{B}$ must then give a lower conditional average value (conditioning on subgraph $\mathcal{B}$) to each $i\in\widetilde{N}_1^{s}$ and a higher conditional average value to each $j\in\widetilde{N}_0^{ns}$. Then, averaging over all $\mathcal{B}_{i'}$ and subgraphs $\mathcal{B}$ gives the result: expectations $\sum_{\widetilde{\mathcal{B}}}b(\mathcal{B}|\widetilde{N}_0^{ns},\widetilde{N}_1^{s})\widetilde{\mathcal{V}}_{0,N_1^t}^{ns}$ and $\sum_{\widetilde{\mathcal{B}}}b(\mathcal{B}|\widetilde{N}_0^{ns},\widetilde{N}_1^{s})\widetilde{\mathcal{V}}_{1,N_1^t}^{s}$ can only increase and decrease, respectively, as we enlarge $\widetilde{N}_1^{s}$. To show that these expectations are decreasing and increasing in $\widetilde{N}_0^{ns}$, respectively, use a similar argument adding some j' to $\widetilde{N}_0^{ns}$. $\square$

Establishing that monotonicity holds given the buyer-optimal or seller-optimal assignments from the interdealer market means that any violation of monotonicity must be the result of some arbitrary change in some degree of freedom in the set of stable assignments, as a result enlarging or contracting $\widetilde{N}_1^{s}$ or $\widetilde{N}_0^{ns}$. Taking Figure 3 for example, this would entail a change in s' after the inclusion of an entrant that is neither connected to i' nor to

j'. Such arbitrary changes could only be consequences of forces outside of the model. For simplicity, the following assumes such forces away.

The only random variable that $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}$ are now left to take expectations over is the particular stable assignment, if the set of stable assignments in not a singleton. With Proposition 2 in hand, we should feel comfortable assuming that any such selection process preserves monotonicity in expectation, barring the types of above mentioned forces. Regardless, in proceeding to explore the implication of this monotonicity, Proposition 2 at least assures that all results hold assuming either only buyer-optimal or seller-optimal stable assignments in the interdealer market.

Next I show how this monotonicity in $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}$ carries through to $\mathcal{V}^{ns}_{0,N_1^t}$ and $\mathcal{V}^{s}_{1,N_1^t}$. With Assumption 2, monotonicity in $\mathcal{V}^{ns}_{0,N_1^t}$ and $\mathcal{V}^{s}_{1,N_1^t}$ is nearly obtained except for the fact that the value set Δ depends on the continuation values $E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]$ and $E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]$. We can transform these values appropriately so that the joint values in Δ are normalized to one:

Definition 3.

$$\lambda \ : \ = \frac{\mathcal{V}^{s}_{0,N_1^t} - \delta E_t\left[V_{0,N_1^{t+1}}\right]}{\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}} \tag{14a}$$

$$\mu \ : \ = \frac{\mathcal{V}^{ns}_{1,N_1^t} - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}{\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}}. \tag{14b}$$

Proposition 3. *λ is increasing in N_1^t and q , and μ is decreasing in N_1^t and q.*

PROOF: With the values in $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ normalized, this leaves $\sum_{\widetilde{\mathcal{B}}} b(\widetilde{\mathcal{B}}|\widetilde{N}_0^{ns}, \widetilde{N}_1^{s})\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ independent of $E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]$ and $E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]$. Thus, writing:

$$\mathcal{V}^{ns}_{0,N_1^t} = \sum_{\widetilde{N}_1^s \geq 0}^{N_1^t} \sum_{\widetilde{N}_0^{ns} \geq 1}^{N-N_1^t} g_1\left(\widetilde{N}_1^s \,\middle|\, N_1^t\right) h_0^B\left(\widetilde{N}_0^{ns} \,\middle|\, N_1^t\right) \sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \,\middle|\, \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}^{ns}_{0,N_1^t} \tag{15a}$$

$$\mathcal{V}^{s}_{1,N_1^t} = \sum_{\widetilde{N}_1^s \geq 1}^{N_1^t} \sum_{\widetilde{N}_0^{ns} \geq 0}^{N-N_1^t} g_1^S\left(\widetilde{N}_1^s \,\middle|\, N_1^t\right) h_0\left(\widetilde{N}_0^{ns} \,\middle|\, N_1^t\right) \sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \,\middle|\, \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}^{s}_{1,N_1^t}, \tag{15b}$$

$\mathcal{V}^{ns}_{0,N_1^t}$ and $\mathcal{V}^{s}_{1,N_1^t}$ are left dependent on N_1^t only through the distributions of $\widetilde{N}_0^{ns}$ and $\widetilde{N}_1^s$. The result then follows immediately using Assumption 2 and Proposition 2, as expectations over increasing [decreasing] functions weakly increase with first order stochastic shifts up [down]. □

At this stage, we can rewrite the definitions of λ and μ as follows:

$$\mathcal{V}^{ns}_{0,N_1^t} = \lambda\left(\delta E_t\left[V_{1,N_1^{t+1}}\right] - \underline{v}\right) + (1-\lambda)\,\delta E_t\left[V_{0,N_1^{t+1}}\right] \tag{16a}$$

$$\mathcal{V}^{s}_{1,N_1^t} = \mu\delta E_t\left[V_{1,N_1^{t+1}}\right] + (1-\mu)\left(\delta E_t\left[V_{0,N_1^{t+1}}\right] + \underline{v}\right). \tag{16b}$$

λ and μ will lie in $[0,1]$ as weakly monotonic increasing and decreasing functions of q and N_1^t, respectively. These functions will also depend on the parameters of the model p, δ, N, as well as $b(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$. And with each period t value functions $(\mathcal{V}^{ns}_{0,N_1^t}, \mathcal{V}^{s}_{1,N_1^t})$ rationally living in Δ, we can reinterpret λ and μ as the weighting that buyers and sellers place on capturing gains in the interdealer market.

The essential properties of these functions are summarized as follows:

$$1.\ \lambda = 0 \text{ for } N_1^t = 0 \text{ and } \mu \simeq 1 \text{ for } N_1^t = 1 \tag{17a}$$

$$2.\ \mu = 0 \text{ for } N_1^t = N, \tag{17b}$$

$$3. \begin{cases} \lambda \text{ is weakly increasing, and} \\ \quad \mu \text{ is weakly decreasing} \end{cases} \text{in } N_1^t \ \big(\text{for each } N_1^t\big) \tag{17c}$$

$$4.\ \lambda = 0 \text{ at } q = 0 \ \big(\text{for each } N_1^t\big) \tag{17d}$$

$$5.\ \frac{\partial}{\partial q}\lambda \geq 0, \text{ and } \frac{\partial}{\partial q}\mu \leq 0 \ \big(\text{for each } N_1^t\big) \tag{17e}$$

$$6.\ \frac{\partial}{\partial q}\big(\lambda(N_1^t+1) - \lambda(N_1^t)\big) \geq 0 \text{ for } q \simeq 0 \text{ (for each } N_1^t < N). \tag{17f}$$

Properties 3 and 5 simply restates Proposition 3, with properties 1 and 2 following from Assumption 6. Properties 4 and 6 are more formally shown for our two demand setups in Appendix C. Property 6 will be of technical use in Proposition 5 of the following section.

[Figure 4]

Figure 4 provides estimates of the functions λ and μ, assuming each link connecting any asset holders/vacants pair forms with probability $m > 0$.[19] The arrows show the directions that these functions move as q is increased. Here, we see asset holders losing expectation over captured gains as N_1^t increases, while vacants enjoy increasingly more expected gains. Similarly, asset holders lose expected gains as q rises (for each N_1^t), while vacants collect additional expected rents from supplying liquidity. We should expect these functions to take on similar form for a general set of demand and buyer network processes, always preserving their monotonic dependence on q and N_1^t. For any shocked asset holder entering the interdealer market, beliefs over the number of available unshocked vacants $\widetilde{N}_0^{ns}$ will decrease (shift left) as q or N_1^t increases. Similarly, beliefs over the number of competing asset holders $\widetilde{N}_1^s$ in the interdealer market will increase (shift right) as q or N_1^t increase. Both of these forces work in conjunction to give the dependence depicted in Figure 4, as entering asset holders place smaller likelihoods on capturing the majority of gains in any interdealer transaction.

[19]Estimates are for the segmented demand setup; dotted lines give $q = .1$, dashed lines $q = .2$, and solid lines $q = .3$. The corresponding figure for the Co-demand setup takes on similar characteristics. Estimates were bootstrapped setting $m = .25$ and drawing random buyer networks for each $(\widetilde{N}_0^{ns}, \widetilde{N}_1^s)$ pair. Stable assignments were solved for using the basic linear programing techniques. p is set to .25 to yield a likelihood for each asset holder entering the market, $(1-p)q$, of roughly .2 or less. This is seen as a conservative range for the likelihood (placing limited emphasis on the role of interdealer trade), with roughly 20% of transactions observed between dealers in the data.

On the other side of the market, these forces work to benefit entering vacants. Unshocked vacants entering the interdealer market benefit from low $\widetilde{N}_0^{ns}$ (competing vacants) and high $\widetilde{N}_1^{s}$ (liquidating asset holders). Thus, their expectations over capturing high gains in the interdealer market only increase with q and N_1^t. Exactly how high λ reaches as market inventories N_1^t reaches its maximum capacity N depend on the sizes of q and N as well as the process $b(\cdot|\widetilde{N}_0^{ns}, \widetilde{N}_1^{s})$.

With λ and μ defined, we can reduce (3) as follows:

$$V_{0,N_1^t} = \left(\begin{array}{c} (1-q) \left(\begin{array}{c} \lambda \max \left\{ \delta E_t \left[V_{0,N_1^{t+1}} \right], \delta E_t \left[V_{1,N_1^{t+1}} \right] - \underline{v} \right\} \\ + (1-\lambda) \, \delta E_t \left[V_{0,N_1^{t+1}} \right] \end{array} \right) \\ + q \delta E_t \left[V_{0,N_1^{t+1}} \right] \end{array} \right) \tag{18a}$$

$$V_{1,N_1^t} = \left(\begin{array}{c} (1-p) \left(\begin{array}{c} (1-q) \max \left\{ \underline{v} + \delta E_t \left[V_{0,N_1^{t+1}} \right], \delta E_t \left[V_{1,N_1^{t+1}} \right] \right\} \\ +q \left(\begin{array}{c} \mu \max \left\{ \underline{v} + \delta E_t \left[V_{0,N_1^{t+1}} \right], \delta E_t \left[V_{1,N_1^{t+1}} \right] \right\} \\ + (1-\mu) \left(\underline{v} + \delta E_t \left[V_{0,N_1^{t+1}} \right] \right) \end{array} \right) \end{array} \right) \\ +p \left(v + \delta E_t \left[V_{0,N_1^{t+1}} \right] \right) \end{array} \right). \tag{18b}$$

Equations (18a) and (18b) embed the essential information needed towards characterizing steady state gains from equilibrium intermediation, which is addressed on the next section.

12 Equilibrium Market Behavior

In deriving the results in this section, I approximate the value gap by $G^t(N_1^t) \approx V^t_{1,N_1^t} - V^t_{0,N_1^t}$, and correspondingly redefine the market clearing condition (5a)-(5b). For N large this comes as a minor adjustment to the model[20], while yielding the useful recursive form to the value gap:

$$G\left(N_1^t\right) = \left(\begin{array}{c} pv + (1-p)\,\underline{v} \\ + \max\left\{0, \delta E_t\left[G\left(N_1^{t+1}\right)\right] - \underline{v}\right\} \left(\begin{array}{c} (1-p)(1-q) \\ +(1-p)\,q\mu - (1-q)\,\lambda \end{array} \right) \end{array} \right). \tag{19}$$

This section's first result characterizes how the choice to carry inventories depends on continuation values and market inventory. The analysis then moves to describe the role of the interdealer market in shaping the dependence of prices on aggregate inventory. The main descriptive results of the paper are summarized with two corollaries. Finally, the role of interdealer prices in driving these forces is characterized in Section 12.2.

12.1 Market Inventory and Asset Prices

Though we've established unique steady state values with Proposition 1, uniqueness of market equilibrium prices and inventories requires that the value gap, given with (18), be monotonically decreasing in N_1^t. This constitutes a law of demand for the market. Proposition 5 below establish sufficient conditions for a law of demand for moderate to low liquidity risk. Before doing so, I first move to more completely characterize equilibrium market behavior, taking as given a downward sloping value gap. This step will later be crucial when formally establishing the link between market inventory and prices.

The value to holding an asset will naturally increase with the arrival likelihood p. Even more, we have $G(N_1^t) \to V_{1,N_1^t} = v$ as p approaches its upper bound of one. For a given upstream price $P < v$ and with p sufficiently close to 1, we should generally expect all dealers to demand and hold assets ($N_1^t = N$) as long as possible until either high demand arrives or they are first hit by a capacity shock, in which case they are forced to sell and collect $\underline{v}$. At the other extreme, for p sufficiently below one, we should expect N_1^t to drop below N. However, if expectations over the arrival of future demand are particularly bleak, dealers may lose the incentive to hold inventories, opting instead to instantaneously convey assets to whichever downstream demand arrives (i.e. funnel assets). Which of these cases are obtained depends on the various parameters of the model. The following proposition formally distinguished between these equilibrium behaviors.

Proposition 4. *Each of the following (exhaustive) cases describe a symmetric Nash equilibrium strategy of asset holders[21]:*

[20] The descriptive results obtained below have been verified using numerical solutions (e.g. see Figures 7(a) and 7(b)) that do not employ this approximation.

[21] When gains to interdealer trade become zero, shocked assets holders are left indifferent between trading with vacant and funneling. In the proof of Proposition 4, I maintain these trades for continuity. Given this assumption, this obtains the unique symmetric Nash strategy by the monotonicity given with Proposition 3 part 1, below.

1. *There exists some integer* $\overline{N}_1 \leq N$ *such that* $\underline{v} = \delta E_t\left[G(N_1^{t+1})\right]$ *for each* $N_1^t > \overline{N}_1$ *and* $\underline{v} \leq \delta E_t\left[G(N_1^{t+1})\right]$ *for each* $N_1^t \leq \overline{N}_1$*. For each* $N_1^t < \overline{N}_1$*, assets holders never funnel. For each* $N_1^t \geq \overline{N}_1$*, unconstrained asset holders unable to sell in period* t *for price* v *mix on funneling (liquidating for* $\underline{v}$*) with some probability* $\theta\left(N_1^t\right) < 1$ *and carrying inventory into period* $t+1$ *with probability* $1 - \theta\left(N_1^t\right)$*. Further,* $\theta\left(N_1^t\right)$ *is weakly increasing in* N_1^t*.*
2. $\underline{v} > \delta E_t\left[G(N_1^{t+1})\right]$ *for all* $N_1^t \in \{0, \ldots, N\}$*. Then, asset holders never funnel assets.*

PROOF: Appendix D.

In case 1, when $N_1^t > \overline{N}_1$ asset holders unable to sell for v are left indifferent between liquidating for $\underline{v}$ and carrying inventory to period $t+1$, in which case asset funneling will be observed in equilibrium. In this case I term the market as being *saturated*. The intuition is natural: as asset holders crowd the market (as N_1^t increases) and the hedge that interdealer trade provides against liquidity risk degrades, the choice not to carry inventories becomes increasingly attractive. Then, when market inventory surpasses the threshold $\overline{N}_1$, residual asset holders must start shaving assets (liquidating for value $\underline{v}$) at a given rate so as to collectively maintain their indifference between liquidating (funneling) and carrying inventories.

It is important to recognize that this mixed strategy equilibrium will be the unique Nash equilibrium symmetric across asset holders. This is because when all shocked asset holders are expected to funnel, the unique optimal action to asset holders unable to sell downstream or in the interdealer market is to carry the asset as inventory, as expectations over future inventories will be low (and thus continuation values will be high). If instead all shocked asset holders are expected to carry inventories, then it remains optimal to funnel in these cases. Only when these asset holders collectively mix on funneling with probability $\theta(N_1^t)$ can they simultaneously choose to funnel. That $\theta(N_1^t)$ is increasing in N_1^t bares a natural economic interpretation: as greater market inventories are realized the incentive to carry assets across periods decays.

In case 2 asset funneling is never rational, leaving gains to interdealer trade strictly positive. In such markets dealers never entirely crowd each other out. Instead, the prospects of future high valued sales remain attractive enough that dealers remain willing to carry inventories, even when $N_1^t = N$ and interdealer trade is excluded. Accordingly, we can set $\overline{N}_1 > N$ in these cases.

In total, Proposition 4 gives conditions under which the market fails to exhibit beneficial interdealer trade. In case 1 when $N_1^t > \overline{N}_1$, dealers crowd each other out and rationally funnel assets rather than actively trading them between each other[22]. If instead N_1^t rises to N in case 2, all dealers hold assets and wait for high demand to arrive as long as possible. In both cases, the interdealer market is contemporaneously closed, leaving any vacants with their expected continuation payoff $V_{0,N_1^t} = \delta E_t\left[V_{0,N_1^{t+1}}\right]$.

[Figure 5]

[22]Again, I maintain interdealer trade in these scenarios only for analytical ease. What is essential is that gains to interdealer trade in these cases fall to zero.

Cases 1 and 2 are each depicted with Figure 5. in Figure 5(a) the market becomes saturated for $N_1^t \geq \overline{N}_1$ where asset holders can simultaneously be observed funnelling assets and carrying inventories in equilibrium. Figure 5(b) gives case 2 without saturation.

Next, as promised, the following proposition more formally establishes our law of demand for small q, as maintained above. Equally as important, part 2 will establish the role of the interdealer market in shaping equilibrium intermediation. In this and the following corollaries, I assume the following weak conidtion on the market. This first condition requires that the interdealer market becomes 'active' at some point. The second is more technical, and will be satisfied provided the dealer market is not too large relative to the asset's total issuance.

Assumption 8.

1. *Either:*

$$\textit{(i)} \qquad \left|\mu(N_1^t+1)-\mu(N_1^t)\right| > 0, \textit{ or} \tag{20a}$$

$$\textit{(ii)} \qquad \left|\frac{\partial}{\partial q}(\lambda(N_1^t+1)-\lambda(N_1^t))\right| > 0, \tag{20b}$$

and,

2. $$N_1^t\left(\hat{N}_1^{t-1}+1\right) \leq N_1^t\left(\hat{N}_1^{t-1}\right)+1. \tag{20c}$$

for each $N_1^t \in \{0,\dots,N-1\}$.[23]

Proposition 5. *There exist* $\overline{q}^1 \geq \overline{q}^2 > 0$ *and some bound* $\overline{dp}^1 > 0$ *such that:*

1. *for each* $q \in [0,\overline{q}^1]$, $G(N_1)$ *is weakly decreasing in* N_1.
2. *for each* $q \in [0,\overline{q}^2]$, *If* $|p(N_1+1)-p(N_1)| < \overline{dp}^1$ *for each* $N_1 \in \{0,\dots,N-1\}$ *then:*

$$\textit{(i)} \qquad \frac{\partial}{\partial q}\left[G(N_1+1)-G(N_1)\right] < 0, \textit{ and} \tag{21a}$$

$$\textit{(ii)} \qquad \frac{\partial}{\partial q}G(N_1) < 0, \tag{21b}$$

for each $N_1 \in \{0,\dots,N-1\}$.

PROOF: Appendix D.

The monotonicity in part 1 is driven by the monotonicity derived from Proposition 3: as the number of assets in total inventory rises, asset holders crowd each other out and suppress the value to holding an asset, V_{1,N_1^t}, while inflating the option value to serving as a vacant in the market, V_{0,N_1^t}. As shown in the Appendix, Proposition 5 part 1 does not require p to be a strictly decreasing function of N_1^t. Instead, part 1 is derived solely through forces derived

[23]Condition 2 requires that an increase residual assets $\hat{N}_1^{t-1}$ increases the subsequent period's market inventory by at most one. A sufficient condition for this is for $\Phi\left(N_1^t;\hat{N}_1^{t-1}\right)$ to be stationary.

from the interdealer market, with the value gap decreasing in N_1^t, as conditions increasingly favor vacants entering the interdealer market.

Part 2 (i) shows that when the influence of N_1^t on demand uncertainty is bounded, the dependence of the value gap (and thus of the equilibrium bid-ask spread) on market inventory increases with liquidity risk. The necessary conditions (a) and (b) ensure that the interdealer market takes center stage in driving the dependence of prices on market demand. If the decay in p is too significant, then the effects of increasing q on equilibrium inventory in upstream primary markets may dominate the effect on demand. Determining in which direction the gap decay moves in response to increasing q away from zero will depend on the particulars of the market (e.g. the function p and process generating $\widetilde{\mathcal{B}}$). But with p near constant and the interdealer market driving the dependence of the value gap on N_1^t, the positive dependence between market inventory and bid-ask spread only strengthens as the interdealer market takes on a more important role hedging the market against liquidity risk.

[Figure 6]

As implied by Proposition 5's part 1, Figure 6 depicts a downward sloping value gap. The fast decay in expected interdealer prices is addressed below in Section 12.2. If the decay in the value is primarily driven by the interdealer market, then an increase in q magnifies this link. Also shown in the proof of Proposition 5, the size of $\frac{\partial}{\partial q}G(N_1^t)$ will increase with N_1^t. Here we see the interdealer market at work, effectively hedging the market against added risk. When N_1^t is small and conditions in the interdealer market are most benevolent for entering asset holders, an increase q yields only a limited effect on the value gap. If instead N_1^t is large and conditions in the interdealer market move toward favoring vacants, the functionality of the interdealer market as a hedge against liquidity risk deteriorates. As such, the effect of this increase in q on the inventory-price link becomes more pronounced.

We can recast part 2 to characterize the role of the interdealer market in pinning each period t's equilibrium market inventory.

Corollary 1. *(*equilibrium liquidity provision*) For each $q \in [0, \overline{q}^2)$ and $s > 0$, if $|p(N_1+1) - p(N_1)| < \overline{dp}^1$ for each $N_1 \in \{0, \dots, N-1\}$ then an increase in q gives a first order stochastic shift down in the distribution of N_1^{t+s}, given N_1^t in period t, and*

PROOF: Appendix D.

Here, we see that liquidity risk applies upward pressure on the equilibrium provision of liquidity from dealers in a persistent way. As we increase q and depart from the efficient benchmark, expectations over the number of future vacants increases. Liquidity risk and interdealer trade play an increasingly important role. Asset intermediation (measured with N_1^t) is persistently substituted for market liquidity (measued with N_0^t) as an endogenous response to the added liquidity risk.

12.2 The role of Interdealer Prices and Search Frictions

With prices providing the basic empirical unit of analysis, this section recasts the above forces in terms of expected interdealer prices. This will also distinguish the role of prices from that of interdealer search frictions. First, decompose λ and μ as follows:

$$\lambda = (1 - E_t[s]) \Pr(\chi_B), \tag{22a}$$

$$\mu = E_t[s] \Pr(\chi_S). \tag{22b}$$

where $E_t[s]$ gives the expected seller share in the interdealer market:

$$E_t[s] := E_t\left[\frac{\sum s_x}{\widetilde{X}}\right]. \tag{23}$$

Again, $\{s_x\}$ gives the set of seller shares in transactions $\widetilde{X}$ given by $\Lambda_{x,\widetilde{\mathcal{B}}}$. Expectations are taken over $\widetilde{N}_0^{ns}$, $\widetilde{N}_1^s$, and $\widetilde{\mathcal{B}}$ (and the particular stable assignment, if needed), while conditioning on N_1^t. Correspondingly, $\Pr(\chi_B) := \lambda / (1 - E_t[s])$ and $\Pr(\chi_S) := \mu / E_t[s]$. We can interpret these values as the likelihoods that buyers and sellers place on transacting in the interdealer market.

In conjunction with (17a)-(17f), the following properties (also shown in the Appendix C) will hold:

1. $E_t[s] \approx 1$ for N_1^t near 0, (24a)
2. $E_t[s] \approx 0$ for N_1^t near N, (24b)
3. $E_t[s]$ is weakly decreasing in N_1^t $\left(\text{for each } N_1^t\right)$. (24c)

These properties highlight the roles that asset holders and vacants fill at both extremes of the market. For N_1^t near 0 with few assets in circulation, shocked asset holders entering the secondary market effectively serve as monopolists, capturing the majority of gains from interdealer transactions. At the other extreme for N_1^t near N, vacants entering the interdealer market serve as monopsonists, assigning high probability to the event of capturing the majority value. These market forces work in conjunction, together pushing in the direction of (17a)-(17f).

Next, denote the set of interdealer prices by:

$$\widetilde{\mathcal{P}}^D := \left\{ s_x \delta E_t\left[G(N_1^{t+1})\right] + (1 - s_x)\, \underline{v} | x \in \widetilde{X} \right\}.$$

Writing the expectation over these prices $E_t\left[\widetilde{\mathcal{P}}^D\right]$, we have:

$$
\begin{aligned}
E_t\left[\widetilde{\mathcal{P}}^D\right] \quad : \quad &= E_t\left[\frac{\sum_{x\in\widetilde{X}}\left(s_x\delta E_t\left[G(N_1^{t+1})\right]+(1-s_x)\,\underline{v}\right)}{\left|\widetilde{X}\right|}\right] \\
&= E_t\left[\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right)\frac{\sum_{x\in\widetilde{X}}(s_x)}{\left|\widetilde{X}\right|}+\underline{v}\right] \\
&= \left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right)E_t\left[s\right]+\underline{v}, \qquad (25)
\end{aligned}
$$

As shown in the Appendix in the proof of Proposition 5, $\delta E_t\left[G(N_1^{t+1})\right]$ is decreasing with N_1^t for $q\in[0,\overline{q}^1]$. And with $E_t\left[s\right]$ decreasing in N_1^t by condition (17c), expected interdealer price $E_t\left[\widetilde{\mathcal{P}}^D\right]$ will exhibit a quick decay with N_1^t. That is, interdealer shares transition from being primarily captured by sellers to being captured by buyers as N_1^t is increased. This transition can be observed in Figures 5 through 7(b).

Using (22a) and (22b), the following expressions relating $E_t\left[\widetilde{\mathcal{P}}^D\right]$ to λ and μ also hold:

$$
\begin{aligned}
\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right)\lambda &= \left(\delta E_t\left[G(N_1^{t+1})\right]-E_t\left[\mathcal{P}^D\right]\right)\Pr\left(\chi_B\right) \qquad (26a)\\
\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right)\mu &= \left(E_t\left[\mathcal{P}^D\right]-\underline{v}\right)\Pr\left(\chi_S\right). \qquad (26b)
\end{aligned}
$$

Substituting for λ and μ, this yields the revealing formulation for the value gap:

$$
G\left(N_1^t\right)=\begin{pmatrix}(v-\underline{v})\,p+\underline{v}+(1-p)\,(1-q)\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right)\\ +(1-p)\,q\left(E_t\left[\mathcal{P}^D\right]-\underline{v}\right)\Pr\left(\chi_S\right)\\ -(1-q)\left(\delta E_t\left[G^{t+1}\right]-E_t\left[\mathcal{P}^D\right]\right)\Pr\left(\chi_B\right)\end{pmatrix}. \qquad (27)
$$

To interpret (27), we can compare it with the analogous expression that excludes the interdealer market:

$$
V_{1,N_1^t}^t=(v-\underline{v})\,p+\underline{v}+(1-p)\,(1-q)\left(\delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]-\underline{v}\right). \qquad (28)
$$

The top term of the right hand side of (27) gives the same term had the interdealer market been shutdown (where $\delta E_t\left[G(N_1^{t+1})\right]=\delta E_t\left[V_{1,N_1^{t+1}}^{t+1}\right]$). Without an interdealer market, the bottom term becomes zero, as shocked asset holders are forced to collect value $\underline{v}$ and vacants are unable to gain from interdealer trade. Introducing the interdealer market effects the value gap in two opposing directions. If asset holders face favorable conditions in the interdealer market with high expected price and transaction likelihood, this added value works to expand the value gap. If instead prices are low with vacants expecting to transact with high likelihood, the effect on the option value to supplying liquidity to the market works to suppress the gap. Therefore, we observe a switch in the net-effect of interdealer trade on the upstream market price $G(N_1^t)$, as expectations over interdealer prices move from favoring sellers to favoring buyers.

If we again bound the dependence of p on market inventory, it is easy to show that the bottom term in (27) is monotonically decreasing in N_1^t. Even more, this dependence will increase with q. As the expected interdealer price decays from its upper bound $\delta E_t\left[G(N_1^{t+1})\right]$ when N_1^t is near 0 to its lower bound $\underline{v}$ when N_1^t is near N, interdealer trade works to magnify the inverse dependence of the value gap on N_1^t. Liquidity risk and active interdealer trade work inseparably to create this link from market inventory to upstream price.

This inventory-price link is simultaneously derived on two margins. First, the likelihood of finding a vacant dealer depends on the total number of assets in circulation in the dealer market. As N_1^t increases and more dealers hold inventories, the frequency of interdealer transactions will rise and then fall non-monotonically. However, upon conditioning on being a seller in the market, this probability– given with $\Pr(\chi_S)$– strictly decreases with N_1^t. The probability of transacting after conditioning on being a buyer, $\Pr(\chi_B)$, strictly rises. These divergent conditional likelihoods unambiguously push in the direction of the link. On the second margin, the prospect to any shocked asset holder of finding herself able to trade on favorable terms– as captured with $E_t\left[s\right]$– unambiguously degrades with N_1^t. That is, expectations over the bargaining positions of dealers move toward favoring buyers as the market becomes more congested with assets. Together, these two margins work simultaneously to influence upstream prices.

Together, expectations over interdealer prices and the likelihood of transacting in the interdealer market work to establish the link from market inventories to asset prices. Such dependence on market inventory, net of external influences through $p(N_1^t)$ or $\Phi(N_1^t;\hat{N}_1^{t-1})$ but working solely through interdealer markets, establishes the market's equilibrium response to the setting's basic frictions to asset intermediation.

13 Conclusion: discussion and extensions

With the above model market inventory is shown to proxy for the effectiveness of the interdealer market as a hedge on dealers' private risks. Interdealer trade effectively boosts prices in the upstream supply market. As interdealer activity is constrained trading costs expand and the ability of the dealer market to allocate assets to the most valuing demanders degrades. Interdealer prices and search frictions work together to drive these forces. The potency of these forces intimately depends on the equilibrium provision of liquidity from dealers, and in turn, on contemporaneous market inventory.

The reader may be concerned with the robustness of the above descriptive results upon loosening some of the various assumptions in the model. In turn, I consider both the simplifying assumptions of a binary inventory space and of the given binary structure of the setting's demand and liquidity risk.

Though the above discussion interprets each dealer to be a separate entity, there is little preventing us from allowing a subset of nodes to constitute a 'firm'. In such an extension to the model, the set of N nodes is partitioned into an integer number F of firms[24], and to include an integer number $I := N/F$ of possible inventories for each firm. An additional assumption on $b(\cdot|\widetilde{N}_0^{ns},\widetilde{N}_1^{s})$ would be needed to remove frictions within each firm, and 'entrants' would be taken I at a time. Then, each dealer i's demand in period t would depend on current private inventory $I_i \in \{1,\ldots,I\}$ as well as on inventories held by other dealers.

[24] This partitioned should come as a division if to maintain ex-ante symmetry across firms

The choice to liquidate or trade some portion of I_i in the interdealer market would then depend on the joint realizations of demand and liquidity shocks through the following period. While such an extension would allow for a richer treatment of the interdealer market, the basic link from market inventories to prices will continue to hold.

One can also consider a straight forward extension allowing for a richer form of liquidity risk in which dealers realize variable costs to carrying inventories. Interdealer trade would then create value as a function of the transacting dealers' relative inventory costs. The universe of assignment games is certainly rich enough to apply to such a setting with analogous entry results holding. Such a model would entail a considerable enrichment of an already high dimensional Bellman system. As such, the above model should be taken as a parsimonious attempt to characterize the role of market inventory and interdealer trade, without indulging in enriched forms of liquidity risk.

Allowing for variable liquidity risk does promise an exciting avenue for future work. Endogenizing the search and matching process to account for assortative matching as a function of the relative liquidity needs of the traders– movng beyond the above model's simple binary setup– could yield a more descriptive characterization of trading behavior in these markets (e.g. see Shimer and Smith (2000) [28])[25]. Negative assortative matching should naturally obtain under regular conditions, as trades between dealers with increasingly divergent liquidity needs create greater value. Concurrent variability over interdealer prices would be driven by heterogeneity in inventory costs as opposed to (or potentially, in addition to) the asymmetric bargain positions that dealers hold in the above setting.

Other natural extensions to the model could discriminate between dealers in the formation of links and allow persistence in dealer relationships. Future work could study games that endogenize link formation, with dealers investing in and maintaining links for future use (e.g. see Elliott (2013) [9]). However, allowing for persistence in the interdealer network departs from this paper's assumption of ex-ante symmetry over dealers. Such complications, along with the many rich topics addressed by the growing financial networks literature, are left outside the scope of this paper.

Finally, this work motivates an agenda for future empirical work. First, the extent to which market inventory can explain asset prices beyond alternative factors, such as asymmetric information, processing costs, and inventory risk sharing is left as an open empirical question. The answer will likely depend on market structure. Incorporating market inventories into price impact regressions offers a straight forward approach (e.g. see Foucault et. al. (2013) [11]). Though, more sophisticated approaches could include controls for market segmentation and correlation in demand. Finally, a suitable proxy for q, or more generally for the degree of heterogeneity in dealers' trading constraints, would be ideal. Only with such a tool can the role of interdealer trade in driving the above market inventory link be aptly assessed.

[25] I accredit Bill Zame for this suggestion.

List of Variables

1. $V_{a,N_1^t} \in \mathbb{R}_+$: net present (expected discounted) value to dealer holding $a \in \{0,1\}$ assets. $a = 0$ termed *vacant*. $a = 1$ termed *asset holder*.
2. $N_1^t \in \{0, N\}$: market inventory, equal to total number of asset holders.
3. $N_0^t = N - N_1^t$: number of vacants (liquidity suppliers).
4. $\hat{N}_1^{t-1} \in \{0, N_1^t\}$: market inventory carried between periods $t-1$ and t.
5. $\mathcal{V}_0^{ns} \in \mathbb{R}_+$: expected gains to unshocked vacant entering interdealer market.
6. $\mathcal{V}_1^{s} \in \mathbb{R}_+$: expected gains to shocked asset holder entering interdealer market.
7. $\widetilde{D} \in \mathbb{Z}_+$: number of high valuing demanders willing to purchase asset at price v.
8. $\Phi\left(N_1^t; \hat{N}_1^{t-1}\right)$: single period inverse asset supply.
9. $\widetilde{N}_1 = \max\left\{0, N_1^t - \widetilde{D}\right\}$: number of asset holders unable to sell for v.
10. $\widetilde{N}_1^s \in \left\{0, \ldots, \widetilde{N}_1\right\}$: number of shocked asset holders unable to sell for v.
11. $\widetilde{N}_0^{ns} \leq \{0, \ldots, N_0^t\}$: number of unshocked vacants.
12. $\widetilde{\mathcal{B}} \in 2^{\widetilde{N}_1^s \widetilde{N}_0^{ns}}$: interdealer market (bipartite network).
13. $\widetilde{X} \in \left\{0, \ldots, \min\left\{\widetilde{N}_1^s, \widetilde{N}_0^{ns}\right\}\right\}$: number of interdealer transactions.
14. $\lambda \in \mathbb{R}_+^{N+1}$: likelihood asset holder captures gains to trade in interdealer market, for each market inventory N_1^t.
15. $\mu \in \mathbb{R}_+^{N+1}$: likelihood vacant captures gains to trade in interdealer market, for each market inventory N_1^t.
16. $G\left(N_1^t\right) = V_{1,N_1^t} - V_{0,N_1^t}$: value gap.

References

[1] Allen, Franklin; Gale, Douglas . "Limited Market Participation and Volatility of Asset Prices". *The American Economic Review*, 84(4):933–55, September 1994.

[2] Amihud, Yakov; Mendelson, Haim. "Dealership market: Market-making with inventory". *Journal of Financial Economics*, 8(1):31–53, March 1980.

[3] Amihud, Yakov; Mendelson, Haim. "Asset pricing and the bid-ask spread". *Journal of Financial Economics*, 17(2):223–49, December 1986.

[4] Biais, Bruno. "Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets". *Journal of Finance*, 48(1):157–85, March 1993.

[5] Brunnermeier, Markus; Pedersen, Lasse Heje. "Market Liquidity and Funding Liquidity". *The Review of Financial Studies*, 22(6):2201–38, June 2009.

[6] Butler, Alexander; Grullon, Gustavo; Weston, James. "Stock Market Liquidity and the Cost of Issuing Equity". *Journal of Financial and Quantititave Analysis*, 40(2):33148, June 2005.

[7] Comerton-Forde, Carole; Hendershott, Terrence; Jones, Charles; Moulton, Pamela; Seasholes, Mark. "Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues". *Journal of Finance*, 65(1):295–331, February 2010.

[8] Duffie, Darrell; Grleanu, Nicolae; Pedersen, Lasse Heje. "Over-the-Counter Markets". *Econometrica*, 73(6):1815–47, November 2005.

[9] Elliott, Matt. "Inefficiencies in Networked Markets". *mimeo*, 2013.

[10] Fama, Eugene; French, Kenneth. "Financing decisions: who issues stock?". *Journal of Financial Economics*, 76(3):54982, June 2005.

[11] Foucault, Theirry; Pagano, Marco; Röell, Alisa. "Market Liquidity: Theory, Evidence, and Policy". *Oxford University Press*, 2013.

[12] Glosten, Lawrence; Milgrom, Paul. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders". *Journal of Financial Economics*, 14(1):71–100, March 1985.

[13] Grossman, Sanford; Miller, Merton. "Liquidity and Market Structure". *Journal of Finance*, 43(3):617–33, July 1988.

[14] Habermeier, Karl; Kirilenko, Andrei. "Securities Transaction Taxes and Financial Markets". *IMF Staff Papers*, 50, Special Issue:165–80, 2003.

[15] Hansch, Oliver; Naik, Narayan; Viswanathan, S. "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange". *Journal of Finance*, 53(5):1623–56, October 1998.

[16] Hasbrouk, Joel. "Trades, Quotes, Inventories, and Information". *Journal of Financial Economics*, 22(2):229–252, December 1988.

[17] Hendershott, Terrence; Madhavan, Ananth. "Click or Call? Auction versus Search in the Over-the-Counter Market". *mimeo*, 2013.

[18] Ho, Thomas; Stoll, Hans. "The Dynamics of Dealer Markets under Competition". *Journal of Finance*, 38(4):1053–74, September 1983.

[19] Hong, Gwangheon; Warga, Arthur. "An Empirical Study of Bond Market Transactions". *Financial Analysts Journal*, 56(2):32–46, March 2000.

[20] Huang, Roger; Cai, Jun; Wang, Xiaozu. "Information-Based Trading in the Treasury Note Interdealer Broker Market". *Journal of Financial Intermediation*, 11(3):269296, July 2002.

[21] Kotowski, Maciej; Leister, C. Matthew. "Trading Networks and Equilibrium Intermediation". *mimeo*, 2014.

[22] Kranton, Rachel; Minehart, Deborah. "A Theory of Buyer-Seller Networks". *American Economic Review*, 91(3):485–508, June 2001.

[23] Manaster, Steven; Mann, Steven. "Life in the Pits: Competitive Market Making and Inventory Control". *Review of Financial Studies*, 11(3):953–75, 1996.

[24] Reiss, Peter; Werner, Ingrid. "Does Risk Sharing Motivaite Interdealer Trading?". *Journal of Finance*, 53(5):1657–1703, October 1998.

[25] Roth, Alvin; Sotomayor, Marilda. "Two-sided Matching". *Handbook of Game Theory with Economic Applications*, 1(1), 1992.

[26] Rubinstein, Ariel; Wolinsky, Asher. "Middlemen". *Quarterly Journal of Economics*, 102(3):581–93, August 1987.

[27] Seasholes, Mark; Hendershott, Terrence. "Market Maker Inventories and Stock Prices". *American Economic Review*, 97(2):210–214, May 2007.

[28] Shimer, Robert; Smith, Lones. "Assortative Matching and Search". *Econometrica*, 68(2):343–369, 2000.

[29] Viswanatha, S.; Wang, James. "Inter-Dealer Trading in Financial Markets". *Journal of Business*, 77(4):987–1040, October 2004.

[30] Vogler, Karl-Hubert. "Risk Allocation and Inter-dealer Trading". *European Economic Review*, 41(8):1615–34, August 1997.

Appendixes

C Appendix: Existence and Uniqueness, and Inter-dealer Market Properties

C.1 Preliminary results

Proposition 6. *If for each transaction x resulting from each network $\widetilde{\mathcal{B}}$, $\Lambda_{x,\widetilde{\mathcal{B}}}(V_0, V_1)$ gives share $s_x \in [0,1]$ of the gains from trade from Δ to the seller, then $\Lambda_{x,\widetilde{\mathcal{B}}}(V_0, V_1)$ is non-expansionary (for any metric $\|\cdot\|$).*

PROOF: Let $\delta V_1 - \delta V_0 - \underline{v}$ give the length and height of the value set. Writing vectors (y,z) is column form $\binom{y}{z}$, then $\left\|\Lambda_{x,\widetilde{\mathcal{B}}}(V_0, V_1) - \Lambda_{x,\widetilde{\mathcal{B}}}(V_0', V_1')\right\|$ is given by:

$$
\begin{aligned}
& \left\|\begin{pmatrix}(1-s_x)(\delta V_1 - \delta V_0 - \underline{v}) + \delta V_0 \\ s_x(\delta V_1 - \delta V_0 - \underline{v}) + \delta V_0 + \underline{v}\end{pmatrix} - \begin{pmatrix}(1-s_x)(\delta V_1' - \delta V_0' - \underline{v}) + \delta V_0' \\ s_x(\delta V_1' - \delta V_0' - \underline{v}) + \delta V_0' + \underline{v}\end{pmatrix}\right\| \\
= & \left\|\begin{pmatrix} s_x(\delta V_0 - \delta V_0') + (1-s_x)(\delta V_1 - \delta V_1') \\ s_x(\delta V_1 - \delta V_1') + (1-s_x)(\delta V_0 - \delta V_0')\end{pmatrix}\right\| \\
= & \left\| s_x \begin{pmatrix}\delta V_0 - \delta V_0' \\ \delta V_1 - \delta V_1'\end{pmatrix} + (1-s_x)\begin{pmatrix}\delta V_1 - \delta V_1' \\ \delta V_0 - \delta V_0'\end{pmatrix}\right\| \\
\leq & \; s_x \left\|\begin{pmatrix}\delta V_0 - \delta V_0' \\ \delta V_1 - \delta V_1'\end{pmatrix}\right\| + (1-s_x)\left\|\begin{pmatrix}\delta V_1 - \delta V_1' \\ \delta V_0 - \delta V_0'\end{pmatrix}\right\| \\
= & \left\|\begin{pmatrix}\delta V_0 - \delta V_0' \\ \delta V_1 - \delta V_1'\end{pmatrix}\right\|,
\end{aligned}
$$

the inequality holding as a triangle inequality. $\square$

Lemma 3. *For any vectors $\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}$ and $\binom{V'_{0,N_1}}{V'_{1,N_1}}_{N_1}$ in $[0,\delta v]^{N+1} \times [\underline{v}, v]^{N+1}$, take the function $\overrightarrow{\mathcal{V}} : [0,\delta v]^{N+1} \times [\underline{v}, v]^{N+1} \longmapsto \Delta$ defined by:*

$$
\overrightarrow{\mathcal{V}}\left(\begin{pmatrix}V_{0,N_1} \\ V_{1,N_1}\end{pmatrix}_{N_1}\right) = \begin{pmatrix}\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right) \\ \widetilde{\mathcal{V}}^{s}_{1,N_1^t}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right)\end{pmatrix},
$$

where each $\widetilde{\mathcal{V}}^{ns}_{a,N_1^t}\left(\binom{V_{0,N_1}}{V_{1,N_1}}\right)$ $(a = 0,1)$ gives values in Δ taking $E_t\left[V^{t+1}_{a,N_1^{t+1}}\right]$ as convex combinations of $\binom{V_{0,N_1}}{V_{1,N_1}}$ with weights $\Pr\left[N_1|N_1^t\right]$[26]. Then, $\overrightarrow{\mathcal{V}}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right)$ satisfies:

$$
\left\|\overrightarrow{\mathcal{V}}\left(\begin{pmatrix}V_{0,N_1} \\ V_{1,N_1}\end{pmatrix}_{N_1}\right) - \overrightarrow{\mathcal{V}}\left(\begin{pmatrix}V'_{0,N_1} \\ V'_{1,N_1}\end{pmatrix}_{N_1}\right)\right\|_\infty \leq \left\|\begin{pmatrix}V_{0,N_1} \\ V_{1,N_1}\end{pmatrix}_{N_1} - \begin{pmatrix}V'_{0,N_1} \\ V'_{1,N_1}\end{pmatrix}_{N_1}\right\|_\infty . \tag{29}
$$

[26]That is, the probability that $N_1^{t+1} = N_1$ conditioning on period t market inventory N_1^t.

PROOF: $\widetilde{\mathcal{V}}^{ns}_{0,N_1}$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1}$ are each convex combination of (ie. expectation over) the respective intermediary's disagreement value (i.e. $\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]$ for vacants and $\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]+\underline{v}$ for asset holders) and the respective value shares as determined by points in Δ (from transaction outcomes). By Assumption 4, the set of transactions is indepedent of $\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]$ and $\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]$ and thus the values from disagreements (multiplied by the likelihood of those outcomes) cancel. By Assumption 5 the value shares from transactions are non-expansionary, and thus $\left\|\overrightarrow{\mathcal{V}}\left(\binom{V_{0,N_1}}{V_{1,N_1}}\right)-\overrightarrow{\mathcal{V}}\left(\binom{V'_{0,N_1}}{V'_{1,N_1}}\right)\right\|_\infty$ is no greater than $\left\|\binom{V_0}{V_1}-\binom{V'_0}{V'_1}\right\|_\infty$ multiplied by the likelihood of a transaction occurring. $\square$

Note that (29) will continue to hold upon setting either $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right)$ or $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right)$ in the function $\overrightarrow{\mathcal{V}}\left(\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}\right)$ to the vector of zeros. This fact is used in the following proof.

C.2 Proof of Proposition 1

First, (3) can be expanded to:

$$V^t_{0,N_1^t}=\begin{bmatrix} (1-q)\left[\sum_{\widetilde{N}_1^s\geq 0,\widetilde{N}_0^{ns}\geq 1} g_1\left(\widetilde{N}_1^s\middle|N_1^t\right)h_0^B\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}\right] \\ +q\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]\end{bmatrix}$$
$$V^t_{1,N_1^t}=\begin{bmatrix} (1-p)\begin{bmatrix} (1-q)\max\left\{\underline{v}+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]\right\} \\ +q\left[\sum_{\widetilde{N}_1^s\geq 1,\widetilde{N}_0^{ns}\geq 0} g_1^S\left(\widetilde{N}_1^s\middle|N_1^t\right)h_0\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\widetilde{\mathcal{V}}^{s}_{1,N_1^t}\right]\end{bmatrix} \\ +p\left(v+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]\right)\end{bmatrix}, \tag{30}$$

where the first sums are taken over $\widetilde{N}_1^s\leq N_1^t$ and $\widetilde{N}_0^{ns}\leq N-N_1^t$, and the second sums are over all networks $\widetilde{\mathcal{B}}$. I prove that the function $[0,\delta v]^{N+1}\times[\underline{v},v]^{N+1}\longmapsto[0,\delta v]^{N+1}\times[\underline{v},v]^{N+1}$ given with the set of $2(N+1)$ functions (two for each $N_1^t=0,\ldots,N$) from the right hand side of (30) gives a contraction mapping in the metric space $\left(\mathbb{R}_+^{2(N+1)},\|\cdot\|_\infty\right)$. Again, we take $E_t\left[V^{t+1}_{a,N_1^{t+1}}\right]$ $(a=1,2)$ as a convex combination of $\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}$ with weights $\Pr\left[N_1|N_1^t\right]$.

Take $\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}$ and $\binom{V'_{0,N_1}}{V'_{1,N_1}}_{N_1}$ in $[0,\delta v]^{N+1}\times[\underline{v},v]^{N+1}$. We suppress the arguments of $\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}$ and $\widetilde{\mathcal{V}}^{s}_{1,N_1^t}$ and prime these functions when evaluated at $\binom{V'_{0,N_1}}{V'_{1,N_1}}_{N_1}$:

$$\widetilde{\mathcal{V}}^{s\prime}_{a,N_1^t}:=\widetilde{\mathcal{V}}^{s}_{a,N_1^t}\left(\begin{pmatrix}V'_{0,N_1}\\ V'_{1,N_1}\end{pmatrix}_{N_1}\right).$$

We also use $\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty$ to denote $\left\|\binom{V_{0,N_1}}{V_{1,N_1}}_{N_1}-\binom{V'_{0,N_1}}{V'_{1,N_1}}_{N_1}\right\|_\infty$.

Expanding, differencing, and taking the sup-norm of the right hand side of (30) gives:

$$\left\|\begin{pmatrix} \begin{bmatrix} (1-q)\left[\sum_{\widetilde{N}_1^s\geq 0,\widetilde{N}_0^{ns}\geq 1} g_1^B\left(\widetilde{N}_1^s\middle|N_1^t\right) h_0^B\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\left(\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}-\widetilde{\mathcal{V}}^{ns\prime}_{0,N_1^t}\right)\right] \\ +q\delta\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \end{bmatrix} \\ \begin{bmatrix} (1-p)\begin{bmatrix} (1-q)\left(\max\left\{\underline{v}+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]\right\}-\max\left\{\underline{v}+\delta E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]\right\}\right) \\ +q\left[\sum_{\widetilde{N}_1^s\geq 1,\widetilde{N}_0^{ns}\geq 0} g_1^S\left(\widetilde{N}_1^s\middle|N_1^t\right) h_0^S\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\left(\widetilde{\mathcal{V}}^{s}_{1,N_1^t}-\widetilde{\mathcal{V}}^{s\prime}_{1,N_1^t}\right)\right] \end{bmatrix} \\ +p\delta\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \end{bmatrix} \end{pmatrix}\right\|_\infty .$$

Take the two cases seperately:

$$\begin{aligned} (1)\ \max\left\{\underline{v}+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]\right\} &= \underline{v}+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right],\text{ and} \\ (2)\ \max\left\{\underline{v}+\delta E_t\left[V^{t+1}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]\right\} &= \delta E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]. \end{aligned}$$

Then, arbitrarily choosing a term from the second (i.e. subtracted) maximand will at most increase the bottom term of the above expression. Thus, we can set $\max\left\{\underline{v}+\delta E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right],\delta E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]\right\}$ to $\underline{v}+\delta E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]$ in case (1) and to $\delta E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]$ in case (2), establishing the following upper bound:

$$\begin{aligned} &\leq \left\|\begin{pmatrix} \begin{bmatrix} (1-q)\left[\sum_{\widetilde{N}_1^s\geq 0,\widetilde{N}_0^{ns}\geq 1} g_1^B\left(\widetilde{N}_1^s\middle|N_1^t\right) h_0^B\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\delta\left(\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}-\widetilde{\mathcal{V}}^{ns\prime}_{0,N_1^t}\right)\right] \\ +q\delta\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \end{bmatrix} \\ \begin{bmatrix} (1-p)\begin{bmatrix} (1-q)\,\delta\left(\max\left\{E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right],E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]\right\}\right) \\ +q\left[\sum_{\widetilde{N}_1^s\geq 1,\widetilde{N}_0^{ns}\geq 0} g_1^S\left(\widetilde{N}_1^s\middle|N_1^t\right) h_0^S\left(\widetilde{N}_0^{ns}\middle|N_1^t\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\delta\left(\widetilde{\mathcal{V}}^{s}_{1,N_1^t}-\widetilde{\mathcal{V}}^{s\prime}_{1,N_1^t}\right)\right] \end{bmatrix} \\ +p\delta\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \end{bmatrix} \end{pmatrix}\right\|_\infty \\ &\leq \delta\left\|\begin{pmatrix} (1-q)\left(\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}-\widetilde{\mathcal{V}}^{ns\prime}_{0,N_1^t}\right)+q\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \\ \begin{bmatrix} (1-p)\begin{bmatrix} (1-q)\max\left\{E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right],E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]\right\} \\ +q\left(\widetilde{\mathcal{V}}^{s}_{1,N_1^t}-\widetilde{\mathcal{V}}^{s\prime}_{1,N_1^t}\right) \end{bmatrix} \\ +p\left(E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right]\right) \end{bmatrix} \end{pmatrix}\right\|_\infty \\ &= \delta\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty \left\|\begin{pmatrix} (1-q)\dfrac{\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}-\widetilde{\mathcal{V}}^{ns\prime}_{0,N_1^t}}{\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty}+q\dfrac{E_t\left[V^{t+1}_{0,N_1^{t+1}}-V^{t+1\prime}_{0,N_1^{t+1}}\right]}{\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty} \\ \begin{bmatrix} (1-p)\begin{bmatrix} (1-q)\dfrac{\max\left\{E_t\left[V^{t+1}_{0,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{0,N_1^{t+1}}\right],E_t\left[V^{t+1}_{1,N_1^{t+1}}\right]-E_t\left[V^{t+1\prime}_{1,N_1^{t+1}}\right]\right\}}{\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty} \\ +q\dfrac{\widetilde{\mathcal{V}}^{s}_{1,N_1^t}-\widetilde{\mathcal{V}}^{s\prime}_{1,N_1^t}}{\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty} \end{bmatrix} \\ +p\dfrac{E_t\left[V^{t+1}_{0,N_1^{t+1}}-V^{t+1\prime}_{0,N_1^{t+1}}\right]}{\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty} \end{bmatrix} \end{pmatrix}\right\|_\infty \\ &\leq \delta\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty \left\|\begin{pmatrix} ((1-q)+q)\,\mathbf{1}_{N_1^t} \\ ((1-p)((1-q)+q)+p)\,\mathbf{1}_{N_1^t} \end{pmatrix}\right\|_\infty = \delta\left\|\overrightarrow{V}-\overrightarrow{V}'\right\|_\infty, \end{aligned}$$

where $1_{N_1^t}$ gives the vector of ones in R_+^{N+1}. The final inequality follows from Lemma 3. With $\delta<1$ we have a contraction mapping, and by Banach's fixed point theorem (30) must yield a unique fixed-point. $\square$

C.3 Conditions (17d) and (17f):

I expand λ and μ as follows:

$$\lambda = \frac{\mathcal{V}^{ns}_{0,N_1^t} - \delta E_t\left[V_{0,N_1^{t+1}}\right]}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}$$

$$= \frac{\left[\begin{array}{c}\left(\sum_{\widetilde{N}_1^s\geq 0}^{N_1^t}\sum_{\widetilde{N}_0^{ns}\geq 1}^{N-N_1^t} g_1^B\left(\widetilde{N}_1^s\left|N_1^t\right.\right) h_0^B\left(\widetilde{N}_0^{ns}\left|N_1^t\right.\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\widetilde{\mathcal{V}}^{ns}_{0,N_1^t}\right)\\ -\delta E_t\left[V_{0,N_1^{t+1}}\right]\end{array}\right]}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}, \tag{31}$$

$$\mu = \frac{\mathcal{V}^{s}_{1,N_1^t} - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}$$

$$= \frac{\left[\begin{array}{c}\left(\sum_{\widetilde{N}_1^s\geq 0}^{N_1^t}\sum_{\widetilde{N}_0^{ns}\geq 1}^{N-N_1^t} g_1^S\left(\widetilde{N}_1^s\left|N_1^t\right.\right) h_0^S\left(\widetilde{N}_0^{ns}\left|N_1^t\right.\right)\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\middle|\widetilde{N}_0^{ns},\widetilde{N}_1^s\right)\widetilde{\mathcal{V}}^{s}_{1,N_1^t}\right)\\ -\delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}\end{array}\right]}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}, \tag{32}$$

(17d).: With Assumption 2.3:

$$g_1^B\left(0\left|N_1^t\right.\right) = h_0^B\left(N_0^t\left|N_1^t\right.\right) = 1 \text{ when } q = 0. \tag{33}$$

Evaluating (31), one obtains for $q = 0$:

$$\begin{aligned}\lambda &= \frac{\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}}\left|N_0^t,0\right.\right)\widetilde{\mathcal{V}}^{ns}_{0,N_1} - \delta E_t\left[V_{0,N_1^{t+1}}\right]}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}}\\ &= \frac{\left(\delta E_t\left[V_{0,N_1^{t+1}}\right]\right) - \delta E_t\left[V_{0,N_1^{t+1}}\right]}{\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v}} = 0,\end{aligned}$$

for each N_1^t.

(17f).: By the above, at $q = 0$ we must have:

$$\nabla_{N_1^t}\lambda = 0.$$

Then, by Proposition 3, $\nabla_{N_1^t}\lambda \geq 0$, and by continuity in q, $\frac{\partial}{\partial q}\nabla_{N_1^t}\lambda \geq 0$ in some neighborhood of $q = 0$. $\square$

C.4 Conditions (24c).3

I first derive an analogous property to Proposition 2:

Lemma 4. $\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \frac{\sum_x s}{\left|\widetilde{X}\right|}$ *is decreasing in* $\widetilde{N}_1^s$ *and increasing in* $\widetilde{N}_0^{ns}$.

PROOF: The proof mirrors that of Proposition 2. Upon including an additional seller i' to $\widetilde{N}_1^s$ (obtaining seller set $i' \cup \widetilde{N}_1^s$), the likelihood of subgraph B (i.e. by excluding all links to i') occuring after averaging over all $B_{i'}$ realizations must equal $b\left(\mathcal{B} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right)$ by link independence. Further, given any realized $B_{i'}$, the value to each $i \in \widetilde{N}_1^s$ can only go down and the value to each $j \in \widetilde{N}_0^{ns}$ can only go up upon entry of i' (conditioning on supgraph B) by Lemma 2. The second fact implies that the number of trasactions $\left|\widetilde{X}\right|$ can only go up, so that any additional transaction x' must give $s_{x'} = 0$. This in turn implies $\frac{\sum_x s}{\left|\widetilde{X}\right|}$ can only go down, and that the expectation $\sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \frac{\sum_x s}{\left|\widetilde{X}\right|}$ is weakly decreasing with $\widetilde{N}_1^s$. To show that $\sum_{\widetilde{\mathcal{B}}} b\left(\mathcal{B} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \frac{\sum_x s}{\left|\widetilde{X}\right|}$ is increasing in $\widetilde{N}_0^{ns}$, use a similar argument by adding some j' to $\widetilde{N}_0^{ns}$. $\square$

Now to show Condition (24c).3, expand $E_t[s]$:

$$E_t[s] = \sum_{\widetilde{N}_1^s \geq 0}^{N_1^t} \sum_{\widetilde{N}_0^{ns} \geq 0}^{N-N_1^t} g_1\left(\widetilde{N}_1^s \middle| N_1^t\right) h_0\left(\widetilde{N}_0^{ns} \middle| N_1^t\right) \sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \frac{\sum_x s}{\left|\widetilde{X}\right|}$$

the result follows from the fact that expectations over monotonically decreasing functions decrease [increase] upon first order shifts up [down] in the underlining distribution. $\square$

D Appendix: Equilibrium Market Behavior

To ease notation, I define the following operator:

Definition 4. *For each $N_1^t \in \{0, \ldots, N-1\}$, the difference operator $\nabla_{N_1^t}$ is defined by $\nabla_{N_1^t} F\left(N_1^t\right) := F\left(N_1^t + 1\right) - F\left(N_1^t\right)$ for any function $F : \{0, \ldots, N\} \to \mathbb{R}$.*

For any two function of market inventory $F_{N_1^t}$ and $H_{N_1^t}$, we can derive the following product rule as follows:

$$\begin{aligned} \nabla_{N_1^t}\left(F^t H^t\right) \quad &: \quad = F_{N_1^{t+1}} H_{N_1^{t+1}} - F_{N_1^t} H_{N_1^t} \\ &\approx \quad H_{N_1^t}\left(F_{N_1^{t+1}} - F_{N_1^t}\right) + F_{N_1^t}\left(H_{N_1^{t+1}} - H_{N_1^t}\right). \end{aligned} \tag{34}$$

In what follows, we will assume that cross difference terms $(F_{N_1^{t+1}} - F_{N_1^t})(H_{N_1^{t+1}} - H_{N_1^t})$ will be negligible, entailing an approximate product rule for $\nabla_{N_1^t}$. This approximation will become increasingly negligable for large values of N.

D.1 Expectations over the Interdealer Market

Under the co-demand setup, the interim probability mass function[27] of $\widetilde{N}_1^s$- for buyers and for sellers entering the interdealer market, respectively- is given by the following conditional binomial distributions:

$$g_1^B\left(\widetilde{N}_1^s \middle| N_1^t\right) \quad : \quad = \sum_{\widetilde{N}_1 \geq \widetilde{N}_1^s}^{N_1^t}\left(f_{\widetilde{N}_1;N_1^t}\left(\widetilde{N}_1; N_1^t\right)\binom{\widetilde{N}_1}{\widetilde{N}_1^s} q^{\widetilde{N}_1^s}\left(1-q\right)^{\widetilde{N}_1 - \widetilde{N}_1^s}\right), \text{ and} \tag{35}$$

$$g_1^S\left(\widetilde{N}_1^s \middle| N_1^t\right) \quad : \quad = \sum_{\widetilde{N}_1 \geq \widetilde{N}_1^s}^{N_1^t}\left(\frac{f_{\widetilde{N}_1;N_1^t}\left(\widetilde{N}_1; N_1^t\right)\binom{\widetilde{N}_1 - 1}{\widetilde{N}_1^s - 1} q^{\widetilde{N}_1^s - 1}\left(1-q\right)^{\widetilde{N}_1 - \widetilde{N}_1^s}}{\sum_{\widetilde{N}_1' \geq 1}^{N_1^t} f_{\widetilde{N}_1;N_1^t}\left(\widetilde{N}_1'; N_1^t\right)}\right), \tag{36}$$

where $f_{\widetilde{N}_1;N_1^t}$ gives the probability mass density function of $\widetilde{N}_1 = N_1^t - \min\{\widetilde{D}, N_1^t\}$, derived from f as:

$$f_{\widetilde{N}_1;N_1^t}\left(\widetilde{N}_1; N_1^t\right) = \begin{cases} f\left(N_1^t - \widetilde{N}_1\right) & \text{if } \widetilde{N}_1 > 0 \\ \sum_{\widetilde{D} \geq N_1^t}^{\infty} f\left(\widetilde{D}\right) & \text{if } \widetilde{N}_1 = 0 \end{cases}. \tag{37}$$

Each term of (35) and (36) give the likelihood that exactly $\widetilde{N}_1$ asset holders are unable to sell to high demanders, conditioning on $\widetilde{N}_1^s$ asset holders having entered the interdealer market. The expression for g_1^S in particular allows for sellers to update their priors over $\widetilde{N}_1^s$ by conditioning on their own inability to sell downstream. Under the segmented demand

[27] For sellers, this means conditioning on realizing a demand and capacity shock but prior to learning $\widetilde{N}_0^{ns}$ and $\widetilde{N}_1^s > 1$. For buyers, this means not realizing a capacity shock but prior to learning $\widetilde{N}_0^{ns} > 1$ and $\widetilde{N}_1^s$.

model, these distributions will reduce to simple binary distributions:

$$g_1^B\left(\widetilde{N}_1^s\left|N_1^t\right.\right) \;:=\; \binom{N_1^t}{\widetilde{N}_1^s}\left(q\left(1-p\right)\right)^{\widetilde{N}_1^s}\left(1-q\left(1-p\right)\right)^{N_1^t-\widetilde{N}_1^s}, \text{ and} \tag{38}$$

$$g_1^S\left(\widetilde{N}_1^s\left|N_1^t\right.\right) \;:=\; \binom{N_1^t-1}{\widetilde{N}_1^s-1}\left(q\left(1-p\right)\right)^{\widetilde{N}_1^s-1}\left(1-q\left(1-p\right)\right)^{N_1^t-\widetilde{N}_1^s}. \tag{39}$$

And $\widetilde{N}_0^{ns}$, under both demand models, follows:

$$h_0^S\left(\widetilde{N}_0^{ns}\left|N_1^t\right.\right) \;:=\; \binom{N_0^t}{\widetilde{N}_0^{ns}}\left(1-q\right)^{\widetilde{N}_0^{ns}}q^{N_0^t-\widetilde{N}_0^{ns}}, \text{ and} \tag{40}$$

$$h_0^B\left(\widetilde{N}_0^{ns}\left|N_1^t\right.\right) \;:=\; \binom{N_0^t-1}{\widetilde{N}_0^{ns}-1}\left(1-q\right)^{\widetilde{N}_0^{ns}-1}q^{N_0^t-\widetilde{N}_0^{ns}} \tag{41}$$

D.2 Proof of Proposition 4

First, given any strategy θ an identical argument given in the proof of Proposition 1 establishes a unique set of stationary values $(V_{0,N_1^t}, V_{1,N_1^t})$. Thus, it remains to derive a symmetric Nash equilibrium strategy θ given values $(V_{0,N_1^t}, V_{1,N_1^t})$, which take θ as given.

We assume a downwards sloping value gap, or that $q \in [0, \overline{q}^1]$ from Proposition 5 (derived below). From the proof of Proposition 5 this implies that an increase in N_1^t (to N_1^t+1) implies a first order shift up (right) in the distribution of residual asset holders $\hat{N}_1^t$, denoted below $\check{N}_1^t$.[28]

For case 1, when $N_1^t > \overline{N}_1$ shocked asset holders are indifferent between liquidating and trading at price $\underline{v}$ in the interdealer market if they can. We assume they trade with found vacants to maintain consistency in the distribution over residual assets $\check{N}_1^t$. Again, denote the number of asset holders able to carry inventories $\check{N}_1^t$. $\check{N}_1^t$ also gives the number of assets holders are left to mix on funneling and on holding their assets as carried inventory. Note, when $N_1^t \le \overline{N}_1$, we have $\check{N}_1^t = \hat{N}_1^t$. Then, use $\hat{N}_1^t$ to denote the number of assets left after asset funneling that are carried into period $t+1$: market clearing condition 5a-5b for period $t+1$ gives subsequent market inventory N_1^{t+1}. Then, the distribution of $\hat{N}_1^t$ conditioning on (1) residual assets (prior to funneling) $\check{N}_1^t$, (2) the original period t market inventory N_1^t, and (3) the strategy of residual asset holders θ is given by:

$$\Pr\left(\hat{N}_1^t\left|\check{N}_1^t, N_1^t, \theta\right.\right) = \binom{\hat{N}_1^t}{\check{N}_1^t}\left(1-\theta\left(N_1^t\right)\right)^{\hat{N}_1^t}\theta\left(N_1^t\right)^{\check{N}_1^t-\hat{N}_1^t}. \tag{42}$$

Using the distribution of $\check{N}_1^t$ conditioning on N_1^t given by $\Pr(\hat{N}_1^t|N_1^t) = \Pr(\hat{N}_1^t|\check{N}_1^t, N_1^t, 0)$ (no funneling), we can derive an updated distribution for carried inventory $\hat{N}_1^t$:

$$\Pr\left(\hat{N}_1^t\left|N_1^t, \theta\right.\right) = \sum_{\check{N}_1^t}\Pr\left(\hat{N}_1^t\left|\check{N}_1^t, N_1^t, \theta\right.\right)\Pr\left(\check{N}_1^t\left|N_1^t\right.\right). \tag{43}$$

[28] We take this relabeling to leave $\hat{N}_1^t$ to denote the post-funneling number of residual asset holders. This avoids a restatement of upstream market clearing condition (5a)-(5b).

$\Pr(\hat{N}_1^t|N_1^t,\theta)$ will exhibit first order shifts left with greater $\theta(N_1^t)$ (for each N_1^t): the independent probability of each asset holder in $\check{N}_1^t$ liquidating as opposed to carrying inventory is given by $\theta(N_1^t)$. Then, with $G(N_1^t)$ a weakly decreasing function of N_1^t we set $\theta(N_1^t)$ to solve:

$$\begin{aligned} \underline{v} &= \delta\left(E_t\left[G\left(N_1^{t+1}\right)|\theta\right]\right) \\ &= \delta\sum_{\hat{N}_1^t}\Pr\left(\hat{N}_1^t\left|N_1^t,\theta\right.\right)G\left(N_1^{t+1}\left(\hat{N}_1^t\right)\right) \end{aligned} \tag{44}$$

so that the discounted gap when N_1^t is above $\overline{N}_1$ is always equal to $\underline{v}$, maintaining indifference over funneling and carrying inventory. Thus, we've established the given strategy to be optimal.

To construct G and show that such a strategy θ can be found, first take $G^{(0)}(N_1^t)$ to give the value gap when mixing on funneling is restricted for all N_1^t ($\theta(N_1^t)=0$). Define $\overline{N}_1^{(0)}$ to give the market inventory so that $\underline{v}\leq\delta E_t\left[G^{(0)}(N_1^{t+1})\right]$ for each $N_1^t\leq\overline{N}_1^{(0)}$ while $\underline{v}>\delta E_t\left[G^{(0)}(N_1^{t+1})\right]$ for each $N_1^t>\overline{N}_1^{(0)}$. If there is no such $\overline{N}_1^{(0)}$ then we are in case 2. Also assume that at least one value for the value gap lies above $\underline{v}$, else $N_1^t>\overline{N}_1^{(0)}$ with probability one for all t giving a case in which interdealer trade is never rational. For each $N_1^t>\overline{N}_1^{(0)}$ define $\theta^{(0)}(N_1^t)$ to give the probability that residual asset holders $\check{N}_1^t$ mix on funneling. The induced distribution $\Pr(\hat{N}_1^t|N_1^t,\theta^{(0)})$ will exhibit a first order shift left as $\theta^{(0)}(N_1^t)$ is increased, yielding a continuously decreasing expectation $\delta E_t\left[G^{(0)}(N_1^{t+1})\right]$ by continuity of (42) with respect to $\theta(N_1^t)$.

We have $\delta E_t\left[G^{(0)}(N_1^{t+1})\right]<\underline{v}$ for $\theta^{(0)}(N_1^t)=0$ by assumption. Ans for $\theta^{(0)}(N_1^t)=1$, $\delta E_t\left[G^{(0)}(N_1^{t+1})\right]=\delta E_t\left[G(N_1^{t+1})\middle|\hat{N}_1^t=0\right]>\underline{v}$ (with all asset holders funneling, leaving $\hat{N}_1^t=0$). By the Intermediate-value Theorem there is some $\theta^{(0)}(N_1^t)$ that gives $\delta E_t\left[G^{(0)}(N_1^{t+1})\right]=\underline{v}$.

Now take $G^{(1)}(N_1^t)$ given with (19) where the discounted value gap is given by the function $\max\left\{0,\delta E_t\left[G^{(0)}(N_1^{t+1})\right]\right\}$. By assumption $q\in[0,\overline{q}^1]$ so that

$$\left(\left(1-p\right)\left(1-q\right)+\left(\left(1-p\right)q\mu-\left(1-q\right)\lambda\right)\right)\geq 0 \tag{45}$$

and (19) is increasing in the discounted expected value gap, and thus $G^{(1)}(N_1^t)\geq G^{(0)}(N_1^t)$ for each N_1^t. Similarly by monotonicity of $\max\{0,\delta E_t\left[G^{(0)}(N_1^{t+1})\right]\}$, $G^{(1)}(N_1^t)$ is also monotonic (ie. decreasing in N_1^t). Accordingly, if the funneling strategy is given by $\theta^{(0)}$ then $\delta E_t\left[G^{(1)}(N_1^{t+1})\right]\geq 0$ for each N_1^t with the discounted value gap strictly above zero for some $N_1^t>\overline{N}_1^{(0)}$. We construct $\theta^{(1)}$ from $\theta^{(0)}$ by decreasing each $\theta^{(0)}(N_1^t)$ for $N_1^t>\overline{N}_1^{(0)}$ until either $\theta^{(1)}(N_1^t)=0$ or $\delta E_t\left[G^{(1)}(N_1^{t+1})\right]=0$ (whichever comes first). Formally,

$$\theta^{(1)}(N_1^t)=\max\left\{0,\left\{\theta\left|\delta E_t\left[G^{(1)}(N_1^{t+1})\right]=0\right.\right\}\right\}. \tag{46}$$

By monotonicity of $G^{(1)}(N_1)$ and again because the induced distribution $\Pr(\hat{N}_1^t|N_1^t,\theta^{(1)})$ exhibits a first order shift left with $\theta^{(1)}(N_1^t)$, if $\theta^{(1)}(N_1^t)=0$ then $\theta^{(1)}(N_1^{t\prime})=0$ for each $N_1^{t\prime}>N_1^t$.

Finally, define $\overline{N}_1^{(1)}$ to give the market inventory so that $\underline{v} \leq \delta E_t \left[G^{(1)}(N_1^{t+1}) \right]$ for each $N_1^t \leq \overline{N}_1^{(1)}$ while $\underline{v} > \delta E_t \left[G^{(1)}(N_1^{t+1}) \right]$ for each $N_1^t > \overline{N}_1^{(1)}$, and given that mixing on funneling is given by $\theta^{(1)}$. In total, $G^{(1)}(N_1^t) \geq G^{(0)}(N_1^t)$ for each N_1^t, $\overline{N}_1^{(1)} \geq \overline{N}_1^{(0)}$, and $\theta^{(1)}(N_1^t) \leq \theta^{(0)}(N_1^t)$ for each $N_1^t > \overline{N}_1^{(1)}$ by construction.

Recursively construct $G^{(n)}$ from $G^{(n-1)}$, $\theta^{(n)}$ from $\theta^{(n-1)}$, and $\overline{N}_1^{(n)}$ from $\overline{N}_1^{(n-1)}$ in a similar fashion. We obtain a weakly increasing sequence $G^{(n)}(N_1^t)$ that is weakly increasing for each N_1^t and bounded above by v. Similarly, we obtain monotone sequences $\theta^{(n)}$ and $\overline{N}_1^{(n)}$ bounded by 0 (below) and N (above), respectively. By the monotone convergence theorem, $G^{(n)} \to G$ and $\theta^{(n)} \to \theta$ (pointwise) with $\overline{N}_1^{(n)} \to \overline{N}_1$, with $\theta(N_1^t)$ optimal for each $N_1^t \in \{0, \ldots, N\}$ given G.

Finally, $\theta(N_1^t)$ (as well as each $\theta^{(n)}$) is increasing in N_1^t by the fact that G and $\delta E_t \left[G(N_1^{t+1}) \right]$ are decreasing functions of N_1^t and $\Pr(\check{N}_1^t | N_1^t)$ exhibits a first order stochastic shift up (right) with N_1^t: a greater mixing on funneling is required to obtain $E_t \left[G(N_1^{t+1}) \right] = \underline{v}$ at larger N_1^t.

For case 2, the gains to carrying assets as inventory between periods is maintained for all $N_1^t \in \{0, \ldots, N\}$, and thus asset funneling is never rational. $\square$

D.3 Proof of Proposition 5

With the market clearing condition (5a)-(5b) pinning $N_1^t(\hat{N}_1^{t-1})$, (20c) is implied by the following restriction on Φ relative to $V_{1,N_1^t}^t - V_{0,N_1^t}^t$:

$$\Phi\left(N_1^t\left(\hat{N}_1^{t-1}\right)+n;\hat{N}_1^{t-1}\right) > G^t\left(N_1^t\left(\hat{N}_1^{t-1}\right)+n\right). \tag{47}$$

for each integer $n \geq 2$. That is, the inverse supply lies strictly above the value gap moving beyond two assets greater than the original market clearing inventory $N_1^t(\hat{N}_1^{t-1})$. We see then that the condition (20c) can be recast as a lower bound on the increase (slope) in $\Phi(N_1^t;\hat{N}_1^{t-1})$ as N_1^t increases. With the value gap given in equilibrium, such conditions can only be verified case-by-case. None the less, we will maintain this as a weak assumption on supply.[29]

Taking condition (20c), the proof proceeds as follows. For the case of asset funneling, the value gap reduces to

$$G\left(N_1^t\right) = \underline{v} + p\left(v-\underline{v}\right) \tag{48}$$

so that:

$$\nabla_{N_1^t} G\left(N_1^t\right) = \left(v-\underline{v}\right)\nabla_{N_1^t} p \leq 0, \tag{49}$$

with $\nabla_{N_1^t} p \leq 0$ by assumption. And that $\nabla_{N_1^t}(V_{1,N_1^t} - V_{0,N_1^t}) \leq 0$ around $N_1^t = \overline{N}_1$ will follow from continuity and monotonicity of the max function.

Now assume that carrying inventories is optimal ($\delta E_t\left[V_{1,N_1^{t+1}}\right] - \delta E_t\left[V_{0,N_1^{t+1}}\right] - \underline{v} \geq 0$). The proof heavily utilizes the following form for the value gap:

$$G\left(N_1^t\right) = \left(\begin{array}{c} \underline{v} + \left(v-\underline{v}\right)p \\ +\left(1-p\right)\left(1-q\right)\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right) \\ +\left(\left(1-p\right)q\mu - \left(1-q\right)\lambda\right)\left(\delta E_t\left[G(N_1^{t+1})\right]-\underline{v}\right) \end{array}\right). \tag{50}$$

Using (50) we can apply $\nabla_{N_1^t}$ and define the following:

1. $$\nabla_{N_1^t} G\left(N_1^t\right) := G\left(N_1^t+1\right) - G\left(N_1^t\right), \text{ and} \tag{51}$$

2. writing:
$$E_t\left[G\left(N_1^{t+1}\right)\right] = \sum_{\widehat{N}_1^t}\left(G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\right)\Pr\left(\widehat{N}_1^t\,|N_1^t\right), \tag{52}$$

   also define:

$$\nabla_{N_1^t}\left(E_t\left[G\left(N_1^{t+1}\right)\right]\right) := \left(\begin{array}{c} \sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\Pr\left(\widehat{N}_1^t\,|N_1^t\right) \\ -\sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\Pr\left(\widehat{N}_1^t\,|N_1^t+1\right) \end{array}\right). \tag{53}$$

[29]Note that the upward pressure on $\Phi(N_1^t;\hat{N}_1^{t-1})$ imposed by residual inventory $\hat{N}_1^{t-1}$ (motivating Assumption 1) works in the direction of this condition. Still, with full generality maintained in $p(N_1^t)$, we must rule out more extreme demand processes violating condition (20c).

The proof proceeds by first establishing the following necessary property and subsequent lemmas.

Definition 5. $\Pr\left(\widehat{N}_1^t \,|N_1^t\right)$ *exhibits Marginal First Order Stochastic Dominance (MFOSD) in* N_1^t *if increasing* N_1^t *by one gives a first order stochastic dominance shift up in the distribution of* $\widehat{N}_1^t$*, in such a way that there are* $\phi\left(\widehat{N}_1^t\right) \leq \Pr\left(\widehat{N}_1^t \,|N_1^t\right)$ *giving the probability shifted from* $\Pr\left(\widehat{N}_1^t \,|N_1^t\right)$ *to* $\Pr\left(\widehat{N}_1^t + 1 \,|N_1^t + 1\right)$ *for each* $\widehat{N}_1^t \in \{1, \ldots, N-1\}$.

When $\Pr(\widehat{N}_1^t|N_1^t)$ satisfies MFOSD in N_1^t, we can write:

$$\Pr\left(\widehat{N}_1^t \,|N_1^t + 1\right) = \Pr\left(\widehat{N}_1^t \,|N_1^t\right) - \phi\left(\widehat{N}_1^t\right) + \phi\left(\widehat{N}_1^t - 1\right). \tag{54}$$

The above property is stronger than first order stochastic dominance, in that each $\phi(\widehat{N}_1^t)$ is required to be bounded above by $\Pr(\widehat{N}_1^t|N_1^t)$. This requires that the shift in the process $\Pr(\widehat{N}_1^t|N_1^t)$ be sufficiently subtle. This property may be violated with a highly active interdealer market, particularly when q and N_1^t are high. This is true because extreme market inventories begin to deprive the interdealer market of the liquidity needed to absorb assets from shocked asset holders, and effectively suppress expectations over $\widehat{N}_1^t$. However, if q is below the following bound, Lemma 5 below shows that the MFOSD property is satisfied.

Definition 6.

$$\overline{q}^1 := \min\left\{ q|\Pr\left(\chi_B\right) = 1 - p, N_1^t = 0, \ldots, N\right\}. \tag{55}$$

where $\Pr(\chi_B)$, the probability of an unshocked vacant buying in the interdealer market, is evaluated at the given q in the above definition. Note that this definition will give $\overline{q}^1$ that solves $\Pr(\chi_B) = 1 - p(N)$ by property 3 in (17a)-(17f).

Lemma 5. *When* $q \in [0, \overline{q}^1]$ *and waiting is optimal,* $\Pr\left(\widehat{N}_1^t \,|N_1^t\right)$ *satisfies the MFOSD property.*

PROOF: As the precise form of $\Pr(\widehat{N}_1^t|N_1^t)$ depends on the particular demand process yielding $p(N_1^t)$, which the model avoids for the sake of generality, I provide a proof outline suitable to apply to any given demand process. First, the following inductive approximation will hold:

$$\Pr\left(\widehat{N}_1^t \,|N_1^t + 1\right) \approx \begin{pmatrix} \Pr\left(\widehat{N}_1^t - 1 \,|N_1^t, N-1\right)(1-p)\left((1-q) + q\Pr\left(\chi_S\right)\right) \\ +\Pr\left(\widehat{N}_1^t \,|N_1^t, N-1\right)\left(p + (1-p)\,q\left(1 - \Pr\left(\chi_S\right)\right)\right) \end{pmatrix} \tag{56}$$

where the probabilities on the right condition on N_1^t and $N - 1 - N_1^t$ vacants. That is, the likelihood that carried market inventory $\widehat{N}_1^t$ is realized is approximately equal to the sum of probabilities of all events in a market of size $N-1$ such that inclusion of an additional asset holder yields carried inventory $\widehat{N}_1^t$. This expression is only approximate precisely because it does not correct for correlation between the demand uncertainty of the additional asset holder (given by p) and in the random variables $\widehat{N}_1^t - 1$ and $\widehat{N}_1^t$. Similarly, we can approximate:

$$\Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right) \approx \left(\begin{array}{c} \Pr\left(\widehat{N}_1^t - 1 \left| N_1^t, N-1\right.\right)(1-q)\Pr\left(\chi_B\right) \\ +\Pr\left(\widehat{N}_1^t \left| N_1^t, N-1\right.\right)\left(q+(1-q)\left(1-\Pr\left(\chi_B\right)\right)\right) \end{array}\right), \tag{57}$$

where the additional dealer enters as a vacant. This yields:

$$\Pr\left(\widehat{N}_1^t \left| N_1^t+1\right.\right) - \Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right) \approx$$

$$\left(\begin{array}{c} \Pr\left(\widehat{N}_1^t - 1 \left| N_1^t, N-1\right.\right)\left(\begin{array}{c} (1-q)\left((1-p)-\Pr\left(\chi_B\right)\right) \\ +q(1-p)\Pr\left(\chi_S\right) \end{array}\right) \\ -\Pr\left(\widehat{N}_1^t \left| N_1^t, N-1\right.\right)\left(\begin{array}{c} (1-q)\left((1-p)-\Pr\left(\chi_B\right)\right) \\ +q(1-p)\Pr\left(\chi_S\right) \end{array}\right) \end{array}\right). \tag{58}$$

Then, with q small, $\Pr(\widehat{N}_1^t - 1 | N_1^t, N-1)$ and $\Pr(\widehat{N}_1^t | N_1^t, N-1)$ will approximate to $\Pr(\widehat{N}_1^t - 1 | N_1^t, N)$ and $\Pr(\widehat{N}_1^t | N_1^t, N)$, respectively. That is, the inclusion of an additional vacant will shift up the distribution of carried inventories only slightly.[30] The above together with the fact that $\Pr(\chi_B) \leq (1-p)$ by our choice of q, yield the approximate expression:

$$\phi\left(\widehat{N}_1^t\right) \approx \Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right)\left(\begin{array}{c} (1-q)\left((1-p)-\Pr\left(\chi_B\right)\right) \\ +q(1-p)\Pr\left(\chi_S\right) \end{array}\right) < \Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right). \tag{59}$$

□

Lemma 6. *If the process on $\widehat{N}_1^t$ given by $\Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right)$ satisfies the MFOSD property, then:*

$$\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \leq \max_{N_1} \nabla_{N_1} G\left(N_1\right). \tag{60}$$

.

PROOF: Assume that increasing N_1^t by one gives a first order stochastic dominance shift up in the distribution of $\widehat{N}_1^t$, and let $\phi(\widehat{N}_1^t) \leq \Pr(\widehat{N}_1^t | N_1^t)$ give the probability shifted from $\Pr(\widehat{N}_1^t | N_1^t)$ to $\Pr(\widehat{N}_1^t + 1 | N_1^t + 1)$ for each $\widehat{N}_1^t \in \{1, \ldots, N-1\}$, so that

$$\Pr\left(\widehat{N}_1^t \left| N_1^t+1\right.\right) = \Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right) - \phi\left(\widehat{N}_1^t\right) + \phi\left(\widehat{N}_1^t - 1\right), \tag{61}$$

and where $\Delta p(0) = \Delta p(N) = 0$. We then have:

$$\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] = \left(\begin{array}{c} \sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\Pr\left(\widehat{N}_1^t \left| N_1^t+1\right.\right) \\ -\sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\Pr\left(\widehat{N}_1^t \left| N_1^t\right.\right) \end{array}\right)$$

[30] With N sufficiently large, we should expect that any error in this approximation will be small enough that we can adjust our expression for $\phi(\widehat{N}_1^t)$ (below) appropriately.

$$
\begin{aligned}
&= \left(\begin{array}{c} \sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\left(\Pr\left(\widehat{N}_1^t\,|N_1^t\right) - \phi\left(\widehat{N}_1^t\right) + \phi\left(\widehat{N}_1^t - 1\right)\right) \\ -\sum_{\widehat{N}_1^t} G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\Pr\left(\widehat{N}_1^t\,|N_1^t\right) \end{array} \right) \\
&= \sum_{\widehat{N}_1^t}\left(-\phi\left(\widehat{N}_1^t\right) + \phi\left(\widehat{N}_1^t - 1\right)\right) G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right) \\
&= \sum_{\widehat{N}_1^t < N} \phi\left(\widehat{N}_1^t\right)\left(G\left(N_1^{t+1}\left(\widehat{N}_1^t + 1\right)\right) - G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\right) \\
&= \sum_{\widehat{N}_1^t < N} \chi\left(\widehat{N}_1^t\right)\phi\left(\widehat{N}_1^t\right)\left(G\left(N_1^{t+1}\left(\widehat{N}_1^t\right) + 1\right) - G\left(N_1^{t+1}\left(\widehat{N}_1^t\right)\right)\right) \\
&= \sum_{\widehat{N}_1^t < N} \chi\left(\widehat{N}_1^t\right)\phi\left(\widehat{N}_1^t\right)\nabla_{N_1^t\left(\widehat{N}_1^t\right)}\left(V_{1,N_1^t\left(\widehat{N}_1^t\right)} - V_{0,N_1^t\left(\widehat{N}_1^t\right)}\right), \qquad (62)
\end{aligned}
$$

where:

$$
\chi\left(\widehat{N}_1^t\right) = \begin{cases} 1 & \text{if } N_1^{t+1}\left(\widehat{N}_1^t + 1\right) = N_1^{t+1}\left(\widehat{N}_1^t\right) + 1 \\ 0 & \text{if } N_1^{t+1}\left(\widehat{N}_1^t + 1\right) = N_1^{t+1}\left(\widehat{N}_1^t\right) \end{cases} \qquad (63)
$$

for each $\widehat{N}_1^t < N$ (which is well defined on $\{0, \ldots, N-1\}$ by condition (20c)). Then with $\chi\left(\widehat{N}_1^t\right)\phi\left(\widehat{N}_1^t\right) \in [0,1]$ (by the MFOSD) we have that $\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]$ is a convex combination of $\nabla_{N_1^t}\left(V_{1,N_1^t} - V_{0,N_1^t}\right)$, giving the result. □

With the above bound established, the proof of part 1 proceeds as follows. First, derive the difference function of the value gap:

$$
\nabla_{N_1^t}\left(V_{1,N_1^t}^t - V_{0,N_1^t}^t\right) =
$$

$$
\left(\begin{array}{c} \delta\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]\left((1-p)(1-q) + (1-p)\,q\mu - (1-q)\,\lambda\right) \\ + \left(\begin{array}{c} \left(\nabla_{N_1^t} p\right)\left(v - \underline{v} - \left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right)\left((1-q) + q\mu\right)\right) \\ + \left(q\,(1-p)\,\nabla_{N_1^t}\mu - (1-q)\,\nabla_{N_1^t}\lambda\right)\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right) \end{array} \right) \end{array} \right). \qquad (64)
$$

With $\delta E_t\left[G(N_1^{t+1})\right] - \underline{v} \geq 0$, $\nabla_{N_1^t}\mu \leq 0 \leq \nabla_{N_1^t}\lambda$, and $\nabla_{N_1^t} p \leq 0$ the bottom term is non-positive. Then, with $\lambda \leq \Pr(\chi_B)$ by expression (22a), and our choice of $q \in [0, \overline{q}^1]$, we have:

$$
\left((1-p)(1-q) + (1-p)\,q\mu - (1-q)\,\lambda\right) \geq 0. \qquad (65)
$$

so the right hand side of (64) is increasing in $\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]$. By Lemmas 5 and 6:

$$
\nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \leq \max_{N_1} \nabla_{N_1}\left(V_{1,N_1} - V_{0,N_1}\right). \qquad (66)
$$

Let N_1^* give any solution the the right hand problem. We can now establish an upper bound to $\nabla_{N_1^t}(V_{1,N_1^t}^t - V_{0,N_1^t}^t)$ by forcing the process generating N_1^{t+1} from $N_1^t = N_1^*$ (defined by

$\Pr(\widehat{N}_1^t\, N_1^t))$ to place probability one on N_1^* (ie. assume N_1^* is a sink). This is precisely because the right hand side of (64) is increasing in $\nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]$ for $q \in [0, \overline{q}^1]$. This gives:

$$\nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]\Big|_{N_1^t = N_1^*} \leq \nabla_{N_1^t} \left(V_{1,N_1^*} - V_{0,N_1^*}\right), \tag{67}$$

giving:

$$\begin{aligned} &\nabla_{N_1^t} \left(V_{1,N_1^*} - V_{0,N_1^*}\right) \leq \\ &\frac{\begin{pmatrix} \left(\nabla_{N_1^t} p\right) \left(v - \underline{v} - \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) \left((1-q) + q\mu\right)\right) \\ + \left(q\,(1-p)\, \nabla_{N_1^t} \mu - (1-q)\, \nabla_{N_1^t} \lambda\right) \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) \end{pmatrix}}{1 - \delta\left((1-p)\,(1-q) + (1-p)\, q\mu - (1-q)\,\lambda\right)} \\ &< \; 0. \end{aligned} \tag{68}$$

establishing part 1.

For part 2., I derive:

$$\frac{\partial}{\partial q} \nabla_{N_1^t} \left(V_{1,N_1^t} - V_{0,N_1^t}\right) =$$

$$\begin{pmatrix} \delta \frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \left((1-p)\,(1-q) + (1-p)\, q\mu - (1-q)\,\lambda\right) \\ + \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \begin{pmatrix} (1-p)\,(\mu - 1) + \lambda \\ + (1-p)\, q \frac{\partial}{\partial q} \mu - (1-q)\, \frac{\partial}{\partial q} \lambda \end{pmatrix} \\ + \begin{pmatrix} \nabla_{N_1^t} p \begin{pmatrix} -\delta \frac{\partial}{\partial q} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \left((1-q) + q\mu\right) \\ - \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) \left((\mu - 1) + q \frac{\partial}{\partial q} \mu\right) \end{pmatrix} \\ + \delta \frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \left(q\,(1-p)\, \nabla_{N_1^t} \mu - (1-q)\, \nabla_{N_1^t} \lambda\right) \\ + \begin{pmatrix} (1-p)\, \nabla_{N_1^t} \mu + \nabla_{N_1^t} \lambda \\ + \left(q\,(1-p)\, \frac{\partial}{\partial q} \nabla_{N_1^t} \mu - (1-q)\, \frac{\partial}{\partial q} \nabla_{N_1^t} \lambda\right) \end{pmatrix} \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) \end{pmatrix} \end{pmatrix}. \tag{69}$$

At $q = 0$, we have $\lambda = \nabla_{N_1^t} \lambda = 0$, so that this expression reduces to:

$$\frac{\partial}{\partial q} \nabla_{N_1^t} \left(V_{1,N_1^t} - V_{0,N_1^t}\right) =$$

$$\begin{pmatrix} \delta\,(1-p)\, \frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \\ + \begin{pmatrix} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \left((1-p)\,(\mu - 1) - \frac{\partial}{\partial q} \lambda\right) \\ - \left(\nabla_{N_1^t} p\right) \left(\delta \frac{\partial}{\partial q} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] + \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) (\mu - 1)\right) \\ - \left(\delta E_t \left[G(N_1^{t+1})\right] - \underline{v}\right) \left(-(1-p)\, \nabla_{N_1^t} \mu + \frac{\partial}{\partial q} \nabla_{N_1^t} \lambda\right) \end{pmatrix} \end{pmatrix} \tag{70}$$

Now, when $q \in [0, \overline{q}^0]$ and an increase in N_1^t gives a first order stochastic shift up in the distribution of $\widehat{N}_1^t$ (ie. the MFOSD property is satisfied), $\nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] < 0$ because $V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}$ is weakly decreasing by part 1. And with $\delta E_t \left[G(N_1^{t+1})\right] - \underline{v} > 0$,

$\nabla_{N_1^t}\mu \leq 0$, and $\frac{\partial}{\partial q}\lambda \geq \frac{\partial}{\partial q}\nabla_{N_1^t}\lambda \geq 0$ by 5 and 6 of the conditions in (17a)-(17f), the second term of the bottom triplet in (70) will be non-negative and the third term will be non-positive. The relative sizes of these terms will depend on the functions p, μ and λ, and in particular their dependence on N_1^t and sensitivity to q. In what follows, I'll denote the bottom triplet:

$$\Gamma := \left(\begin{array}{c} \nabla_{N_1^t} E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]\left((1-p)(\mu-1) - \frac{\partial}{\partial q}\lambda\right) \\ -\left(\nabla_{N_1^t}p\right)\left(\delta\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] + \left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right)(\mu-1)\right) \\ -\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right)\left(-(1-p)\nabla_{N_1^t}\mu + \frac{\partial}{\partial q}\nabla_{N_1^t}\lambda\right) \end{array} \right). \tag{71}$$

I next expand $\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]$ using the product rule:

$$\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] = \frac{\partial}{\partial q}\sum_{\widehat{N}_1^t}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\Pr\left(\widehat{N}_1^t\,|N_1^t\right)$$

$$\begin{aligned} &= \left(\begin{array}{c} \sum_{\widehat{N}_1^t}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\frac{\partial}{\partial q}\Pr\left(\widehat{N}_1^t\,|N_1^t\right) \\ +\sum_{\widehat{N}_1^t}\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\Pr\left(\widehat{N}_1^t\,|N_1^t\right) \end{array} \right) \\ &= \left(\begin{array}{c} \sum_{\widehat{N}_1^t}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\frac{\partial}{\partial q}\Pr\left(\widehat{N}_1^t\,|N_1^t\right) \\ E_t\left[\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\right] \end{array} \right), \end{aligned} \tag{72}$$

and I derive:

$$\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) =$$

$$\left(\begin{array}{c} \left(((1-p)(1-q) + (1-p)q\mu - (1-q)\lambda)\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]\right) \\ +\left(\begin{array}{c} (1-p)(\mu-1)+\lambda \\ (1-p)q\frac{\partial}{\partial q}\mu - (1-q)\frac{\partial}{\partial q}\lambda \end{array} \right)\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right) \end{array} \right), \tag{73}$$

which reduces at $q = 0$ to:

$$\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) = \left(\begin{array}{c} \delta(1-p)\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] \\ +\left((1-p)(\mu-1) - \frac{\partial}{\partial q}\lambda\right)\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right) \end{array} \right). \tag{74}$$

With $0 \leq \frac{\partial}{\partial q}\lambda$ and $\mu \leq 1$ the bottom term is non-positive. The following evaluates the sign of $\frac{\partial}{\partial q}(V_{1,N_1^t} - V_{0,N_1^t})$ and $\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right]$ for each case.

For part 2, I take to case $\nabla_{N_1^t}p = 0$ for each N_1^t (the segmented market setup gives a special case of this), causing the second term of Γ to drop out. And with $q = 0$, the gap $V_{1,N_1} - V_{0,N_1}$ is independent of market inventory N_1 at each time, causing the first term of

Γ to drop out. This yields:

$$\Gamma = -\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right)\left(-(1-p)\nabla_{N_1^t}\mu + \frac{\partial}{\partial q}\nabla_{N_1^t}\lambda\right) \leq 0. \tag{75}$$

With $V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}$ in the top term of (72) a constant for each $\widehat{N}_1^t$ (because $\nabla_{N_1^t}p = 0$), and because $\sum_{\widehat{N}_1^t}\Pr(\widehat{N}_1^t|N_1^t) = 1$, it must be that $\sum_{\widehat{N}_1^t}\frac{\partial}{\partial q}\Pr(\widehat{N}_1^t|N_1^{t)}$ equals zero, so that the top term of (72) drops out[31]. This then yields an updated (74):

$$\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) = \begin{pmatrix} \delta(1-p)E_t\left[\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\right] \\ +\left((1-p)(\mu-1) - \frac{\partial}{\partial q}\lambda\right)\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right) \end{pmatrix}, \tag{76}$$

yielding an upper bound:

$$\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) \leq \frac{\left((1-p)(\mu-1) - \frac{\partial}{\partial q}\lambda\right)\left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right)}{1-\delta(1-p)}, \tag{77}$$

which is unambiguously negative. Taking expectations over N_1^t will preserve this negativity:

$$E_t\left[\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right)\right] = \frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] < 0, \tag{78}$$

with the equality holding again at $q = 0$. Thus, we've shown $\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) \leq 0$ with $\Omega \leq 0$ for $\nabla_{N_1^t}p = 0$. With $\mu < 1$ and $\nabla_{N_1^t}\mu < 0$ (by Assumption 8) the above inequalities bounding $\frac{\partial}{\partial q}\left(V_{1,N_1^t} - V_{0,N_1^t}\right)$ and Ω become strict. And by continuity of all relevant terms in $q \in [0,1]$ and $\nabla_{N_1^t}p \in [0,1]^N$, this will also hold for some neighborhood of $\left(q, \nabla_{N_1^t}p\right) = (0,0) \in [0,1] \times [0,1]^N$. Take $\overline{q}' = B'$ to be the diameter of largest ball included in this neighborhood.

I next show that given (70) the sign of $\frac{\partial}{\partial q}\nabla_{N_1^t}(V_{1,N_1^t} - V_{0,N_1^t})$ is determined by the sign of Γ when $\nabla_{N_1^t}p$ is near the origin. We can write

$$\begin{aligned} \frac{\partial}{\partial q}\nabla_{N_1^t}\left(V_{1,N_1^t} - V_{0,N_1^t}\right) &= \nabla_{N_1^t}\frac{\partial}{dq}\left(V_{1,N_1^t} - V_{0,N_1^t}\right), && (79) \\ \frac{\partial}{\partial q}\nabla_{N_1^t}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right] &= \nabla_{N_1^t}\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}\right], && (80) \end{aligned}$$

[31] This step is crucial because an increase in q causes a stochastic shift down in the distribution of $\widehat{N}_1^t$; with a strictly downward sloping value gap this force would work against $\frac{\partial}{\partial q}(V_{1,N_1^t} - V_{0,N_1^t}) < 0$. When this effect is small (for small $\left|\nabla_{N_1^t}p\right|$ and q), this force is dominated by those derived from the interdealer marker, as shown above.

with $\nabla_{N_1^t}$ constituting a linear operator. Plugging in (72) into the bottom equation gives:

$$\frac{\partial}{\partial q}\nabla_{N_1^t}E_t\left[V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right]=$$

$$\nabla_{N_1^t}\left(\begin{array}{c}\sum_{\widehat{N}_1^t}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\frac{\partial}{\partial q}\Pr\left(\widehat{N}_1^t|N_1^t\right)\\+E_t\left[\frac{d}{dq}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\right]\end{array}\right)$$

$$=\left(\begin{array}{c}\nabla_{N_1^t}\sum_{\widehat{N}_1^t}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\frac{\partial}{\partial q}\Pr\left(\widehat{N}_1^t|N_1^t\right)\\+\nabla_{N_1^t}E_t\left[\frac{d}{dq}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\right]\end{array}\right). \tag{81}$$

As argued above, when $\nabla_{N_1^t}p=0$ the term $\sum_{\widehat{N}_1^t}(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}})\frac{\partial}{\partial q}\Pr(\widehat{N}_1^t|N_1^t)$ will equal zero[32], leaving:

$$\nabla_{N_1^t}\frac{\partial}{\partial q}E_t\left[V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right]=\nabla_{N_1^t}E_t\left[\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\right], \tag{82}$$

and updating (70) to:

$$\nabla_{N_1^t}\frac{\partial}{\partial q}\left(V_{1,N_1^t}-V_{0,N_1^t}\right)=$$

$$\left(\delta\left(1-p\right)\nabla_{N_1^t}E_t\left[\frac{\partial}{\partial q}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\right]+\Gamma\right). \tag{83}$$

With an increase in N_1^t to N_1^{t+1} yielding a first order shift up in $\Pr(\widehat{N}_1^t|N_1^t)$ when $q\in[0,\overline{q}^1]$, and because (i) $\nabla_{N_1^t}(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}})\leq 0$ by part 1 and
(ii) $\nabla_{N_1^t}E_t\left[V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right]\leq\sum_{\widehat{N}_1^t}\nabla_{N_1^t}(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}})\phi(\widehat{N}_1^t)$ by the proof of Lemma 6, we can write:

$$\nabla_{N_1^t}E_t\left[V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right]=\sum_{\widehat{N}_1^t}\nabla_{N_1^t}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\varphi\left(\widehat{N}_1^t|q\right) \tag{84}$$

for some set of probabilities $\{\varphi(\widehat{N}_1^t|q)\}_{\widehat{N}_1^t}$, with $\sum_{\widehat{N}_1^t}\varphi(\widehat{N}_1^t|q)\in[0,1]$. Then,

$$\frac{\partial}{\partial q}\nabla_{N_1^t}E_t\left[V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right]=$$

$$\sum_{\widehat{N}_1^t}\left(\begin{array}{c}\frac{\partial}{\partial q}\nabla_{N_1^t}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\varphi\left(\widehat{N}_1^t|q\right)\\+\nabla_{N_1^t}\left(V_{1,N_1^{t+1}}-V_{0,N_1^{t+1}}\right)\frac{\partial}{\partial q}\varphi\left(\widehat{N}_1^t|q\right)\end{array}\right). \tag{85}$$

[32]One can show that for small $\nabla_{N_1^t}p>0$ this term is negative, working in the direction of our result. This is also precisely why we can not formally state the analogous property $\frac{d}{dq}\left[(V_{1,N_1+1}-V_{0,N_1+1})-(V_{1,N_1}-V_{0,N_1})\right]>0$ for part 3, with this term taking the opposite sign to Γ.

Again, with $\nabla_{N_1^t} p = 0$, $\nabla_{N_1^t}(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}}) = 0$ leaving:

$$\frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right] = \sum_{\widehat{N}_1^t} \frac{\partial}{\partial q} \nabla_{N_1^t} \left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right) \varphi \left(\widehat{N}_1^t \,| q \right). \tag{86}$$

Thus,

$$\frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right] \leq \max_{\widehat{N}_1^t} \frac{\partial}{\partial q} \nabla_{N_1^t} \left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right). \tag{87}$$

With the right hand side of (83) increasing in $\frac{\partial}{\partial q} \nabla_{N_1^t} E_t \left[V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right]$, this implies the following upper bound on $\frac{\partial}{\partial q} \nabla_{N_1^t}(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}})$:

$$\frac{\partial}{\partial q} \nabla_{N_1^t} \left(V_{1,N_1^{t+1}} - V_{0,N_1^{t+1}} \right) \leq \frac{\Gamma}{1 - \delta (1 - p)} < 0, \tag{88}$$

with the second inequality becoming strict when either $\nabla_{N_1^t} \mu < 0$ or $\frac{\partial}{\partial q} \nabla_{N_1^t} \lambda > 0$. Finally, by continuity of all relevant terms in q and $\nabla_{N_1^t} p$, this will also hold for some neighborhood of $(q, \nabla_{N_1^t} p) = (0, 0) \in [0, 1] \times [0, 1]^N$. Take $\overline{q}'' = B''$ to be the diameter of largest ball included in this neighborhood, and define $\overline{q}^2 := \min(\overline{q}', \overline{q}'')$ and $B := \min(B', B'')$. This concludes the proof of Proposition 5. $\square$

D.4 Proof of Corollary 1

With $\frac{\partial}{\partial q} \left(V_{1,N_1^t} - V_{0,N_1^t} \right)$ negative for each $q \in [0, \overline{q}^2)$, integrating $V_{1,N_1^t} - V_{0,N_1^t}$ over $[0, q) \subseteq [0, \overline{q}^2)$ yields a value gap $V_{1,N_1^t} - V_{0,N_1^t}$ decreasing in q for each N_1^t. Then, given N_1^t, the distribution of carried inventories $\widehat{N}_1^t$ shifts left with q, as the frequency of asset holders forced to liquidate (having experienced a liquidity shock and unable to sell for a price above $\underline{v}$) increases. And in addition, residual asset holders may begin to funnel if $\overline{N}_1$ drops below N_1^t. Thus, with the distribution of $\widehat{N}_1^t$ shifted left and each $V_{1,N_1^t} - V_{0,N_1^t}$ shifting down, the distribution of N_1^{t+1} (conditioning on N_1^t) unambiguously shifts left. This will also hold at any $t + s$ given N_1^{t+s} for $s > 0$, by stationarity. Thus, as a convolution of distributions– each having experienced a stochastic shift left– the distribution of N_1^{t+s} conditional on N_1^t must shift left as q increases.

$\square$

E Appendix: Numerical Solutions

Examples (a) and (b) use InterdealerSegD.m for the segmented demand setup with geometrically distributed $\widetilde{D}$, InterdealerCoD.m for the co-demand setup, with both using VTIL.m to estimate expectated shares in the interdealer market for a Poisson random network (with links between dealer and between vacants excluded). All other subfunctions are included. For each state dependent value $Y_{N_1^t}$ we use $Y \in \mathbb{R}_+^{N+1}$ to denote the column vector of values.

For Geometric Co-demand setup:

$$
\begin{aligned}
p\left(N_1^t\right) &= 1 - F\left(N_1^t - 1; N_1^t\right) + \sum_{\widetilde{D}=0}^{N_1^t-1} \frac{\widetilde{D}}{N_1^t} f\left(\widetilde{D}; N_1^t\right) \\
&= 1 - \left(1 - (1-s)^{N_1^t}\right) + \frac{1}{N_1^t} \sum_{\widetilde{D}=0}^{N_1^t-1} \widetilde{D}\,(1-s)^{\widetilde{D}}\, s \\
&= (1-s)^{N_1^t} + \frac{s\,(1-s)}{N_1^t} \frac{1 - N_1^t\,(1-s)^{N_1^t-1} + \left(N_1^t-1\right)(1-s)^{N_1^t}}{s^2} \\
&= \frac{1-s}{s} \frac{1}{N_1^t} \left(1 - (1-s)^{N_1^t}\right).
\end{aligned}
$$

Step 1: Bootstrap values for Poisson random network

$\widetilde{\mathcal{V}}_{0,N_1^t}^{ns}$ and $\widetilde{\mathcal{V}}_{1,N_1^t}^{s}$, probability mass function of $\widetilde{X}$, $f_{\widetilde{X};\left(\widetilde{N}_1^s,\widetilde{N}_0^{ns}\right)}$, and expected seller shares $E\left[s_1\right]$, given each interdealer market dimension $\left(\widetilde{N}_1^s, \widetilde{N}_0^{ns}\right)$ are bootstrapped using the function VTIL.m. VTIL.m takes values (N, m), where $m \in (0,1)$ gives the probability of any buyer and seller being matched. For each bipartite network $\widetilde{\mathcal{B}}$, VTIL.m solves for a stable assignment:

$$
\begin{gathered}
\min_{(s,t)\in\mathbb{R}^{\widetilde{N}_1^s}\times\mathbb{R}^{\widetilde{N}_0^{ns}}} \sum_{i\in\widetilde{N}_1^s} s_i + \sum_{j\in\widetilde{N}_0^{ns}} t_j \\
s.t\ s_i \geq 0, t_j \geq 0 \\
s.t\ s_i + t_j \geq \alpha_{ij}.^{33}
\end{gathered}
$$

[33] See Roth and Sotomayor (2001), page 200.

Step 2: Derive expectations over interdealer market: λ and μ.

Subfunctions Interdealer_lammuSegD.m and Interdealer_lammuCoD.m calculate λ and μ as functions of N_1^t using:

$$
\begin{aligned}
\lambda &= \sum_{\widetilde{N}_1^s \geq 0}^{N_1^t} \sum_{\widetilde{N}_0^{ns} \geq 1}^{N-N_1^t} g_1\left(\widetilde{N}_1^s \left| N_1^t\right.\right) h_0^B\left(\widetilde{N}_0^{ns} \left| N_1^t\right.\right) \sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}_{0,N_1^t}^{ns} \\
\mu &= \sum_{\widetilde{N}_1^s \geq 1}^{N_1^t} \sum_{\widetilde{N}_0^{ns} \geq 0}^{N-N_1^t} g_1^S\left(\widetilde{N}_1^s \left| N_1^t\right.\right) h_0\left(\widetilde{N}_0^{ns} \left| N_1^t\right.\right) \sum_{\widetilde{\mathcal{B}}} b\left(\widetilde{\mathcal{B}} \middle| \widetilde{N}_0^{ns}, \widetilde{N}_1^s\right) \widetilde{\mathcal{V}}_{1,N_1^t}^{s},
\end{aligned}
$$

using $\widetilde{\mathcal{V}}_{0,N_1^t}^{ns}$ and $\widetilde{\mathcal{V}}_{1,N_1^t}^{s}$ from step 1.

Step 3: Derive process generating carried inventories: $\Pr\left(\widehat{N}_1^t; N_1^t\right)$.

Step 3.1:

This takes on seperate forms for the co-demand and segmented demand setups. For the co-demand setup, we have $\widetilde{N}_1 = \max\left\{0, N_1^t - \widetilde{D}\right\}$. Thus,

$$
\begin{aligned}
f_{\widetilde{N}_1;N_1^t}(x) &= \begin{cases} 0 & \text{if } x \notin \{0,\ldots,N_1^t\} \\ f\left(N_1^t - x\right) & \text{if } x \in \{0,\ldots,N_1^t\} \end{cases} \\
F_{\widetilde{N}_1;N_1^t}(x) &= \begin{cases} 0 & \text{if } x < 0 \\ 1 - F\left(N_1^t - x - 1; N_1^t\right) & \text{if } x \geq 0 \end{cases}.
\end{aligned}
$$

Applying $f(x) = (1-s)^x s$ for the geometric distribution with parameter s, we derive:

$$
f_{\widetilde{N}_1;N_1^t}(x) = (1-s)^{N_1^t - x} s.
$$

And for the segmented demand setup $f_{\widetilde{N}_1}$ is given by:

$$
f_{\widetilde{N}_1;N_1^t}(x) = \binom{N_1^t}{x} p^{N_1^t - x} (1-p)^x,
$$

with support $x \in \{0,\ldots,N_1^t\}$.

Step 3.2: Given $\widetilde{N}_1$ (and N_1^t):

$$
\begin{aligned}
f_{\widetilde{N}_1^s;\widetilde{N}_1,N_1^t}(x) &= \binom{\widetilde{N}_1}{x} q^x (1-q)^{\widetilde{N}_1 - x} \\
f_{\widetilde{N}_0^{ns};N_1^t}(x) &= \binom{N - N_1^t}{x} q^{N - N_1^t - x} (1-q)^x
\end{aligned}
$$

Step 3.3: $\widetilde{N}_1 - \widetilde{N}_1^s$ gives the number of asset holders unable to sell at high price but able to carry inventories. Taking $f_{\widetilde{X};\left(\widetilde{N}_1^s,\widetilde{N}_0^{ns}\right)}$ from the interdealer market, we have $\widehat{N}_1^t = \left(\widetilde{N}_1 - \widetilde{N}_1^s\right) + \widetilde{X}$, so that:

$$f_{\widehat{N}_1^t;\left(\widetilde{N}_1^s,\widetilde{N}_0^{ns}\right)}(x) = f_{\widetilde{X};\left(\widetilde{N}_1^s,\widetilde{N}_0^{ns}\right)}\left(x - \left(\widetilde{N}_1 - \widetilde{N}_1^s\right)\right).$$

Then, for each $\widehat{N}_1^t \in \left\{0, \ldots, \widetilde{N}_1\right\}$:

$$f_{\widehat{N}_1^t;\left(\widetilde{N}_1,N_1^t\right)}(x) = \sum_{\substack{\widetilde{N}_1^s \in \left\{0,\ldots,\widetilde{N}_1\right\}, \\ \widetilde{N}_0^{ns} \in \left\{\widehat{N}_1^t - \left(\widetilde{N}_1 - \widetilde{N}_1^s\right),\ldots,N-N_1^t\right\}}} \left(\begin{array}{c} f_{\widetilde{N}_1^s;\widetilde{N}_1}(\widetilde{N}_1^s) f_{\widetilde{N}_0^{ns};N_1^t}(\widetilde{N}_0^{ns}) \times \\ f_{\widetilde{X};\left(\widetilde{N}_1^s,\widetilde{N}_0^{ns}\right)}\left(x - \left(\widetilde{N}_1 - \widetilde{N}_1^s\right)\right) \end{array} \right).$$

Step 3.4: Finally, we have:

$$\Pr\left(\widehat{N}_1^t; N_1^t\right) := f_{\widehat{N}_1^t;N_1^t}(x) = \sum_{\widetilde{N}_1 \in \left\{\widehat{N}_1^t,\ldots,N_1^t\right\}} f_{\widetilde{N}_1;N_1^t}(\widetilde{N}_1) f_{\widehat{N}_1^t;\left(\widetilde{N}_1,N_1^t\right)}(x).$$

Step 4: derive stationary value gap

We construct fixed point processes to converge on value vector V_0, V_1, and

$$G\left(N_1^t\right) := \begin{cases} V_{1,N_1^t} - V_{0,N_1^t-1} & \text{if } N_1^t > 0 \\ V_{1,N_1^t} - \delta E_t\left[V_{0,N_1^{t+1}}\right] & \text{if } N_1^t = 0 \end{cases}.$$

Start with initial values $V_0^{(0)} = [0]_{N_1^t}$ and $V_1^{(0)} = [pv + (1-p)\,\underline{v}]_{N_1^t}$. I assume supply can be written $\Phi\left(N_1^t - \hat{N}_1^{t-1}\right)$; that N be sufficiently small relative to total issuance. The following process $\left(\Pr^{(n)}, V_0^{(n)}, V_0^{(n)}\right)$ will converge to the state process function $\Pr\left(N_1^t; N_1^{t-1}\right)$ and upwards to the stationary values V_0 and V_1as $n \to \infty$:

Step 4+n.1: derive $\Pr^{(n)}\left(N_1^t; N_1^{t-1}\right)$ from $\Pr\left(\widehat{N}_1^t; N_1^{t-1}\right)$ and

$$G^{(n-1)}\left(N_1^t\right) := \begin{cases} V_{1,N_1^t}^{(n-1)} - V_{0,N_1^t-1}^{(n-1)} & \text{if } N_1^t > 0 \\ V_{1,N_1^t}^{(n-1)} - \delta E_t\left[V_{0,N_1^{t+1}}^{(n-1)}\right] & \text{if } N_1^t = 0 \end{cases}$$

using

$$\begin{aligned} \Phi\left(N_1^t - \hat{N}_1^{t-1}\right) &\leq G^{(n-1)}\left(N_1^t\right), \text{ with} \\ \Phi\left(N_1^t + 1 - \hat{N}_1^{t-1}\right) &> G^{(n-1)}\left(N_1^t + 1\right). \end{aligned} \tag{89}$$

Precisely, for each $N_1^t \in \{0, \ldots, N\}$ take $\widehat{\mathcal{N}}_{1,N_1^t}^{(n-1)} \subseteq \{0, \ldots, N_1^t\}$ to be the set of $\hat{N}_1^{t-1}$ that solves (89), then:

$$\Pr^{(n)}\left(N_1^t; N_1^{t-1}\right) = \sum_{\widehat{N}_1^{t-1} \in \widehat{\mathcal{N}}_{1,N_1^t}^{(n-1)}} \Pr\left(\widehat{N}_1^{t-1}; N_1^{t-1}\right).$$

Then, construct $\Omega^{(n)} := \left[\Pr^{(n)}\left(N_1^t; N_1^{t-1}\right)\right]_{N_1^{t-1}, N_1^t}$.

Step 4+n.2:

Given values $\left(V_0^{(n-1)}, V_0^{(n-1)}\right)$ and $\Pr^{(n)}$ from step 4+n.1, take

$$\begin{aligned}
V_{0,N_1^t}^{(n)} &= \left[\begin{array}{c} (1-q)\left[\left(\begin{array}{c} \lambda\left(\delta E_t\left[V_{1,N_1^{t+1}}\right] - \underline{v}\right) \\ +(1-\lambda)\,\delta E_t\left[V_{0,N_1^{t+1}}\right] \end{array}\right)\right] \\ +q\delta E_t\left[V_{0,N_1^{t+1}}\right] \end{array}\right] \\
V_{1,N_1^t}^{(n)} &= \left[\begin{array}{c} (1-p)\left[\begin{array}{c} (1-q)\,\delta E_t\left[V_{1,N_1^{t+1}}\right] \\ +q\left(\begin{array}{c} \mu\delta E_t\left[V_{1,N_1^{t+1}}\right] \\ +(1-\mu)\left(\underline{v} + \delta E_t\left[V_{0,N_1^{t+1}}\right]\right) \end{array}\right) \end{array}\right] \\ +p\left(v + \delta E_t\left[V_{0,N_1^{t+1}}\right]\right) \end{array}\right]
\end{aligned}$$

where $E_t\left[V_{0,N_1^{t+1}}\right]$ and $E_t\left[V_{1,N_1^{t+1}}\right]$ are given by the expectation $V_{0,N_1^t}^{(n)}$ and $V_{1,N_1^t}^{(n)}$ (respectively) given probability mass function $\Pr^{(n)}\left(N_1^t; N_1^{t-1}\right)$ derived in Step 4+n.1:

$$\begin{aligned}
\left[E_t\left[V_{0,N_1^{t+1}}\right]\right]_{N_1^t} &= \Omega^{(n)} V_0^{(n)}, \\
\left[E_t\left[V_{1,N_1^{t+1}}\right]\right]_{N_1^t} &= \Omega^{(n)} V_1^{(n)}.
\end{aligned}$$

Outputs

Expected interdealer price given N_1^t is calculated as:

$$E_t\left[\widetilde{\mathcal{P}}^D\right] = \left(\delta E_t\left[G(N_1^{t+1})\right] - \underline{v}\right) E_t[s] + \underline{v},$$

using steady state $\delta E_t\left[G\left(N_1^{t+1}\right)\right]$ from above and $E_t[s]$. Value gaps, expected future value gaps, and expected interdealer prices are plotted for $N_1^t \in \{1, \ldots, N-1\}$, as μ is not defined at $N_1^t = 0$ and λ is not defined for $N_1^t = N$.

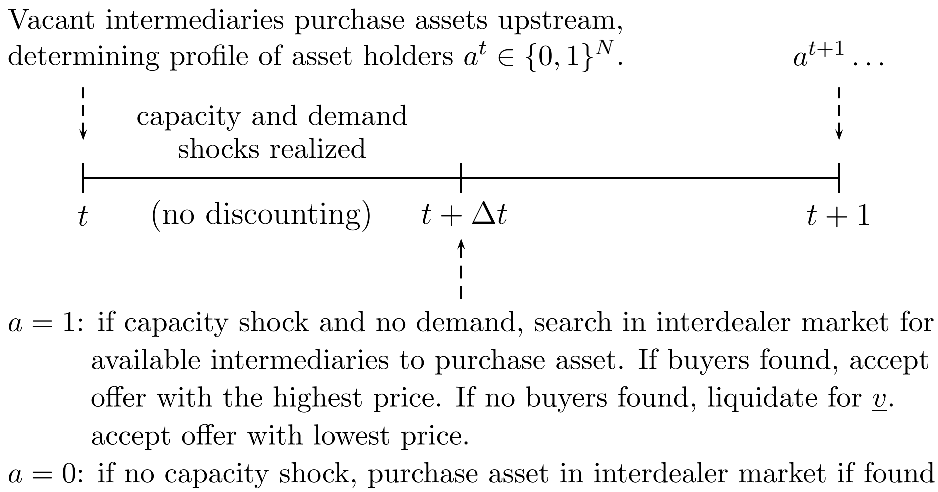

Figure 1: Period t timeline

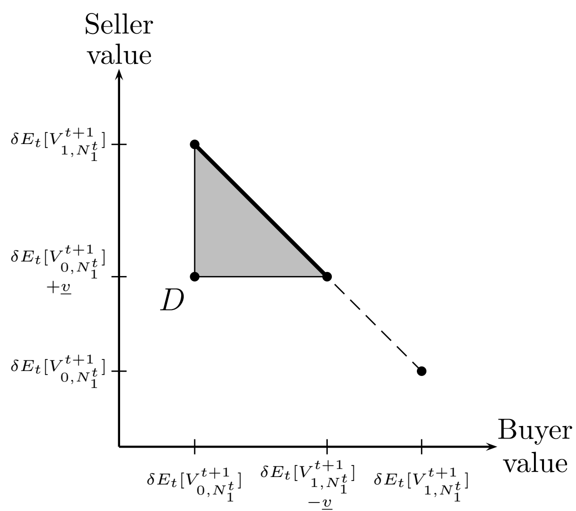

Figure 2: Interdealer value set

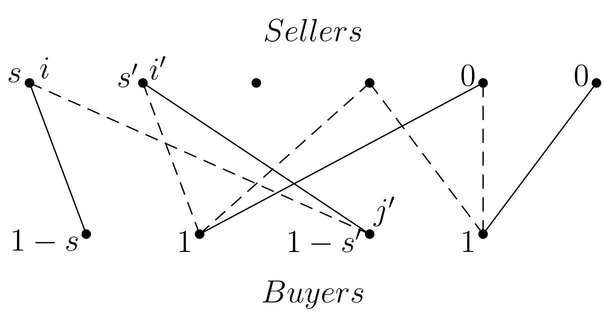

Figure 3: Interdealer market buyer network

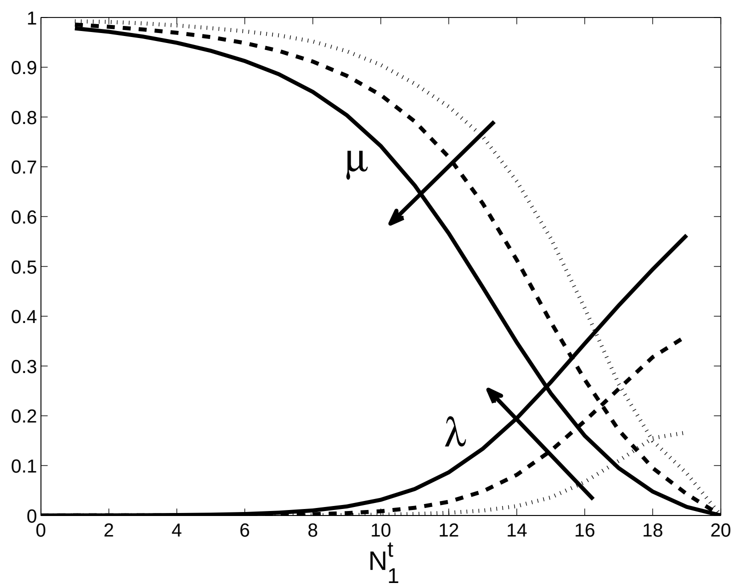

Figure 4: Estimates of λ and μ as functions of N_1^t

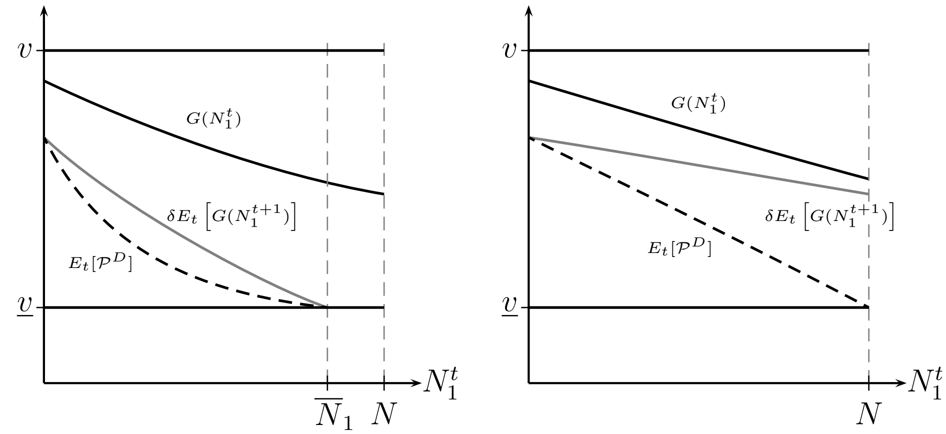

(a) Case 1: $\overline{N}_1 \leq N$

(b) Case 2: $\overline{N}_1 > N$

Figure 5: Proposition 4

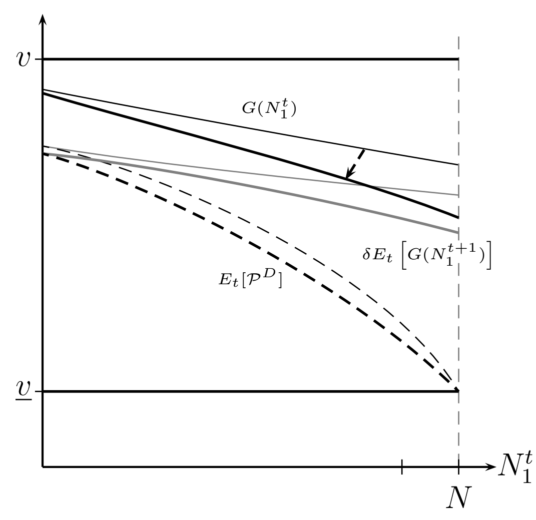

Figure 6: Increase to liquidity risk q and the price-inventory link.

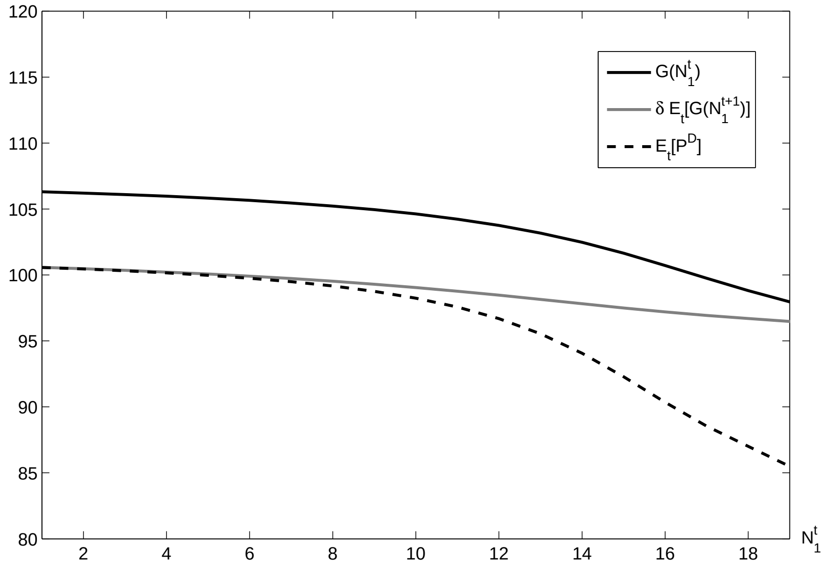

(a) Segmented Demand: $N = 20$, $p = .3$, $q = .2$, linear supply, $m = .2$ (i.i.d. links)

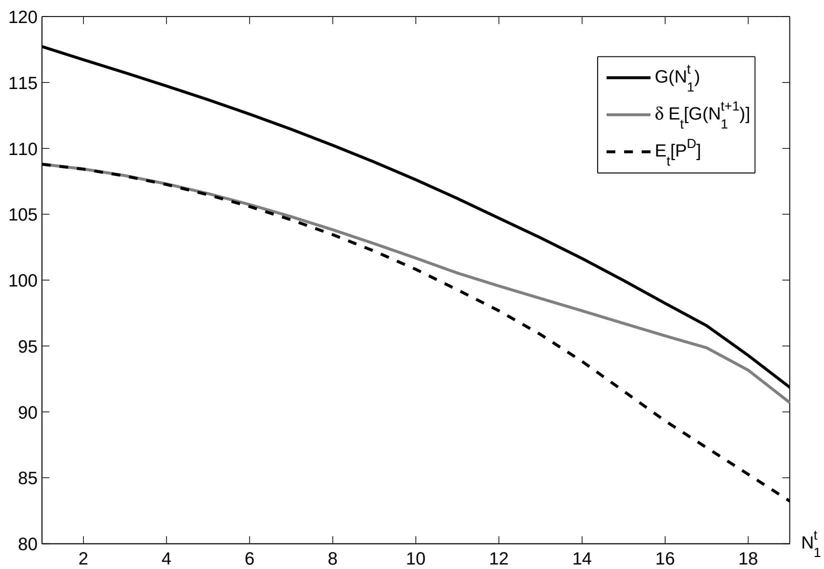

(b) Co-Demand: $N = 20$, $\widetilde{D} \sim$ Geometric(.3), $q = .2$, linear supply, $m = .2$ (i.i.d. links)

Figure 7: Numerical examples

Information Acquisition and Response in Peer-effects Networks

C. Matthew Leister*

Abstract

Recently the network games literature turned its attention to strategic interactions under incomplete and asymmetric information. The canonical assumption takes players' signal qualities to be exogenous. This paper allows players to privately acquire costly payoff-relevant information prior to simultaneous play. The presence of peer effects, captured by the economy's network structure, implies that information externalities assume a rich and nuanced form. When pairwise peer effects are symmetric, asymmetries in acquired information are inefficiently low relative to the utilitarian benchmark. And with information privately acquired, all players face strictly positive gains to overstating their informativeness as to strategically influence the beliefs and behaviors of neighbors. If strategic substitutes in actions are present and significant, low centrality players move against their signals in anticipation of their neighbors' actions. A blueprint for optimal policy design is developed. Applications to market efficiency in financial crises and two-sided markets are discussed.

Keywords: Network Games, Coordination Games, Endogenous Information, Peer Effects, Network Centrality
JEL: C72, D21, D82, D85, G14

*University of California, Berkeley. E-mail: leister@berkeley.edu

[1]I'd like to thank David Ahn, Ivan Balbuzanov, Haluk Ergin, Satoshi Fukuda, Nicolae Gârleanu, Sanjeev Goyal, Brett Green, Shachar Kariv, Sanket Korgaonkar, Maciej Kotowski, Kaushik Krishnan, Sheisha Kulkarni, Natalia Lazzati, Martin Lettau, Raymond Leung, Sheng Li, Gustavo Manso, Pau Milan, David Minarsch, Michèle Muller, Carl Nadler, Paulo Natenzon, Marcus Opp, Chris Shannon, Philipp Strack, Xavier Vives, and many seminar attendees for helpful comments. All errors are my own.

14 Introduction

Many environments involve individuals acquiring and using information toward both learning more about the world and inferring the information of others. This ubiquitous dual role of information plays out in financial markets, labor markets, social networks and trends, as well as in professional communities. Relevant to each of these examples, individuals commonly face asymmetric incentives to invest in costly information depending on their identity, market position, and social ties. And when information is used to infer the observations of influential players, the strategic response to signals establishes a crucial component to the private value of information. This paper studies the role of such peer effects in shaping the incentives to acquire and strategically respond to information. It examines both the positive and normative implications of the resulting disparities in acquired information qualities.

An example embodying this duality while in the presence of directed peer effects is given with the following vignette. At some point in time, an independent research institute develops and patents a novel drill technology. The new drill potentially means that a large, previously untapped field of deep-sea oil deposits can now be safely resourced. The institute advertises the promise of the drill, and is willing to lease out the rights to operate the technology on a per-drill basis. A three-player market consists of two petroleum firms (firms A and B) comprising a competitive duopoly and a lobbyist for the petroleum industry. The three players pursue their own due diligence as to confirm or refute the drill's value. We can capture the resulting network of relationships with the following figure.

[Figure 6]

The sign and direction of links emanating from each individual capture the competitive and supportive influences that others' investment choices have on their private incentive to adopt the drill.

Peer effects in technology adoption feed into the incentives to acquire and respond to private information in the following ways. With both firms simultaneously researching the technology,[1] any acquired information regarding the drill's value brings with it the knowledge of greater competition. That is, a firm that learns the drill is effective also learns that they are likely to face stiff competition when drilling. This is precisely because information regarding the drill's efficacy can also be used to infer the competition's observations and subsequent investment in the drill. The lobbyist, on the other hand, will decide whether to utilize her resources promoting subsidies toward the employment of the drill technology or focus her efforts elsewhere. Choosing the optimal agenda to pursue requires her own due diligence. And as a function of the connections that she has with firms A and B, her incentives to acquire information will depend on how informed she can expect the firms will be. This is because upon learning of the drill's value, the more she subsequently promotes the drill the more she will need the firms to follow suit and utilize the technology. For the firms, the more efficient the drill appears the more likely they can expect subsidization in the near future – *if* the lobbyist is also expected to do her research.

[1] Hendricks and Porter (1996) [37] provide evidence of non-cooperative exploration in these industries.

Crucially, the clarity in any individual's inference of others' observations depends on the equilibrium extent of research undergone by the others in the market. Those who learn of the technology's value also learn that other highly informed individuals observe its value. Put succinctly, the collective incentives of firms A, B and the lobbyist to acquire information intricately depend on each individual's expectation of the information acquisitions, observations and subsequent actions of the others. The in-equilibrium incentives to invest in information will ultimately depend on the strategic interdependencies that each player's market position entails. Those in highly competitive positions in the market (e.g. competitive firms) will, ceteris paribus, face less value to information than those in supported or complimented market roles (e.g. lobbyists and experts).

With weighted, directed, and signed peer effects pushing and pulling equilibrium incentives, what are the welfare implications of equilibrium information acquisition? Precisely, who over invests and who under invests in information relative to the utilitarian benchmark? And, do players carry incentives to distort other's beliefs regarding their acquired information qualities? While a rich literature studying coordination games with endogenous information[2] broadly focusing on symmetric beauty-contests has offered a number of results relevant to these questions,[3] the following network setup offers a novel platform toward assessing inefficiencies in more diverse economies.

First, the essential structural property that drives the direction of inefficiencies is the extent of symmetry in pairwise relationships. Symmetric networks, in which pairwise peer effects are identical, provide generalizations to many features obtained in the coordination games with endogenous information literature. For example, in symmetric beauty contests under/over acquisition of information in equilibrium has been shown to accompany strategic complements/substitutes in the second stage. In symmetric *networks*, a more general *bunching* in acquired information qualities obtains. For example, those facing a majority of strategic complements acquire the most information in equilibrium but also *under* acquire relative to the utilitarian benchmark. Those facing a majority of strategic substitutes acquire the least but *over* acquire. Departing from these results, the direction of inefficiencies *reverse* upon introducing sufficient anti-symmetry in pairwise relationships. Precisely, when pairwise peer effects exhibit opposing signs, acquired information exhibits inefficient *spreading* in equilibrium.

A second novelty unique to network settings is the introduction of players strategically moving against their signals. Under sufficient network irregularity and for players occupying adverse positions in the network (i.e. facing significant strategic substitutes), the endogenous choice to invest in costly information and strategically move against signal realizations arises. Put crudely, players may *short the network*. Inefficiencies naturally arise with this behavior, with the direction of these inefficiencies continuing to be driven by pairwise symmetry. In

[2] This literature is commonly referred to as "*global games with endogenous information*".

[3] To list a few examples, Morris and Shin (2002) [47] and later Myatt and Wallace (2009) [49] illustrate how strategic effects in actions can influence information choice. Vives (1988) [61], (2008) [62] and Hellwig and Veldkamp (2009) [34] show how strategic complements (substitutes) can directly spill into complements (substitutes) in information acquisition, and in turn derive inefficient under (over) acquisition in equilibrium. And, Colombo et al. (2014) [19] provide an encompassing analysis of the inefficiencies that arise from the strategic use of private and public information, casting equilibrium play against both the efficient acquisition *and* efficient use of information. Section 18.3 further discusses relation to this literature.

symmetric networks, the equilibrium extent to which these players acquire and move against their signals is inefficiently *low*. Precisely, the rationality in this equilibrium behavior is valued by the very neighbors that invoke it. And consistent with the preceding message, this value reverses when peer effects are anti-symmetric. That is, those moving against their signals impose a net cost on those they influence.

An important question arises when considering such environments comprised of a finite number of strategically informed players. What would happen if players could influence others' beliefs? With signal qualities privately acquired, players face a marginal cost due to their inability to directly influence others' perceptions of their expertise. In reality, firms in an array of industries are commonly observed marketing the qualities of their research departments. Lobbyists are found promoting the extent of their expertise in their given industry or interest. While such marketing may serve a number of goals, this paper taps into a common impetus for this behavior, found within equilibrium information acquisition and response. Once again, the strength and direction of this force ultimately depends on the network's extent of symmetry among pairwise relationships.

Elaborating on this, our three-player petroleum market is seen to display symmetry in each pair's peer effects. In this environment, the marginal value derived from the strategic use of information takes on a uniformly-positive orientation, regardless of players' positions in the network. If firm A, for example, is able to influence firm B's beliefs by acquiring additional information, this discourages firm B's strategic responsiveness. For the lobbyist, firms A's additional informativeness only encourages her corresponding behavior. Both of these effects work in firm A's favor. A similar story holds for firm B. For the lobbyist, her additional informativeness encourages the actions of both firms. And if the firms consequently acquire additional information, the value that the lobbyist obtains from her own research, which allows her to infer the observations and subsequent actions of the firms, only increases. In other words, *everyone* carries the incentive to exaggerate the quality of their acquired information. As will be seen, the extent of *connectedness* to others in the network drives the magnitude of the strategic incentives to information acquisition.

To study these heterogeneous environments in a reduced form while maintaining scope, the following model employs a familiar quadratic-payoffs setup under the general linear peer-effects pioneered by Ballester et al. (2006) [5]. Incorporating incomplete information, the model captures players' information investments in an initial stage. Signals are observed, informing players of their marginal values to second-stage action. When correlation between payoffs is introduced, signals begin to inform of the likely observations of neighbors. In line with the above vignette, the clarity of this inference is a function of the signal's quality as well as the qualities that neighbors are expected to acquire. An *information-response game* is derived and characterized, played on the same network of peer effects but transformed by the equilibrium correlation in signals. Here, players choose the extent to which their strategies respond to their information. The resulting equilibrium profile of strategic responses defines players' *informational centralities* in the game.

As a function of the unique linear equilibrium of the information-response game, the incentives to acquire information across players are derived. Marginal values to information are shown to scale with the square of each player's responsiveness. The scaling of marginal values with absolute informational centralities carries with it the potential for players moving against their signals. As such, information acquisition takes on a U-shaped non-monotonicity in networks. Acquisition at the bottom *decreases* with centrality in the information-response game, with the least central players investing in high levels of information as to move against the anticipated actions of neighbors.

After characterizing equilibrium behaviors and addressing the welfare and strategic implications of information acquisition, we turn to optimal policy design. A hypothetical *neutral player* is designated. Though an active member in the network, this player behaves as though she is in isolation, without peer influences. Then given a symmetric network, players that respond more so than the neutral player under-acquire information. Those responding less so but positively to their signal realizations over acquire information. And those moving against their signals under acquire information. With positive strategic values to information throughout the network, allowing players to publicly observe the information investments of the most central players as well as those moving against their signals increases aggregate welfare. As these players internalize the strategic value to information acquisition, the network collectively adjusts information investments efficiently. Importantly, this alignment in strategic values and informational externalities for these two sets of players persists in anti-symmetric networks. Thus together, the origin (i.e. no information acquisition) and the extent of acquisition and response of the neutral player provide a normalized yardstick useful for designing optimal transparency-based interventions, portable across network structures.

Applications of the model are then considered. The incorporation of both strategic substitutes and complements into the analysis affords a high level of flexibility and scope. Our three player network of firms A, B and our lobbyist provides one industrial organization incorporating both strategic substitutes and complements. Supply chains may also embody an array of both positive peer effects (e.g. between vertically positioned firms) and negative peer effects (e.g. between horizontal competing firms).[4] Section 18.1 further explores two more applications: financial markets under liquidity crises and two-sided markets. Both of these examples call on networks with both positive and negative links, with the former also exhibiting anti-symmetric relationships.

The implications for markets in crises are as follows. In liquidity flush markets, with traders unconstrained in their asset positions, strategic substitutes in asset demand implies strategic substitutes in information acquisition. Market crowding between firms' information investments parallels the market's informational inefficiency derived in rational expectations, as in the seminal work of Grossman and Stiglitz (1980) [30]. From a welfare perspective, the strategic use of costly information implies *over* investment of information in the market. The application then move beyond competitive markets to explore the implications of firms facing severe funding constraints during a liquidity crises. À la the type of liquidity spirals studied in Brunnermeier and Pedersen (2009) [11], a subset of firms are assumed to exhibit upward sloping demands, with high market prices allowing them to retain inventories and avoid unwanted liquidations. As the proportion of constrained firms to unconstrained firms grows large, firms throughout the market under acquire information. Constrained firms impose positive externalities on each other as they collect information, and aim to coordinate on high market liquidity outcomes. While unconstrained firms impose negative externalities on each other, they fail to internalize the sizable value that their information investments

[4] Ostrovsky (2008) [54] and Kotowski and Leister (2014) [40] study the tension between vertical strategic complements and horizontal strategic substitutes in competitive supply chains.

provide to constrained firms.

Taking job-search networks as a tangible example of a two-sided market, industry insiders and workers researching job opportunities compete with those within their group while complimenting the investment choices in the counterpart group. With these networks exhibiting extensive symmetry amongst pairwise relationships, the shorter, less competitive side of the market (insiders, commonly) under invests in information. The longer, competitive side of the market (workers) over acquires information. Here, insiders fail to internalize the value that their expertise endows workers, while workers over exert themselves researching job opportunities.

The organization of the paper is as follows. Section 15 provides the model's setup and discusses the optimal information acquisition and response problem of a single, isolated player. Section 16 then defines and characterizes equilibria in general networks. It discusses and derives the information-response game, and corresponding ex ante incentives to invest in information. It then offers a number of revealing examples describing the potential for equilibrium multiplicities and negative signal responses. Section 17 formalizes the welfare and strategic considerations discussed above under moderately sized peer effects. Welfare and strategic information acquisition for players moving against their signals are then addressed. A more general analysis of optimal policy design is then developed. Finally, Section 18 discusses applications, covers basic extensions of the model, returns to related literature, and concludes. A Supplemental Section 19 below more closely explores the relationship between network structure and information costs.

15 Model Setup

Time is discrete with two periods $t = 1, 2$. Period $t = 1$ gives the information acquisition game (*first stage*). Period $t = 2$ gives a Bayesian game in which N players simultaneously act in response to their information (*second stage*). For the second stage we adopt the bilinear payoffs studied by Ballester et al. (2006) [5]. We extend their setup to incorporate incomplete information regarding the marginal benefits of action $x_i \in \mathbb{R}$ for each player $i \in \{1, \dots, N\}$.

The following notation is used. Each player i directly cares about state $\tilde{\omega}_i := \gamma\omega + \sqrt{1-\gamma^2}\omega_i$, a mixture of a player-specific state ω_i with a common (shared) state ω, each drawn from $\Omega \subset \mathbb{R}$.[5] The loading $\sqrt{1-\gamma^2}$ on ω_i merely normalizes the variance of $\tilde{\omega}_i$, simplifying the following analysis. A more general treatment is addressed in Section 18.2 with minor modification to the following. The respective state pairs (ω, ω_i) for each i and (ω_i, ω_j) for each i and $j \neq i$ are taken as jointly independent. Together, γ and ω scale the public alignment in preference shocks. $\gamma\omega$ should be interpreted as a publicly-shared but commonly-unknown component to the marginal value to adopting some technology in the second stage. $\sqrt{1-\gamma^2}\omega_i$ gives the corresponding idiosyncratic component.

[5]While here $\boldsymbol{\omega}$ will denote the vector of states $(\omega, (\omega_i)_{i=1}^N)$, bold symbols will generally be used to denote profiles (vectors) of respective parameters and variables, with components for each $i \in \{1, \dots, N\}$. We can consider ω and ω_i to follow standard normal distributions, though the more general properties in players' expectations required in the analysis are given below with E1-E4.

All information is learned after the second stage, with each player i realizing her payoff:

$$u_i(\mathbf{x}|\omega,\omega_i) = (a_i + \tilde{\omega}_i)\, x_i - \frac{1}{2}\sigma_{ii}x_i^2 + \sum_{j\neq i}\sigma_{ij}x_ix_j.$$

a_i scales i's publicly-known average marginal gain to x_i, or her expected predisposition for second-stage action. It incorporates her average marginal value to action x_i, leaving residual uncertainty to be captured by the state $\tilde{\omega}_i$. σ_{ii} gives a positive constant scaling the concavity in her utility, capturing diminishing returns to x_i. σ_{ij} measures the influence that j's action x_j has on i's marginal gain to x_i (j's *peer effect* on i) and takes values in $\mathbb{R}$. Positive σ_{ij} will correspond to strategic complements, negative values to strategic substitutes, with $\sigma_{ij} = 0$ designating that j lies outside of i's *neighborhood*. Σ will be used to denote the square matrix $[\sigma_{ij}]$ with 0's along the diagonal. The sizes of the elements a_i and σ_{ij} for each $j \neq i$ relative to σ_{ii} determine the responsiveness of i's ideal action to the second-stage actions of her neighbors.

Each player i does not directly observe any component of $\boldsymbol{\omega}$. However, at $t = 2$ i does receive information (θ_i, e_i), giving signal realization $\theta_i \in \Theta \subset \mathbb{R}$ of quality $e_i \in [0,1]$ informing her of $\tilde{\omega}_i$. i does not observe (θ_j, e_j) for each $j \neq i$. Thus, i is free to choose private information-contingent second-stage strategy $X_i(\cdot|\cdot) : \Theta \times [0,1] \to \mathbb{R}$ mapping privately observed signal θ_i to an action in $\mathbb{R}$ given her quality e_i.

In the first stage each i privately invests in the signal quality e_i. The cost of quality (i.e. information acquisition effort) is given by the convex function $\kappa \in \mathcal{C}^2$ satisfying: $\kappa(0) = 0$ and $\kappa'(e_i), \kappa''(e_i) \geq 0$ for each $e_i \in [0,1]$. Beyond these standard conditions we assume the following:

Assumption 1. *$\kappa \in \mathcal{C}^3$ satisfies: $\kappa'(0) = 0$, $\kappa'''(e_i) \geq 0$ for every $e_i \in [0,1]$, and there exists an unique $e^\dagger \in (0,1)$ solving $e^\dagger = \kappa'\left(e^\dagger\right)$.*

$\kappa'(0) = 0$ implies that the marginal cost to the lowest quality information is negligible. $\kappa'''(e_i) \geq 0$ implies that the convexity in information qualities are non-decreasing, and primary serves as a technical condition sufficing for existence of a first-stage equilibrium. Uniqueness of a solution to $e^\dagger = \kappa'\left(e^\dagger\right)$ will be seen to yield a unique interior solution to any isolated player's information acquisition problem.

Without significant loss of generality we normalize $\sigma_{ii} = 1$ for each i and scale other terms as needed.[6] Again, Section 18.2 discusses extensions incorporating heterogeneous σ_{ii}, individual costs functions κ_i, as well as additional idiosyncrasies into $\tilde{\omega}_i$. All of these extensions preserve the following analysis and results.

Together, the couple (e_i, X_i) defines a pure strategy for each i in the two-stage game. As players do not directly observe quality investments of others, $\mu_{ij} : [0,1] \to \mathbb{R}_+$ will denote the $t = 2$ belief held by player i regarding j's first-stage quality investment e_j. Thus, the initial period $t = 2$ expected payoff of player i as a function of the vector of other players'

[6] Setting $\sigma_{ii} = \sigma_{jj}$ for each i and j does carry the implication that all players face common total variation in their payoffs. This allows the network of peer effects to drive all variation in equilibrium information acquisition.

strategies $\mathbf{X}_{-i}$, private information (θ_i, e_i), and beliefs μ_i can be written:

$$u_i\left(x_i, \mathbf{X}_{-i}|\theta_i, e_i, \mu_i\right) = \left(a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i\right]\right) x_i - \frac{1}{2}x_i^2 + \sum_{j \neq i} \sigma_{ij} x_i \mathbb{E}_i\left[X_j\left(\theta_j|e_j\right)|\theta_i, e_i, \mu_i, \right]. \quad (1)$$

This yields a second-stage linear best response:

$$BR_i\left(\mathbf{X}_{-i}|\theta_i, e_i, \mu_i\right) = a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i\right] + \sum_{j \neq i} \sigma_{ij} \mathbb{E}_i\left[X_j\left(\theta_j|e_j\right)|\theta_i, e_i, \mu_i\right]. \quad (2)$$

That is, each i responds to her conditional expectation of $\tilde{\omega}_i$ and to what her information informs her of the observations and actions of neighbors.

States and signals may be taken to be joint-normally distributed. The following requires only that priors be centered about the origin and posteriors be linear-in-qualities:

E1. $\mathbb{E}_i\left[\omega\right] = \mathbb{E}_i\left[\omega_i\right] = \mathbb{E}_i\left[\theta_i\right] = 0$,

E2. $\mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i\right] = e_i\theta_i$ for each $e_i \in [0, 1]$,

E3. $\mathbb{E}_i\left[\theta_i^2|e_i\right] = 1$ for each $e_i \in [0, 1]$, and

E4. $\mathbb{E}_i\left[\theta_j|\theta_i, e_i, \mu_i\right] = \int_{[0,1]} \mu_{ij}\left(e_j\right) \gamma^2 e_j e_i \theta_i de_j$.

As is common to model information investment as a number of costly draws of a normally distributed signal of given precision, Appendix F.1 applies this particular structure to derive properties E1-E4 directly. Information structures with two states also easily satisfy E1-E4.[7] Together, these give the essential properties used through the following analysis.

Conditions E1 and E2 together imply $\theta_i = \tilde{\omega}_i$ at $e_i = 1$. Condition E3 requires a normalization obtained by the appropriate increasing affine transformation to signals. The factor $\gamma^2 e_j e_i$ in condition E4 gives the correlation of the signals θ_i and θ_j.[8] Noting that any strictly-monotonic transformation does not change the informational content of signals,[9] conditions E1-E3 merely simplify the following analysis. Loss of generality does come with the linear-multiplicative separability of condition E4. The following analysis and results hinge only on multiplicative separability, however. All qualitative properties remain intact under the more general (non-linear) extension $\mathbb{E}_i\left[\theta_j|\theta_i, \mu_i, e_i\right] = \int_{[0,1]} \mu_{ij}\left(e_j\right) \gamma^2 \eta(e_j)\eta(e_i)\theta_i de_j$ for any non-negative and strictly monotone $\eta \in \mathcal{C}^1$. Finally, as the following will consider pure first-stage strategy profiles $\mathbf{e} \in [0, 1]^N$, sequential rationality in beliefs requires $\mu_{ij}^*\left(e_j\right) = 1$ for each i and j. Therefore, condition E4 reduces in equilibrium to $\mathbb{E}_i\left[\theta_j|\theta_i, e_i, \mu_i^*\right] = \gamma^2 e_j e_i \theta_i$.

Though this paper's focus is on the role of general peer effects in equilibrium information acquisition, to help fix ideas the following example solves the information acquisition and optimal response problems of a single, isolated player.

[7] Additional examples incorporating an arbitrary finite number of states can be constructed.

[8] With $\kappa(\cdot)$ a function of the quality of information that is used directly to infer $\tilde{\omega}_i$, we can interpret the efforts of i to be focused toward information sources most relevant to her particular qualities or tastes. For example, a firm's inference of the value of a production technology requires acquiring information of the technology's particular attributes most consequential to the firm's marginal product. The most important attributes should depend on the firm's specific qualities, preexisting input profile, and compatibility between coexisting technologies. Thus, κ should be interpreted as a general cost to research.

[9] Precisely, such a transformation merely rescales the value of θ_i for each $\tilde{\omega}_i$.

Example. [isolated player's problem] *Consider the information response problem of a single player i having chosen quality e_i in period $t = 1$, and now maximizing the following period $t = 2$ objective:*

$$u_i\left(x_i|\theta_i, e_i\right) = \left(a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i\right]\right)x_i - \frac{1}{2}x_i^2 = \left(a_i + e_i\theta_i\right)x_i - \frac{1}{2}x_i^2.$$

The first order condition to her problem, conditioning on information (θ_i, e_i), yields:

$$\frac{\partial}{\partial x}u_i\left(x_i|\theta_i, e_i\right) = \left(a_i + e_i\theta_i\right) - x_i = 0,$$

which gives:

$$X^*\left(\theta_i|e_i\right) = a_i + e_i\theta_i.$$

That is, i responds to her realized signal by an amount equal to the qualities of the signal, e_i. This yields period $t = 1$ expected (indirect) utility:

$$\mathbb{E}_i\left[u_i\left(X^*\left(\theta_i|e_i\right)|\theta_i, e_i\right)|e_i\right] = \mathbb{E}_i\left[\left(a_i + e_i\theta_i\right)\left(a_i + e_i\theta_i\right) - \frac{1}{2}\left(a_i + e_i\theta_i\right)^2|e_i\right] = \frac{1}{2}\left(a_i^2 + e_i^2\right),$$

which uses condition E3: $\mathbb{E}_\theta\left[\theta_i^2|e_i\right] = 1$. Then, the period $t = 1$ first-order condition for any interior $e^\dagger \in (0, 1)$ is given with:

$$e^\dagger = \kappa'\left(e^\dagger\right). \tag{3}$$

Under Assumption 1, a unique $e^\dagger \in (0, 1)$ solving (3) obtains. Further, as the above holds for all values of a_i, we see that without peer effects the isolated player (i) acquires a nonzero amount of information and (ii) responds positively to her information ($\frac{\partial}{\partial\theta_i}X^*\left(\theta_i|e_i\right) \geq 0$).

As seen in the example, the value of information exhibits a natural convexity, even when a player i acts in isolation at $t = 2$. This is because more precise information increases i's posterior belief that her response to her signal is in the optimal direction, while holding the size of her response fixed. Then, additionally allowing her to optimally increase the size of her response provides additional value. These two effects multiply each other, yielding an increasing marginal value to signal quality.

16 Equilibrium information acquisition and response

16.1 Equilibrium definitions

The following equilibrium notions are presented backward inductively.

Definition 1. **[second-stage equilibrium]** *Given profile of qualities* $\mathbf{e}$ *and beliefs* $\boldsymbol{\mu}$*, an* information response equilibrium *(IRE) is a profile of strategies* $\mathbf{X}^* := (X_1^*, \ldots, X_N^*)$ *given as a Bayesian Nash equilibrium of the second stage game:*

$$X_i^*(\theta_i|e_i) \in \arg\max_{x\in\mathbb{R}} \mathbb{E}_i \left[u_i \left(x, \left(X_j^*(\theta_j|e_j) \right)_{j\neq i} |\omega, \omega_i \right) \middle| \theta_i, e_i, \mu_i \right],$$

for each $\theta_i \in \Theta$ *and* $i \in N$*. Expectation* $\mathbb{E}_i$ *is taken over* $\tilde{\omega}_i$*,* $\boldsymbol{\theta}_{-i}$ *and* $\mathbf{e}_{-i}$ *using beliefs* μ_i*, taking other players' strategies* $\mathbf{X}_{-i}^*$ *as given.*

Given private information (θ_i, e_i), each player i best responds to her signal by investing in her action, taking the profile of all other players' actions $\mathbf{X}_{-i}$ as fixed. Her information is relevant to learning about both $\tilde{\omega}_i$ and what other players observe and do at $t = 2$.

The first-stage equilibrium for given second-stage equilibrium $\mathbf{X}^*$ and beliefs $\boldsymbol{\mu}$ is defined as follows.

Definition 2. **[first-stage equilibrium]** *Given IRE* $\mathbf{X}^*$ *and beliefs* $\boldsymbol{\mu}$*, an* information acquisition equilibrium *(IAE) is a profile of qualities* $\mathbf{e}^* := (e_1^*, \ldots, e_n^*)$ *given as a Nash equilibrium of the first stage game:*

$$e_i^* \in \arg\max_{e\in[0,1]} \mathbb{E}_i \left[u_i \left(X_i^*(\theta_i|e), \left(X_j^*(\theta_j|e_j) \right)_{j\neq i} |\omega, \omega_i \right) \middle| e, \mu_i \right] - \kappa(e),$$

for each i*, where expectation* $\mathbb{E}_i$ *is taken over* $\tilde{\omega}_i$*,* $\boldsymbol{\theta}$ *and* $\mathbf{e}_{-i}$ *using beliefs* μ_i*, taking strategies* $\mathbf{X}^*$ *as given.*

That is, each player i optimally invests in the quality of her signal θ_i at cost $\kappa(e_i^*)$, anticipating second-stage play as given by $\mathbf{X}^*$. Together, IAE $\mathbf{e}^*$, IRE $\mathbf{X}^*$ and sequentially rational beliefs $\boldsymbol{\mu}^*$ define a weak perfect Bayesian equilibrium of the two-stage game.

The following begins by characterizing equilibrium information acquisition and response under our general network setting. Section 16.3 then provides a number of examples exploring the breadth of equilibrium behaviors.

16.2 Equilibrium characterizations

Here, we characterize IRE and interior IAE of the two-stage game. As displayed below, an ex ante expected equilibrium $\boldsymbol{\alpha}^* \in \mathbb{R}^N$ can be obtained to yield average second-stage actions by averaging over realized signals.[10] A key innovation, however, is that in addition to this expected game played on a network, players play an information-response game on the same network. However, the network of peer effects is *transformed by the correlation in signals*, which is *induced by qualities acquired in the first stage*. Information now tells players not only about their marginal gain to action (i.e. the relevant state of the world $\tilde{\omega}_i$) but also about what to expect neighbors will see and do at $t = 2$. Accordingly, the relative responsiveness of each player i's strategy to their signal θ_i will depend not only on their quality of information e_i, but also on each neighbor j's equilibrium information investment and corresponding strategic responsiveness to their own signal, θ_j. Crucially, the resulting intricate interdependence of information responses is introduced precisely when players' payoffs are correlated through the common state ω: when $\gamma > 0$.

Formalizing the discussion, define the correlation adjusted adjacency matrix as:

$$\begin{aligned} \Sigma^c &:= \left[\gamma^2 e_i \sigma_{ij} e_j\right]_{i,j;i\neq j} \\ &= \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}, \end{aligned} \tag{4}$$

where $\mathbf{I}_{\boldsymbol{\phi}}$ denotes the diagonal matrix with entries given by (generic) vector $\boldsymbol{\phi}$.[11] Then, when $(\mathbf{I}-\Sigma)$ and $(\mathbf{I}-\Sigma^c)$ are invertible[12] the following unique linear second-stage solution obtains.

Theorem 1. [linear IRE] *For any* $\mathbf{e}$ *and sequentially rational* $\boldsymbol{\mu}^*$ *there exists a unique linear IRE of the form:*

$$\begin{aligned} \mathbf{X}^* &= (\mathbf{I}-\Sigma)^{-1}\mathbf{a} + \mathbf{I}_{\boldsymbol{\theta}}(\mathbf{I}-\Sigma^c)^{-1}\mathbf{e} \\ &= \left[\alpha_i^* + \beta_i^* \theta_i\right], \end{aligned} \tag{5}$$

denoting:

$$\begin{aligned} \boldsymbol{\alpha}^* &:= (\mathbf{I}-\Sigma)^{-1}\mathbf{a}, \\ \boldsymbol{\beta}^* &:= (\mathbf{I}-\Sigma^c)^{-1}\mathbf{e}. \end{aligned}$$

Note that $\boldsymbol{\alpha}^*$ is independent of $\mathbf{e}$, while $\boldsymbol{\beta}^*$ is a function of the vector of qualities chosen in the first stage. As shown in Appendix F.2 with the theorem's proof, IRE $\mathbf{X}^*$ is the unique equilibrium in a broad class of strategies that yield convergent higher-order expectations across players in the network.

A valuable interpretation of Theorem 1 utilizes the notion of *weighted Bonacich centrality* (Bonacich (1987) [7]). Formally, for given $N \times N$ graph $\mathbf{G} := [g_{ij}]$ of interaction terms (with zero diagonal) and weighting vector $\boldsymbol{\phi} \in \mathbb{R}^n$, the weighted Bonacich centrality measure is

[10] $\boldsymbol{\alpha}^*$ corresponds to the solution of Ballester et al. (2006) [5] but in expectation.

[11] $\mathbf{I_e}\Sigma\mathbf{I_e}$ is referred to as a Hadamard product of $[e_i e_j]$ with Σ, named after Jacques Salomon Hadamard (1865 –1963).

[12] Assumption F1 in Appendix F.2 provides a weak sufficient condition for this to hold.

defined as:

$$\begin{aligned}\mathbf{b}(\mathbf{G},\boldsymbol{\phi}) &:= (\mathbf{I}-\mathbf{G})^{-1}\boldsymbol{\phi} \\ &= \sum_{\tau=0}^{\infty}\mathbf{G}^{\tau}\boldsymbol{\phi}.\end{aligned}$$

This measure is well defined provided $(\mathbf{I}-\mathbf{G})$ is invertible. Each i'th component of $\mathbf{b}(\mathbf{G},\boldsymbol{\phi})$ gives an aggregation of the total number of paths starting from player i, with sub-paths emanating from each player j weighted by ϕ_j.[13] While the matrix $\mathbf{G}$ provides a benchmark network of bilateral relationships, the components of $\boldsymbol{\phi}$ capture each player's relative prominence within the network.

Placing this centrality concept into the context of Theorem 1, we see that ex ante expected actions are proportional to players' Bonacich centrality on Σ weighted by the vector of constants $\mathbf{a}$, $\mathbf{b}(\Sigma,\mathbf{a})$. The strategic response of each player's strategy to her private information also depends on her centrality. However, the centrality of interest is now (i) adjusted for the correlation of players' signals, and (ii) weighted by the vector of signal qualities $\mathbf{e}$. For the former, scaling down links by signal correlations adjusts for the inference of neighbors' second-stage actions. For the latter, weighting the resulting Bonacich centrality measure by $\mathbf{e}$ accounts for the value that information carries toward directly inferring the payoff-relevant state, $\tilde{\omega}_i$. The resulting alternative measure of centrality, or *informational centrality*, resonates with the unweighted Bonacich centrality $\mathbf{b}(\Sigma,\mathbf{1})$ directly derived from the network Σ. $\mathbf{b}(\Sigma^c,\mathbf{e})$ instead offers an adjusted measure of player position in the information-response game.

In light of Theorem 1 and as a technical note, scaling γ is analytically equivalent to uniformly scaling each term in Σ, via the product $\gamma^2 e_i\sigma_{ij}e_j$ in (4). Much of the following analysis will consider small or bounded values of γ. Thus, with γ directly scaling links in Σ^c, this can be aptly interpreted as taking moderately sized peer effects in the information-response game.

As seen with (5), the ways in which players respond to their information in the unique linear IRE depends in an intricate way on the network of peer effects and on the acquired signal qualities of neighbors.[14] The following begins to characterize the collective incentives to acquire information, providing a necessary condition for any interior first-stage strategy.

Theorem 2. **[IAE and information response]** *Given signal quality profile $\mathbf{e}_{-i}$, player i's private marginal gain to signal quality e_i is given by β_i^2/e_i, yielding the necessary condition for any IAE $\mathbf{e}^*$:*

$$\frac{\beta_i^{*2}}{e_i^*} = \kappa'(e_i^*), \tag{6}$$

for each i with $e_i^* \in (0,1)$.[15]

With $e_i\kappa(e_i)$ an increasing function in e_i, $\mathbf{e}^*$ is thus ordered with respect to the size of

[13] Other variations of this centrality measure are defined with weighted walks starting from neighbors (see Jackson (2008) [39]), while this definition's weighting begins at the originating node.

[14] Or more precisely, the sequentially rational beliefs regarding the signal qualities of others.

[15] The existence of an IAE is established with Proposition S.1 in Supplemental Section 19.

players' informational centrality, $|\beta^*|$. Intuitively, we should expect the responsiveness of each player's strategy to their signal to be proportional –in some way– to the quality of her information, regardless of the presence of peer effects. Theorem 2 affirms this intuition.

Next, Corollary 1 ties player degree with their incentives to acquire information under moderate peer effects. It describes the speeds and directions that players diverge from $e^\dagger$ as peer effects are introduced.

Corollary 1. *Under Assumption 1, the following limit obtains:*

$$\lim_{\gamma \to^+ 0} \frac{\partial e_i^*}{\partial(\gamma^2)} = \frac{e^{\dagger 2} \sum_{k \neq i} \sigma_{ik}}{\kappa'(e^\dagger) - 1}. \tag{7}$$

As γ departs from zero, or as peer effects are introduced, players with the highest degree depart upward away from quality $e^\dagger$ relatively faster than those with lower degree. The speed at which players adjust their qualities decreases in the concavity of $\kappa(e^\dagger)$, which measures the sensitivity in marginal gains to information around $e^\dagger$. This speed increases in $e^{\dagger 2}$, which measures the initial marginal informational content that signals provide toward inferring neighbors' second-stage observations.

As will be observed in Sections 16.3, the degree-wise ordering in $\mathbf{e}^*$ that is implied by Corollary 1 may not persist as γ is further increased. That is, while degree describes players' initial incentives to acquire quality, it does not fully determine these incentives when peer effects are more pronounced. The ease and extent to which this ordering may be violated will intimately depend on both the network's structure and the shape of κ. Supplemental Section 19 further explores this relationship, and develops network properties that allow player degree to persistently order the equilibrium extent of information acquisition.

16.3 Examples

The following examples illustrate the breadth of equilibrium properties in this setting. The first example illustrates the potential for multiple IAE, even under the unique IRE given with Theorem 1. Multiplicity can arise under either strategic complements or strategic substitutes.

Example 1. *For this and subsequent examples we consider the following strictly convex information cost function:*

$$\kappa(e) = K \frac{e^{\eta_1}}{(1-e^2)^{\eta_2}},$$

where $\eta_1 \geq 2$ and $\eta_2 > 0$. It can be shown that Assumption 1 is satisfied under these bounds, yielding isolation quality $e^\dagger$. This functional form provides a standard family of convex costs functions that asymptote as $e_i \to^- 1$. It also allows for a broad range of convexities. Crudely, increasing η_1 increases convexity at higher values of e_i while increasing η shifts convexity toward lower values of e_i.

First, consider any regular network in which each player i is connected to four other players with symmetric peer effects $\sigma_{ij} = \sigma_{ji} = p > 0$ (for neighbor j). Normalize $\gamma = 1$, and set $\eta_1 = 3$, $\eta_2 = .5$ and $K = .312$. Figure 2(a) provides the set of symmetric equilibria of the information acquisition game as we increase the size of p above. In these examples, *information responses are given by the increasing relationship* $\beta_i^* = \sqrt{e_i^* \kappa'(e_i^*)}$ *from Theorem 2.*

Under second stage strategic complements, first stage information acquisition reinforces itself. The added convexity in the value of information introduces potential coordination problems. For values of p *above* $5/22$ *the players can coordinate on low, medium, or high levels of information. Higher levels of information further incentivize acquisition as the players' signals correlate with each other.*

[Figure 2]

Next, consider the two player network of players 1 *and* 2*, with symmetric negative peer effect* $\sigma_{12} = \sigma_{21} = -p \leq 0$. *Set* $\eta_1 = 2$, $\eta_2 = 1$, *with* $K = .03$.[16] *Now, the propensity for an asymmetric equilibrium arises, as seen in Figure 2(b). For values of* p *below* .85 *the symmetric equilibrium of the information acquisition game gives the unique IAE (solid line). Information responses are again given using* $\beta_i^* = \sqrt{e_i^* \kappa'(e_i^*)}$. *Above* $p = .85$ *there also exists an asymmetric equilibrium in which one player acquires a highly precise signal while the other acquires an imprecise signal (dashed lines). It can be verified that in this equilibrium the low-quality player rather prefers the symmetric equilibrium, while the high-quality player strictly prefers her equilibrium informational advantage.*

The next example exhibits the potential for players to move against their information given substantial negative peer effects. Precisely, for players facing enough negative influence from others, the incentives to acquire information may increase with the size of these influences, but with these players moving against their signals in anticipation of their neighbors' actions. Strikingly, this non-monotonicity can be quite significant, with the incentives to acquire information quickly falling to zero and abruptly restoring itself at more extreme influences.

Example 2. *Again, normalize* $\gamma = 1$. *Take the wheel and spoke network with center player* 1 *and peripheral players* $i \in \{2, 3, 4\}$, *as depicted in Figure 3(left).*

[Figure 3]

Each player imposes a symmetric negative externality on 1*:* $\sigma_{1i} = \sigma_{i1} = -p \leq 0$ *for each* $i \in \{2, 3, 4\}$. *Finally, the peripheral players are symmetrically linked in a circle with weights* $\sigma_{ij} = 1/3$ *for each pair* $i, j \in \{2, 3, 4\}$. *Take the information cost function given in Example 1, setting* $\eta_1 = 2$ *and* $\eta_2 = 1$ *with* $K = .03$.

Figure 3(right) plots qualities e_i^* *and responses* β_i^* *in the unique equilibrium of the information-response game symmetric across the peripheral players 2-4 ('per.'), over a range of* p *values. As* p *departs from zero, information acquisition drops slightly for the center,*

[16] K is adjusted down with the new values of η_1 and η_2 to obtain interior solutions, with the latter set so that the qualitative properties of the equilibrium are well displayed.

as the negative externalities between the peripheral clique and the center increase. For values of p between 0.111 *and* 0.183 *the center acquires no information. Then, for values of p above* 0.183 *the incentives to acquire information are quickly restored. However, now the center moves against the network, with* 1 *acquiring information and moving opposite to her signal (i.e.* $\beta_1^* < 0$*) in anticipation of the actions of the periphery. As negative externalities become more acute, further information is incentivized, with 1's behavior further reinforcing information acquisition throughout the network.*

Example 2 highlights the potential for non-monotonicity in information acquisition. The responsiveness of each player to her private information in an IRE $\mathbf{X}^*$ remains unbounded at the origin. Players with intermediate centrality in the information-response game face moderate incentives to acquire information in period $t = 1$. And as illustrated in Figure 3, such non-monotonicity need not be gradual, but rather the incentives to acquire information can quickly vanish for players with particularly low centrality. Then, for ever lower levels of centrality the incentives to acquire information may be restored to great extent, but now with information used to infer and move against neighbors' second-stage actions.

Both multiplicity and negative signal responses come as interesting equilibrium incarnations. None the less, the following establishes sufficient conditions for the exclusion of these cases. Under moderate peer effects, a unique equilibrium in which players move in the direction of their signals always obtains.

Proposition 1. *Under Assumption 1:*

1. *there is some* $\gamma^u > 0$ *such that if* $\gamma \in [0, \gamma^u)$ *a unique IAE exists, and*
2. *there is some* $\gamma^p > 0$ *such that if* $\gamma \in [0, \gamma^p)$ *all players acquire qualities in* $(0, 1)$ *and respond positively to their signals (i.e.* $\beta_i^* > 0$ *for each i) in equilibrium.*

17 Equilibrium welfare and the strategic value to information

The welfare analysis takes the following approach. First, we will see that when allowing players' to respond optimally to signal realizations, correlation in signals is necessary for the presence of inefficiencies in information acquisition. That is, externalties and strategic motives arise in the first stage only when players can use their information in the second stage to infer the observations of neighbors. Departing from the case of zero correlation (i.e. $\gamma = 0$), strong welfare statements are derived given moderately sized peer effects (i.e. small γ). These results address the directions of both (i) the equilibrium profile of information qualities when information investments are publicly observed and (ii) the utilitarian solution relative to the equilibrium described in Theorem 2. We then turn to more significantly sized peer effects, incorporating the welfare implications of players moving against their signals and under the potential for multiple information acquisition equilibria.

In the following welfare benchmark, we take the second-stage information response equilibrium $\mathbf{X}^*$ –a function of qualities $\mathbf{e}$– as given. Further, given quality profile $\mathbf{e}$ we impose sequential rationality in beliefs throughout: $\mu_i^{j*}(e_j) = 1$ for each i and j. That is, the planner is free to publicly announce the information qualities that she prescribes. This prevents inefficiencies derived from inconsistent beliefs. Incorporating these elements into our benchmark leaves first-stage behavior as the sole endogenous (potential) source of inefficiencies.

As a function of the realized quality profile $\mathbf{e}$, and given IRE $\mathbf{X}^*$ and sequentially rational beliefs $\boldsymbol{\mu}^*$, players' ex ante values reduce as follows:

$$\begin{aligned}\boldsymbol{\nu}\left(\mathbf{X}^*|\mathbf{e}\right) &:= \left[\mathbb{E}_i[u_i\left(\mathbf{X}^*|\theta_i, e_i, \mu_i^*\right)|e_i, \mu_i^*] - \kappa\left(e_i\right)\right] \\ &= \frac{1}{2}\left(\mathbf{I}_{\boldsymbol{\alpha}^*}\boldsymbol{\alpha}^* + \mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^*\right) - \kappa\left(\mathbf{e}\right).\end{aligned} \tag{8}$$

This reduction to quadratic payoffs is easily shown in Appendix F.3. Then taking $\nu\left(\mathbf{X}^*|\mathbf{e}\right)$ we can define the following Pareto problem:

$$\max_{\mathbf{e}\in[0,1]^N} \sum_k \lambda_k \nu_k\left(\mathbf{X}^*|\mathbf{e}\right), \tag{9}$$

for non-negative Pareto weights $\boldsymbol{\lambda}$ taken from the $(N-1)$-simplex. First order conditions yielding the planner's solution $\mathbf{e}^{po}\left(\boldsymbol{\lambda}\right)$ are given for each $i \in N$ by:

$$\sum_k \lambda_k \frac{\partial}{\partial e_i}\nu_k\left(\mathbf{X}^*|\mathbf{e}\right) = 0. \tag{10}$$

The following establishes correlation in the players' payoffs as necessary for any equilibrium inefficiencies that may arise.

Proposition 2. *At $\gamma = 0$ we have $e_i^* = e_i^{po}\left(\boldsymbol{\lambda}\right)$, and $e_i^* = e^\dagger$ under Assumption 1, for each i.*

Proof. With $\boldsymbol{\alpha}^*$ independent of $\mathbf{e}$, (10) can be written:

$$\begin{aligned}\sum_k \lambda_k \frac{\partial}{\partial e_i}\nu_k\left(\mathbf{X}^*|\mathbf{e}\right) &= \lambda_i \frac{\partial}{\partial e_i}\nu_i\left(\mathbf{X}^*|\mathbf{e}\right) + \sum_{k\neq i}\lambda_k \frac{\partial}{\partial e_i}\nu_k\left(\mathbf{X}^*|\mathbf{e}\right) \\ &= \lambda_i\left(\left(\frac{\partial}{\partial e_i}\mathbb{E}_i[u_i\left(\mathbf{X}^*|\theta_i, e_i, \mu_i^*\right)|e_i, \mu_i^*] - \kappa'\left(e_i\right)\right) + \beta_i^* \sum_{k\neq i}\gamma^2 e_i e_k \sigma_{ik}\frac{\partial \beta_k^*}{\partial e_i}\right) + \sum_{k\neq i}\lambda_k \beta_k^* \frac{\partial \beta_k^*}{\partial e_i},\end{aligned}$$

where the first term in brackets takes $\boldsymbol{\beta}_{-i}^*$ fixed. The first order condition of i's IAE problem is given by setting this term to zero. Now, $\boldsymbol{\beta}^* = \mathbf{b}\left([0], \mathbf{e}\right) = \mathbf{e}$ when $\gamma = 0$, and thus $\frac{\partial}{\partial e_i}\beta_k^* = 0$ for each $k \neq i$. Thus, when $\gamma = 0$ the Pareto optimal and IAE solutions align. Finally, $e_i^* = e^\dagger$ for each i under Assumption 1 follows from $\boldsymbol{\beta}^* = \mathbf{e}$ at $\gamma = 0$ and Theorem 2. $\square$

Under our general treatment of peer effects, it does not come surprisingly that equilibria are not generally Pareto efficient. We next begin to more completely describe the nature of inefficiency in the model. The following measures for the strategic value to information and informational externalities are required.

First, a loss in value to information is realized by each player i who, in equilibrium, is unable to directly influence others' beliefs regarding her information investment. Precisely, with qualities privately chosen at $t = 1$, incentive compatibility constrains each i when weighing the costs and benefits of acquiring information quality. If instead i could publicly invest in quality e_i and directly influence others' beliefs, she may derive additional value from acquiring more or less quality than in equilibrium (holding $\mathbf{e}^*_{-i}$ fixed). Informational externalities, on the other hand, are directly imposed on i's neighbors. Also derived from the influences that e_i has on neighbors' responses, these externalities are instead measured by the effect that e_i has on *neighbors'* welfare.

Formalizing this, consider the following utilitarian problem, given from (9) by setting $\boldsymbol{\lambda} = \frac{1}{N}\mathbf{1}$:

$$\max_{\mathbf{e}\in[0,1]^N} \sum_k \nu_k \left(\mathbf{X}^*|\mathbf{e}\right). \tag{11}$$

The partial derivative of aggregate welfare with respect i's quality is given by:

$$\begin{aligned} \frac{\partial}{\partial e_i} \sum_k \nu_k \left(\mathbf{X}^*|\mathbf{e}\right) &= \beta_i^* \frac{\partial \beta_i^*}{\partial e_i} + \sum_{k\neq i} \beta_k^* \frac{\partial \beta_k^*}{\partial e_i} - \kappa'\left(e_i\right) \\ &= \underbrace{\underbrace{\underbrace{\left(\frac{\beta_i^{*2}}{e_i} - \kappa'\left(e_i\right)\right)}_{=0 \text{ in IAE } \mathbf{e}^* \text{ f.o.c.}} + \beta_i^* \sum_{k\neq i} \gamma^2 e_i e_k \sigma_{ik} \frac{\partial}{\partial e_i}\beta_k^*}_{=0 \text{ in public acquisition equilibrium } \mathbf{e}^{pb} \text{ f.o.c.}} + \sum_{k\neq i} \beta_k^* \frac{\partial}{\partial e_i}\beta_k^*}_{=0 \text{ in planner's solution } \mathbf{e}^{pl} \text{ f.o.c.}} . \end{aligned} \tag{12}$$

In IAE $\mathbf{e}^*$, where qualities are privately acquired, the first term in brackets is set to zero by i in her optimization problem.[17] That is, the term β_i^{*2}/e_i is given by β_i^* multiplied by the marginal influence of i's quality on her own response $\partial\beta_i^*/\partial e_i$, while setting the marginal influence of i's information quality on others' responses $\partial\beta_k^*/\partial e_i$ to zero. This corresponds with Theorem 2. The middle sum adjusts for the marginal effect that i's information quality imposes on each k's response in IRE $\mathbf{X}^*$, *when* the acquisition of e_i is directly observed by each $k \neq i$. This term is excluded in IAE under i's incentive compatibility constraint, again where e_i^* is chosen fixing μ_k and in turn β_k^* for each $k \neq i$. If instead i were free to publicly choose her signal quality in the first stage she would internalize this influence. The term captures i's *strategic* incentive toward influencing her neighbors' information responses. Setting these terms in (12) to zero for each i yields the *public* information acquisition equilibrium $\mathbf{e}^{pb}$. Finally, in the planner's problem the direct marginal influence that i's quality carries for *others'* payoffs is also accounted for. That is, when total marginal gains to welfare from i's quality is set to zero, we obtain one of N first order conditions that determine the *planner's* solution $\mathbf{e}^{pl} := \mathbf{e}^{po}(\frac{1}{N}\mathbf{1})$.[18] Together, the final two terms adjust for the effect that i's quality has on welfare that is not internalized by i in the first stage.

[17] Note that the private information acquisition benchmark is equivalent to a one-stage game in which players simultaneously choose information qualities and information contingent strategies.

[18] This planner's benchmark $\mathbf{e}^{pl}$ is commonly referred to as the "*second-best team*" solution, with the "*first-best team*" or "*team-efficient*" solution determined when the planner can also control how players use their information. See Burguet and Vives (2000) [12] or Vives (2008) [62] chapter 6 for discussions.

Thus, we can define the following measures:

$$\begin{aligned}
\xi_i^{st}\left(\mathbf{e},\mathbf{X}^*\right) &:= \beta_i^* \sum_{k\neq i} \gamma^2 e_i e_k \sigma_{ik} \frac{\partial}{\partial e_i}\beta_k^*, && (13)\\
\xi_i^{ex}\left(\mathbf{e},\mathbf{X}^*\right) &:= \sum_{k\neq i} \beta_k^* \frac{\partial}{\partial e_i}\beta_k^*, && (14)\\
\xi_i\left(\mathbf{e},\mathbf{X}^*\right) &:= \xi_i^{st}\left(\mathbf{e},\mathbf{X}^*\right) + \xi_i^{ex}\left(\mathbf{e},\mathbf{X}^*\right). && (15)
\end{aligned}$$

We refer to $\xi_i^{st}\left(\mathbf{e},\mathbf{X}^*\right)$ as i's marginal *strategic value* to information at quality profile $\mathbf{e}$. At IAE $\mathbf{e}^*$, it informs us in which direction i would deviate in the first stage *if* her quality investment were publicly observed by $t = 2$. In addition, it tells us how much i would be willing to pay (in utils) per unit of quality if she could directly influence her neighbors' beliefs of e_i. We refer to $\xi_i^{ex}\left(\mathbf{e},\mathbf{X}^*\right)$ as i's marginal *informational externalities* at quality profile $\mathbf{e}$. This marginal cost would not be internalized in the event that i could publicly choose her information quality. Taking these marginal costs together, $\xi_i\left(\mathbf{e},\mathbf{X}^*\right)$ gives the sum of i's marginal strategic value and informational externalities. We term this the marginal *public value* from quality e_i (at $\mathbf{e}$). When evaluated at IAE $\mathbf{e}^*$, the vector $\boldsymbol{\xi}\left(\mathbf{e}^*,\mathbf{X}^*\right)$ evaluates the equilibrium *gradient* of $\sum_k \nu_k\left(\mathbf{X}^*|\mathbf{e}\right)$, pointing in the direction of the social planner's optimal deviation from the equilibrium quality profile $\mathbf{e}^*$.

Closed forms of these measures are derived in Appendix F.3. In accordance with Proposition 2, all equate to zero at $\gamma = 0$, when private signals are uninformative of others' second-stage observations and responses. When $\gamma > 0$ this no longer holds. For moderate peer effects, $\xi_i^{st}\left(\mathbf{e},\mathbf{X}^*\right)$ and $\xi_i^{ex}\left(\mathbf{e},\mathbf{X}^*\right)$ are proportional to the following network measures.

Lemma 1. [limiting marginal inefficiencies] *The following limits obtain:*

$$\begin{aligned}
\lim_{\gamma\to^+0} \frac{\partial \xi_i^{st}\left(\mathbf{e}^*,\mathbf{X}^*\right)}{\partial(\gamma^4)} &= 2e^{\dagger 5}\sum_{k\neq i}\sigma_{ik}\sigma_{ki}, && (16)\\
\lim_{\gamma\to^+0} \frac{\partial \xi_i^{ex}\left(\mathbf{e}^*,\mathbf{X}^*\right)}{\partial(\gamma^2)} &= 2e^{\dagger 3}\sum_{k\neq i}\sigma_{ki}. && (17)
\end{aligned}$$

That is, the rate of increase of i's marginal strategic value as γ^4 departs from zero is proportional to $(e^\dagger)^5$ multiplied by the sum-of-products of i's peer effects (i.e. the sum of her out-links multiplied by their respective in-links). The rate-of-increase of i's marginal externalities as γ^2 departs from zero is approximately $(e^\dagger)^3$ multiplied by i's in-degree.

The intuition behind these limits goes as follows. Each i's marginal strategic value as the network of peer effects is pronounced in the information-response game depends on both i's outward and inward directed links. Outward links measure the extent to which i *cares* about each of her neighbor's second-stage actions. Inward links measure the influence that i's quality has on each neighbors' payoffs. Together, the neighbor-wise product of links scale i's marginal strategic value to her information. Marginal externalities, on the other hand, depend only on the influence that i's quality has on others. Precisely, marginal externalities scale by the sum of influences imposed on the network. For moderately size peer effects, this is propotional to i's in-degree.

As in (7), the sizes of these limiting derivatives depend on the initial level of information acquisition, $e^{\dagger}$. $e^{\dagger}$ scales the initial strategic responsiveness of strategies, as well as the initial extent to which players can infer others' signal realizations from their own signals. With strategic values involving the additional inference by neighbors of i's signal realization, (16) scales with an additional factor of $e^{\dagger 2}$.

To explore the broader implications of Lemma 1, we next focus in on the set of symmetric networks. This family of network architectures offers a broad and flexible class of familiar environments.

17.1 Symmetric pairwise peer effects and welfare

Here we further describe the nature of inefficient information acquisition by focusing on symmetric network structures. This is primarily done as symmetry is commonly observed in many real-world peer-effects environments. Be them competitive or cooperative, most relationships in society tend to be reflexive, in both direction and size. Competitors tend to be competitors, while collaborators can find a sometimes delicate balance of cooperative synergies. Symmetric peer-effects networks represent environments in which individual pairs can be either competitive or cooperative, and at various extents. As will be shown, such networks carry with them a natural tendency for strategic information acquisition.

First, we show that marginal strategic values borne by players interacting under symmetric and moderate peer effects are positive. This holds regardless of other network details. When influences between player pairs balance with each other, revealing and even exaggerating one's signal quality (if this were feasible) unambiguously increases private payoff. Remarkably, both positive and negative links reinforce this effect.

Secondly, we show that in these environments, the equilibrium response to the network of peer effects is *weak* relative to the utilitarian benchmark. This is manifested as an inefficient dispersion in $\mathbf{e}^*$. When both positive and negative links are present, this can imply that the most informed players under acquire information while the least informed players over acquire.

Formally, we consider the following family of network structures.

Assumption 2A. *Σ is symmetric: $\sigma_{ij} = \sigma_{ji}$ for each $i \neq j$.*

Taking (16) under the symmetry of Assumption 2A, each i's marginal strategic value positively scales with her *sum-of-squared* degree: $\sum_k \sigma_{ik}^2$. As such, $\xi_i^{st}(\mathbf{e}^*, \mathbf{X}^*)$ is strictly positive in symmetric, connected networks under moderate peer effects. Both positive and negative links reinforce the size of i's marginal strategic value to information. And with (17), each i's marginal externalities positively scale with her in-degree, which under Assumption 2A equates with her out-degree. With these measures taking on clear directions under symmetric, moderately-sized peer effects, the following can be shown.

Proposition 3A. **[symmetric, moderate peer effects]** *For symmetric Σ, there exists some γ^w with $0 < \gamma^w \leq \min\{\gamma^m, \gamma^s\}$ such that if $\gamma \in (0, \gamma^w)$ and for $\mathbf{e}^*$, $\mathbf{e}^{pb}$ and $\mathbf{e}^{pl}$ we have*[19]*:*

[19] For 1., we assume Σ to have no isolated players: $\sigma_{ij} \neq 0$ for some j for every i.

1. $\mathbf{e}^{pb} > \mathbf{e}^*$, *with* $({e_i}^{pb} - {e_i}^*) > ({e_j}^{pb} - {e_j}^*) > 0$ *for each* i *and* j *with* $\sum_{k \neq i} \sigma_{ik}^2 > \sum_{k \neq j} \sigma_{jk}^2$,
2. $e_i^* > e_j^*$ *and* $({e_i}^{pl} - {e_i}^*) > ({e_j}^{pl} - {e_j}^*)$ *for each* i *and* j *with* $\sum_{k \neq i} \sigma_{ik} > \sum_{k \neq j} \sigma_{jk}$.

From 1., players are *disincentivized* to acquire information as a result of incentive compatibility constraints. If players' could convincingly persuade others of their first-stage actions, they would always exaggerate their informativeness. The relative strength of this incentive scales with each player's sum-of-squared degree. With 2., the planner's optimal deviation from IAE $\mathbf{e}^*$ entails increases to signal qualities to higher degree players that are no less than increases prescribed to players with lesser degree. With equilibrium qualities similarly ordered according to player degree (for moderate peer effects), the *asymmetry in acquired information qualities are inefficiently low* as a result of marginal externalities. That is, the players' equilibrium information qualities are "*bunched*". And if all links in the network of peer effects are non-negative (non-positive), then $\xi_i^{ex}(\mathbf{e}^*, \mathbf{X}^*)$ will be non-negative (non-positive) with the most informed players imposing the greatest externalities. When both strategic complements and substitutes exist in the network, the ordering provided in Proposition 3A.2 establishes the more general result.

The economic interpretation of parts 1. and 2. in Proposition 3A are more broadly described as follows. For part 1., consider a player i with both positive and negative links with other players. If i could publicly acquire additional quality, this would encourage the responsiveness of positively linked neighbors, and simultaneously discourage the responsiveness of her negatively linked neighbors. Such directed influences are precisely due to the correlation in signals: learning that ω is likely high also informs i's neighbors that i likely observes similar information and will respond accordingly. These directional influences strictly work in i's favor regardless of the sign of her link with j. The symmetry in each pair's relationship implies a clear direction in these incentives. Thus, a player's *connectedness* in a symmetric network determines the size of the marginal strategic value to her information.[20]

With part 2., the network of peer effects can more broadly be interpreted as simultaneously quantifying the sizes and directions of externalities in the economy (in-links), as well as the sizes and signs of network effects imposed on each player (out-links). Externalities and network effects balance in symmetric networks. Thus, those that respond most positively to the network –through their information investments– are precisely those that endow the most value upon others from acquiring their information. And those that respond most negatively are precisely those that impose the most negative externalities upon others. Thus with respect to the utilitarian benchmark, players collectively under respond to a symmetric network of peer effects.

All of the above equilibrium properties are illustrated with the following example.

[20] Hauk and Hurkens (2001) [33] obtain a similar under acquisition in homogenous Cournot markets. In the network setting, a player i's connectedness –sum-of-squared degree– scales the size of her under acquisition arising from the privacy of e_i.

[Figure 4]

[Figure 5]

Example 3. *Take the network structure given in Figure 4, having three classes of players comprised of the center triad (class "x"), outer triad (class "z"), and three players bridging the two triads (class "y"). A general definition of player classes is provided in Supplemental Section 19. Here, γ is set to* 1.

Taking the cost function from Example 2 with $\eta_1 = 2$, $\eta_2 = 1$, and $K = .1$, we consider the unique equilibria symmetric across players within each class. Equilibrium qualities e_c^*, differences $(e_c^{pl} - e_c^*)$, and marginal strategic value $\xi_c^{st}(\mathbf{e}^*, \mathbf{X}^*)$ are provided in Figure 5 for each class $c \in \{x, y, z\}$ over a range of p values. At $p = 0$ peer effects include only complements. Accordingly, externalities remain positive for all classes over a range of small p. As competition between classes y and z heightens, class z's $(e_z^{pl} - e_z^*)$ drops below zero. Marginal strategic value, on the other hand, unambiguously rises for classes y and z as these players place additional weight on each other.

With negative links representing inter-player competition, the incentives of low informational centrality players to distort the beliefs of more central neighbors –as to discourage their information responses– only heighten with great inter-class competition. While marginal externalities derive the majority of the marginal public value to e_i, marginal strategic value continues to capture and describe the incentives to distort beliefs. If strategic substitutes are significant for some players, the miss-orientation between the planner's and these players' preferences magnifies with greater competition.

17.2 Network asymmetries and welfare

Here we further explore the ramifications of network symmetry. We first consider analogues of the above results in networks with anti-symmetric pairwise peer effects. These *anti-symmetric* networks provide the opposite extreme to symmetric networks. As illustrated with the application to financial markets in liquidity crises of Section 18.1, such anti-symmetry in pairwise relationships may pervade a market when traders face asymmetric constraints in the second stage.

Formally, consider the following condition on Σ:

Assumption 2B. *Σ is anti-symmetric: $\sigma_{ij} = -\sigma_{ji}$ for each $i \neq j$.*

That is, for each peer effect the opposite-pointing effect gives the opposite-signed relationship. We refer to these pairwise relationships as *anti-symmetric*. Here, the natural analogue to Proposition 3A obtains.

Proposition 3B. [anti-symmetric, moderate peer effects] *For anti-symmetric Σ, there exists some γ^w with $0 < \gamma^w \leq \min\{\gamma^m, \gamma^s\}$ such that if $\gamma \in (0, \gamma^w)$ and for $\mathbf{e}^*$, $\mathbf{e}^{pb}$ and $\mathbf{e}^{pl}$ we have:*

1. *$\mathbf{e}^{pb} < \mathbf{e}^*$, with $({e_i}^{pb} - {e_i}^*) < ({e_j}^{pb} - {e_j}^*) < 0$ for each i and j with $\sum_{k \neq i} \sigma_{ik}^2 > \sum_{k \neq j} \sigma_{jk}^2$,*
2. *$e_i^* > e_j^*$ and $({e_i}^{pl} - {e_i}^*) < ({e_j}^{pl} - {e_j}^*)$ for each i and j with $\sum_{k \neq i} \sigma_{ik} > \sum_{k \neq j} \sigma_{jk}$.*

In this setting, players face opposite strategic incentives. They now face the incentives to *understate* their informativeness: to *"play dumb"*. IAE now exhibit over-dispersion under moderate peer effects. In these networks, the most informed players will tend to over acquires while the least informed players will under acquire.

But, what if the network is neither purely symmetric nor anti-symmetric? With the strategic use of information taking extremes under symmetric and anti-symmetric networks, their manifestation in networks with both symmetric and anti-symmetric relationships may be less pronounced. The following example explores this more general setting.

Example 4. *First consider the two-player directed network where player 1 faces strategic substitutes in 2's action, $\sigma_{12} = -p < 0$, while player 2 faces strategic complements in 1's action of equal size, $\sigma_{21} = p$. Then, one can derive an exact expression for marginal strategic values:*

$$\xi_i^{st}(\mathbf{e}^*, \mathbf{X}^*) = -\gamma^4 2 \frac{\beta_i^{2*}}{e_i^*} p^2 e_1^{*2} e_2^{*2},$$

for $i = 1, 2$. That is, both players face the incentive to understate their information investment, in accordance with Proposition 3B.1. Precisely, player 1 has the incentive to understate her quality as to encourage 2's information investment. On the other hand, player 2 faces a similar incentive, but rather in order to discourage player 1's information investment.

Now consider the extended network in which player 2 is positively and symmetrically influenced by a player 3: $\sigma_{23} = \sigma_{32} = q > 0$. The structure of peer effects is offered in Figure 6. One can similarly derive:

$$\begin{aligned}
\xi_1^{st}(\mathbf{e}^*, \mathbf{X}^*) &= -\gamma^4 2 \beta_1^{2*} p^2 e_1^* e_2^{*2}, \\
\xi_2^{st}(\mathbf{e}^*, \mathbf{X}^*) &= \gamma^4 2 \beta_2^{2*} e_2^* \left(q^2 e_3^{*2} - p^2 e_1^{*2}\right), \\
\xi_3^{st}(\mathbf{e}^*, \mathbf{X}^*) &= \gamma^4 2 \beta_3^{2*} q^2 e_3^* e_2^{*2}.
\end{aligned}$$

Thus, player 2 may no longer face significant marginal strategic value to her acquired information if $q^2 e_3^{*2} \approx p^2 e_1^{*2}$ in IAE $\mathbf{e}^*$.

[Figure 6]

We see that environments that couple symmetric and anti-symmetric relationships carry ambiguous strategic motives. When positive strategic values induced by symmetric relationships counterbalance negative strategic values induced by asymmetric relationships, players may be left without a unidirectional motive to influence others' beliefs. The private investment of information simultaneously imposes positive and negative strategic motives behind information acquisition. The net result is left as a function of each particular player's position in the networks of directed peer effects.

Next we address welfare implications when peer effects are large, incorporating negative information responses and multiple equilibria.

17.3 General peer effects and welfare

This section extends our welfare analysis to include more significant peer effects, and incorporates the potential for negative signal responses and multiple equilibria. As we will see, the observed U-shaped non-monotonicity in the incentives to invest in information carries over to externalities. As suggested throughout the preceding sections, the essential structural property driving the direction of the utilitarian optimum will be the extent of symmetry or anti-symmetry in pairwise relationships. We continue by taking Assumptions 2A and 2B as extremal benchmarks to pairwise symmetry and anti-symmetry (resp.) through the network. While clearly most real-world networks may not align exactly with one of these two cases, the following welfare properties can be applied by considering the extent of symmetry at a local level for sub-components of an observed peer-effects network.

To derive Lemma 1, Appendix F.3 takes the geometric expansions of the closed forms of $\boldsymbol{\xi}^{st}$ and $\boldsymbol{\xi}^{ex}$, respectively. Then taking their leading terms –which dominate their respective sums for small γ– the limits (16) and (17) are established. While affording formal proofs of Propositions 3A and 3B under moderate peer effects, these leading terms remain useful in assessing the directions of informational externalities and strategic values in the network. As derived in Appendix F.3, the approximations to $\boldsymbol{\xi}^{st}$ and $\boldsymbol{\xi}^{ex}$ for symmetric networks are given as:

$$\xi_i^{st}(\mathbf{e}^*, \mathbf{X}^*) \approx 2\frac{\beta_i^{*2}}{e_i^*}\gamma^4 \sum_{k \neq i} e_i^{*2} \sigma_{ik}^2 e_k^{*2}, \tag{18}$$

$$\xi_i^{ex}(\mathbf{e}^*, \mathbf{X}^*) \approx 2\frac{\beta_i^*}{e_i^*}\left(\beta_i^* - e_i^*\right), \tag{19}$$

for each i. And for anti-symmetric networks the negations of these corresponding approximations obtain.

In symmetric networks, we see that (18) is strictly positive for $\beta_i^* \neq 0$, consistent with Proposition 3A.1. We can assess (19), on the other hand, using Figure 7(a). In the top panel e_i^* is graphed against β_i^*. The exact form of this relationship is implicitly defined with expression (6) of Theorem 2. For any given Σ and in any IAE $\mathbf{e}^*$, the players will be spread across the domain at various points, yielding each i's e_i^*. Below this, the approximation (19) is plotted. With the exact form of marginal externalities scaling with signal response β_i^*/e_i^*, these marginal costs always pass through the origin. When $\beta_i^* = e_i^* = e^\dagger$, (19) again obtains a value of zero.

[Figure 7]

Thus, we obtain a *reversal* in the sign of marginal externalities when players move against their information. Non-monotonicity in the private value of information extends to the public value of information. For $\beta_i^* < 0$, the second-stage optimality of i's negative response implies that the value she derives from strategically moving against her signal outweighs the value from inferring and responding with her expectation of $\tilde{\omega}_i$. This is precisely because in IRE $\mathbf{X}^*$, the network imposes significant cost to i if she moves in the direction of her information. In symmetric networks and when $\beta_i^* < 0$, this cost translates to value imparted to i's competitors: to each j with $\sigma_{ji} < 0$ *and* $\beta_j^* > 0$. And with i failing to internalize this positive externality, she *under* acquires information relative to the efficient benchmark.

This reversal in the direction of the utilitarian solution relative to e_i^* can be illustrated with Example 2. p again gives the size of the negative links connecting the center player 1 to the peripheral players $\{2, 3, 4\}$. For p values below 0.111 player 1 moves in the direction of her signal realization, for values between 0.111 and 0.183 she acquires no information, and for values above 0.183 she moves against her signal in anticipation of the periphery's second-stage actions.

[Figure 8]

Figure 8 provides the planner's solution e_1^{pl} for the center (dashed line) along with IAE e_1^*. Below $p = 0.111$ player 1 over acquires information while facing positive marginal strategic values, as consistent with Proposition 3A. Internalizing marginal externalities on the periphery (as well as 1's marginal strategic values), the planner sends i's quality to zero *early*. Then for $p > .145$, $(e_1^{pl} - e_1^*)$ becomes positive with the planner setting β_1^* to be negative. Thus, player 1 under acquires information, and moves against her signal *late*. When player 1 finally starts moving against her signal (for $p > 0.183$) the gap between the planner's solution e_1^{pl} and e_1^* drops. Thus, the reversal in ξ_1^{ex} as β_1^* crosses the origin translates to a leftward horizontal shift in e_1^{pl}. While the corresponding figures for the periphery are omitted, $(e_{per.}^{pl} - e_{per.}^*)$ and $\xi_{per.}^{st}$ remain strictly positive and vary only mildly over the range of p values shown.

The economic message is noteworthy. When players acquire and move against their information in symmetric networks, the direction of this strategic behavior is socially efficient from a utilitarian perspective. But, the equilibrium extent to which these players invest in information is inefficiently *low.* The periphery now benefits from 1's negative response, and is only further encouraged to respond positively to their own private information. The value that such players create for others by acquiring and moving against their information is, once again, not internalized in equilibrium.

Returning to (18) and Figure 7(a), if ξ_i^{ex}'s leading term plays a dominant role in its sum, the exact form will shadow the depicted quadratic form. Inclusion of second order terms, or of marginal strategic values giving ξ_i –both of which will be positive away from the origin– give a more accurate approximation to the gradient of the utilitarian function. This higher-order approximation will (i) continue to cross the origin, with higher-order terms also scaled by β_i^*/e_i, and (ii) cross the β_i^*-axis (again) at some point to the left of $\beta_i^* = e^\dagger$.

The exact point at which ξ_i crosses the β_i^*-axis to the right of the origin designates the lower bound of the set of players that exhibit positive margin public values to e_i^*, while setting $\beta_i^* > 0$. This includes all i that set $\beta_i^* > e^\dagger$: region (III) in the figure. Players setting $\beta_i^* \in (0, e_i^*)$ in region (II) face negative marginal externalities up to some β_i^* left of $e^\dagger$. Finally, for players moving against their information, for $\beta_i^* < 0$ giving region (I), marginal externalities once again switch positive.[21]

For the hypothetical "*knife-edge*" player that sets $\beta_i^* = e_i^* = e^\dagger$, such an i must satisfy the equilibrium condition: $\sum_{k \neq i} \sigma_{ik} e_k^* \beta_k^* = 0$. That is, the sum of i's inferred network effects –the weighted sum of expected neighbors' responses– equals zero. Responding as an informed player within the network, such an i continues to use her information to infer the actions of neighbors. However, on net, i's incentives to strategically respond by adjusting her signal response upward or downward from $e^\dagger$ *perfectly* balance. From the outside observer, i behaves as though she acts in isolation. But in actuality, the net influence that the network imposes on her behavior equates with zero. And given the symmetry of the network, so must her total externality imposed on others. We term such an i the "*neutral player*".

Figure 7(b) provides the corresponding functions under an anti-symmetric network (Assumption 2B). While the equilibrium relationship providing e_i^* as a function of β_i^* remains unaltered, the corresponding approximations to marginal externalities and marginal strategic values invert. Now, players face negative marginal strategic values. The resulting influence –either up or down– on others' incentives to acquire information from understating their informativeness always works in their favor. In region (IV) we now see players that significantly move against their information imposing negative externality on the network. Their strategic behavior only reduces the incentives of more central players to acquire information. Players in region (V), moving in the direction of their signals but less so than the neutral player, under acquire information. The very peer influences that induce them to respond less to their information add value, on net, to the network. Those in regions (VI), to the right of the neutral player, face additional incentive to acquire information, which translates to negative marginal externalities.

[21] While marginal externalities and marginal strategic values are zero at $\beta_i^* = 0$, we see from Figure 8 that the planner's solution can depart from zero information. Though the gradient of the utilitarian function is fixed at zero at the origin, this does not imply that the planner and IAE solutions align: $\beta_i^* = 0$ may give an inflection point to the planner's objective.

17.4 Policy design

A number of policies could be implemented that nudge the economy in the direction of an efficient outcome. A tax and transfer policy gives an invasive but effective approach. If ξ_i is negative for at least some i and positive for others, a revenue neutral plan taxing the information investments of the former while subsidizing the latter could be at least partially effective. When all links are non-positive or non-negative, subsidy-only and tax-only plans, respectively, would be required.

A less invasive policy geared toward acquisition transparency provides an alternative design. Public certification of the information investments of targeted individuals give one example. Centralized verification and publication of research, or policies that physically display the efforts of individuals within the network give others. All of these examples involve targeting selected positions within the information-response game.

The preceding section suggests a more descriptive design for each of these policy types. Let us focus on symmetric networks, leaving the natural analogue for anti-symmetric networks. For tax and transfer policies, players moving against their information or who set $\beta_i^* > e^\dagger$ should be incentivized (subsidized) to acquire additional information, while those with $\beta_i^* \in (0, e^\dagger)$ should be discouraged (taxed). If links are non-negative and the network resides in region (III), then a natural policy multiplier is realized. Each dollar publicly offered to encourage the acquisition activities of highly central players in the information-response game feeds through to influence the acquisition activities of less central players. Upon introducing negative peer effects as well, such an incentive scheme continues to feed through to others' incentives. Less central players exhibiting $\xi_i < 0$ will be discouraged from acquiring information: an aggregate welfare enhancing effect. And conversely, taxing the acquisition activities of these low-centrality players will tend to encourage the acquisitions of the most central players.

For policies enhancing first-stage transparency, players with $\beta_i^* > e^\dagger$ or $\beta_i^* < 0$ should be targeted. Under only positive links or when negative links are also present, such policies again exhibit a natural multiplier. The incorporation of marginal strategic values into the objectives of the targeted players further encourages others in regions (I) and (III), and discourages those in region (II). In both Figures 7(a) and 7(b), we see a preservation of the property that players to the left of the origin and right of the neutral player tend to exhibit marginal strategic values that are aligned with their marginal externalities, while for those just right of the origin these measures miss-align. Thus, increasing transparency of players with responses outside of the interval $(0, e^\dagger)$ remains a robust and simple rule-of-thumb for these designs, regardless of the extent of symmetry or anti-symmetry in pairwise peer effects.[22]

Certification-based policies will be most feasible in symmetric networks for the following reason. Implementation for players facing positive marginal strategic values requires only a one time certification of their information investments. For those facing negative marginal strategic values (i.e. anti-symmetric peer effects), nothing prevents these players

[22] An even more precise design to the above proposals would target players with positive marginal public values to their information, $\xi_i > 0$, which includes some players in region (II). However, with second-order effects and marginal strategic values shifting the intercept to the left, targeting all i with $\beta_i^* > e^\dagger$ can be taken as a conservative design.

from acquiring additional information subsequent to certification. With other players rationally anticipating this behavior, one-time certifications in anti-symmetric networks may be unimplementable.[23]

A few empirical challenges must also be addressed in any of these designs. First, unless data on information responses in the market can be obtained, retrieval of the peer effects network Σ will be necessary in order to derive equilibrium $\boldsymbol{\beta}^*$. Further, understanding of the costs of information κ is needed to elicit the value of $e^{\dagger}$. Players may also face their own idiosyncratic information costs in real-world peer-effects environments. Section 18.2 addresses extensions that incorporate heterogeneous information costs.

18 Discussion

In this section we explore applications to financial markets in crises and to two-sided markets. Then, basic extensions of the model incorporating further heterogeneity across players' preferences are developed. The model's broader relation to the literature is discussed before concluding.

18.1 Applications

MARKET EFFICIENCY and CRISES. Here we apply the above setup to financial markets and crises. The above welfare properties are cast against the Efficient Market Hypothesis, and applied toward equilibrium information acquisition during financial crises. For the latter, this will require a mixture of both symmetric and anti-symmetric pairwise peer effects.

First consider the following stylized model of a competitive, liquidity-flush market. N traders comprise nontrivial shares of a market for a risky asset. The market price in the second stage is an increasing function of the total of their chosen holdings: $\phi(\bar{x}) = A + B\bar{x}$, where x_i gives i's holding of the asset, $\bar{x} := \sum_{i=1}^{N} x_i$, and $A, B > 0$. Then, as a function of the asset's fundamental value ω, each i's payoff is given by:

$$\begin{aligned} u_i(\mathbf{x}|\omega) &= (\omega + p_i\phi(\bar{x}))\, x_i - x_i^2 \\ &= (p_i A + \omega)\, x_i + (p_i B - 1)\, x_i^2 + x_i p_i B \sum_{k \neq i} x_k, \end{aligned} \tag{20}$$

where here we set $p_i < 0$ capturing a downward sloping demand from each trader. Then, the $t = 2$ expectation $\mathbb{E}_i[\omega + p_i\phi(\bar{x})|\theta_i, e_i]$ in i's best response (2) captures her long-term expected return to her investment, a decreasing function of the expected market price at which assets are purchased. The quadratic term $-x_i^2$ captures decreasing returns to holding inventory, derived from the opportunity costs of funds.

The market price $\phi(\bar{x})$, which here traders do not condition on when choosing second-stage actions, is meant to capture the strategic value that players derive from private information in the market. We can think of each trader i's final holding x_i to be realized by i placing some market order (buy or sell) in the second stage, without complete knowledge

[23] When feasible, continuous monitoring of players below 0 and above $e^{\dagger}$ could insure policy compliance.

of the transaction price ultimately realized.[24] As seen in (20), the sensitivity of the asset's price to others' demands scales by B, which will depend on the total size of the N traders in proportion to the broader market. The larger is B and p_i (in size) the more i cares about the short-term demands of the other $N-1$ traders in the market. And the larger the size of $p_i B < 0$, the more i will strategically avoid highly demanded assets. Thus in reduced form, this stylized setup captures the strategic uses of private information in financial markets under monopolistic competition.

The application can be placed against the Efficiency Market Hypothesis, as follows. As first characterized by Grossman and Stiglitz (1980) [30], when prices are observed and used to infer the private information acquired by others, the asset's price can not both fully and rationally reveal all information of the asset's underlying value.[25] Precisely, if costly private information is fully transmitted through observation of the asset's market-clearing price, then the ex ante incentives to acquire the information are compromised. Here, through the strategic use of information, the shear presence of competing traders similarly reduces the incentives to acquire private information. This is now due to the inference of others' observations and equilibrium actions: privately observing that the asset has high long-term value also informs traders of high short-term market prices. The traders continue to crowd each other as they compete for valuable assets.

The application elicits an important distinction between the informational efficiency versus the welfare efficiency of the market. While the incentives to acquire information display strategic substitutes, the extent of crowding out that ensues is *inefficient*. With each peer effect taking a negative value, each i will obtain $\xi_i < 0$ with $e_i^* > e^{pl}$. In other words, the market will reside in region (II) of Figure 7(a). The informational *in*efficiency of the market is *efficient* from a utilitarian perspective, but to an inefficient extent. In other words, the traders over exert themselves in competition as they research the asset's long-term value.[26]

[Figure 9]

We can further apply the model to yield similar statements on inefficiencies during financial crises. Consider some subset of the traders undergoing severe funding constraints. Precisely, while these traders carry asset holdings prior to the second stage, their abilities to retain their inventories will depend on the market price $\phi(\bar{x})$. If liquidity is sufficiently thin amongst these traders, liquidity spirals may ensue.

[24]This is akin to Kyle (1984a) [41] and (1985) [42], where an insider's market order is a function only of the asset's value and not the market-clearing price. In rational-expectations equilibrium, Kyle (1989) [43] allows traders to submit demand schedules over market prices. The strategic use of information comes in the form of inference of *market depth*: each informed trader i submits her demand schedule given her information, inferring (i) the private observations of other informed traders, and thus (ii) their submitted demand schedules and the extent of noise traders in the market, and ultimately (iii) the sensitivity of the asset's price to i's demand.

[25]Hellwig and Veldkamp (2009) [34] also highlight a similar kinship with Grossman and Stiglitz (1980) [30].

[26]Sanford Grossman and Joseph Stiglitz [30] close with an open question of '*whether it is socially optimal to have 'informationally efficient markets'*,". The above model thus provides one answer, and that is "*no*". When price discovery is introduced, complementarity in information acquisition may also arise, pushing in the opposite direction of over acquisition while reinforcing the under acquisition during crises by all N traders, as described below.

Brunnermeier and Pedersen (2009) [11] provide a theoretical framework of liquidity spirals during crises. Their model captures the dynamic interdependence of market prices and traders' funding constraints. With an initial fall in the asset's perceived fundamental value, speculators' funding constraints force liquidity-starved traders to sell off inventory. As the market price drops, margin calls force these traders to further liquidate, causing a further drop in the asset's market price. This only further constrains the traders, and the spiral worsens.[27]

In effect, these severely constrained traders' demand functions exhibit an upward sloping form.[28] As a reduced representation, we capture this by setting $p_i > 0$ for each of these traders. How exactly would the market look? Figure 9 illustrates the network architecture for the N traders. Liquidity constrained traders, facing positive and directed peer effects, will be the most central players in the information-response game.[29] With respect to Figure 7, unconstrained traders will lie to the left of the neutral player, while highly constrained traders will lie to the right.

The stakes are high for constrained traders. If ω is high, this implies both that the traders can expect a large returns on their holdings, but more importantly, that the current market price will remain high. This is crucial, as the availability of market liquidity is necessary for them to maintain their holdings without the burden of funding constraints.[30]

[Figure 10]

As illustrated in Figure 10 for a market of eight traders, equilibrium welfare exhibits a paradigm shift as the extent of liquidity through the market declines. This shift is driven by the orientation (i.e. symmetry or anti-symmetry) in the local peer effects that each market participant faces. First, when most traders are unconstrained (i.e. "normal" times) the market takes on one similar to the competitive market described above. For the few constrained traders in the market, the majority of their relationships will be anti-symmetric. Residing in region (VI) of Figure 7(b), these traders will over invest in equilibrium. Responding intensely to the news of a high ω, their impact on the market price only crowds the market activities of unconstrained traders. Then, as the proportion of constrained traders grows, these traders face more symmetric and positive peer effects while unconstrained traders face more anti-symmetric relationships. When liquidity problems pervade the market, the constrained traders enter region (III) in Figure 7(a), with unconstrained traders moving to region (V)

[27] Here, the strategic component of information to constrained firms is even more evident as private information may allow them to forecast market prices and infer the potential for constraints to bind over the short term.

[28] See also Gennotte and Leland (1990) [26], Angeletos and Werning (2006) [2] and Gárleanu et al. (2014) [25] for models with inverted equilibrium demand functions of constrained traders in crises.

[29] In the language of Supplemental Section 19, these firms' weighting functions lie strictly above those of unconstrained traders, and thus there will always exist an equilibrium in which they acquire more information. While marginal values to information may be higher for these traders, so too may their marginal costs if the opportunity costs of funds to these traders are large. This can be captured using idiosyncratic κ_i: see Section 18.2.

[30] One can either model information as directly acquired by the traders' funders, or by the banks but with signal realizations verifiable to the funders.

of Figure 7(b). When the number of constrained traders grows to three or more, *all* traders under invest in information. Those under significant funding constraints face positive externalities from the information investments of others. Their informativeness allows the constrained traders to coordinate on asset retention in high market-liquidity outcomes (i.e. high $\phi(\bar{x})$), which tend to occur when the asset is "good" (i.e. high ω).

A striking set of equilibrium behaviors arise among the few unconstrained traders during a crisis. Their acquisition activities impose positive externalities on constrained traders. Moving with their information in region (V) in Figure 7(b), unconstrained traders' informativeness further aids constrained traders to coordinate on high market-liquidity outcomes. They thus under acquire in equilibrium. When six or seven constrained traders pervade the market, the few unconstrained traders acquire zero information in the planner's solution. When the number of constrained traders rises to seven, however, the lone unconstrained trader moves to region (IV) and finds additional value to acquiring information, inferring and moving against the actions of others in the market.

Observation. *As the extent of funding constraints across traders increases, the market equilibrium shifts from being over informed to under informed from a welfare perspective. Crucially, in severe liquidity crises as constrained traders attempt to coordinate on high market-liquidity outcomes, both constrained and unconstrained traders do not internalize the positive externalities their information imposes on the constrained side of the market. In extreme crises, unconstrained firms acquiring and moving against the market do so at a cost to aggregate welfare.*

One can also introduce additional network irregularity by applying this framework to over-the-counter markets. As in Babus and Kondor (2014), if a network designating feasible trades constrains the market, and with each bilateral transaction assigned its own clearing price, Figure 9 would take on a more incomplete network structure. Only trading pairs would be linked in the corresponding peer-effects network, with the sign of out-links determined by the extent of available liquidity to the corresponding trader. The above observation extends. Precisely, highly connected traders that are liquidity deprived may significantly over acquire information relative to the utilitarian benchmark *if* their neighbors are generally unconstrained. Traders that are unconstrained but have many constrained neighbors will, again, under acquire in equilibrium.

Finally, the above policy discussion applies to the application as follows. Competition amongst firms in normal times suggests that certification-based policies may be unimplementable. During crises, however, constrained banks face positive strategic values. Stress tests, coupled with the certification of identified constrained firms, offers a simple and implementable policy intervention. Constrained banks' anticipation of being identified and certified encourages them to internalize their strategic use to information. As they acquire additional information, the market is pushed in the direction of the utilitarian solution.

With marginal strategic values to these firms scaling with their quality-weighted sum-of-squared degree (see expression (18) above), one can verify that such transparency-based policies will be most effective in incomplete network structures under large pairwise peer effects (e.g. in over-the-counter markets). In these networks, marginal strategic values can be sizable in proportion to marginal externalities.[31] With a few constrained neighbors

[31] To formalize the statement in the context of Section 18.2, trader i will have large strategic values relative

imposing large positive externalities on each trader, and vice versa, the effect of internalized strategic values moves the market farther toward the utilitarian solution than in completely but weakly connected network structures.

TWO-SIDED NETWORKS. As an example of a two sided network, consider a job-search market with network structure depicted in Figure 11. Here, a pool of insiders, which may include head hunters or industry professionals, link to workers searching for a job. A particular firm to whom each insider has ties posts a number of open positions. The quality of the firm as an employer (culture, benefits, job security and growth, etc.) are captured by an unknown common state ω. At $t = 2$, each insider i exerts resources x_i toward filling the firm's vacancies with workers they know. Each worker j invests time and effort x_j tailoring their resumes to align with the firm's qualifications and formally applying to relevant positions through their acquainted insiders. The optimal second-stage actions of each player will depend on the expected quality of the firm as an employer, as well as the anticipated actions of neighbors. The workers compete with each other to fill job vacancies, while the insiders compete with each other to connect the workers with the firm and collect value in the form of commission or gained social capital.

[Figure 11]

Abstracting away from variability in the size of counterpart links, the network will generally be symmetric with the welfare properties depicted in Figure 7(a) applying. If the insiders out number workers, facing more positive links across the two groups than negative links amongst other competing insiders, then they will generally reside in region (III). This places the workers in region (II). In this case, insiders face greater incentive to research the firm and will under acquire information in equilibrium. This is because they fail to internalize the positive externalities that their expertise provide their clientele. The less informed workers will over acquire information and over exert themselves researching .

Strategic substitutes within each side and complements across each side of the market introduce clear network irregularities. However, additional heterogeneity may exist across peer effects within either side of the market. As seen in Figure 11, insider i enjoys only two links with workers, which pushes against her centrality in the information-response game. However, i faces the luxury of being the only insider connected to worker j. On the other hand, while the other two insiders enjoy high connectivity with workers, they face stiff competition between each other as their clienteles highly overlap. More generally, the insiders most central in the information-response game will be those that strike an ideal balance between their connectedness (i.e. degree) and the centrality of the workers they connect with (here, client exclusivity).[32]

While this example lends itself well to job-search networks, an array of two-sided markets should adopt similar welfare properties. Entrepreneurs and venture capital investing in new platforms or technologies, Hospitals and pharmaceutical sales firms investing in new medicines or medical technologies, or any other two-sided market in which all players acquire information regarding a fundamental common state will apply.

to externalities when $\gamma_i\sigma_{ij}$ is large relative to σ_{ii} for each neighbor j.

[32]Supplemental Section 19 discusses this further.

Observation. *Two-sided markets in which the shorter side of competing insiders matches competing workers with value-creating transactions exhibit under acquisition amongst experts and over acquisition amongst workers. Experts fail to internalize the positive externalities that their information impose on workers, and workers fail to internalize the negative externalities their information impose on other workers.*

18.2 Basic Extensions

A number of generalizations to the basic model can be considered. As suggested by footnote 6, setting $\sigma_{ii} = \sigma_{jj} = 1$ comes with loss in generality in the degree of variation in players' payoffs. To account for idiosyncrasies in this variation, one can rather define $\tilde{\omega}_i = \gamma_i \omega + \iota_i \omega_i$ for $\gamma_i, \iota_i \in \mathbb{R}_+$, and take the ex post payoff structure:

$$u_i(\mathbf{x}|\omega, \omega_i) = (a_i + \tilde{\omega}_i)\, x_i - \frac{1}{2}\sigma_{ii} x_i^2 + \rho \sum_{j \neq i} \sigma_{ij} x_i x_j,$$

where $\rho \in \mathbb{R}_+$ directly scales the size of peer effects. The corresponding second-stage linear best response is:

$$BR_i(\mathbf{X}_{-i}|\theta_i, e_i, \mu_i) = \frac{a_i + \mathbb{E}_i[\tilde{\omega}_i|\theta_i, e_i]}{\sigma_{ii}} + \rho \sum_{j \neq i} \frac{\sigma_{ij}}{\sigma_{ii}} \mathbb{E}_i[X_j(\theta_j|e_j)\,|\theta_i, e_i, \mu_i].$$

Such a generalization comes with only two necessary modifications to the model's primitives, made to conditions E2 and E4 to give $\mathbb{E}_i[\tilde{\omega}_i|\theta_i, e_i] = v_i e_i \theta_i$ and $\mathbb{E}_i[\theta_j|\theta_i, \mu_i, e_i] = \int_{[0,1]} \mu_{ij}(e_j)\, \gamma_j e_j \gamma_i e_i \theta_i de_j$, respectively, where $v_i := \sqrt{\gamma_i^2 + \iota_i^2}$ gives the variance in i's relevant state $\tilde{\omega}_i$.

Inline with these generalizations, an updated correlation-adjusted adjacency matrix $\Sigma^c := [\gamma_i e_i \sigma_{ij} \gamma_j e_j]$ can be defined. With v_i scaling i's interim expectation of $\tilde{\omega}_i$, informational centralities are now further weighted by the extent of variation in players' relevant states:

$$\beta^* = (\mathbf{I} - \Sigma^c)^{-1}\, \mathbf{I}_{\mathbf{v}} \mathbf{e}.$$

The analogue to network symmetry incorporates an adjustment to peer effects:

Assumption 2C. $\mathbf{I}_{\boldsymbol{\sigma}}^{-1} \mathbf{I}_{\gamma} \Sigma \mathbf{I}_{\gamma}$ *where* $\boldsymbol{\sigma} := [\sigma_{ii}]$ *is symmetric:* $\frac{\gamma_i}{\sigma_{ii}} \sigma_{ij} = \frac{\gamma_j}{\sigma_{jj}} \sigma_{ji}$ *for each* $i \neq j$.

This generalization of Assumption 2A comes with a natural interpretation. Players with low σ_{ii} possess relatively high propensities to choose high actions in the second stage, on average, as well as to acquire information and respond highly to their signal realizations, ceteris paribus. These are exactly the players that have significant influence in the information-response game. Thus, Assumption 2C requires that these influential players have proportionally greater impact on the preferences of those with less influence.[33] The weighting by

[33] Taking the inter-bank network as an example, The Bank of America's expected extent of information acquisition should carry proportionally greater influence on the preferences of smaller banks than do the information investments of these banks on the incentives of The Bank of America.

γ_i and γ_j adjusts for the loading that each player places on the shared state ω. That is, the players' impacts scale directly with the extent that their preferences correlate with others' preferences. The corresponding assumption for anti-symmetric networks can also be defined and applied in a similar way.

Finally, one is free to introduce further idiosyncrasies through the curvature of information costs by providing $\kappa_i(e_i)$ for each i. With these extensions, all of the above results are preserved. With all of these modifications, we obtain the identical expression to (6): $\beta_i^{*2}/e_i^* = \kappa_i'(e_i^*)$ for any interior e_i^*. Players that are most "*central*" in the information-response game are now those with the right combination of being (i) central in the updated network Σ^c, and (ii) having a natural propensity to acquire information, as determined by the relative sizes of $\sqrt{\gamma_i^2 + \iota_i^2}$ and σ_{ii} and the extent of convexity in κ_i.

Analogous limit results are easily obtained, with partials taken with respect to ρ rather than γ^2, and by sending $\rho \to^+ 0$. The corresponding expressions to (7), (16), and (17) are respectively:

$$\begin{aligned}
\lim_{\rho \to^+ 0} \frac{\partial e_i^*}{\partial \rho} &= \frac{\gamma_i e_i^\dagger}{\kappa_i'(e_i^\dagger) - 1} \sum_{k \neq i} \gamma_k e_k^\dagger \frac{\sigma_{ik}}{\sigma_{ii}}, \\
\lim_{\rho \to^+ 0} \frac{\partial \xi_i^{st}(\mathbf{e}^*, \mathbf{X}^*)}{\partial (\rho^2)} &= e_i^\dagger (\gamma_i e_i^\dagger)^2 \sum_{k \neq i} (\gamma_k e_k^\dagger)^2 \frac{\sigma_{ik}}{\sigma_{ii}} \frac{\sigma_{ki}}{\sigma_{kk}}, \quad \text{and} \\
\lim_{\rho \to^+ 0} \frac{\partial \xi_i^{ex}(\mathbf{e}^*, \mathbf{X}^*)}{\partial \rho} &= \gamma_i e_i^{\dagger 2} \sum_{k \neq i} \gamma_k e_k^\dagger \frac{\sigma_{ki}}{\sigma_{kk}}.
\end{aligned}$$

As one may anticipate Proposition 3A maintains, but with $\sum_{k \neq i} \gamma_k e_k^\dagger \frac{\sigma_{ik}}{\sigma_{ii}}$ defining each player i's *effective* degree for moderate peer effects:[34]

Proposition 3C. [symmetric, moderate peer effects] *For symmetric Σ, there exists some γ^w with $0 < \gamma^w \leq \min\{\gamma^m, \gamma^s\}$ such that if $\gamma \in (0, \gamma^w)$ and for $\mathbf{e}^*$, $\mathbf{e}^{pb}$ and $\mathbf{e}^{pl}$ we have:*

1. $\mathbf{e}^{pb} > \mathbf{e}^*$, *with* $({e_i}^{pb} - {e_i}^*) > ({e_j}^{pb} - {e_j}^*) > 0$ *for each i and j with* $\sum_{k \neq i} \left(\gamma_k e_k^\dagger \frac{\sigma_{ik}}{\sigma_{ii}}\right)^2 > \sum_{k \neq j} \left(\gamma_k e_k^\dagger \frac{\sigma_{jk}}{\sigma_{jj}}\right)^2$,
2. $e_i^* > e_j^*$ *and* $({e_i}^{pl} - {e_i}^*) > ({e_j}^{pl} - {e_j}^*)$ *for each i and j with* $\sum_{k \neq i} \gamma_k e_k^\dagger \frac{\sigma_{ik}}{\sigma_{ii}} > \sum_{k \neq j} \gamma_k e_k^\dagger \frac{\sigma_{jk}}{\sigma_{jj}}$.

And, the analogues to Assumption 2B and Proposition 3B that incorporate these extensions can similarly be constructed.

Crucially, the basic message of Figure 7 will continue to hold. Approximations to ξ_i^{st} and ξ_i^{ex} are derived with γ_i and $1/\sigma_{ii}$ scaling each peer effect σ_{ij}. Thus, the corresponding figures can be produced for each individual player. Under Assumption 1, exactly where each i falls on their β_i^*-axis relative to the origin and their respective $e_i^\dagger$ continues to be driven by the network of peer effects. In symmetric networks (now, Assumption 2C), those in their corresponding region (III) choosing β_i^* above $e_i^\dagger$ underinvest, those to the right of the origin

[34] The results of Supplemental Section 19 also maintain, with our notions of degree centrality and weighting function defined in terms of normalized peer effects $\frac{\gamma_i}{\sigma_{ii}}\sigma_{ij}$ for each i and j.

falling in region (II) tend to overinvest, and players moving against their signals in region (I) again underinvest. All players to the right of the origin moving in the direction of their signal realizations continue to face positive marginal strategic values.

Thus, one can view the above model's homogenous setup in the first-stage information-acquisition game –outside of network effects– as simplifying the analysis, allowing the network to "speak clearly". None the less, the extent of symmetry in pairwise peer effects coupled with players' informational centralities continue to play crucial roles shaping equilibrium inefficiencies in a much broader set of economies.

18.3 Related Literature

Here related papers are discussed, along with a number of potential avenues for future research. The paper relates to a family of papers studying information games with communication on networks. In Calvó-Armengol and de Martí (2007) [13], (2009) [14], and Calvó-Armengol et al. (2009) [15] the network is defined by the exogenous correlation matrix between signals. In Calvó-Armengol et al. (2011) [16] this network is endogenized through a communications device, and the authors study the relative extent of active and passive communication in equilibrium (i.e. "speaking" and "listening", resp.). Thus, each player's communication quality is endogenously directed to each of her neighbors. Beyond the above setup's treatment of acquired information as pertaining directly to fundamental payoffs, these papers have a number of model elements that distinguish them from this one. As a closest comparative, Calvó-Armengol et al. (2011) [16]'s first-stage signal qualities are fully observed in the second stage. Thus, strategic values are fully internalized in their model. And as only strategic complements are considered in their formal analysis, the characterization of negative signal responses and corresponding welfare implications discussed above are not considered. The authors find underinvestment in communication, which relates with the under acquisition above when $\sigma_{ij} \geq 0$ for each i and j. In the above, however, states are global rather than local via common state ω, and thus inefficient acquisition under strategic compliments is driven by the correlation in and strategic use of information in the second stage. Further, players facing negative peer effects and with signal responses β_i^* in some interior subset of $(0, e^\dagger)$ over acquire in symmetric networks.

As discussed in the introduction, this paper is closely related to the coordination games with costly information acquisition literature. The above case of symmetric networks provides a generalization to many related results found in the literature, while incorporating strategic substitutes and complements simultaneously through the network. Hellwig and Veldkamp (2009) [34] study costly acquisition of signals chosen from a subset of available signals of various qualities and correlation profiles. The authors Proposition 1 offers a closest analogue to the above Theorems 1 and 2, which establish the feed through of strategic complementarity and substitutability (separately) into first-stage information values.[35] Also reminiscent of their findings, strategic complementarity can imply multiple symmetric equilibria. However, the type of multiplicity of equilibria illustrated above in Example 2 are derived solely from strategic complements, rather than discreteness in the signal technology.

[35] Vives (1988) [61] together with Vives (2008)'s [62] exercise 8.15 similarly establish this feed through for strategic substitutes and compliments, respectively.

In contrast, equilibrium uniqueness under strategic complements and continuous signals is derived in Hellwig and Veldkamp (2009) [34] as well as Myatt and Wallace (2009) [49]. In beauty contest games with a continuum of agents, the extents of complementarities are implicitly bounded. In the above network setting, strategic complements can be more pronounced while the set of convex cost functions considered are less constrained,[36] thus yielding the observed multiplicity under significant positive peer effects.

This literature also offers an exciting research agenda studying the effects of public information on the equilibrium actions and welfare in a general network setting. Morris and Shin (2002) [47] first highlighted the potential adverse effects of public information, showing that players may coordinate on less precise public announcements. In an information acquisition setting, Colombo et al. (2014) [19][37] show that public information crowds out private information[38], while acquired private information is inefficiently low *if and only if* the equilibrium degree of coordination falls short of the efficient benchmark (see Colombo et al. (2014) [19] Corollary 1 and Proposition 5 (ii), resp.). In a network setting under both strategic complements and substitutes, the efficiency of equilibrium coordination depends on each player's informational centrality (e.g. Proposition 3A, above). The effects of public information on both the positive and normative properties of equilibrium information acquisition in these settings are left as open questions.

The above coordination games literature assumes agents to reside on a continuum, and thus the strategic values studied here are not realized. In a network setting, a continuum of players is clearly inapplicable. With the exception of Hauk and Hurkens (2001) for a competitive Cournot production market, the welfare implication of incentive compatibility in information acquisition are novel. The extent of player *connectedness* in the network as driving the size of her strategic values provides a network characterization. Further, the symmetry in the coordination games with endogenous information that have thus far been studied plays an important part in driving inefficiencies in information acquisition. While symmetric networks offers analogous welfare results to many seen in these papers, the fact that the direction of the utilitarian solution inverts under anti-symmetric networks suggests caution when applying these welfare properties in settings that incorporate anti-symmetric relationships. And as the above application of Section 18 suggests, anti-symmetric relationships may be common in environments with a subset of constrained players.

A number of oligopoly models have studied information acquisition outside of a network setting. Novshek and Sonnenschein (1982) [52] and Vives (1983) [60] study the effects of private information when firms face an uncertain demand function. Taking the extent of information acquisition exogenously, the authors' consider comparative statics of equilibrium production and welfare with respect to signal qualities. Their Lemma 1 establishes a direct

[36]To see this, here qualities are chosen from $[0, 1]$ while in these and most of the coordination games literature they are taken from $[0, \infty)$. Appendix F.1 provides a mapping from the accumulated i.i.d. normal draws setup to information qualities. Note that a constant marginal cost to these draws excludes the possibility of initial positive gains as players search for and locate the most efficient sources of informative signals.

[37]They allow for both strategic complements and substitutes (though not simultaneously) in their setting. Their welfare benchmark that corresponds to that taken here involves *not* allowing the planner to enforce the efficient use of information.

[38]Myatt and Wallace (2009) [49] find a similar result, with the publicity of information endogenously determined.

dependence in the slope of equilibrium strategies to signal quality, as a function of the extent of complementarity between firms' goods. Similar equilibrium properties obtain under the more general network treatment of Theorems 1 and 2 above, upon homogenizing the size and direction of links. Beyond this close similarity at a positive level, the papers' welfare analyses depart from each other with the consumer side of the market excluded in network games.

Related to transparency, a number of papers have addressed information transmission in a network settings similar to that taken here, but without endogenizing information qualities. Hagenbach and Koessler (2010) [31] and Galeotti et al. (2013) [24] study cheap talk in networks, taking exogenous biases as common information amongst the players. Kondor and Babus (2014) [3] study information diffusion and trade between traders connected through a network.[39] The authors define a "conditional guessing game", which solves for transmission of information in rational expectations, as a function of observed prices and the network structure. And taking an extreme to transparent play, Hagenbach et al. (2014) [32] study full disclosure under certifiable pre-play communication. In the above setting, these authors' acyclicity condition is satisfied.[40]

The above model's exclusion of information transmission provides a first benchmark to information acquisition in a network setting, while maintaining reach in its applications. Studying the incentives to acquire the information that agents carry when also faced with particular transmission mechanisms offers an exciting avenue for research. Both the positive and normative implications under full information disclosure offers a promising starting point.

Finally, Bramoulle et al. (2014) [9] study the set of network games equivalent to potential games. The authors characterize both the presence of multiple equilibria and of equilibria that involve players taking zero action (i.e. a corner solution in their setup) using the size of the lowest eigenvalue for the network's adjacency matrix. In the above, corner solutions in the information-response game play an important role when incorporating the possibility of players moving against their signals, as illustrated with Example 2. Bramoulle et al. (2014) [9]'s eigenvalue characterization of corner equilibria provides a valuable tool to designate the presence of players moving against their information. Here, the second-stage game can be characterized as a potential game if the network is symmetric.[41]

18.4 Conclusion

A flexible framework for studying information acquisition in linear peer-effects networks is developed. An intuitive characterization of the equilibrium strategic responsiveness of players to their private signals is derived. Using this characterization, marginal information values are derived in equilibrium. Scaling with the square of this responsiveness, marginal values to information take on a U-shaped dependence on network centrality in the information-response game. Under significant strategic substitutes, the least central players find additional use from information through inferring and moving against the actions of neighbors.

[39] That is, their network captures the set of feasible trades that can occur.

[40] Precisely, where the worst type in finite set $S_i \subseteq \Theta$ is given by the lowest element *if* $\beta_i^* > 0$, and the highest element if $\beta_i^* < 0$.

[41] The information-response game is no longer equivalent to a potential games when Σ is not symmetric.

Equilibrium welfare and the strategic motives behind information acquisition are analyzed. The extent of symmetry in pairwise relationships drives the direction of inefficiencies, both when players move in the direction of or against their information. Under moderate and symmetric relationships, players under respond to the network of peer effects. With both strategic complements and substitutes present, the most informed players under acquire information and the least informed players over acquire. Incorporating players moving against their signals, the extend of information acquired by these players is inefficiently low. Thus, the U-shaped non-monotonicity in the incentives to acquire information in networks carries over to welfare.

All of these welfare properties reverse when the presence of anti-symmetric relationships pervade the network structure. As our example of a market in crisis illustrates, anti-symmetric relationships may play an important role when a nontrivial set of traders in the market face liquidity constraints and thus value high market prices. When liquidity becomes scarce through the market, the few unconstrained firms fail to internalize the externalities they impose on the constrained side of the market. When these unconstrained firms move in the direction of their signal realizations, they under acquire information. If they instead move against their signals, their strategic acquisition of information quality exceeds that of the social planner's prescription. The flexibility in peer-effects networks is essential when assessing the welfare implications in these heterogeneous settings, capturing an ray of equilibrium behaviors.

Marginal strategic values take on unambiguous and opposing directions in symmetric and anti-symmetric networks. Players face clear incentives to overstate and understate their informativeness in these respective settings. The size of these incentives are proportional to players' sum-of-squared degrees. Thus, player connectedness characterizes the size of marginal strategic values to information, while symmetry in pairwise relationships continues to capture its direction. The analysis elicits a transparency based policy with a simplistic implementation: certify the information acquired by the experts: the most central players in the information response game. And when possible, certify the information investments of those moving against their signals: the least central players in the information response game.

In summary, the above network setting offers a flexible framework to both extend and assess the robustness of many results offered by the coordination games with endogenous information literature. While moderate, symmetric networks offer a natural extension to heterogeneous environments, anti-symmetry and the incorporation of players moving against their information offer both positive and normative properties unattained in symmetric settings. The role of observable prices determined in market clearing under rational expectations, as well as to other forms of information transmission are left as important open questions for future work.

References

[1] Angeletos, George-Marios; Pavan, Alessandro. "Efficient Use of Information and Social Value of Information". *Econometrica*, 74(4):1103–1142, 2007.

[2] Angeletos, George-Marios; Werning, Iván. "Crises and Prices: Information Aggregation, Multiplicity, and Volatility". *American Economic Review*, 96(5):1720–1736, December 2006.

[3] Babus, Ana; Kondor, Peter. "Trading and Information Diffusion in Over-the-Counter Markets". *mimeo*, 2014.

[4] Bala, Venkatesh; Goyal, Sanjeev. "A Noncooperative Model of Network Formation". *Econometrica*, 68(5):1181 1229, September 2000.

[5] Ballester, Coralio; Calvó-Armengol, Antoni; Zenou, Yves. "Who's Who in Networks. Wanted: The Key Player". *Econometrica*, 74(5):1403–1427, September 2006.

[6] Blume, Lawrence; Brock, William; Durlauf, Steven; Jayaraman, Rajshri. "Linear Interactions Models". *NBER Working Paper Series*, Working Paper 19212, 2013.

[7] Bonacich, Phillip. "Power and Centrality: A Family of Measures". *American Journal of Sociology*, 92(5):1170–1182, March 1987.

[8] Bramoullé, Yann; Kranton, Rachel. "Public Goods in Networks". *Journal of Economic Theory*, 135:478–494, 2007.

[9] Bramoullé, Yann; Kranton, Rachel; D'Amours, Martin. "Strategic Interaction and Networks". *American Economic Review*, 104(3):898930, June 2014.

[10] Brunnermeier, Markus; Pedersen, Lasse Heje. "Predatory Trading". *The Journal of Finance*, 60(4):1825–1863, August 2005.

[11] Brunnermeier, Markus; Pedersen, Lasse Heje. "Market Liquidity and Funding Liquidity". *The Review of Financial Studies*, 22(6):2201–38, June 2009.

[12] Burguet, Roberto; Vives, Xavier. "Social learning and costly information acquisition". *Economic Theory*, 15:185–205, 2000.

[13] Calvó-Armengol, Antoni; de Martí, Joan. "Communication Networks: Knowledge and Decisions". *The American Economic Review*, 97(2):86–91, May 2007.

[14] Calvó-Armengol, Antoni; de Martí, Joan. "Information Gathering in Organizations: Equilibrium, Welfare and Optimal Network Structure". *mimeo*, 2009.

[15] Calvó-Armengol, Antoni; de Martí, Joan; Prat, Andrea. "Endogenous Communication in Complex Organizations". *mimeo*, 2009.

[16] Calvó-Armengol, Antoni; de Martí, Joan; Prat, Andrea. "Communication and Influence". *mimeo*, 2011.

[17] Calvó-Armengol, Antoni; Jackson, Matt. "The Effects of Social Networks on Employment and Inequality". *American Economic Review*, 94(3):426–454, 2004.

[18] Carlsson, Hans; van Damme, Eric. "Global Games and Equilibrium Selection". *Econometrica*, 61(5):989–1018, September 1993.

[19] Colombo, Luca; Femminis, Gianluca; Pavan, Alessandro. "Information Acquisition and Welfare". *The Review of Economic Studies*, forthcoming, 2014.

[20] DeMarzo, Peter; Vayanos, Dimitri; Zwiebel, Jeffrey. "Persuasion Bias, Social Influence, and Unidimensional Opinions". *The Quarterly Journal of Economics*, 118(3):909–968, 2003.

[21] Diego García; Strobl, Günter. "Relative Wealth Concerns and Complementarities in Information Acquisition". *mimeo*, July 2010.

[22] Elliott, Matthew; Golub, Benjamin. "A Network Approach to Public Goods". *mimeo*, 2014.

[23] Fainmesser, Itay; Galeotti, Andrea. "The Value of Network Information". *mimeo*, 2013.

[24] Galeotti, Andrea; Ghiglino, Christian; Squintani, Francesco. "Strategic Information Transmission Networks". *Journal of Economic Theory*, 148(5):1751–1769, 2013.

[25] Gárleanu, Nicolae; Panageas, Stavros; Yu, Jianfeng. "Financial Entanglement: A Theory of Incomplete Integration, Leverage, Crashes, and Contagion". *forthcoming, American Economic Review*, 2013.

[26] Gennotte, Gerard; Leland, Hayne. "Market Liquidity, Hedging, and Crashes". *American Economic Review*, 80(5):999–1021, December 1990.

[27] Golosov, Michael; Lorenzoni, Guido; Tsyvinski, Aleh. "Decentralized trading with private information". *Econometrica*, 82(3):1055–1091, July 2014.

[28] Goyal, Sanjeev. "Connections: An Introduction to the Economics of Networks". *Princeton University Press*, 2007.

[29] Grossman, Sanford. "On the Efficiency of Competitive Stock Markets Where Trades Have Diverse Information". *The Journal of Finance*, 31(2):573–585, May 1976.

[30] Grossman, Sanford; Stiglitz, Jospeh. "On the Impossibility of Informationally Efficient Markets". *American Economic Review*, 70(3):393–408, June 1980.

[31] Hagenbach, Jeanne; Koessler, Frédéric. "Strategic Communication Networks". *The Review of Economic Studies*, 77:10721099, 2010.

[32] Hagenbach, Jeanne; Koessler, Frédéric; Perez-Richet, Eduardo. "Certifiable Pre-Play Communication: Full Disclosure". *mimeo; forthcoming Econometrica*, 2014.

[33] Hauk, Esther; Hurkens, Sjaak. "Secret Information Acquisition in Cournot Markets". *Economic Theory*, 18:661–681, 2007.

[34] Hellwig, Christian; Veldkamp, Laura. "Knowing What Others Know: Coordination Motives in Information Acquisition". *The Review of Economic Studies*, 76:223–251, 2009.

[35] Hellwig, Martin. "On the Aggregation of Information in Competitive Markets". *Journal of Economic Theory*, 22:477–498, 1980.

[36] Helsleya, Robert; Zenou, Yves. "Social Networks and Interactions in Cities". *Journal of Economic Theory*, 150:426–466, 2014.

[37] Hendricks, Kenneth; Porter, Robert. "The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts". *American Economic Review*, 86(3):388–407, June 1996.

[38] Hiller, Timo. "Peer Effects in Endogenous Networks". *mimeo*, 2013.

[39] Jackson, Matt. "Social and Economics Networks". *Princeton University Press*, 2008.

[40] Kotowski, Maciej; Leister, C. Matthew. "Trading Networks and Equilibrium Intermediation". *mimeo*, 2014.

[41] Kyle, Albert. "Market Structure, Information, Futures Markets, and Price Formation". *Storey, G. G. Schmitz, A. and Sarris, A. H. (eds), International Agriculurla Trade: Advanced Readings in Price Formation, Market Structure, and Price Stability*, 53:45–64, 1984a.

[42] Kyle, Albert. "Continuous Auctions and Insider Trading". *Econometrica*, 53(6):1315–1335, 1985.

[43] Kyle, Albert. "Informed Speculation with Imperfect Competition". *The Review of Economic Studies*, 56(3):317–355, July 1989.

[44] Li, Lode; McKelvey, Richard. "Optimal Research for Cournot Oligopolists". *Journal of Economic Theory*, 42:140–166, 1987.

[45] Malamud, Semyon; Rostek, Marzena. "Decentralized Exchange". *mimeo*, 2012.

[46] Marschak, Thomas; Reichelstein, Stefan. "Network Mechanisms, Informational Efficiency, and Hierarchies". *Journal of Economic Theory*, 79:106–141, 1998.

[47] Morris, Stephen; Shin, Hyun Song. "The Social Value of Public Information". *American Economic Review*, 92(5):1521–1534, December 2002.

[48] Morris, Stephen; Shin, Hyun Song. "Coordination Risk and the Price of Debt". *European Economic Review*, 48(1):133–153, 2003.

[49] Myatt, David; Wallace, Chris. "Endogenous Information Acquisition in Coordination Games". *The Review of Economics Studies*, 79:340–374, 2012.

[50] Myatt, David; Wallace, Chris. "Cournot Competition and the Social Value of Information". *mimeo*, 2013.

[51] Nikitin, Maxim; Smith, Todd. "Information Acquisition, Coordination, and Fundamentals in a Financial Crisis". *Journal of Banking & Finance*, 32:907–914, 2008.

[52] Novshek, William; Sonnenschein, Hugo. "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release". *The Bell Journal of Economics*, 13(1):214–218, Spring 1982.

[53] Oliveros; Santiago. "Abstention, Ideology and Information Acquisition". *Journal of Economic Theory*, 148(3):871–902, May 2013.

[54] Ostrovsky; Michael. "Stability in Supply Chain Networks". *American Economic Review*, 98(3):897–923, 2008.

[55] Persico, Nicola. "Information Acquisition in Auctions". *Econometrica*, 68(1):135–148, January 2000.

[56] Szkupy, Michal; Trevino, Isabel. "Information Acquisition and Transparency in Global Games". *mimeo*, 2014.

[57] Takayama, Akira. "Mathematical Economics". *Cambridge University Press*, 2nd edition, 1985.

[58] Veldkamp, Laura. "Information Choice in Macroeconomics and Finance". *Princeton University Press*, 2011.

[59] Verracchiq, Robert. "Information Acquisition in a Noisy Rational Expectations Economy". *Econometrica*, 50(6):1415–1430, 1982.

[60] Vives, Xavier. “Duopoly Information Equilibrium: Cournot and Bertrand”. *Journal Of Economic Theory*, 34:71–94, 1983.

[61] Vives, Xavier. “Aggregation of Information in Large Cournot Markets”. *Econometrica*, 56(4):851–876, July 1988.

[62] Vives, Xavier. “Information and Learning in Markets: The Impact of Market Microstructure”. *Princeton University Press*, 2008.

[63] Vives, Xavier. “Strategic Supply Function Competition With Private Information”. *Econometrica*, 6(76):19191966, November 2011.

[64] Yang, Ming. “Coordination with Flexible Information Acquisition”. *mimeo*, 2014.

[65] Zhou, Junjie; Chen, Ying-Ju. “Targeted Information Release in Social Networks”. *NET Institute working paper series*, (13-4), September 2013.

F Appendix

F.1 Linear-in-qualities expectations: examples

Two states. As the most basic example of an information structure embodying Conditions E1-E4, consider the case of two aggregate states $\omega \in \{-1, 1\}$ with $\gamma = 1$ and priors $\Pr(\omega = 1) = \Pr(\omega = -1) = 1/2$. Then, player i's quality e_i gives the probability of the signal being correct, $\Pr(\theta_i = \omega) = \frac{e_i+1}{2}$. The conditional expectation $\mathbb{E}_i[\theta_j|\theta_i, e_i, \mu_i^*] = e_i e_j \theta_i$ for each $j \neq i$ can be derived as the correlation in the players i and j's signals, $e_i e_j$, multiplied by i's signal realization θ_i. In this case, ω is naturally interpreted as giving 'high' ($\omega = 1$) and 'low' ($\omega = -1$) marginal gains to action x_i, for each player i.

Multiple normal draws. Considering the more general definition of $\tilde{\omega}_i$ provided Section 18.2, one can also consider a normally distributed states and signals setup with normal errors. Assume $\omega \sim N(0, 1)$, $\omega_i \sim N(0, 1)$, and thus $\tilde{\omega}_i \sim N(0, v_0)$ where $v_0 := \gamma_i^2 + \iota_i^2$. Now, consider player i who draws $S_i \in \mathbb{Z}_+$ signals $\{\vartheta_i^s\}_{s=1}^{S_i}$ taking values $\vartheta_i^s = \tilde{\omega}_i + \varepsilon_i^s$ with error $\varepsilon_i^s \sim N(0, v_1)$; that is, each ϑ_i^s has precision v_1^{-1}. Clearly E1 is satisfied. Player i can then use her signals to infer $\tilde{\omega}_i$ by the usual Bayesian updating rule:

$$\mathbb{E}_i\left[\tilde{\omega}_i \middle| \{\vartheta_i^s\}_{s=1}^{S_i}\right] = \frac{v_1^{-1} \sum_{s=1}^{S_i} \vartheta_i^s}{v_0^{-1} + S_i v_1^{-1}}.$$

Define the aggregate signal and information quality:

$$\begin{aligned}
\theta_i &:= \frac{1}{\sqrt{v_0 + \frac{v_1}{S_i}}} \frac{1}{S_i} \sum_{s=1}^{S_i} \vartheta_i^s, \\
e_i &:= \frac{1}{\sqrt{v_0}} \frac{v_1^{-1}}{v_0^{-1} + S_i v_1^{-1}} S_i \sqrt{v_0 + \frac{v_1}{S_i}} = \sqrt{\frac{v_0}{v_0 + \frac{v_1}{S_i}}}.
\end{aligned}$$

The average $\frac{1}{S_i}\sum_{s=1}^{S_i} \vartheta_i^s$ will have precision $S_i v_1^{-1}$. It is then straight forward to show that (the extended version of) E2 and E3 are satisfied:

$$\begin{aligned}
\mathbb{E}_i[\tilde{\omega}_i|\theta_i, e_i] &= \sqrt{\gamma_i^2 + \iota_i^2} e_i \theta_i, \\
\mathbb{E}_i\left[\theta_i^2|e_i\right] &= 1.
\end{aligned}$$

Now consider player j who draws $S_j \in \mathbb{Z}_+$ signals $\{\vartheta_j^s\}_{s=1}^{S_i}$ taking values $\vartheta_j^s = \tilde{\omega}_j + \varepsilon_j^s$ with error $\varepsilon_j^s \sim N(0, v_1)$. Then $\mathbb{E}_i[\theta_j|\theta_i, e_i, \mu_i^*]$ is derived from simple linear regression of θ_j

on θ_i:

$$\begin{aligned}
\mathbb{E}_i\left[\theta_j|\theta_i, e_i, \mu_i^*\right] &= \frac{Cov\left(\theta_i, \theta_j\right)}{Sd\left(\theta_j\right)}\theta_i \\
&= \frac{\gamma_i\gamma_j v_0}{\sqrt{v_0 + \frac{v_1}{S_i}}\sqrt{v_0 + \frac{v_1}{S_j}}}\theta_i \\
&= \gamma_i\gamma_j e_i e_j \theta_i,
\end{aligned}$$

establishing (the extended version of) condition E4 under sequentially rational μ_i^*.

F.2 Section 16 proofs: Equilibrium information acquisition and response

Existence of a second-stage equilibrium is only ensured if the size of peer effects are suitably constrained. This motivates the following assumption, maintained throughout.

Assumption F1. $\left(\mathbf{I} - [p_{ij}\sigma_{ij}]\right)^{-1}$ *is well defined for every* $\mathbf{p} \in [0,1]^{N(N-1)}$.

Assumption F1 is a strengthening of the condition $p > \lambda\mu_i(\mathbf{G})$ in Bellester et al. (2006) [5] Theorem 1 bounding the spectral radius of the relevant diagonally dominant matrix under complete information. Assumption F1 implies that the relevant diagonally-dominant matrix in the second stage's information-response game remains invertible for all first-stage outcomes. Primarily a technical condition, this suffices for existence and uniqueness of a pure linear Bayesian equilibrium at period $t = 2$.

Proof of Theorem 1. For all purposes, $\mathbf{I}$ will denote the $n \times n$ identity matrix. Linearity of the ex-post best responses allows us to take expectations of (2) and obtain i's first order condition of her information response problem. This gives optimal action:

$$\begin{aligned}
X_i\left(\theta_i, e_i\right) &= \left(a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i\right]\right) + \sum_{k\neq i}\sigma_{ik}\mathbb{E}_i\left[X_k\left(\theta_k, e_k\right)|\theta_i, e_i, \mu_i^*\right] \\
&= \left(a_i + e_i\theta_i\right) + \sum_{k\neq i}\sigma_{ik}\mathbb{E}_i\left[X_k\left(\theta_k, e_k\right)|\theta_i, e_i, \mu_i^*\right].
\end{aligned} \tag{21}$$

Next, we are free to take expectations of (21) over realizations of player i's signal θ_i. Denoting the vector of expected stage two actions $\boldsymbol{\alpha}^*$, this gives:

$$\boldsymbol{\alpha}^* = \mathbf{a} + \Sigma\boldsymbol{\alpha}^*. \tag{22}$$

This can easily be solved to give the following expectations equilibrium:

$$\boldsymbol{\alpha}^* = \left[\mathbf{I} - \Sigma\right]^{-1}\mathbf{a}. \tag{23}$$

Note that $\boldsymbol{\alpha}^*$ does not depend on $\mathbf{e}$.[42]

[42] Expression (23) is analogous to expression (4) in Ballester et al. (2006) [5], but now in expectations.

Next, we derive the information responses given in (5). Linearity of this expression is derived by the linearity in best responses (2) and in expectations. Consider the following profile of strategies:

$$\mathbf{X}^{*}(\boldsymbol{\theta})=\boldsymbol{\alpha}^{*}+\mathbf{I}_{\boldsymbol{\theta}}\boldsymbol{\beta}^{*},$$

with $\beta_i^* \in \mathbb{R}$ denoting each player i's responsiveness to her signal. For each component i taking $\boldsymbol{\beta}_{-i}^*$ as above we verify that i plays a linear strategy. Taking differences of (21) at θ_i and $\theta_i' < \theta_i$ then gives[43]:

$$\begin{aligned}
X_i^*(\theta_i|e_i) - X_i^*(\theta_i'|e_i) &= \begin{pmatrix} \mathbb{E}_i[\tilde{\omega}_i|\theta_i, e_i, \mu_i^*] - \mathbb{E}_i[\tilde{\omega}_i|\theta_i', e_i, \mu_i^*] \\ +\sum_{k\neq i}\sigma_{ik}\left(\mathbb{E}_i[X_k^*(\theta_k|e_k)|\theta_i, e_i, \mu_i^*]\right. \\ \left.-\mathbb{E}_i[X_k^*(\theta_k|e_k)|\theta_i', e_i, \mu_i^*]\right. \end{pmatrix} \\
&= \begin{pmatrix} e_i\theta_i - e_i\theta_i' \\ +\sum_{k\neq i}\sigma_{ik}\left(\mathbb{E}_i[\alpha_k^* + \theta_k\beta_k^*|\theta_i, e_i, \mu_i^*]\right. \\ \left.-\mathbb{E}_i[\alpha_k^* + \theta_k\beta_k^*|\theta_i', e_i, \mu_i^*]\right. \end{pmatrix} \\
&= e_i(\theta_i - \theta_i') + \sum_{k\neq i}\sigma_{ik}\left(\left(\mathbb{E}_i[\theta_k|\theta_i, e_i, \mu_i^*] - \mathbb{E}_i[\theta_k|\theta_i', e_i, \mu_i^*]\right)\beta_k^*\right) \\
&= e_i(\theta_i - \theta_i') + \sum_{k\neq i}\sigma_{ij}\left(\left(\gamma^2 e_i e_k\theta_i - \gamma^2 e_i e_k\theta_i'\right)\beta_k^*\right) \\
&= (\theta_i - \theta_i')\left(e_i + \sum_{k\neq i}\sigma_{ik}\gamma^2 e_i e_k\beta_k^*\right).
\end{aligned}$$

With $\frac{X_i^*(\theta_i|e_i) - X_i^*(\theta_i'|e_i)}{\theta_i - \theta_i'}$ independent of the choice of θ_i and θ_i', player i also plays a linear strategy, with (optimal) responsiveness

$$\beta_i^* = \frac{X_i^*(\theta_i|e_i) - X_i^*(\theta_i'|e_i)}{\theta_i - \theta_i'} = e_i + \sum_{k\neq i}\sigma_{ik}\gamma^2 e_i e_k\beta_k^*. \tag{24}$$

We thus have:

$$\boldsymbol{\beta}^* = \mathbf{e} + \gamma^2\mathbf{I}_{\mathbf{e}}\Sigma\mathbf{I}_{\mathbf{e}}\boldsymbol{\beta}^*. \tag{25}$$

With $(\mathbf{I} - \gamma^2\mathbf{I}_{\mathbf{e}}\Sigma\mathbf{I}_{\mathbf{e}})^{-1}$ well defined by Assumption F1, solving for $\boldsymbol{\beta}^*$ gives the unique linear information response equilibrium:

$$\boldsymbol{\beta}^* = \left(\mathbf{I} - \gamma^2\mathbf{I}_{\mathbf{e}}\Sigma\mathbf{I}_{\mathbf{e}}\right)^{-1}\mathbf{e}.$$

Finally, we can easily write:

$$X_i^*(\theta_i|e_i) = \alpha_i^* + \theta_i\beta_i^*,$$

for each i, giving the $t = 2$ IRE strategy seen in (5).

To establish the stronger uniqueness claim succeeding Theorem 1, the following establishes a similar result to that shown in Dewan and Myatt (2008), adapted to our network

[43] One can always find such a signal pair, else signals are never informative.

setting. The second stage best response function of any i given first stage outcome $\mathbf{e}$ (and correct beliefs μ_i^* regarding $\mathbf{e}_{-i}$) is again:

$$BR_i\left(\mathbf{X}_{-i}|\theta_i, e_i, \mu_i^*\right) = a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i, e_i, \mu_i^*\right] + \sum_{k\neq i} \sigma_{ik}\mathbb{E}_i\left[X_k^*\left(\theta_k|e_k\right)|\theta_i, e_i, \mu_i^*\right].$$

Suppressing the (e_i, μ_i^*) conditionals, the composition of $BR_i\left(\mathbf{X}_{-i}|\theta_i\right)$ with $BR_j\left(\mathbf{X}_{-j}|\theta_j\right)$ for each $j \neq i$ gives:

$$\begin{aligned}
BR_i^2\left(\cdot|\theta_i\right) &= a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i\right] + \sum_{k\neq i}\sigma_{ik}\mathbb{E}_i\left[a_k + \mathbb{E}_k\left[\tilde{\omega}_k|\theta_k\right] + \sum_{k'\neq k}\sigma_{kk'}\mathbb{E}_k\left[\cdot|\theta_k\right]\middle|\theta_i\right] \\
&= a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i\right] + \sum_{k\neq i}\sigma_{ik}\left(a_k + \mathbb{E}_i\left[\mathbb{E}_k\left[\tilde{\omega}_k|\theta_k\right]|\theta_i\right] + \sum_{k'\neq k}\sigma_{kk'}\mathbb{E}_i\left[\mathbb{E}_k\left[\cdot|\theta_k\right]|\theta_i\right]\right) \\
&= \left(\begin{array}{c} a_i + \mathbb{E}_i\left[\tilde{\omega}_i|\theta_i\right] + \sum_{k\neq i}\sigma_{ik}a_k + \sum_{k\neq i}\sigma_{ik}\mathbb{E}_i\left[\mathbb{E}_k\left[\tilde{\omega}_k|\theta_k\right]|\theta_i\right] \\ + \sum_{k\neq i}\sum_{k'\neq k}\sigma_{ik}\sigma_{kk'}\mathbb{E}_i\left[\mathbb{E}_k\left[\cdot|\theta_k\right]|\theta_i\right] \end{array}\right) \\
&= a_i + \sum_{k\neq i}\sigma_{ik}a_j + e_i\theta_i + \sum_{k\neq i}\sigma_{ik}\gamma^2 e_k^2 e_i\theta_i + \sum_{k\neq i}\sum_{k'\neq j}\sigma_{ik}\sigma_{kk'}\mathbb{E}_i\left[\mathbb{E}_k\left[\cdot|\theta_k\right]|\theta_i\right].
\end{aligned}$$

In vector form[44]:

$$BR^2\left(\cdot|\boldsymbol{\theta}\right) = \left(\begin{array}{c} \mathbf{a} + \Sigma\mathbf{a} + \mathbf{I}_{\boldsymbol{\theta}}\mathbf{e} + \mathbf{I}_{\boldsymbol{\theta}}\gamma^2\mathbf{I}_{\mathbf{e}}\Sigma\mathbf{I}_{\mathbf{e}}\mathbf{e} \\ + \left[\sum_{k\neq i}\sum_{k'\neq k}\sigma_{ik}\sigma_{kk'}\mathbb{E}_i\left[\mathbb{E}_k\left[\cdot|\theta_k\right]|\theta_i\right]\right] \end{array}\right).$$

We can iterate this to yield the τ'th best-response dynamic $BR^\tau\left(\cdot|\boldsymbol{\theta}\right)$:

$$BR^\tau\left(\cdot|\boldsymbol{\theta}\right) = \left(\begin{array}{c} \left(\mathbf{I} + \sum_{t=1}^{\tau}\Sigma^{t-1}\right)\mathbf{a} + \mathbf{I}_{\boldsymbol{\theta}}\left(\mathbf{I} + \sum_{t=1}^{\tau}\gamma^2\left(\mathbf{I}_{\mathbf{e}}\Sigma\mathbf{I}_{\mathbf{e}}\right)^{t-1}\right)\mathbf{e} \\ + \left[\sum_{k\neq i}\cdots\sum_{h\neq j}\sigma_{ik}\cdots\sigma_{jh}\mathbb{E}_i\left[\cdots\mathbb{E}_j\left[\cdot|\theta_j\right]|\theta_i\right]\right] \end{array}\right).$$

When each $|\sigma_{ij}| < 1$ the bottom term will converge to zero provided strategies are bounded. More generally, we require the following property to hold.

Definition 3 (non-explosive expectations)**.** *For any sequence of players $(i_1, i_2, \ldots)$ with $i_t \neq i_{t+1}$ and each $i_t \in \{1, \ldots, N\}$ and $t \in \mathbb{N}$, the operator $\overline{\mathbb{E}}_{i_t}[\cdot]$ is defined inductively as $\overline{\mathbb{E}}_{i_t}[\cdot] := \sigma_{i_{t-1}i_t}\overline{\mathbb{E}}_{i_{t-1}}[\mathbb{E}_{i_t}[\cdot]]$, with $\overline{\mathbb{E}}_{i_1}[\cdot] = \mathbb{E}_{i_1}[\cdot]$. Then for any given (potentially non-linear) IRE $\mathbf{X}^*$ and quality profile $\mathbf{e}$, expectations over the network are* non-explosive *if* $\lim_{t\to\infty}\overline{\mathbb{E}}_{i_t}\left[X_{i_t}\left(\theta_{i_t}|e_{i_t}\right)\right] = 0$.

[44] $\mathbf{I}_{\boldsymbol{\phi}}$ gives the diagonal matrix with elements from generic vector $\boldsymbol{\phi}$.

Given expectations are non-explosive, we then obtain:

$$
\begin{aligned}
\lim_{\tau\to\infty} BR^{\tau}\left(\cdot|\boldsymbol{\theta}\right) &= \lim_{\tau\to\infty}\left(\mathbf{I}+\sum_{\tau=1}^{k}\Sigma^{\tau}\right)\mathbf{a}+\mathbf{I}_{\boldsymbol{\theta}}\left(\mathbf{I}+\sum_{\tau=1}^{k}\gamma^{2}\left(\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{\tau}\right)\mathbf{e} \\
&= \left(\mathbf{I}-\Sigma\right)^{-1}\mathbf{a}+\mathbf{I}_{\boldsymbol{\theta}}\left(\mathbf{I}-\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{e} \\
&= \boldsymbol{\alpha}^{*}+\mathbf{I}_{\boldsymbol{\theta}}\boldsymbol{\beta}^{*}=:\mathbf{X}^{*}
\end{aligned}
$$

which gives the unique linear information response equilibrium of Theorem 1. Thus, any equilibrium in which expectations are non-explosive must be $\mathbf{X}^{*}$.

□

Proof of Theorem 2. Writing each player $k \neq i$'s information response strategy as $X_k^{*}\left(\theta_k|e_k\right)=\alpha_k^{*}+\theta_k\beta_k^{*}$:

$$
\begin{aligned}
u_i\left(x_i,\mathbf{X}^{*}_{-i}|\theta_i,e_i,\mu_i\right) &= \left(a_i+\mathbb{E}_i\left[\tilde{\omega}_i|\theta_i,e_i\right]\right)x_i-\frac{1}{2}x_i^{2}+\sum_{k\neq i}\sigma_{ik}x_i\mathbb{E}_i\left[X_k^{*}\left(\theta_k|e_k\right)|\theta_i,e_i,\mu_i\right] \\
&= \left(a_i+e_i\theta_i\right)x_i-\frac{1}{2}x_i^{2}+\sum_{k\neq i}\sigma_{ik}x_i\left(\alpha_k^{*}+\beta_k^{*}\mathbb{E}_i\left[\theta_k|\theta_i,e_i,\mu_i\right]\right) \\
&= \left(a_i+e_i\theta_i\right)x_i-\frac{1}{2}x_i^{2}+\sum_{k\neq i}\sigma_{ik}x_i\left(\alpha_k^{*}+\beta_k^{*}\gamma^{2}e_ie_k\theta_i\right).
\end{aligned}
$$

By the optimality of $\mathbf{X}_i^{*}$ in stage two, we can apply the envelope theorem:

$$
\frac{\partial}{\partial\beta_i^{*}}\mathbb{E}_i\left[u_i\left(\mathbf{X}^{*}\left(\boldsymbol{\theta}|\mathbf{e}\right)|\omega,\omega_i\right)|\theta_i,e_i,\mu_i\right]=0.
$$

Further, as information acquisition is unobserved by others in $t=2$, incentive compatibility of e_i^* requires that the response of β_j^* to shifting e_i be set to zero: $\frac{\partial}{\partial e_i}\beta_j^* = 0$. Thus we obtain:

$$\begin{aligned}
& \frac{\partial}{\partial e_i}\mathbb{E}_i\left[u_i\left(\mathbf{X}^*\left(\boldsymbol{\theta}|\mathbf{e}\right)|\omega,\omega_i\right)|e_i,\mu_i\right] \\
= & \frac{\partial}{\partial e_i}\mathbb{E}_i\left[\mathbb{E}_i\left[u_i\left(\mathbf{X}^*\left(\boldsymbol{\theta}|\mathbf{e}\right)|\omega,\omega_i\right)|\theta_i,e_i,\mu_i\right]\right] \\
= & \frac{\partial}{\partial e_i}\mathbb{E}_i\left[\begin{array}{c}(a_i+e_i\theta_i)\left(\alpha_i^*+\theta_i\beta_i^*\right)-\frac{1}{2}\left(\alpha_i^*+\theta_i\beta_i^*\right)^2 \\ +\sum_{k\neq i}\sigma_{ik}\left(\alpha_i^*+\theta_i\beta_i^*\right)\left(\alpha_k^*+\beta_k^*\gamma^2 e_i e_k\theta_i\right)\end{array}\right] \\
= & \frac{\partial}{\partial e_i}\mathbb{E}_i\left[\begin{array}{c}\left(\beta_i^* e_i-\frac{1}{2}\beta_i^{*2}+\gamma^2\sum_{k\neq i}\sigma_{ik}e_i e_k\beta_i^*\beta_k^*\right)\theta_i^2 \\ +\text{const}_0+\text{const}_1\cdot\theta_i\end{array}\right] \\
= & \mathbb{E}_i\left[\frac{\partial}{\partial e_i}\left(\begin{array}{c}\left(\beta_i^* e_i-\frac{1}{2}\beta_i^{*2}+\gamma^2\sum_{k\neq i}\sigma_{ik}e_i e_k\beta_i^*\beta_k^*\right)\theta_i^2 \\ +\text{const}_0\end{array}\right)\right] \\
= & \left(\begin{array}{c}\beta_i^*\left(1+\gamma^2\sum_{k\neq i}\sigma_{ik}\beta_k^* e_k^*\right)\mathbb{E}_{\theta_i}\left[\theta_i^2|e_i\right] \\ +\left(\beta_i^* e_i-\frac{1}{2}\beta_i^{*2}+\gamma^2\sum_{k\neq i}\sigma_{ik}\beta_i^*\beta_k^* e_i e_k\right)\frac{\partial}{\partial e_i}\mathbb{E}_{\theta_i}\left[\theta_i^2|e_i\right]\end{array}\right) \\
= & \beta_i^*\left(1+\gamma^2\sum_{k\neq i}\sigma_{ik}e_k\beta_k^*\right),
\end{aligned}$$

with $\mathbb{E}_{\theta_i}\left[\theta_i^2|e_i\right]=1$ and $\frac{\partial}{\partial e_i}\mathbb{E}_{\theta_i}\left[\theta_i^2|e_i\right]=0$ by condition E3. This yields i's period $t=1$ marginal gains to information acquisition:

$$\frac{\partial}{\partial e_i}u_i\left(\mathbf{X}_i^*|e_i,\mathbf{e}_{-i}\right)=\beta_i^*\left(1+\gamma^2\sum_{k\neq i}\sigma_{ik}e_k\beta_k^*\right) \tag{26}$$

Thus, the period $t=1$ vector of marginal gains to quality is given by:

$$\left[\frac{\partial}{\partial e_i}u_i\left(\mathbf{X}_i^*|e_i,\mathbf{e}_{-i}\right)\right]=\gamma^2\mathbf{I}_{\boldsymbol{\beta}^*}\Sigma\mathbf{I}_{\mathbf{e}}\boldsymbol{\beta}^*+\boldsymbol{\beta}^*. \tag{27}$$

When $\gamma=0$ then (25) reduces to $\boldsymbol{\beta}^*(\mathbf{e})=\mathbf{e}$. Equating marginal gains to marginal costs of quality in IAE gives $\mathbf{e}^*=\kappa'(\mathbf{e}^*)$, which corresponds to expression (6), and yields $e^\dagger$ from (3) for each i so that each player chooses the quality that the isolated player chooses.

When $\gamma>0$ then (25) can be rearranged as:

$$\gamma^2\Sigma\mathbf{I}_{\mathbf{e}}\boldsymbol{\beta}^*=\mathbf{I}_{\mathbf{e}}^{-1}\left(\boldsymbol{\beta}^*-\mathbf{e}\right). \tag{28}$$

Substituting this into (27) gives the marginal gains to information:

$$\left[\frac{\partial}{\partial e_i}u_i\left(\mathbf{X}_i^*|e_i,\mathbf{e}_{-i}\right)\right]=\mathbf{I}_{\mathbf{e}}^{-1}\mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^*.$$

Equating this with the marginal cost of information then gives the first-stage interior IAE condition (6):

$$\mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^* = \mathbf{I}_{\mathbf{e}^*}\kappa'\left(\mathbf{e}^*\right). \tag{29}$$

□

Proof of Corollary 1. Applying the implicit function theorem to expression (6)[45]:

$$\begin{aligned}
\frac{\partial e_i^*}{\partial\left(\gamma^2\right)} &= -\frac{\frac{\partial\left(\beta_i^{*2}/e_i^*\right)}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^*-\kappa'\left(e_i^*\right)\right)}{\partial e_i}} + \sum_{k\neq i}\frac{\partial e_i^*}{\partial\beta_k^*}\frac{\partial\beta_k^*}{\partial\left(\gamma^2\right)} \\
&= -\frac{2\beta_i^*/e_i^*\frac{\partial\beta_i^*}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^*\right)}{\partial e_i} - \kappa''\left(e_i^*\right)} + \sum_{k\neq i}\frac{\partial e_i^*}{\partial\beta_k^*}\frac{\partial\beta_k^*}{\partial\left(\gamma^2\right)} \\
&= \frac{2\frac{\beta_i^*}{e_i^*}\frac{\partial\beta_i^*}{\partial(\gamma^2)}}{\kappa''\left(e_i^*\right) - \frac{\left(e_i^*2\beta_i^*\frac{\partial\beta_i^*}{\partial e_i}-\beta_i^{*2}\right)}{e_i^2}} + \sum_{k\neq i}\frac{\partial e_i^*}{\partial\beta_k^*}\frac{\partial\beta_k^*}{\partial\left(\gamma^2\right)} \\
&= \frac{2\frac{\beta_i^*}{e_i^*}\sum_{k\neq i}e_i^*\sigma_{ik}e_k^*\beta_k^*}{\kappa''\left(e_i^*\right) - \frac{\left(e_i^*2\beta_i^*\frac{\partial\beta_i^*}{\partial e_i}-\beta_i^{*2}\right)}{e_i^2}} + \sum_{k\neq i}\frac{\partial e_i^*}{\partial\beta_k^*}\frac{\partial\beta_k^*}{\partial\left(\gamma^2\right)}.
\end{aligned}$$

Taking the limit $\gamma\to^+0$ of the expression, and noting that $\lim_{\gamma\to^+0}\frac{\partial\beta_i^*}{\partial e_i}=1$ while $\lim_{\gamma\to^+0}\frac{\partial e_i^*}{\partial\beta_k^*}=0$ for each $k\neq i$, yields:

$$\lim_{\gamma\to^+0}\frac{\partial e_i^*}{\partial\left(\gamma^2\right)} = \frac{2e^{\dagger 3}\sum_{k\neq i}\sigma_{ik}}{\kappa''\left(e^\dagger\right)-1}.$$

Note that $\kappa''\left(e^\dagger\right)-1>0$ by the optimality of $e^\dagger$ at $\gamma=0$ and Assumption 1.

□

Proof of Lemma 1. The existence of the bound γ^m follows from Assumption 1, by continuity in $\boldsymbol{\beta}^*$ and $\mathbf{e}^*$ for each i at $\gamma=0$, and by the implicit function theorem. Precisely, $\boldsymbol{\beta}^*=\mathbf{e}^*:=\left(e^\dagger,\ldots,e^\dagger\right)>\mathbf{0}$ when $\gamma=0$, and thus that marginal gains to quality β_i^{*2}/e_i^* are continuous at $\gamma=0$. Assumption 1 implies a unique $e^\dagger$ solving $\beta_i^{*2}/e_i^*=e^\dagger=\kappa'\left(e^\dagger\right)$ for each i. Further,

$$\begin{aligned}
\det\left(D_{\mathbf{e}}[\beta_i^{*2}/e_i^*-\kappa'\left(e_i^*\right)]\right)\Big|_{\left(\mathbf{e}=\left(e^\dagger,\ldots,e^\dagger\right),\gamma=0\right)} &= \det\left(\left(1-\kappa'\left(e^\dagger\right)\right)\mathbf{I}\right) \\
&= \left(1-\kappa'\left(e^\dagger\right)\right)^N\neq 0,
\end{aligned}$$

and thus by the IFT there exists an open neighborhood $U\subseteq[0,1]^N$ of $\left(e^\dagger,\ldots,e^\dagger\right)$ and $W\subseteq[0,1]$ of $\gamma=0$ such that for every $\gamma\in W$ there is a unique IAE $\mathbf{e}^{*,\gamma}\in U$.

[45] One could employ the multivariate implicit function theorem, noting that changes in e_i^* will result as second-stage β_k^* for each $k\neq i$ adjust with γ^2. We avoid the multivariate implicit function theorem by employing the chain rule, and summing over partials of e_i^* with respect to β_k^* for each $k\neq i$ (last term).

Now, the best response correspondence $\overline{BR}(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*; \gamma)$ (see proof of Proposition S.1) is upper hemicontinuous in $(\mathbf{e},\gamma)$ by continuity of β_i^{*2}/e_i in $\mathbf{e}$ and γ and of $\kappa(\cdot)$ in e_i at $\mathbf{e} = \left(e^\dagger, \ldots, e^\dagger\right)$ and $\gamma = 0$. There must then also exist some neighborhood $V \subseteq [0,1]^N \times [0,1]$ of $(\left(e^\dagger, \ldots, e^\dagger\right), 0)$ such that $\overline{BR}(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*; \gamma) \subseteq U$ for any $(\mathbf{e}, \gamma) \in V$. This then implies that $[0,1]^N \backslash U$ does not contain any IAE for all $\gamma \in W \cap V \subseteq [0,1]$, and thus that $\mathbf{e}^{*,\gamma}$ gives the unique IAE for each $\gamma \in [0, \gamma^m) \subseteq W \cap V$.

We construct the interval $[0, \gamma^s)$ as follows, which incorporates the potential for multiplie equilibria. $\boldsymbol{\beta}^*$ is continuous in $\mathbf{e}$ with $\boldsymbol{\beta}^* = \left(e^\dagger, \ldots, e^\dagger\right) > \mathbf{0}$ at $\gamma = 0$. Thus for each i, there must exist some $\gamma_i^s > 0$ such that if $\gamma < \gamma_i^s$ then $\beta_i^* > 0$ for any $\mathbf{e} \in [0,1]^N$.[46] Then defining $\gamma^s := \min_i \{\gamma_i^s\}$ and by the existence of IAE given with the proof of Proposition S.1 below, we must have that $\boldsymbol{\beta}^* > \mathbf{0}$ provided $\gamma \in [0, \gamma^s)$ and any IAE $\mathbf{e}^*$.

□

F.3 Section 17 proofs: Equilibrium welfare and the strategic value to information

First we derive equilibrium welfare, expression (8) in the text. Restating player i's expected payoff:

$$u_i(x_i, \mathbf{X}_{-i} | \theta_i, \mathbf{e}) = (a_i + e_i\theta_i)\, x_i - \frac{1}{2} x_i^2 + \sum_{k \neq i} \sigma_{ik} x_i \left(\alpha_{kj}^* + \beta_k^* \gamma^2 e_i e_k \theta_i\right).$$

Subtracting information cost $\kappa(e_i)$ and taking expectations over signals $\boldsymbol{\theta}$ gives her period $t = 1$ value:

$$\begin{aligned}
\nu_i(\mathbf{X}_i^* | e_i, \mathbf{e}_{-i}) &= \mathbb{E}_i \left[\left(\begin{array}{c} (a_i + e_i\theta_i)(\alpha_i^* + \theta_i\beta_i^*) - \frac{1}{2}(\alpha_i^* + \theta_i\beta_i^*)^2 \\ + \sum_{k \neq i} \sigma_{ik} (\alpha_i^* + \theta_i\beta_i^*)(\alpha_k^* + \beta_k^* \gamma^2 e_i e_k \theta_i) \end{array} \right) \right] - \kappa(e_i) \\
&= a_i\alpha_i^* + e_i\beta_i^* - \frac{1}{2}\left(\alpha_i^{*2} + \beta_i^{*2}\right) + \sum_{k \neq i} \sigma_{ik} \left(\alpha_i^*\alpha_k^* + \beta_i^*\beta_k^* \gamma^2 e_i e_k\right) - \kappa(e_i).
\end{aligned}$$

Writing this in vector form gives:

$$\nu(\mathbf{X}^* | \mathbf{e}) = \left(\mathbf{I_a}\boldsymbol{\alpha}^* + \mathbf{I_e}\boldsymbol{\beta} - \frac{1}{2}\left(\mathbf{I}_{\boldsymbol{\alpha}^*}\bar{\mathbf{X}}^* + \mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^*\right) + \mathbf{I}_{\boldsymbol{\alpha}^*}\Sigma\boldsymbol{\alpha}^* + \gamma^2 \mathbf{I}_{\boldsymbol{\beta}^*}\mathbf{I_e}\Sigma\mathbf{I_e}\boldsymbol{\beta}^* \right) - \kappa(\mathbf{e}). \quad (30)$$

Next, left multiplying (22) by $\mathbf{I}_{\boldsymbol{\alpha}^*}$ gives:

$$\mathbf{I}_{\boldsymbol{\alpha}^*}\boldsymbol{\alpha}^* = \mathbf{I_a}\boldsymbol{\alpha}^* + \mathbf{I}_{\boldsymbol{\alpha}^*}\Sigma\boldsymbol{\alpha}^*, \quad (31)$$

while rearranging (25) gives:

$$\frac{1}{\gamma^2}(\boldsymbol{\beta}^* - \mathbf{e}) = \mathbf{I_e}\Sigma\mathbf{I_e}\boldsymbol{\beta}^*. \quad (32)$$

[46] This uses Assumption F1 to maintain that $\boldsymbol{\beta}^*$ is well defined for each $\mathbf{e} \in [0,1]^N$.

Substituting (31) and (32) into (30) then gives:

$$\begin{aligned}\nu\left(\mathbf{X}^*|\mathbf{e}\right) &= \left(\mathbf{I}_{\bar{\mathbf{X}}^*}\boldsymbol{\alpha}^* + \mathbf{I}_{\mathbf{e}}\boldsymbol{\beta}^* - \frac{1}{2}\left(\mathbf{I}_{\boldsymbol{\alpha}^*}\boldsymbol{\alpha}^* + \mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^*\right) + \mathbf{I}_{\boldsymbol{\beta}^*}\left(\boldsymbol{\beta}^* - \mathbf{e}\right)\right) - \kappa\left(\mathbf{e}\right) \\ &= \frac{1}{2}\left(\mathbf{I}_{\boldsymbol{\alpha}^*}\boldsymbol{\alpha}^* + \mathbf{I}_{\boldsymbol{\beta}^*}\boldsymbol{\beta}^*\right) - \kappa\left(\mathbf{e}\right),\end{aligned}$$

giving expression (8).

For the proofs of Lemma 1 and Proposition 3A we next derive expressions for partials $\frac{\partial \boldsymbol{\beta}^*}{\partial e_i}$. This yields expressions for $\xi_i^{st}\left(\mathbf{e}, \mathbf{X}^*\right)$ and $\xi_i^{ex}\left(\mathbf{e}, \mathbf{X}^*\right)$ solely in terms of Σ and $\mathbf{e}$

Using u and v for row and column dummies (respectively) the system of equations giving IRE $\boldsymbol{\beta}^*$ can be written as:

$$[u]\text{: } \beta_u^* - e_u\left(1 + \gamma^2 \sum_{k \neq u} \sigma_{uk} e_k \beta_k^*\right) = 0,$$

for each $u \in \{1, \ldots, N\}$. Partial differentiating each $[u]$ by β_v^* gives:

$$[f_{uv}]\text{: } \frac{\partial[u]}{\partial \beta_v^*} = \begin{cases} -\gamma^2 e_u \sigma_{uv} e_v & \text{if } u \neq v \\ 1 & \text{if } u = v \end{cases},$$

for each $u, v \in \{1, \ldots, N\}$. In matrix form this is exactly $\mathbf{I} - \gamma^2 \mathbf{I}_{\mathbf{e}} \Sigma \mathbf{I}_{\mathbf{e}}$. Partial differentiating each $[u]$ by e_i gives:

$$[du]\text{: } \sum_v f_{uv} \frac{\partial \beta_v^*}{\partial e_i} + b_u = 0,$$

for each $du \in \{d1, \ldots, dN\}$, where

$$b_u := \frac{\partial[u]}{\partial e_i} = -\frac{\beta_i^*}{e_i} \cdot \begin{cases} \gamma^2 e_u \sigma_{ui} e_i & \text{if } u \neq i \\ 1 & \text{if } u = i \end{cases}.$$

In vector form $\mathbf{b}$ gives $\frac{\beta_i^*}{e_i}\left(\mathbf{I} - \gamma^2 \mathbf{I}_{\mathbf{e}} \Sigma \mathbf{I}_{\mathbf{e}} - 2\mathbf{I}\right)\mathbf{1}_i$, where $\mathbf{1}_u$ gives the vector of zeros with a one in row u. Solving for $\frac{\partial \beta_u^*}{\partial e_i}$ in matrix form gives the comparative static of $\boldsymbol{\beta}^*$ with respect to

e_i:[47]

$$
\begin{aligned}
\frac{\partial \boldsymbol{\beta}^*}{\partial e_i} &= -F^{-1}\mathbf{b} \\
&= -\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \left(\frac{\beta_i^*}{e_i}\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e} - 2\mathbf{I}\right)\mathbf{1}_i\right) \\
&= -\frac{\beta_i^*}{e_i}\left(\mathbf{I} - 2\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1}\right)\mathbf{1}_i \\
&= \frac{\beta_i^*}{e_i}\left(2\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} - \mathbf{I}\right)\mathbf{1}_i \qquad (33) \\
&= \frac{\beta_i^*}{e_i}\left(\mathbf{I} + \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1}\mathbf{1}_i. \qquad (34)
\end{aligned}
$$

$\xi_i^{st}(\mathbf{e}, \mathbf{X}^*)$ and $\xi_i^{ex}(\mathbf{e}, \mathbf{X}^*)$ can be expressed solely in terms of Σ and $\mathbf{e}$ by substituting (33) into the following expressions:

$$
\begin{aligned}
\xi_i^{st}(\mathbf{e}, \mathbf{X}^*) \quad &: \quad = \beta_i^* \sum_{k \neq i} \gamma^2 e_i e_k \sigma_{ik} \frac{\partial}{\partial e_i} \beta_k^* \\
&= \quad \beta_i^* \mathbf{1}_i' \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e} \frac{\partial \boldsymbol{\beta}^*}{\partial e_i} \\
\xi_i^{ex}(\mathbf{e}, \mathbf{X}^*) \quad &: \quad = \sum_{k \neq i} \beta_k^* \frac{\partial}{\partial e_i} \beta_k^* \\
&= \quad \left(\boldsymbol{\beta}^* - \mathbf{1}_{\beta_i^*}\right)' \frac{\partial \boldsymbol{\beta}^*}{\partial e_i}.
\end{aligned}
$$

[47]An equivalent setup of the above is provided in Takayama (1985) [57], pgs. 403-5.

For $\xi_i^{st}(\mathbf{e},\mathbf{X}^*)$ we have:

$$\begin{aligned}
\xi_i^{st}(\mathbf{e},\mathbf{X}^*) &= \beta_i^* \mathbf{1}_i' \gamma^2 \mathbf{I_e}\Sigma\mathbf{I_e}\frac{\partial\boldsymbol{\beta}^*}{\partial e_i} \\
&= \frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}+\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{1}_i \\
&= \frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}+\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)\left(\sum_{\tau=0}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= \frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}+2\left(\sum_{\tau=1}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\right)\mathbf{1}_i \\
&= \frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_{e^*}}+2\left(\sum_{\tau=2}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\right)\mathbf{1}_i \\
&= 2\frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\left(\sum_{\tau=2}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= 2\frac{\beta_i^{*2}}{e_i}\mathbf{1}_i'\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\mathbf{1}_i.
\end{aligned}$$

For $\xi_i^{ex}(\mathbf{e},\mathbf{X}^*)$ we have:

$$\begin{aligned}
\xi_i^{ex}(\mathbf{e},\mathbf{X}^*) &= \frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\left(\mathbf{I}+\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{1}_i \\
&= \frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\left(\mathbf{I}+\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)\left(\sum_{\tau=0}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= \frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\left(\mathbf{I}+2\sum_{\tau=1}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= \frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\left(2\sum_{\tau=1}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= 2\frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\left(\sum_{\tau=1}\left(\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^\tau\right)\mathbf{1}_i \\
&= 2\frac{\beta_i^*}{e_i}\left(\boldsymbol{\beta}^*-\mathbf{1}_{\beta_i^*}\right)'\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{1}_i.
\end{aligned}$$

Together:

$$\begin{aligned}
\xi_i^{st}(\mathbf{e},\mathbf{X}^*) &= \gamma^4 2\frac{\beta_i^{*2}}{e_i^*}\mathbf{1}_i'\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{I_e}\Sigma\mathbf{I_e}\mathbf{1}_i, \\
\xi_i^{ex}(\mathbf{e},\mathbf{X}^*) &= \gamma^2 2\frac{\beta_i^*}{e_i^*}\left(\boldsymbol{\beta}^*-\beta_i^*\mathbf{1}_i\right)'\mathbf{I_e}\Sigma\mathbf{I_e}\left(\mathbf{I}-\gamma^2\mathbf{I_e}\Sigma\mathbf{I_e}\right)^{-1}\mathbf{1}_i.
\end{aligned}$$

One can also use Theorem 1 to substitute in corresponding expressions for $\boldsymbol{\beta}^*$ and β_i^*, respectively, that are solely in terms of Σ and $\mathbf{e}$.

Lemma 1 is established using the leading term of the Taylor expansion:

$$\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} = \sum_{\tau=0}^{\infty} \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{\tau},$$

which will dominate the sum for small γ. Formal proofs are as follows.

Proof of Lemma 1 and derivations of (18) and (19). We can rewrite the expression for $\xi_i^{st}(\mathbf{X}^*, \mathbf{e})$ by expanding $\left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1}$ as follows:

$$\begin{aligned}
\xi_i^{st}(\mathbf{e}, \mathbf{X}^*) &= \gamma^4 2 \frac{\beta_i^{*2}}{e_i^*} \mathbf{1}_i' \mathbf{I_e} \Sigma \mathbf{I_e} \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \mathbf{I_e} \Sigma \mathbf{I_e} \mathbf{1}_i && (35)\\
&= 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\sum_{\tau=2}^{\infty} \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{\tau} \right) \mathbf{1}_i && (36)\\
&= 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 + \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^3 \sum_{\tau=0}^{\infty} \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{\tau} \right) \mathbf{1}_i && (37)\\
&= \gamma^4 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 \mathbf{1}_i + \gamma^6 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\mathbf{I_{e^*}} \Sigma \mathbf{I_e}\right)^3 \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \mathbf{1}_i. && (38)
\end{aligned}$$

For the second term:

$$\frac{\partial}{\partial(\gamma^4)} \left(\gamma^6 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^3 \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \mathbf{1}_i \right) \to 0,$$

for each i, as $\gamma \to 0$. Thus focusing on the first term:

$$\begin{aligned}
\gamma^4 2 \frac{\beta_i^{*2}}{e_i} \mathbf{1}_i' \left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 \mathbf{1}_i &= \gamma^2 2 \frac{\beta_i^{*}}{e_i} \mathbf{1}_i' \left[\sum_{k \neq i} e_i \sigma_{ik} e_k e_k \sigma_{kj} e_j \right] \mathbf{1}_i \\
&= \gamma^4 2 \frac{\beta_i^{*2}}{e_i} \sum_{k \neq i} e_i \sigma_{ik} e_k e_k \sigma_{ki} e_i. && (39)
\end{aligned}$$

Taking a partial derivative of (38) with respect to γ^2, and with $e_i^* \to e^\dagger$ as $\gamma \to 0$ for each i, we obtain expression (16):

$$\lim_{\gamma \to^+ 0} \frac{\partial \xi_i^{st}(\mathbf{e}^*, \mathbf{X}^*)}{\partial(\gamma^4)} = 2 e^{\dagger 5} \sum_{k \neq i} \sigma_{ik} \sigma_{ki}.$$

For symmetric Σ (Assumption 2A) with $\sigma_{ki} = \sigma_{ik}$, we can rewrite (41) to give expression (18), as well as the corresponding (negated) expression under network anti-symmetry (Assumption 2B).

Next, we can rewrite the expression for $\xi_i^{ex}(\mathbf{X}^*, \mathbf{e})$, again expanding $(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e})^{-1}$:

$$\begin{aligned}
\xi_i^{ex}(\mathbf{e}, \mathbf{X}^*) &= \gamma^2 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \mathbf{I_e} \Sigma \mathbf{I_e} \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \mathbf{1}_i \\
&= 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \left(\sum_{\tau=1}^{\infty} \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{\tau} \right) \mathbf{1}_i \\
&= 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e} + \left(\gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 \sum_{\tau=0}^{\infty} \gamma^{2\tau} \left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^{\tau} \right) \mathbf{1}_i \\
&= \gamma^2 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \left(\mathbf{I_e} \Sigma \mathbf{I_e} + \gamma^2 \left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \right) \mathbf{1}_i. \qquad (40)
\end{aligned}$$

For the second term:

$$\frac{\partial}{\partial(\gamma^2)} \left(\gamma^4 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \left(\left(\mathbf{I_e} \Sigma \mathbf{I_e}\right)^2 \left(\mathbf{I} - \gamma^2 \mathbf{I_e} \Sigma \mathbf{I_e}\right)^{-1} \right) \mathbf{1}_i \right) \to 0,$$

for each i, as $\gamma \to 0$. Focusing again on the first term:

$$\begin{aligned}
\gamma^2 2 \frac{\beta_i^*}{e_i} (\boldsymbol{\beta}^* - \beta_i^* \mathbf{1}_i)' \mathbf{I_e} \Sigma \mathbf{I_e} \mathbf{1}_i &= \gamma^2 2 \frac{\beta_i^*}{e_i} \boldsymbol{\beta}^{*\prime} \mathbf{I_e} \Sigma \mathbf{I_e} \mathbf{1}_i \\
&= \gamma^2 2 \frac{\beta_i^*}{e_i} \left(\left[\sum_{k \neq i} e_k \sigma_{kj} e_j \beta_k^* \right]_{j=1}^{N} \right)' \mathbf{1}_i \\
&= \gamma^2 2 \frac{\beta_i^*}{e_i} \sum_{k \neq i} e_k \sigma_{ki} e_i \beta_k^*. \qquad (41)
\end{aligned}$$

Then:

$$\lim_{\gamma \to {}^+0} \frac{\partial \xi_i^{ex}(\mathbf{e}^*, \mathbf{X}^*)}{\partial(\gamma^2)} = 2 e^{\dagger 3} \sum_{k \neq i} \sigma_{ki},$$

for each i, yielding expression (17).

For symmetric Σ with $\sigma_{ki} = \sigma_{ik}$, we can rewrite (41) to give expression (18):

$$\xi_i^{ex}(\mathbf{e}, \mathbf{X}^*) \approx \gamma^2 2 \frac{\beta_i^*}{e_i} \sum_{k \neq i} e_k \sigma_{ik} e_i \beta_k^* = \gamma^2 2 \frac{\beta_i^*}{e_i} (\beta_i^* - e_i),$$

with the second equality using Theorem 1. This also yields the corresponding (negated) expression under network anti-symmetry (Assumption 2B).

□

Proof of Propositions 3A and 3B. For part 1 of Proposition 3A, apply the implicit function theorem to the difference $(e_i^{pb}-e_i^*)$ to give[48]:

$$\frac{\partial\left(e_i^{pb}-e_i^*\right)}{\partial\left(\gamma^4\right)}=-\frac{\frac{\partial\left(\beta_i^{*2}/e_i^{pb}+\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)\right)}{\partial(\gamma^4)}}{\frac{\partial\left(\beta_i^{*2}/e_i^{pb}+\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)-\kappa'\left(e_i^{pb}\right)\right)}{\partial e_i}}+\frac{\frac{\partial\left(\beta_i^{*2}/e_i^*\right)}{\partial(\gamma^4)}}{\frac{\partial\left(\beta_i^{*2}/e_i^*-\kappa'\left(e_i^*\right)\right)}{\partial e_i}}+\sum_{k\neq i}\left(\frac{\partial e_i^{pb}}{\partial\beta_k^*}-\frac{\partial e_i^*}{\partial\beta_k^*}\right)\frac{\partial\beta_k^*}{\partial\left(\gamma^2\right)}.$$

Taking the limit $\gamma\to^+0$ of the expression, $\lim_{\gamma\to^+0}\frac{\partial\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)}{\partial e_i^{pb}}=0$, because $\xi_i^{ex}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)=0$ at $\gamma=0$, and $\xi_i^{ex}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)$ is $\mathcal{C}^1$ in γ. Thus, the denominators of the first two terms converge to $\kappa''\left(e^\dagger\right)-1$, as in the proof of Corollary 1. With $e_i^*\to e_i^{pb}$ as $\gamma\to^+0$ with both e_i^* and e_i^{pb} $\mathcal{C}^1$ in γ, $\frac{\partial\left(\beta_i^{*2}/e_i^{pb}\right)}{\partial(\gamma^2)}\to\frac{\partial\left(\beta_i^{*2}/e_i^*\right)}{\partial(\gamma^2)}$. Again noting that $\lim_{\gamma\to^+0}\frac{\partial\beta_i^*}{\partial e_i^*}=1$ while $\lim_{\gamma\to^+0}\frac{\partial e_i^*}{\partial\beta_k^*}=\lim_{\gamma\to^+0}\frac{\partial e_i^{pb}}{\partial\beta_k^*}=0$ for each $k\neq i$, implying that the second sum converges to zero as $\gamma\to^+0$, this leaves:

$$\lim_{\gamma\to^+0}\frac{\partial\left(e_i^{pb}-e_i^*\right)}{\partial\left(\gamma^4\right)}=\frac{\frac{\partial\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)}{\partial(\gamma^4)}}{\kappa''\left(e^\dagger\right)-1}=\frac{e^{\dagger5}\sum_{k\neq i}\sigma_{ik}^2}{\kappa''\left(e^\dagger\right)-1}>0.$$

The second equality following from Lemma 1. By continuity of all functions in γ, this positivity must hold for some neighborhood of $\gamma=0$.

A similar expression can be derived for j, giving:

$$\lim_{\gamma\to^+0}\frac{\partial\left(e_i^{pb}-e_i^*\right)}{\partial\left(\gamma^4\right)}-\lim_{\gamma\to^+0}\frac{\partial\left(e_j^{pb}-e_j^*\right)}{\partial\left(\gamma^4\right)}=\frac{e^{\dagger5}\left(\sum_{k\neq i}\sigma_{ik}^2-\sum_{k\neq j}\sigma_{jk}^2\right)}{\kappa''\left(e^\dagger\right)-1}>0,$$

the final inequality following by assumption: $\sum_{k\neq i}\sigma_{ik}^2>\sum_{k\neq j}\sigma_{jk}^2$. Again, by continuity of all functions in γ, this positivity must hold for some neighborhood of $\gamma=0$.

For part 2 of Proposition 3A, again apply the implicit function theorem to the difference $(e_i^{pl}-e_i^*)$ to give:

$$\frac{\partial\left(e_i^{pl}-e_i^*\right)}{\partial\left(\gamma^2\right)}=\left(\begin{array}{c}\left(-\frac{\frac{\partial\left(\beta_i^{*2}/e_i^{pl}+\xi_i^{st}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)+\xi_i^{ex}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)\right)}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^{pl}+\xi_i^{st}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)+\xi_i^{ex}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)-\kappa'\left(e_i^{pl}\right)\right)}{\partial e_i}}+\frac{\frac{\partial\left(\beta_i^{*2}/e_i^*\right)}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^*+\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)-\kappa'\left(e_i^*\right)\right)}{\partial e_i}}\right)\\+\sum_{k\neq i}\left(\frac{\partial e_i^{pl}}{\partial\beta_k^*}-\frac{\partial e_i^*}{\partial\beta_k^*}\right)\frac{\partial\beta_k^*}{\partial(\gamma^2)}\end{array}\right).$$

$$\frac{\partial\left(e_i^{pl}-e_i^*\right)}{\partial\left(\gamma^2\right)}=\left(\begin{array}{c}\left(-\frac{\frac{\partial\left(\beta_i^{*2}/e_i^{pl}+\xi_i^{st}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)+\xi_i^{ex}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)\right)}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^{pl}+\xi_i^{st}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)+\xi_i^{ex}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)-\kappa'\left(e_i^{pl}\right)\right)}{\partial e_i}}+\frac{\frac{\partial\left(\beta_i^{*2}/e_i^*+\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)\right)}{\partial(\gamma^2)}}{\frac{\partial\left(\beta_i^{*2}/e_i^*+\xi_i^{st}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)-\kappa'\left(e_i^*\right)\right)}{\partial e_i}}\right)\\+\sum_{k\neq i}\left(\frac{\partial e_i^{pl}}{\partial\beta_k^*}-\frac{\partial e_i^{pb}}{\partial\beta_k^*}\right)\frac{\partial\beta_k^*}{\partial(\gamma^2)}\end{array}\right).$$

[48] One could employ the multivariate implicit function theorem, noting that changes in e_i^{pb} and e_i^* will result as second-stage β_k^* for each $k\neq i$ adjust with γ^2. We avoid the multivariate implicit function theorem finding the total derivative, summing over partials of e_i^{pb} and e_i^* with respect to β_k^* for each $k\neq i$ (last term).

Taking the limit $\gamma \to^{+} 0$ of the expression, $\lim_{\gamma\to^{+}0} \frac{\partial \xi_i^{ex}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)}{\partial e_i^{pb}} = 0$, $\lim_{\gamma\to^{+}0} \frac{\partial \xi_i^{st}\left(\mathbf{e}^{pl},\mathbf{X}^*\right)}{\partial(\gamma^2)} = 0$ from (38), while $e_i^* \to e_i^{pl}$, along with all of the limits above. This leaves:

$$\lim_{\gamma\to^{+}0} \frac{\partial\left(e_i^{pl} - e_i^*\right)}{\partial\left(\gamma^2\right)} = \frac{\frac{\partial \xi_i^{ex}\left(\mathbf{e}^{pb},\mathbf{X}^*\right)}{\partial(\gamma^2)}}{\kappa''\left(e^{\dagger}\right) - 1} = \frac{e^{\dagger 3}\sum_{k\neq i}\sigma_{ik}}{\kappa''\left(e^{\dagger}\right) - 1} > 0.$$

A similar expression can be derived for j, giving:

$$\lim_{\gamma\to^{+}0} \frac{\partial\left(e_i^{pl} - e_i^*\right)}{\partial\left(\gamma^2\right)} - \lim_{\gamma\to^{+}0} \frac{\partial\left(e_j^{pl} - e_j^*\right)}{\partial\left(\gamma^2\right)} = \frac{e^{\dagger 3}\left(\sum_{k\neq i}\sigma_{ik} - \sum_{k\neq j}\sigma_{jk}\right)}{\kappa''\left(e^{\dagger}\right) - 1} > 0,$$

the final inequality following by assumption: $\sum_{k\neq i}\sigma_{ik} > \sum_{k\neq j}\sigma_{jk}$. By continuity of all functions in γ, this positivity must hold for some neighborhood of $\gamma = 0$.

By Corollary 1 and a similar argument, $e_i^* > e_j^*$ in some neighborhood of $\gamma = 0$. Taking the meet of these two neighborhoods, as well as for each pair i, j with $\sum_{k\neq i}\sigma_{ik} > \sum_{k\neq j}\sigma_{jk}$, gives the result.

The proof of Proposition 3B is analogous to the above.

□

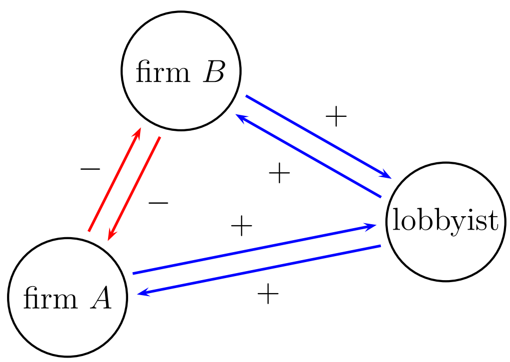

Figure 1: An oil industry and political lobbyist network

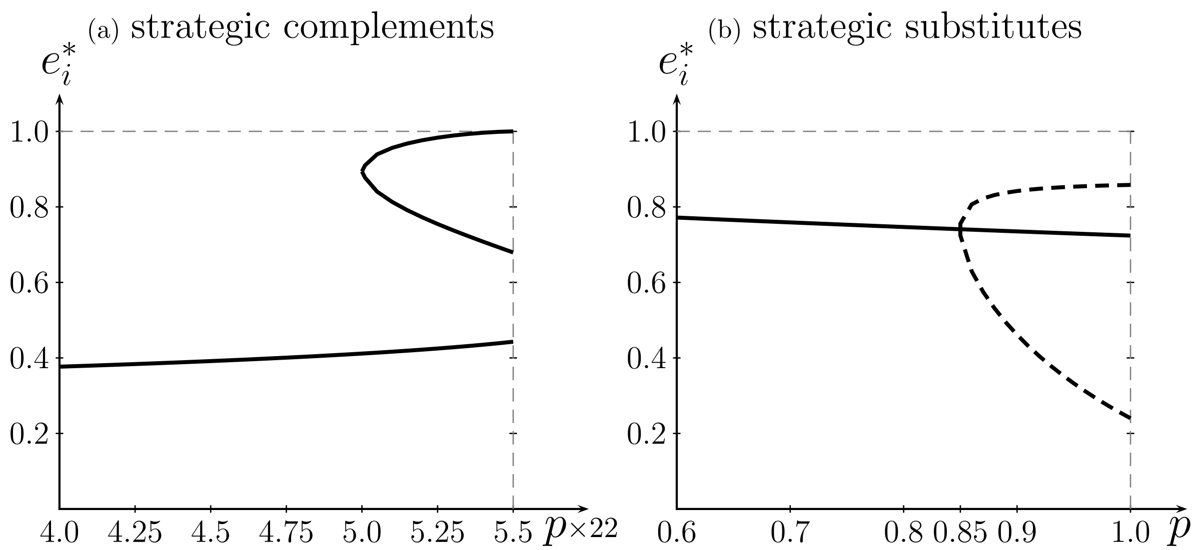

Figure 2: [Example 1] equilibrium multiplicity

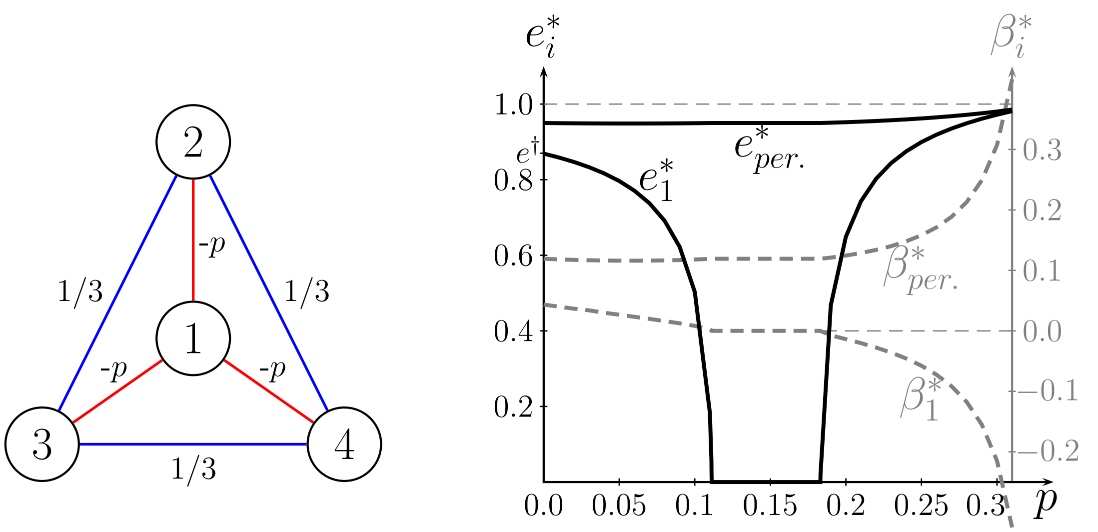

Figure 3: [Example 2] unique equilibrium with negative signal response

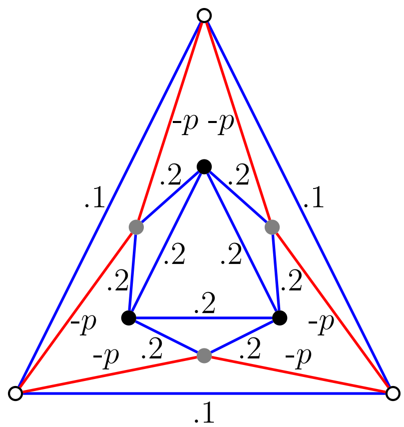

Figure 4: [Example 3] a network with three classes of players. Solid nodes give class x, gray nodes give class y, white nodes give class z.

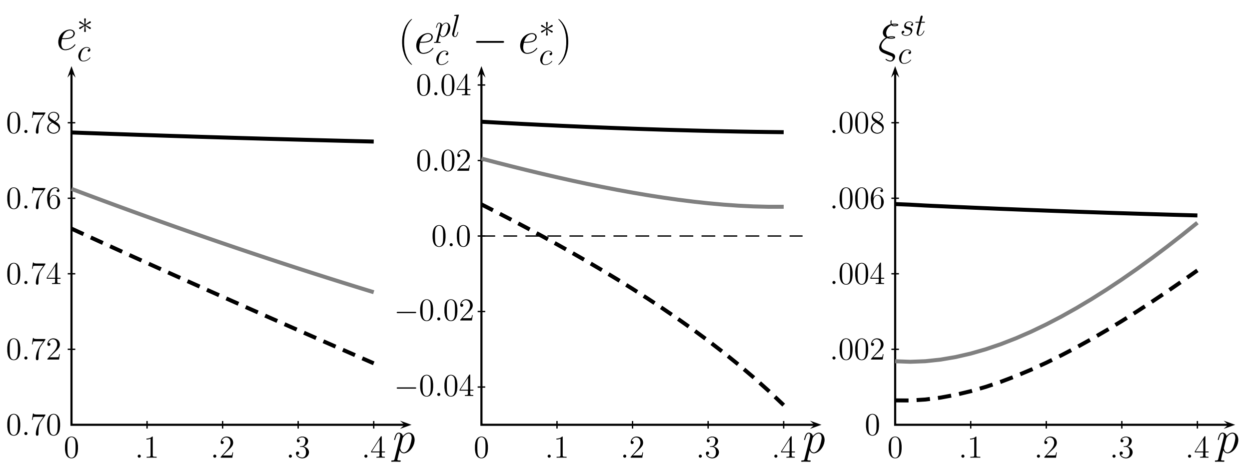

Figure 5: [Example 3] **Left**: equilibrium qualities. **Middle**: absolute welfare difference. **Right**: marginal strategic value. **All**: solid lines give class x, gray lines give class y, dashed lines give class z.

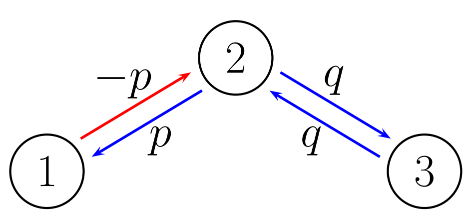

Figure 6: [Example 4] an asymmetric network

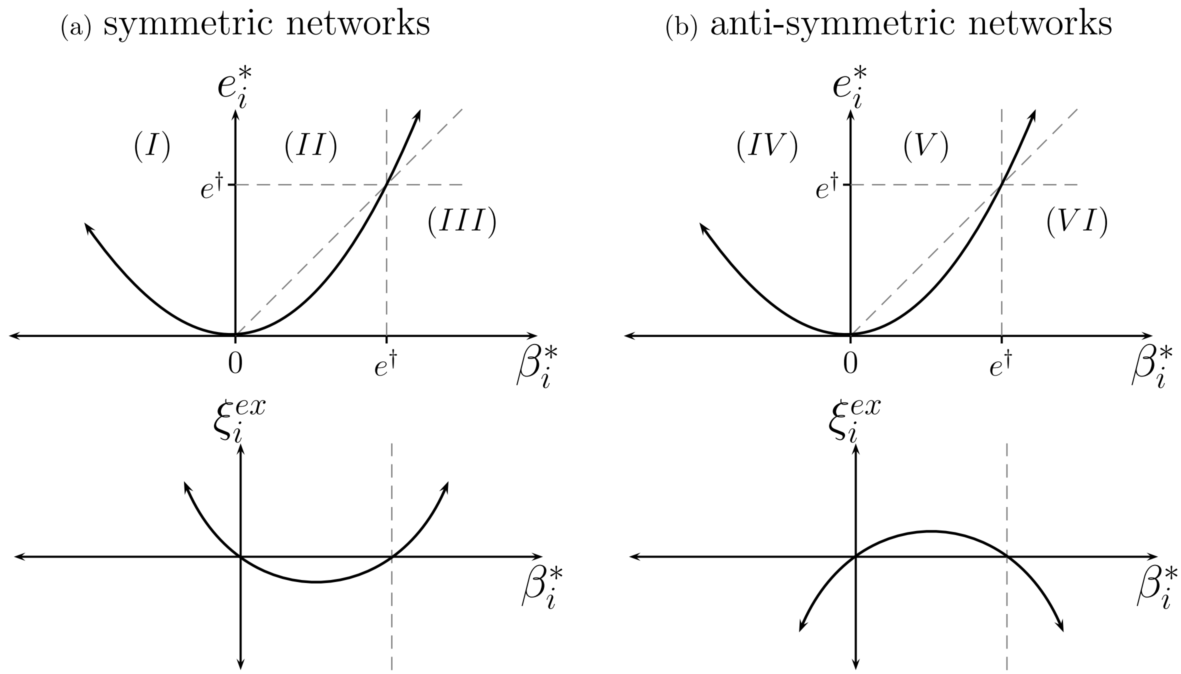

Figure 7: [Directional inefficiencies] leading terms of marginal externalities.

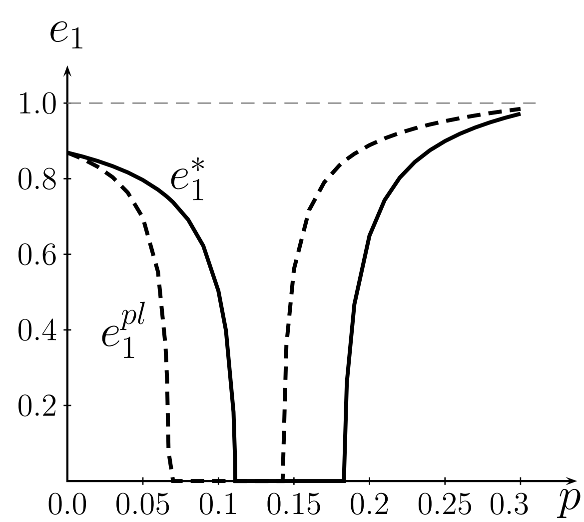

Figure 8: [Example 2] welfare inefficiency of player 1 equilibrium signal quality

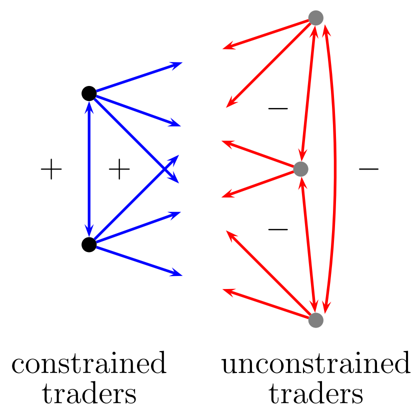

Figure 9: Market with liquidity-constrained traders

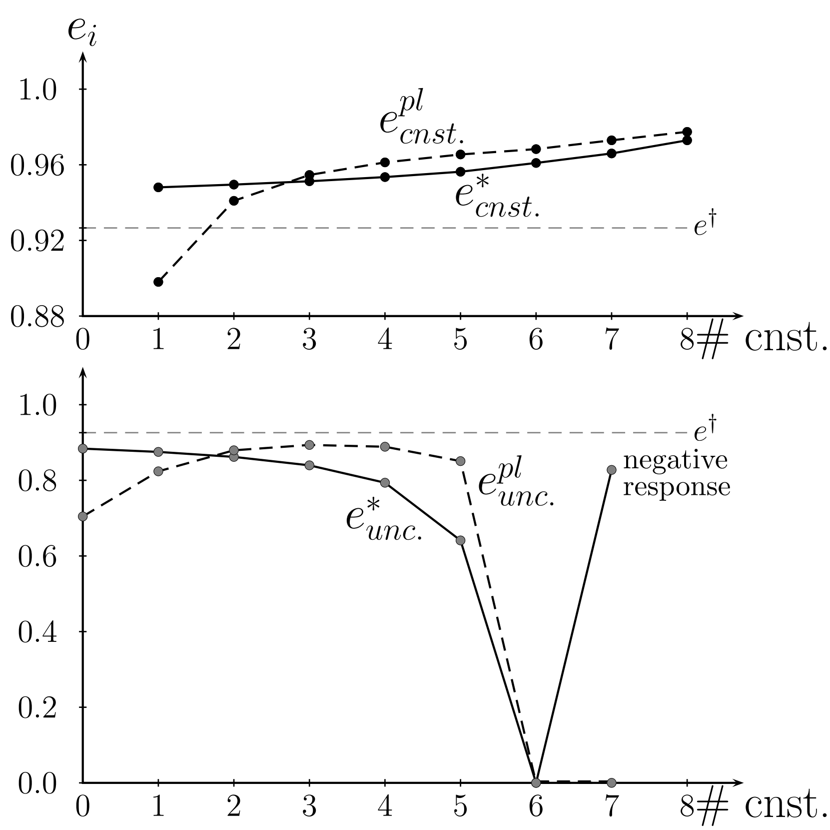

Figure 10: [Efficiency and liquidity crises] Unique equilibrium information qualities versus number of constrained traders (# cnst.) out of eight traders. All links are of size .1, $\eta_1 = 2$, $\eta_2 = 1$, and $K = .01$ for the cost function in Example 1, giving $e^\dagger = 0.927$.

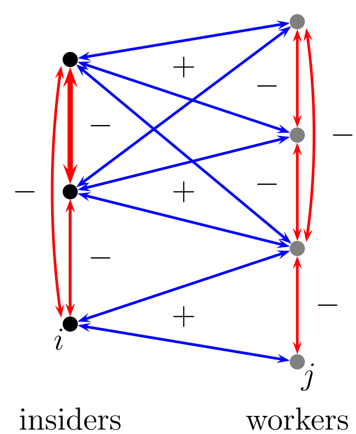

Figure 11: A job-search network

19 Supplemental Section: Who is more Informed?

Theorem 2 offers an important step toward describing information acquisition under general peer effects. However, the fact that $\boldsymbol{\beta}^*$ is endogenously determined as a function of $\mathbf{e}^*$ limits this result from providing a full description of the incentives to acquire information as a function of player-position in the network. Here we reveal a basic challenge in the task of characterizing exactly who acquires more information than others. In light of this fact, we then develop a class of network structures that robustly order the relative extent of information acquisition across players, for all $\gamma > 0$ and over the set of convex κ.[1] All of the results of this Section will refrain from assumptions on the extent (or lack of) symmetry in pairwise peer-effects. Further, we can modify Assumption 1 requiring only the conditions $\kappa'(0) = 0$ and $\kappa''' \geq 0$. As shown in the proof of Proposition S.1, these will suffice for IAE existence for all $\gamma \in [0, 1]$.

Toward better understanding the players' underlining incentives to acquire information, a useful thought experiment is to walk through the best-response dynamic of the period $t = 1$ game. We allow players to simultaneously choose their preferred e_i taking as given their current sequentially rational belief $\mu_{ij}^*(e_j)$ for each $j \neq i$. Start from the profile $\mathbf{e}^{(1)} := (0, \ldots, 0)$, and for this discussion assume Assumption 1 to hold. Here, signals are neither informative of the state nor informative of the actions of neighboring players. However, each player –mindful of the positive direct effect that the state has on their marginal gains to period $t = 2$ action– will prefer to invest in (unique) quality $e^\dagger$ that solves $e^\dagger = \kappa'\left(e^\dagger\right)$ (see Example 15). Then, given positive quality profile $\mathbf{e}^{(2)} = (e^\dagger, \ldots, e^\dagger)$ and updated beliefs $\mu_{ij}^*\left(e^\dagger\right)$, correlation between players' signals is introduced. That is, players' signals now inform them of what others will see and do. Players with high degree will realize an extra *kick* to their marginal benefit to information in the first stage, as additional quality further informs them of their neighbors' $t = 2$ actions. Players with particularly low degree will also obtain information regarding what their neighbor's will see and do. However, the optimal response to "learning neighbors will likely choose high actions" moves against their private response to learning that their marginal gains to action are likely high. Thus, the net responsiveness of these players' strategies to their signals decrease. By Theorem 2, this in turn decays the incentives to acquire precise signals in the first stage.

The direction of the best-response dynamic $\left\{\mathbf{e}^{(n)}\right\}_{n=1}^{\infty}$ from $n = 3$ and on will depend on the structure of the network. Whether or not high degree players will continue to invest more in information than low degree players depends on the relative informativeness of neighbors. Thus, though information acquisition can be ordered with respect to informational centrality $\mathbf{b}(\Sigma^c, \mathbf{e})$, whether the ordering in this measure ultimately aligns with players' degrees in equilibrium depends on both (i) more delicate properties of the network Σ, (ii) the shape of κ, and (iii) the size of γ. The potential for such sensitivity in $\mathbf{e}^*$ is illustrated with the following example.

Example S.1. *Take the six-player star network with center player* 1 *and periphery players* $i \in \{2, \ldots, 6\}$. *We assume center-periphery peer-effects to be undirected:* $\sigma_{i1} = \sigma_{1i} = p > 0$ *(while* $\sigma_{ij} = 0$ *for each pair* $i, j \in \{1, \ldots, 6\}$*). Here, the center player acquires the most information in a unique equilibrium (see Proposition S.2 below).*

[1] For the former, this is provided the second-stage system yields a finite solution.

Now, as depicted in Figure 1, consider adding two more players (7 and 8) with links to the center that are weaker than those of the original periphery players: links of size cp with $c \in [0, 1)$. *However, these additional players enjoy an added positive link of size* $q > 0$ *between each other, reinforcing their behavior. Now, players 2 through 6 are highly influenced by the most central player (player 1), while players 7 and 8 place less weight on the center but together reinforce each other's actions.*

[Figure 1]

Taking $p = 1/5$ *and* $c = 0$*, for example. For any* $q > 1/5$ *players* $\{7, 8\}$ *have greater degree than players* $\{2, \ldots, 6\}$. *As such, players* $\{7, 8\}$ *acquire more information when* γ *is sufficiently small, by Corollary 1. The ordering in a unique* $\mathbf{e}^*$ *when* γ *is large, however, will also depend on the curvature of the cost function* κ. *Thus, take* $\gamma = 1$ *and* $q = 6.9$ *for example, we borrow again the cost function from Example 2 setting* $\zeta = 2$ *and range* η *from* $.5$ *to* 2*, yielding the black and gray cost functions depicted in Figure 2(left).*

When $\eta = .5$ *(low convexity) the marginal cost of information varies mildly over a wide range of small* e_i *values. This results in high dispersion across equilibrium qualities. In this scenario, having access (high influence) to the center player bears heavily on the incentives to acquire a precise signal. As seen in Figure 2(right), players* $\{2, \ldots, 6\}$ *acquire more information than* $\{7, 8\}$. *If instead* $\eta = 2$ *(high convexity) and the marginal cost of information varies quickly over a narrow range of small* e_i*, equilibrium dispersion is more slight:* e_1^* *lies only slightly above the equilibrium qualities of the other players. In this scenario, degree centrality again most encourages information acquisition. As under small gamma,* $\{7, 8\}$ *acquires more information than* $\{2, \ldots, 6\}$.

[Figure 2]

Example S.1 illustrates the tautology that the curvature of information costs and the details of the network structure work in tandem to determine the relative extent of acquired qualities across players. This makes the goal of robustly ordering $\mathbf{e}^*$ over players using some fully portable centrality measure, defined solely over the network structure Σ, unreachable. With intercentrality (Ballester et al. (2006) [5]) and Bonacich centrality measures defined solely on Σ, a one-to-one representation of equilibrium information acquisition and the network structure can not exist. This is true even when the network is undirected and non-negative, as Example S.1 shows.

The following begins to constrain the problem of describing information acquisition in our general network setting. We establish network properties that suffice to order equilibrium qualities. This ordering will be independent over γ and hold over the set of convex κ, for at least one IAE. The properties derived will exclude examples such as the star-with-clique above, and align the essential network properties discussed in Example S.1: degree centrality and neighbors' informational centralities.

First, the following definition and equilibrium notion will help to simplify the task of describing the role of network architecture.

Definition S.1. *For given network* Σ *consider a partition* $\mathcal{P} = \{P_c\}_{c=1}^{C}$ *of* $\{1, \ldots, N\}$*, with subsets ("*classes*") indexed by* $c = 1, \ldots, C \leq N$*, where* $C := |\mathcal{P}|$.[2] *Then, player* i*'s* weighting function $w_i : \{1, \ldots, C\} \to \mathbb{R}$ *with respect to* $\mathcal{P}$ *is defined by:*

$$w_i\,(c) = \sum\nolimits_{k \neq i : k \in P_c} \sigma_{ik}.$$

$\mathcal{P}$ *gives an* equivalence relation *if* $w_i\,(\cdot) = w_j\,(\cdot)$ *for each* $i, j \in P_t$ *and for every* c.

Weighting functions aggregate the weights that a given player places on the individual members of each class. We will use $w_c\,(\cdot)$ to denote the common weighting function of players in equivalence class P_c. Note that an equivalence relation always exists for any network: namely, the *discrete* partition of individual players. One can also find a suitably *coarse* relation that groups all players of equivalent objectives.[3] The goal of partitioning the players in this manner is to discard details of Σ less essential to the problem of information acquisition, while preserving the more germane network properties that drive equilibrium dispersion in $\mathbf{e}^*$. Conducive to this goal, for any equivalence relation $\mathcal{P}$ an equilibrium that is symmetric within classes will always exist.

Proposition S.1. [**class-symmetric IAE**] *For equivalence relation* $\mathcal{P}$ *and any* $\kappa \in \mathcal{C}^3$ *with* $\kappa'(0) = 0$ *and* $\kappa'', \kappa''' \geq 0$*, there exists a class-symmetric equilibrium in which* $\beta_i^* = \beta_j^*$ *and* $e_i^* = e_j^*$ *if* $i, j \in \mathcal{P}_c$ *for* $c \in \{1, \ldots, C\}$.

The second half of Example 1 provides an IAE that violates class symmetry. Precisely, the asymmetric equilibrium violates class symmetry when both players are included within the same class.

Reflecting again on Example S.1, we see that three classes are used to induce sensitivity in the ordering of $\mathbf{e}^*$ to the shape of κ. When players place non-negative aggregate weight on those within their class, this extent of network irregularity (i.e. three classes) is necessary to establish such sensitivity.

Proposition S.2. [**two-class networks**] *For equivalence relation* $\mathcal{P} = \{r, s\}$ *with* $w_r(r), w_s(s) \geq 0$ *and* $w_r\,(r) + w_r\,(s) > w_s\,(r) + w_s\,(s)$*, and any* $\kappa \in \mathcal{C}^3$ *with* $\kappa'(0) = 0$ *and* $\kappa'', \kappa''' \geq 0$*, there exists a class-symmetric equilibrium such that* $e_r^* \geq e_s^*$*, and where if* $e_r^*, e_s^* \in (0, 1)$ *then* $e_r^* > e_s^*$ *with* $\beta_r^* > \beta_s^*$.

Note that given $e_r^* > e_s^*$, $\beta_r^* > \beta_s^*$ in the last statement of the theorem is equivalent to $\beta_r^* > 0$ by Theorem 2. Thus, signal responses are ordered with the highest degree class moving positively with their signal. Allowing for $\beta_s^* < 0$, Proposition S.2 captures a striking equilibrium property. For a class s moving against their information, anticipating the actions of players in P_r, each $j \in P_s$ chooses a quality that is bounded above by e_r^*. With each j's signal used merely to infer the the actions of those in P_r, and with e_r^* intrinsically

[2] That is, $\bigcup \mathcal{P} = N$ with $P_s \cap P_{s'} = \varnothing$ for distinct $s, s' \in \{1, \ldots, C\}$.

[3] Equivalent in the sense that players within a class set place equivalent weights on other classes. Note that given partitions $\mathcal{P}^1$ and $\mathcal{P}^2$ one can construct coarser partition $\mathcal{P}$ by joining elements $P^1 \in \mathcal{P}^1$ and $P^2 \in \mathcal{P}^2$ to give $P^1 \cup P^2 = P \in \mathcal{P}$ when $P^1 \cap P^2 \neq \emptyset$. That is, a coarsesed set of partitions can be obtained, most often being s single coarsesed set pooling interchangeable players.

bounding the extent of this inference, e_r^* provides a natural bound on j's incentives to acquire information.[4] This natural bound can be clearly observed above in Figure 3(right).

Next, the following notions allow for any arbitrary number of classes, and establish alternative conditions on the network structure that suffice for an ordering in $\mathbf{e}^*$, again robust to the relative convexity in κ. This family of *class-ordered* networks will offer a generalization of core-periphery-like structures, incorporating signed, weighted, and directed links. Note that the following ordering in $\mathcal{P}$ is defined solely using properties of the network Σ.

Definition S.2. *We say that class r <u>dominates</u> class s (denoted $r \succsim s$) if the following two conditions hold:*

1. *w_r crosses w_s at most once from below: $w_r(c) \geq w_s(c)$ if $c \geq x$ and $w_r(c) \leq w_s(c)$ if $c \leq x$ for some $x \in \{1, \ldots, C\}$, and*

2. *players in P_r have degree no smaller than players in P_s:*

$$\sum\nolimits_{c=1}^{C} w_r(c) \geq \sum\nolimits_{c=1}^{C} w_s(c). \tag{1}$$

r <u>strictly dominates</u> s (denoted $r \succ s$) if the inequality in (1) is strict.

The cumulative ordering (1) with single crossing in condition 1 imply that more central classes are more influenced by others (have higher degree), and that these classes tend to place relatively more weight on the most central players.

[Figure 3]

From a technical vantage point, dominance gives an appealingly weak condition that suffices for the relative weighting functions to aggregate any non-negative, non-decreasing function f in similar order. That is, and as illustrated in Figure 3 (left), $r \succsim s$ implies that $w_r(c')$ must lie weakly above $w_s(c')$ for the highest classes c' which give the greatest values $f(c')$. Formally, this gives the following lemma.

Lemma S.1. *If $r \succsim s$ then:*

$$\sum\nolimits_{c=1}^{C} f(c)\, w_r(c) \geq \sum\nolimits_{c=1}^{C} f(c)\, w_s(c), \tag{2}$$

for any non-decreasing function f on $\{1, \ldots, C\}$: $f(c') \geq f(c) \geq 0$ for $c' \geq c$. If $r \succ s$ then the inequality in (2) is strict.

The proof of this is simple to obtain and is provided in the appendix. The following class of network structures can now be defined. Note that the ordering in index $\{1, \ldots, C\}$ has thus far been immaterial. Here, however, the ordering in $\mathcal{P}$ plays a more central role.

[4] The qualification $e_r^* < 1$ is needed to exclude equilibria in which the classes coordinate on simultaneously acquiring perfectly precise signals in order to move against them.

Definition S.3. *The network* Σ *is* class ordered *if there is an equivalence relation* $\mathcal{P}$ *such that for each* $r \in \{2, \ldots, C\}$ *we have* $r \succsim r-1$. *The network* Σ *is* strictly class ordered *if* $r \succ r-1$ *for each* $r > 1$.

The class orderedness of a network establish a definitive ordering amongst its classes. The most connected nodes will place proportionally more of their weight on precisely those classes that are most connected in the network. Above in Example 3, Definition S.3 is satisfied under class ordering $x \succsim y \succsim z$ (see Figure 4, above). Each class's weighting function is plotted in Figure 3 (right). w_c exhibits dominance between adjacent classes: w_y single crossing w_z from below for all $p \geq 0$.[5]

Though examples of networks of two classes may come readily (e.g. star, circle-spoke), the range of class-ordered networks may be less obvious to the reader. The following example lends to the scope of class-ordered structures.

Example S.2. *The binary networks given in Figure 4 where each link designates positive peer effect* $\sigma_{ij} = p > 0$ *are class ordered. The most central class (i.e. the "core") are given with solid nodes, with the subsequent ordering over classes designated for representative members. Alternatively, all of these examples are also class ordered for* $p < 0$ *with the ordering over classes reversed.*

[Figure 4]

We see that class-ordered networks encompass a wide range of structures exhibiting a natural ordering over its players. These networks can be viewed as a generalized family of core-periphery like structures, allowing for weighted links that may be positive, negative, or directed. Many hierarchical[6] social settings will embody these properties. And in network formation environments, many related models of investment with endogenous link formation –both under strategic substitutes (Bala and Goyal (2000) [4]) and complements (Hiller (2013) [33])– have been shown to yield core-periphery structures.[7]

We come to the main result of the section. When the network of peer effects takes on the above ordering, the following *class-ordered* equilibria always exist.

Proposition S.3. **[class-ordered equilibria]** *If* Σ *is class ordered, taking* $r, s \in \{1, \ldots, C\}$ *with* $r \succsim s$ *and constrain* $\gamma \in [0, \gamma^s)$. *Then, for any* $\kappa \in \mathcal{C}^3$ *with* $\kappa'(0) = 0$, *and* $\kappa'', \kappa''' \geq 0$, *there exists a class-symmetric equilibrium such that* $e_r^* \geq e_s^*$.

Thus, provided players always move in the direction of their signals, player degree robustly orders signal responsiveness independent of the convexity of κ.

The appealing property of class-ordered networks is that highly central players (here, players with the highest degree) proportionally place more of their weight on players that are

[5] Note that here, $x \succsim z$. Such transitivity need not hold for the network to be class ordered.

[6] I thank Anja Prummer for suggesting the natural application to social hierarchies.

[7] Refer to Calvó-Armengol et al. (2011) [13] Section 5.2 for class of "hierarchical" structures that yield properties similar to class orderedness.

also of high centrality. Definition S.3 provides an ordering underlining such nested weighting. In class-ordered networks, this ordering captures both value to having high degree with the value to being connected to the most informed players. In a class-ordered equilibria, it is precisely the neighbors with greatest degree who are most informed.

Returning the two-sided market application, if highly connected insiders are also those that enjoy exclusivity in their clientele, informational centrality will likely be ordered according to degree, with the network adopting a class-ordered structure. If instead the more connected insiders tend to compete with each other for workers, as in the case of Figure 11, the ultimate informational centralities realized by each insider will more intimately depend on the shape of κ. Akin to Example 3, when κ displays significant elasticity yielding moderate dispersion in $\mathbf{e}^*$, degree centrality will dominate. If instead κ displays moderate elasticity yielding significant dispersion in $\mathbf{e}^*$, exclusivity will drive information centrality. While all insiders on the sufficiently short side of the market under acquire information relative to the utilitarian benchmark, exactly who most acquires and simultaneously most *under* acquires information will depend on the precise properties of Σ and κ.

G Supplemental Appendix

G.1 Section 19 proofs: Class-ordered networks

Proof of Proposition S.1. Take the compact subspace of $[0,1]^N$ comprising all class-symmetric vectors $\mathbf{e}$:

$$\mathcal{E}^s := \left\{\mathbf{e} \in [0,1]^N : e_i = e_j \text{ if } i,j \in P \in \mathcal{P}\right\}.$$

Note that $\mathcal{E}^s$ is a closed subset of a compact space, and is thus compact. Now, take the incentive-compatible first-stage best response correspondence for player i:

$$\begin{aligned} BR_i\left(e_i|\mu_i^*,\boldsymbol{\beta}^*\right) &= \underset{e_i\in[0,1]}{\arg\max}\, \mathbb{E}_i\left[u_i\left(\mathbf{X}^*|\omega,\omega_i\right)|e_i,\mu_i^*\right] - \kappa\left(e_i\right), \\ &= \underset{e_i\in[0,1]}{\arg\max}\, \frac{1}{2}\beta_i^{*2} - \kappa\left(e_i\right), \end{aligned} \tag{3}$$

which holds $\boldsymbol{\beta}_{-i}^*$ and μ_i^* fixed but allows β_i^* to optimally adjust to e_i. The second equality uses expression (8) derived in Section F.3. First, by the compactness of $[0,1]$ and continuity of β_i^* and of $\kappa\left(e_i\right)$ in e_i,[8] $BR_i\left(e_i|\mu_i^*,\boldsymbol{\beta}^*\right)$ is non-empty by the Weierstrass extreme-value theorem.

By construction, the set:

$$\left[BR_i\left(e_i|\mu_i^*,\boldsymbol{\beta}^*\right)\right] \cap \mathcal{E}^s$$

is non-empty, and thus the restriction:

$$\overline{BR}\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right) := \left[BR_i\left(e_i|\mu_i^*,\boldsymbol{\beta}^*\right)\right] \cap \mathcal{E}^s, \tag{4}$$

is a well defined vector-valued mapping from $\mathcal{E}^s \to \mathcal{E}^s$. By continuity of $\boldsymbol{\beta}^*$ and κ in $\mathbf{e} \in [0,1]$ a compact set, and applying the Maximum theorem, $\overline{BR}\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right)$ is upper hemicontinuous. Marginal gains to information are given by:

$$\beta_i^{*2}/e_i = e_i\left(1+\sum_{k\neq 1}\sigma_{ik}e_k\beta_k^*\right)^2$$

by Theorem 1, which is linear in e_i by incentive compatability (μ_k^* and β_k^* are held fixed) and obtains $\beta_i^{*2}/e_i = 0$ at $e_i = 0$. When $\kappa'(0) = 0$ and $\kappa''' \geq 0$, each $\overline{BR}_i\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right)$ is convex valued: if $\beta_i^{*2}/e_i > \kappa'(e_i)$ for some e_i then $\beta_i^{*2}/e_i' > \kappa'(e_i')$ for each $e_i' > e_i$, and if $\beta_i^{*2}/e_i < \kappa'(e_i)$ then $\beta_i^{*2}/e_i' < \kappa'(e_i')$ for each $0 < e_i' < e_i$ (excluding $e_i = 0$ which gives a minimum).[9] $\overline{BR}\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right)$ then gives a convex polyhedron in $[0,1]^N$. Then, by Kakutani's fixed point theorem, $\overline{BR}\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right)$ yields a fixed point in $\mathcal{E}^s$. By construction of $\mathcal{E}^s$, the properties of the fixed point satisfy those of the theorem.

□

[8] Continuity follows from Assumptions F1 and $\kappa \in \mathcal{C}$.

[9] With $\kappa'(0) = 0$ and $\kappa''' \geq 0$, each $\overline{BR}_i\left(\mathbf{e},\boldsymbol{\beta}^*,\boldsymbol{\mu}^*\right)$ will either (i) give a unique value in $[0,1]$ if $\kappa''' > 0$ or (i) give a corner or the entire unit interval if $\kappa''' = 0$ (quadratic κ).

Proof of Proposition S.2. Assuming quality profile $e_r \geq e_r$ we show that there exists a first-stage best response for class r weakly above every best response for class s. Write the system giving the IRE as a function of (e_r, e_s):

$$\begin{array}{l} [1]\ \beta_r^* - \left(e_r + \gamma^2 e_r \left(w_r\left(r\right) e_r \beta_r^* + w_r\left(s\right) e_s \beta_s^*\right)\right) = 0 \\ [2]\ \beta_s^* - \left(e_s + \gamma^2 e_s \left(w_s\left(s\right) e_s \beta_s^* + w_s\left(r\right) e_r \beta_r^*\right)\right) = 0 \end{array}.$$

Together these imply:

$$\beta_r^* = \frac{e_r\left(1 + \gamma^2 e_s^2\left(w_r\left(s\right) - w_s\left(s\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r\right)\left(1 - \gamma^2 w_s\left(s\right) e_s\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^2 e_s^2}, \tag{5}$$

$$\beta_s^* = \frac{e_s\left(1 + \gamma^2 e_r^2\left(w_s\left(r\right) - w_r\left(r\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r\right)\left(1 - \gamma^2 w_s\left(s\right) e_s\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^2 e_s^2}. \tag{6}$$

Multiplying by e_r and e_s, respectively:

$$\begin{aligned} e_r\beta_r^* &= \frac{e_r^2\left(1 + \gamma^2 e_s^2\left(w_r\left(s\right) - w_s\left(s\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r\right)\left(1 - \gamma^2 w_s\left(s\right) e_s\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^2 e_s^2}, \\ e_s\beta_s^* &= \frac{e_s^2\left(1 + \gamma^2 e_r^2\left(w_s\left(r\right) - w_r\left(r\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r\right)\left(1 - \gamma^2 w_s\left(s\right) e_s\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^2 e_s^2}. \end{aligned}$$

With $e_r \geq e_s$, then $e_r\beta_r^* \geq e_s\beta_s^*$ is implied by:

$$\begin{aligned} w_r\left(s\right) - w_s\left(s\right) &> w_s\left(r\right) - w_r\left(r\right) \\ \Leftrightarrow w_r\left(s\right) + w_r\left(r\right) &> w_s\left(r\right) + w_s\left(s\right), \end{aligned}$$

which is assumed.

Now, rewriting the system as:

$$\begin{array}{l} [1]\ \beta_r^{*2} = e_r\beta_r^*\left(1 + \gamma^2\left(w_r\left(r\right) e_r\beta_r^* + w_r\left(s\right) e_s\beta_s^*\right)\right) \\ [2]\ \beta_s^{*2} = e_s\beta_s^*\left(1 + \gamma^2\left(w_s\left(s\right) e_s\beta_s^* + w_s\left(r\right) e_r\beta_r^*\right)\right) \end{array}.$$

$$
\begin{aligned}
\beta_r^{*2} - \ \beta_s^{*2} &= e_r\beta_r^* - e_s\beta_s^* + \gamma^2 \left(\begin{array}{c} e_r\beta_r^* \left(w_r\left(r\right) e_r\beta_r^* + w_r\left(s\right) e_s\beta_s^*\right) \\ -e_s\beta_s^* \left(w_s\left(s\right) e_s\beta_s^* + w_s\left(r\right) e_r\beta_r^*\right) \end{array} \right) \\
&= e_r\beta_r^* - e_s\beta_s^* + \gamma^2 e_s\beta_s^* e_r\beta_r^* \left(\begin{array}{c} \left(w_r\left(r\right) \frac{e_r\beta_r^*}{e_s\beta_s^*} + w_r\left(s\right)\right) \\ -\left(w_s\left(s\right) \frac{e_s\beta_s^*}{e_r\beta_r^*} + w_s\left(r\right)\right) \end{array} \right) \\
&= e_r\beta_r^* - e_s\beta_s^* + \gamma^2 e_s\beta_s^* e_r\beta_r^* \left(\begin{array}{c} w_r\left(r\right) \left(\frac{e_r\beta_r^*}{e_s\beta_s^*} - 1\right) - w_s\left(s\right) \left(\frac{e_s\beta_s^*}{e_r\beta_r^*} - 1\right) \\ +\left(w_r\left(r\right) + w_r\left(s\right)\right) - \left(w_s\left(s\right) + w_s\left(r\right)\right) \end{array} \right) \\
&= e_r\beta_r^* - e_s\beta_s^* + \gamma^2 e_s\beta_s^* e_r\beta_r^* \left(\begin{array}{c} \left(e_r\beta_r^* - e_s\beta_s^*\right) \left(\frac{w_r(r)}{e_s\beta_s^*} + \frac{w_s(s)}{e_r\beta_r^*}\right) \\ +\left(w_r\left(r\right) + w_r\left(s\right)\right) - \left(w_s\left(s\right) + w_s\left(r\right)\right) \end{array} \right) \\
&= \left(\begin{array}{c} \left(e_r\beta_r^* - e_s\beta_s^*\right) \left(1 + \gamma^2 \left(w_r\left(r\right) e_r\beta_r^* + w_s\left(s\right) e_s\beta_s^*\right)\right) \\ +\gamma^2 e_s\beta_s^* e_r\beta_r^* \left(w_r\left(r\right) + w_r\left(s\right)\right) - \left(w_s\left(s\right) + w_s\left(r\right)\right) \end{array} \right).
\end{aligned}
$$

If $w_r(r)$ and $w_s(s)$ are positive, then $\beta_r^{*2} - \ \beta_s^{*2} \geq 0$ is implied by $(w_r(r) + w_r(s)) - (w_s(s) + w_s(r)) > 0$ (with strict inequality when $e_s, e_r > 0$), which are all assumed. Take any class-symmetric $\mathbf{e}$ and $\boldsymbol{\beta}$ that satisfy $e_r \geq e_s$. Again denote $BR(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*) := [BR_i(e_i|\mu_i^*, \boldsymbol{\beta}^*)]$ from the proof of Proposition S.1. For each $j \in P_s$ and any $e_j \in BR_j\left(e_j|\mu_j^*, \boldsymbol{\beta}^*\right)$, by Theorem 2 we must have either $e_j = 1$ with $e_j\kappa'(e_j) < \beta_j^{*2}$ or $e_j < 1$ with $e_j\kappa'(e_j) = \beta_j^{*2}$ for e_j to be a best response. Thus, in either case by $\beta_r^{*2} \geq \beta_s^{*2}$, the marginal gain $\beta_r^{*2}/e_j \geq \beta_s^{*2}/e_j$ when e_i is set to e_j, implying that any $i \in P_r$ would have a profitable deviation up away from e_j. This then implies existence of some $e_i \in BR_i(\mathbf{e}|\mu_i^*, \boldsymbol{\beta}^*) \geq e_j$.

Now take the compact subspace of $[0,1]^2$ that includes all weakly increasing class- symmetric vectors $\mathbf{e}$: $\mathcal{E}^+ := \{\mathbf{e} \in [0,1]^2 : e_i \geq e_j,\ i \in P_r, j \in P_s\}$. Note that $\mathcal{E}^+$ is a closed subset of a compact space, and is thus compact. By the above, $BR(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*) \cap \mathcal{E}^s \cap \mathcal{E}^+$ is non-empty, and thus the restriction:

$$
\overline{\overline{BR}}(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*) := BR(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*) \cap \mathcal{E}^s \cap \mathcal{E}^+,
$$

where $\mathcal{E}^s$ is given by (G.1), is a well defined mapping from $\mathcal{E}^s \cap \mathcal{E}^+ \to \mathcal{E}^s \cap \mathcal{E}^+$. By continuity of $\boldsymbol{\beta}^*$ and κ in $\mathbf{e} \in [0,1]$ a compact set, and applying the Maximum theorem, $\overline{\overline{BR}}(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*)$ is upper-hemicontinuous. $\kappa'(0) = 0$ and $\kappa''' \geq 0$ again suffice for $\overline{BR}(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*)$ to be convex valued (see proof of Theorem S.1). By Kakutani's fixed point theorem, $\overline{\overline{BR}}(\mathbf{e}|\boldsymbol{\mu}^*, \boldsymbol{\beta}^*)$ yields a fixed point in $\mathcal{E}^s \cap \mathcal{E}^+$.

Finally, we show that $e_r^* > e_s^*$ and $\beta_r^* > 0$ when $e_r^*, e_s^* \in (0,1)$. Rewriting (5) and (6) evaluated at IAE with $e_r^* \geq e_s^*$:

$$
\begin{aligned}
\frac{\beta_r^*}{e_r^*} &= \frac{\left(1 + \gamma^2 e_s^{*2} \left(w_r\left(s\right) - w_s\left(s\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r^*\right)\left(1 - \gamma^2 w_s\left(s\right) e_s^*\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^{*2} e_s^{*2}}, \\
\frac{\beta_s^*}{e_s^*} &= \frac{\left(1 + \gamma^2 e_r^{*2} \left(w_s\left(r\right) - w_r\left(r\right)\right)\right)}{\left(1 - \gamma^2 w_r\left(r\right) e_r^*\right)\left(1 - \gamma^2 w_s\left(s\right) e_s^*\right) - \gamma^4 w_r\left(s\right) w_s\left(r\right) e_r^{*2} e_s^{*2}}.
\end{aligned}
$$

If $\beta_r^* < 0$, this implies that $\gamma^2 e_s^{*2}(w_r(s) - w_s(s)) < 1$, which implies also that $\gamma^2 e_r^{*2}(w_s(r) - w_r(r)) < 1$ and $\gamma^2 e_r^{*2}(w_s(r) - w_r(r)) < \gamma^2 e_s^{*2}(w_r(s) - w_s(s))$ by $(w_s(r) - w_r(r)) < (w_r(s) - w_s(s))$ and $e_r^{*2} \geq e_s^{*2}$. Thus, $\beta_s^*/e_s^* < \beta_r^*/e_r^*$, implying that

$\beta_s^{2*}/e_s^* > \beta_r^{*2}/e_r^*$, and thus by Theorem 2 that $e_s^* > e_r^*$ as $e_r^* < 1$, yielding a contradiction. Thus $\beta_r^* > 0$ with $\beta_r^{*2} > \beta_s^{*2}$ by the above, implying that $e_r^* > e_s^*$.

□

Proof of Lemma S.1. Let x be defined as in Definition S.2. Rearranging the second part of Definition S.2 gives:

$$\sum\nolimits_{c \geq x} (w_r(c) - w_s(c)) \geq \sum\nolimits_{c<x} (w_s(c) - w_r(c)). \tag{7}$$

Then, rearranging the result:

$$\begin{aligned} \sum_{c=1}^{C} f(c)\, w_r(c) - \sum_{c=1}^{C} f(c)\, w_s(c) &= \begin{pmatrix} \sum_{c\geq x} f(c)(w_r(c) - w_s(c)) \\ -\sum_{c<x} f(c)(w_s(c) - w_r(c)) \end{pmatrix} \\ &\geq \begin{pmatrix} \sum_{c\geq x} f(x)(w_r(c) - w_s(c)) \\ -\sum_{c<x} f(x)(w_s(c) - w_r(c)) \end{pmatrix} \\ &= f(x) \begin{pmatrix} \sum_{c\geq x} (w_r(c) - w_s(c)) \\ -\sum_{c<x} (w_s(c) - w_r(c)) \end{pmatrix} \\ &\geq 0. \end{aligned}$$

The first inequality follows from $f(\cdot)$ non-decreasing, while the second inequality follows from (7) and $f(x) \geq 0$. The final inequality is strict if $f(c) > 0$ for each c and $r \succ s$.

□

Proof of Proposition S.3. We use class indices for all strategies and weighting functions, when convenient. First, we will need the following definitions and Lemma. Take $i \in P_r \in \mathcal{P}$ and $j \in P_s \in \mathcal{P}\backslash\{P_r\}$. Take any class- symmetric $\mathbf{e}$ that satisfies the conditions of the theorem. The set of class-ordered profiles:

$$\mathcal{E}^+ := \left\{\mathbf{e} \in [0,1]^N : \ e_i \geq e_j \text{ or each } i \in P_r, j \in P_s \text{ with } r \geq s\right\}$$

is a closed, compact subset of $\mathbb{R}^{2N}$. By i's first order condition of the IRE:

$$\beta_i^* = \beta_r^* = e_r \left(1 + \gamma^2 \sum_c w_r(c)\, e_c \beta_c\right). \tag{8}$$

Denote as a function of $\mathbf{e}$ and $\boldsymbol{\beta}$:

$$\Lambda_r := \sum_c w_r(c)\, e_c \beta_c.$$

Λ_r captures the size of the aggregate peer effect on i in $\boldsymbol{\beta}^*$. Analogous expressions can be derived for class s. By our choice of $(\mathbf{e}, \boldsymbol{\beta})$ and with $\gamma \in [0, \gamma^s)$, $e_c\beta_c$ is non-negative and non-decreasing across classes. By the class orderedness of Σ and Lemma S.1, the factor $(1 + \gamma^2 \sum_c w_r(c)\, e_c\beta_c)$ must also be increasing across classes, and by our choice of $\mathbf{e}$ the vector of optimal responses $\boldsymbol{\beta}^*$ must also respect the ordering $\beta_r^* \geq \beta_s^*$ if and only if $r \geq s$. We can now establish the following Lemma.

Lemma S.B.1. *If $\gamma \leq \gamma^s$, Σ is class ordered and $\mathbf{e}$ is class ordered (i.e. weakly increasing across classes), then for $i \in P_r$, $j \in P_{r-1}$, and for every $e_j \in BR_j\left(e_j|\mu_j^*, \boldsymbol{\beta}^*\right)$, there exists $e_i \in BR_i\left(e_i|\mu_i^*, \boldsymbol{\beta}^*\right)$ with $e_i \geq e_j$.*

Proof of Lemma S.B.1. Again use $BR_i\left(e_i|\mu_i^*, \boldsymbol{\beta}^*\right)$ to denote i's first-stage incentive- compatible best response. If $\beta_s^* \geq 0$ then $\beta_r^{*2} \geq \beta_s^{*2}$. For any $e_j \in BR_j\left(e_j|\mu_j^*, \boldsymbol{\beta}^*\right)$, by Theorem 2 we must have either $e_j = 1$ with $e_j\kappa'\left(e_j\right) < \beta_j^{*2}$ or $e_j < 1$ with $e_j\kappa'\left(e_j\right) = \beta_j^{*2}$ for e_j to be a best response. Thus, in either case by $\beta_r^{*2} \geq \beta_s^{*2}$, the marginal gain $\beta_r^{*2}/e_j \geq \beta_s^{*2}/e_j$ when e_i is set to e_j, implying that any $i \in P_r$ would have a (weak) profitable deviation up away from e_j. This then implies existence of some $e_i \in BR_i\left(\mathbf{e}|\mu_i^*, \boldsymbol{\beta}^*\right) \geq e_j$.

□

The proof proceeds analogous to that of Proposition S.1. Take $BR_i\left(e_i|\mu_i^*, \boldsymbol{\beta}^*\right)$ the incentive-compatible best-response correspondence for i in her first-stage problem, holding $\boldsymbol{\mu}_{-i}^*$ and $\boldsymbol{\beta}_{-i}^*$ fixed. The set:

$$[BR_i\left(e_i|\mu_i^*, \boldsymbol{\beta}^*\right)] \cap \mathcal{E}^s \cap \mathcal{E}^+$$

is non-empty by construction,[10] and thus the restriction:

$$\overline{\overline{BR}}\left(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*\right) := [BR_i\left(e_i|\mu_i^*, \boldsymbol{\beta}^*\right)] \cap \mathcal{E}^s \cap \mathcal{E}^+$$

is a well defined vector-valued mapping from $\mathcal{E}^s \cap \mathcal{E}^+ \to \mathcal{E}^s \cap \mathcal{E}^+$. By continuity of $\boldsymbol{\beta}^*$ and κ in $\mathbf{e} \in [0,1]$ a compact set, and applying the Maximum theorem, $\overline{\overline{BR}}\left(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*\right)$ is upper-hemicontinuous. $\kappa'(0) = 0$ and $\kappa''' \geq 0$ again suffice for $\overline{BR}\left(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*\right)$ to be convex valued (see proof of Theorem S.1). By Kakutani's fixed point theorem, $\overline{\overline{BR}}\left(\mathbf{e}, \boldsymbol{\beta}^*, \boldsymbol{\mu}^*\right)$ yields a fixed point in $\mathcal{E}^s \cap \mathcal{E}^+$. By construction of $\mathcal{E}^s$ and $\mathcal{E}^+$, the properties of the fixed point satisfy those of the theorem.

□

[10]To include $r > s+1$, Lemma S.B.1 is used here $r-s$ times.

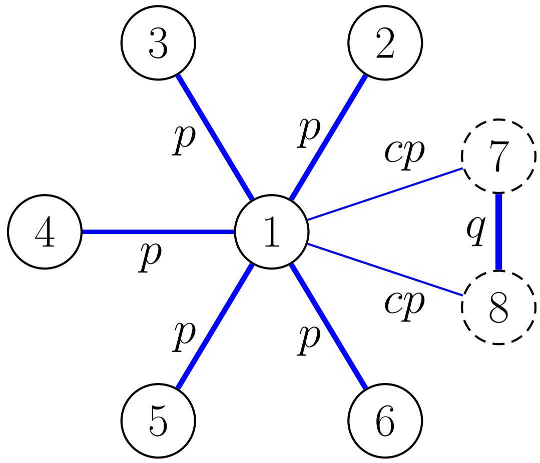

Figure 1: [Example S.1] star with clique

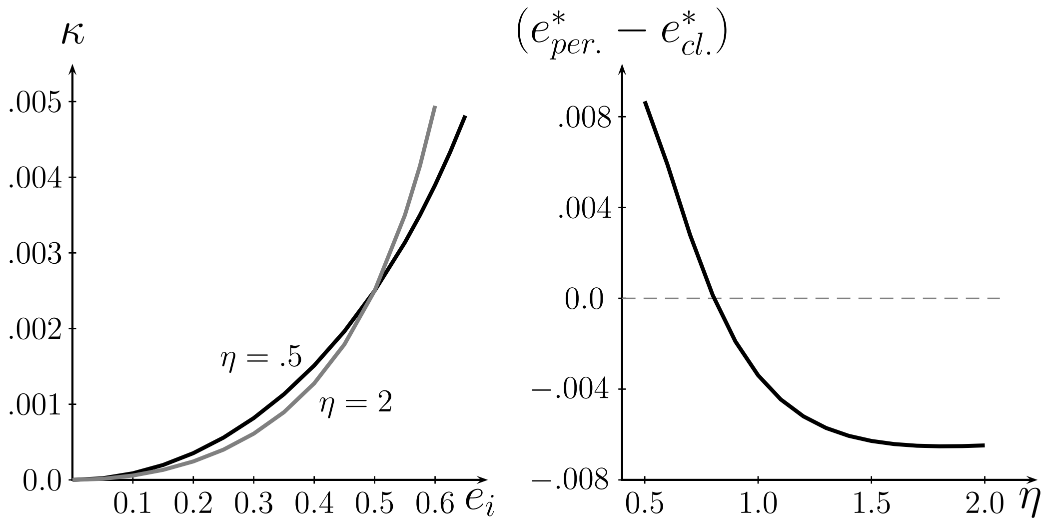

Figure 2: [Example S.1] Sensitivity in $\mathbf{e}^*$ ordering to κ.

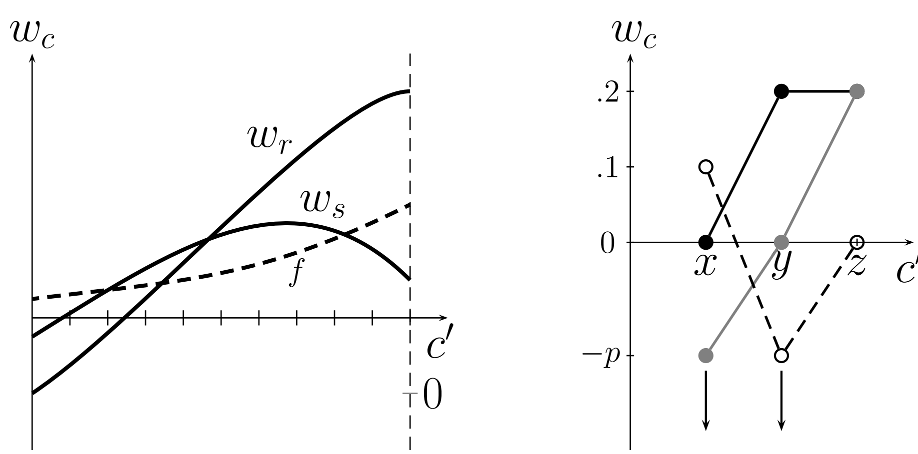

Figure 3: **Left** *Dominance* orders weighting functions to aggregate any non-negative, non-decreasing f in similar order. **Right** The network in Example 3 (Figure 4) is class ordered for all $p > 0$.

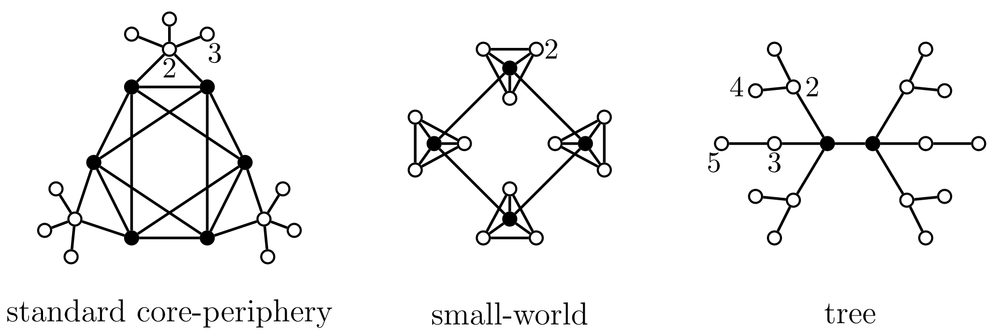

Figure 4: [Example S.2] class-ordered networks