

Bilateral Trading in Networks

DANIELE CONDORELLI

Department of Economics, University of Essex

E-mail: dcond@essex.ac.uk

ANDREA GALEOTTI

Department of Economics, European University Institute and University of Essex

E-mail: andrea.galeotti@eui.eu

LUDOVIC RENO

Department of Economics, University of Essex

E-mail: lrenou@essex.ac.uk

In many markets, goods flow from initial producers to final customers traveling through many layers of intermediaries and information is asymmetric. We study a dynamic model of bargaining in networks that captures these features. We show that the equilibrium price demanded over time is non-monotonic, but the sequence of transaction prices declines over time, with the possible exception of the last period. The price-dynamic is, therefore, reminiscent of fire-sale and hot-potato trade dynamic. Traders who intermediate the object arise endogenously and make a positive profit. The profit-earning intermediaries are not necessarily traders with many connections; for the case of multi-layer networks, they belong to the path that reaches the maximum number of potential buyers using the minimal number of intermediaries. This is not necessarily the path of the network that maximises the probability of consumption by traders who value the most the object (i.e., welfare).

Key words: Asymmetric Information, Bargaining, Bilateral Trading, Networks.

JEL Codes: C78, D82, D85

1. INTRODUCTION

We study strategic intermediation in markets characterised by incompleteness of trading opportunities, dispersed information, and resale. A finite number of risk-neutral traders

are connected in a network. A trader can bargain with another only if there is a link between them, and the absence of a link between two traders subsumes prohibitive trading costs. One arbitrary trader is the producer of a single valuable object. A trader's valuation for the object can be either low or high; it is private information and independent of other traders' valuations. In each round of trade, the current owner of the object either consumes it, in which case the game ends, or makes a take-it-or-leave-it offer to a connected trader that she chooses. The trader who receives the offer either accepts or rejects it. After this decision is taken, a new round of trade starts.

As a prime example, consider markets for agricultural goods in developing countries. Local producers access only a limited number of traders, who operate in geographically close markets; these traders, in turn, access other traders, who operate in markets located farther away from the original producers. Local producers could access such distant markets, but the lack of infrastructure and the difficulties of raising capital make such journeys infeasible. The structure of who can trade directly with whom, the network, is a technology that can be altered with large investments, like building a bridge or a road. However, in the short run and as a first approximation, it can be seen as fixed and known to market participants. Furthermore, negotiations are often bilateral, and products are exchanged for cash between intermediaries en route from local producers. Finally, each trader has precise information about the demand of the local market in which they operate—i.e., they know their valuation for the object. In developing countries, however, markets are dispersed and the communication infrastructures are poor. The implication is that asymmetric information among traders about the state of the demand in each local market is pervasive.¹

Incompleteness of trading opportunities, dispersed market information, and resale are also prominent features of over-the-counter (henceforth, OTC) financial markets. Products such as foreign currencies, swaps, forward rate agreements, and exotic options are often traded via bilateral negotiation in OTC markets. These securities are subject to counter-party risk, and, therefore, bonds of trusts and information flows between

1. This description is a brief summary of the stylised facts of markets in developing countries, which have been documented in several empirical studies, such as Fafchamps and Minten (1999, 2001), Jensen (2007), Aker (2010), Svensson and Yanagizawa (2009), and Allen (2014). These papers focus on a variety of goods, ranging from non-storable goods, such as fish in Jensen (2007), to highly storable goods, such as grain in Aker (2010).

firms/dealers are particularly important. This gives rise to trading costs that can be heterogeneous across pairs of traders. Furthermore, private observability of customer order flow and individual liquidity shocks are sources of asymmetric information among dealers. A large amount of inter-dealer trading is often observed, a phenomenon also referred to as hot-potato trading.²

Our analysis addresses the following questions: How does the underlying network structure affect the way market participants set prices? What are the network locations that provide larger payoffs to traders? What are the main welfare implications of trading in networks under asymmetric information? In Section 3, we answer these questions for the class of multi-layer networks. These networks capture environments in which there is a natural direction of trade from upstream to downstream traders. In the rest of this introduction, we present the results for general networks (contained in Section 4), and, when appropriate, we comment on additional insights that are obtained from the analysis of multi-layer networks.

The set of weak-Markov perfect Bayesian equilibria (henceforth “equilibria”) that we characterise has a simple structure. A high-value trader who acquires the object consumes it. In contrast, a low-value trader engages in a sequence of offers to other connected traders until the object is sold (unless, at some point, her own consumption value is higher than the discounted resale value of the object). All of her offers except the last come at prices that only traders with a high value are willing to accept. We refer to these offers as *consumption offers* because, once accepted, the object is consumed. Refused consumption offers are followed by an offer that low and high-value traders accept. We refer to these offers as *resale offers* because they come at a price equal to the expected revenue that the low-value trader obtains from reselling to other traders (i.e., their *resale value*).

We show that the equilibrium sequence of asking prices is non-monotonic in time. Prices in resale offers are decreasing in time because, as time passes, all traders become more pessimistic about the total expected demand in the network. However, prices in

2. We refer to Lyons (1997) for a seminal paper on hot-potato trading. Li and Schurhoff (2014) study the trade of US municipal bonds in the OTC market. They document that bonds move from the municipality through an average of six inter-dealer trades. They also document that there is systematic price dispersion across dealers, with earlier dealers maintaining systematically larger margins. We refer to Allen and Babus (2009) for a survey of networks in financial markets.

consumption offers spike, as sellers are attempting to exploit their positional power in the network to appropriate the surplus of connected traders with high value. The equilibrium price dynamic, then, is reminiscent of fire-sale and hot-potato trade dynamics: a low-value trader buys the object at a price that equals the expected sale price; if she fails to sell it at a high price, she is “forced” to sell it at a price lower than the one she paid. The fact that the game ends, once consumption offers are accepted, and that prices in resale offer are declining, imply that the realised sequence of transaction prices is also declining, with the possible exception of the last one.

We then investigate how the network location of a trader affects her payoffs. Only traders who receive a resale offer can make a positive profit. We call these traders *dealers*. Low-value dealers break-even in expected value: they buy at a price which equals their resale value. High-value dealers make a positive profit: they acquire the object at a price lower than their consumption value. Dealers arise endogenously, depending on their network location. In multi-layer networks, the dealers are the traders that lie on the path of the network that maximises the number of traders connected to the nodes in that path. In general networks, we show that if a high-value dealer is *essential* in connecting the local producer to another trader, then the former obtains a higher expected payoff.³

Finally we show that the equilibrium trading path does not always maximise ex-ante allocative efficiency. In multi-layer networks, when traders are sufficiently patient, the owner of the object makes consumption offers to all but one of her trading partners and, if all these offers are rejected, she makes a resale offer to the remaining partner (the dealer). The dealer is the trading partner with the highest resale value and, as already mentioned, is located on the path that maximises the number of traders connected to it. However, this path does not necessarily maximises the probability that a high-value trader consumes the object.

A central message in our paper is that the interplay between asymmetric information and strategic intermediation in networks leads to spatial price dispersion (the object is priced differently across network locations), and inefficiency. These predictions are in line with most of the empirical work studying markets in developing countries. Moreover, recent empirical studies have provided compelling evidence that the introduction of information technology in developing countries— e.g., the availability of mobile phones—

3. A trader i is essential to connect j to the initial producer if trader i lies in every path from the initial owner to trader j .

has reduced asymmetric information across markets and, as a consequence, has decreased spatial price dispersion and inefficiencies— see e.g., Jensen (2007), Aker (2010), and Svensson and Yanagizawa (2009). In our framework, we could think of such an innovation as leading to a setting in which either the demand of each local market (the private valuation of the trader) becomes common knowledge among traders, or the initial producer can directly access all other traders. In both cases, the equilibrium predicts constant prices across locations, larger profits to the initial producer, and allocative efficiency.

Understanding how networks affect trade is a central question in economics. A large literature has focused on buyer-seller networks—see Manea (2015) for a survey of this literature. A few recent papers have studied strategic intermediation in networks under complete information. Blume et al. (2009), Choi et al. (2016), Nava (2015), and Gale and Kariv (2007, 2009) focus on markets in which intermediaries post prices. In Kotowski and Leister (2014), negotiation occurs via auctions. Manea (2014) studies a model of bilateral bargaining with random selection of proposer.⁴ We contribute to this literature by examining, for the first time, a dynamic model of trade in the presence of asymmetric information. The prevalence of asymmetric information in decentralised markets motivates this extension, and, as we discussed above, our results provide novel predictions that are consistent with existing empirical evidence.⁵

Our paper also relates to bilateral bargaining models with one-sided asymmetric information. Some classical papers are Ausubel and Deneckere (1989), Fudenberg and Tirole (1983), Fudenberg et al. (1985), Gul et al. (1986) and Hart (1989). Most of this literature focuses on one seller and one buyer; exceptions are Fudenberg et al. (1987) and De Fraja and Muthoo (2000) in which there are multiple buyers. We extend the analysis to allow for resale; indeed, both Hart (1989) and De Fraja and Muthoo (2000) are a special case of our model. The possibility of resale generates new questions. It also creates new difficulties, as it implies that the continuation value of a trader depends on

4. Other related papers are Siedlerek (2015) and Gofman (2011). See Condorelli and Galeotti (2015) for a survey of models of strategic intermediation in networks.

5. Our approach, which combines both incomplete information and an explicit network structure, stands in contrast to recent models of trading that employ the random-matching approach pioneered by Rubinstein and Wolinsky (1985)— e.g., Duffie et al. (2005), Satterthwaite and Shneyerov (2007), and Golosov et al. (2014).

the entire vector of beliefs. As a consequence, some methods of analysis that are used in the classical bilateral bargaining model are not easily transposed to our setting.

Finally, our paper is related to the literature on the effect of resale on market outcomes. Calzolari and Pavan (2006) consider a simple seller-intermediary-buyer architecture. Haile (2003), Garratt and Tröger (2006), Hafalir and Krishna (2008), Jehiel and Moldovanu (1999) and Zheng (2002) analyse a complete network and their focus on auctions and inter-bidder resale. We work with a general network and explore a different set of questions.

2. THE NETWORK TRADING GAME

A *trading network* is a directed graph $G = (N, E)$, where $N = \{1, \dots, n\}$ is the set of *traders*, and $E \subseteq 2^{N \times N}$ is the set of directed edges; edges represent potential trading relationships. If $ij \in E$, we say that j is a *trading partner* of i —i.e., trader i can sell to trader j . The set of i 's trading partners is $N_i = \{j \in N \setminus \{i\} : ij \in E\}$.⁶

There is a single unit of an indivisible good, the *object*, initially owned by one of the traders, s^0 . Each trader i has a private value for the object, $v_i \in \{v_L, v_H\}$, where $0 < v_L < v_H$. The (common) prior probability that $v_i = v_H$ is $\mu_i \in (0, 1)$, and values are independently distributed. Without loss of generality, we restrict attention to networks in which there is at least one (directed) path from s^0 to each $i \neq s^0$.

There is an infinite number of rounds. Starting with $t = 0$, each round of trade is as follows: (i) The current owner of the object, denoted s^t (i.e., the *seller*), chooses to either consume the object or not and, in the latter case, makes a take-it-or-leave-it offer to one of her trading partners, $b^t \in N_{s^t}$ (i.e., the *buyer*), at some price $p_t \in \mathbb{R}_+$; (ii) b^t decides whether to accept the offer or reject it. The game ends if s^t consumes the object. If s^t does not consume, the game proceeds to round $t + 1$ and $s^{t+1} = b^t$ if b^t accepted the offer; otherwise, $s^{t+1} = s^t$. We assume perfect observation—i.e., all traders observe all decisions.

The t -period payoff to the seller is v_{s^t} if she consumes, p_t if her offer is accepted, and 0 otherwise; buyer b^t obtains $-p_t$ if she accepts, and 0 otherwise. All other traders get 0. Traders have a common discount factor $\delta \in (0, 1)$ and maximise their expected discounted sum of payoffs.

6. The trading network is undirected if $ij \in E$ implies that $ji \in E$.

We call the multi-stage extensive-form game with observed actions, independent types, and incomplete information that we have described above the *network trading game*.

A non-terminal history h^t at the beginning of period t is a sequence of offers made and purchasing decisions up to period t , that is, $h^t = ((i_\tau, p_\tau), a_\tau)_{\tau=0}^{t-1}$, where $(i_\tau, p_\tau) \in N \times \mathbb{R}_+$ is the offer made at period τ and $a_\tau \in \{a, na\}$ is the decision of trader i_τ whether to purchase the good at price p_τ (a stands for “accepted” and na for “not accepted”). We omit the reference to consumption decisions since the game ends as soon as a trader consumes the object. We denote by $h^t \oplus (i_t, p_t)$ the history that follows h^t if (i_t, p_t) is played at h^t . Behavioral strategies are defined in the usual way, see Fudenberg and Tirole (1991).

Throughout the paper, we restrict attention to *weak-Markov perfect Bayesian equilibria* (henceforth referred to as wMPBE). We refer to section 8.2.3 of Fudenberg and Tirole (1991) for a formal definition of Perfect Bayesian equilibrium. We impose the additional restriction that degenerate beliefs are never updated. A wMPBE is a perfect Bayesian equilibrium in which the consumption decision and the offer made by seller s^t at history h^t only depend on her private information, the profile of beliefs at h^t , and the offer (i_{t-1}, p_{t-1}) made in the previous round, if made by seller s^t . The acceptance strategy of buyer b^t at history $(h^t, (b^t, p_t))$ only depends on her private information, the profile of beliefs at $(h^t, (b^t, p_t))$, the price asked p_t , and the identity of the seller.⁷

Restricting attention to wMPBE implies that, at the beginning of a period, a seller’s equilibrium continuation payoff depends only on her private information and the profile of beliefs.⁸ We call *resale value* the continuation payoff of a low-value seller i at the beginning of a round of trade, and we denote this quantity as R_i , omitting reference to beliefs. We now present a preliminary result about *all* wMPBE.

7. In a strong-Markov PBE, the seller’s strategy only depends on the profile of beliefs at every history. It is well known that strong-Markov PBE do not always exist in bargaining games with incomplete information. For this reason, the literature has focused on wMPBE, see Fudenberg et al. (1985), and Fudenberg and Tirole (1983).

8. To see this, consider two histories (h^t, \tilde{h}^t) such that the profile of beliefs and the seller are the same at both histories. If $(i_{t-1}, p_{t-1}) = (\tilde{i}_{t-1}, \tilde{p}_{t-1})$, the statement follows directly from the definition of Markov strategies. If $(i_{t-1}, p_{t-1}) \neq (\tilde{i}_{t-1}, \tilde{p}_{t-1})$, the strategies alone do not imply the statement. However, if the continuation payoffs were different, the seller would have a profitable deviation, namely to deviate to the strategy that gives her the highest continuation payoff.

Proposition 1. *In every wMPBE:*

- (1) *A high-value seller consumes the object.*
- (2) *A seller never makes an offer to buyer b at a price strictly below her discounted resale value, δR_b . Hence, a low-value trader, with the exception of the initial seller, makes zero profit.*

To understand (1), note that the continuation payoff of a high-value seller must be at least v_H for her to not consume the object immediately. However, the sum of all continuation payoffs is at most δv_H , since the maximal surplus is v_H in every period. Therefore, at least one trader must have a strictly negative continuation payoff. This is impossible because each trader can always reject all offers and guarantee herself a zero payoff. This observation greatly simplifies our analysis as it implies that only low-value sellers bargain with their trading partners. As for (2), note that sellers have, in equilibrium, correct beliefs about the resale value of their trading partners and, since they hold all the bargaining power, they never offer the object at prices below discounted resale values.

Proposition 1 illustrates a number of properties that hold for *all* wMPBE. To sharpen the equilibrium analysis, however, we need to impose additional assumptions. In Section 3, we restrict attention to a particular class of networks, construct explicitly a wMPBE (thus proving existence), and discuss the insights that emerge from our construction. In Section 4, we show that these insights generalise to arbitrary networks.

3. MULTI-LAYER NETWORKS

In this section, we restrict attention to a particular class of networks, called *multi-layer networks*. A multi-layer network is an acyclic network with the additional property that for every distinct triples (i, j, j') such that $(ij, ij') \in E \times E$, there is no directed path from j to j' or from j' to j . Directed trees rooted at s^0 are examples of multi-layer networks. An example of a multi-layer network that is not a tree is depicted in Figure 1.

There are two main reasons for paying special attention to multi-layer networks.⁹ First, these networks capture environments in which there is a natural direction of

9. Multi-layer networks, or variants of them, are assumed in recent models of trading on networks with resale, e.g., Gale and Kariv (2009), Gofman (2011), Manea (2014) and Kotowski and Leister (2014).




FIGURE 1

A multi-layer network.

trade, from upstream to downstream traders, as is typical in most supply chains in manufacturing, retail, and agriculture. Second, in multi-layer networks, the buyers that a seller can reach either directly or indirectly— i.e., for which there is a path from the seller to the buyer— cannot have received offers in the past. This implies that we can construct wMPBE with the property that the continuation payoff of a low-value buyer upon accepting an offer (i.e., her discounted resale value) is independent of the offers that have been made before she acquires the object.

Throughout the section, we assume that $\mu_i = \mu_j$ for all $(i, j) \in N \times N$, so that the only source of heterogeneity among traders is their location on the trading network. We also focus the analysis on arbitrarily large discount factors and refer the interested reader to the appendix for a complete analysis.

3.1. Construction of an equilibrium

We first define an auxiliary game of bargaining between one seller and multiple buyers and construct an equilibrium (of that game). Second, we return to trading games on multi-layer networks, define a partition of the traders (the layers), and then construct an equilibrium recursively, using the equilibrium strategies of the appropriate auxiliary game.

3.1.1. An auxiliary game: multi-lateral bargaining. There is one seller s , with value v_L , who can bargain with a finite set $N_s = \{1, \dots, m\}$ of buyers. Each buyer

$b \in N_s$ has private valuation $v_b \in \{\delta R_b, \delta v_H\}$, with $R_b \in [v_L, v_H)$, and μ is the common prior that $v_b = \delta v_H$. Without loss of generality, we assume that $R_1 \geq R_2 \geq \dots \geq R_m$. This auxiliary game is similar to a network trading game in which the traders are $\{s\} \cup N_s$ and the set of (directed) edges is $E = \{sb : b \in N_s\}$. Later, we will interpret R_b as the continuation payoff of buyer b (if she were to buy the object).

We consider two separate cases: $R_1 > v_L$ and $R_1 = v_L$. For each case, we present an equilibrium path and leave the details of the equilibrium strategies to the appendix.

Lemma 1. *Assume that $R_1 = v_L$. For all sufficiently large discount factors, there exists an equilibrium with the following equilibrium path: (i) the seller asks δv_H , in sequence, to each buyer, (ii) high-value buyers accept the offer, whereas low-value buyers reject it, and (iii) if all offers at δv_H have been rejected, the seller consumes the object. As $\delta \rightarrow 1$, the seller's expected payoff converges to $(1 - (1 - \mu)^m)v_H + (1 - \mu)^m v_L$.*

The intuition for Lemma 1 is as follows. Since $\delta R_1 < v_L$ for all $\delta < 1$, there are strictly positive gains from trade only if the seller trades with a high-value buyer (i.e., when $\delta v_H > v_L$). As the seller has all the bargaining power, she can then extract the entire expected surplus by offering δv_H , in sequence, to each buyer and consuming the object if all offers have been rejected. A sufficiently large discount factor guarantees that the seller is indeed better off offering to all buyers in sequence rather than consuming immediately.

We now turn our attention to the case $R_1 > v_L$ (so that $\delta R_1 > v_L$ for sufficiently large δ). A special case is when there is a single buyer; this case is known as the ‘‘gap case’’ in the bilateral bargaining literature. In what follows, we say that the seller and a buyer engage in bilateral bargaining if they follow the equilibrium strategies of the bilateral bargaining game, as characterised in Hart (1989). We review Hart's construction later on.

Lemma 2. *Assume that $R_1 > v_L$. For all sufficiently large discount factors, there exists an equilibrium with the following equilibrium path: (i) the seller asks, in sequence, δv_H to all buyers $b > 1$, (ii) all high-value buyers $b > 1$ accept the offer δv_H , whereas low-value buyers $b > 1$ reject it, (iii) if all offers at δv_H have been rejected, the seller and buyer 1 engage in bilateral bargaining. As $\delta \rightarrow 1$, the seller's expected payoff converges to $(1 - (1 - \mu)^{m-1})v_H + (1 - \mu)^{m-1}R_1$.*

The logic behind Lemma 2 is as follows. The seller can guarantee herself δR_1 with an offer at δR_1 to buyer 1, since it is accepted with probability one (see Proposition 1). As a consequence, the seller has no gains to make from trading with low-value buyers, other than buyer 1. The best the seller can do is therefore to offer δv_H to all these buyers. In the event that all these offers are rejected, it becomes common belief that all buyers $b \neq 1$ have a low value and, from that point onward, the seller bargains exclusively with buyer 1.

We conclude this subsection by reviewing Hart's equilibrium construction. A reader familiar with the bargaining literature may skip the review. We already know that the seller never makes an offer to the buyer below δR_1 in equilibrium and, consequently, the low-value buyer accepts all offers up to δR_1 with probability one. Also, the seller's belief about the buyer having a high-value must decrease over time. This is known as the *skimming property* in the bargaining literature and follows from the observation that a high-value buyer has a strict incentive to accept all offers that a low-value buyer would accept. In turn, the skimming property implies that bargaining ends with probability one at or before $T < \infty$. Indeed, if at period T the seller's belief μ_T is less than R_1/v_H , the seller finds it profitable to sell the object at price δR_1 , which is accepted with probability one, rather than offering a higher price, which would be accepted with probability at most μ_T if it is less than δv_H (and with probability zero, otherwise). That bargaining ends in finite time then follows by noting that the seller's beliefs must decrease at a rate uniformly bounded from below and thus be less than μ_T in finite time (See Lemma 3 in Fudenberg et al. (1985) and Proposition 5 for a generalisation to network trading games.)

We now show how to construct an equilibrium. Assume that bargaining ends at or before period T (to be determined later). At the last period, the seller must offer $p_0 = \delta R_1$. Any higher offer would be rejected by the low-value buyer; the seller would then find it profitable to make another offer next period, contradicting that bargaining ends at that period. Any lower offer clearly decreases the seller's payoff. For future reference, let $\mu_0 = 0$ and $V_0(\hat{\mu}) = \delta R_1$ for all $\hat{\mu} \in [\mu_0, 1]$.

At the second to last period, the seller must offer a price p_1 that makes the high-value buyer indifferent between accepting and rejecting the offer, i.e., $\delta v_H - p_1 = \delta(\delta v_H - p_0)$. Any higher offer is rejected, which leads to a payoff of $\delta^2 R_1$ to the seller; the seller would be better off offering δR_1 . Any lower offer clearly decreases the seller's payoff. Moreover, the high-value buyer must accept the offer p_1 with probability one for the

seller's maximisation problem to have a solution. The seller's expected payoff is therefore $\hat{\mu}p_1 + \delta(1 - \hat{\mu})p_0$, where $\hat{\mu}$ is the seller's current belief. Observe that we can rewrite the seller's expected payoff as

$$V_1(\hat{\mu}) = \frac{\hat{\mu} - \mu_0}{1 - \mu_0} p_1 + \frac{1 - \hat{\mu}}{1 - \mu_0} \delta V_0(\mu_0),$$

for all $\hat{\mu} \in [\mu_0, 1]$. Finally, for the seller to indeed have an incentive to offer the sequence (p_1, p_0) instead of (p_0) when her current belief is $\hat{\mu}$, we need that $V_1(\hat{\mu}) \geq V_0(\hat{\mu})$. This is equivalent to $\hat{\mu} \geq \mu_1$, with μ_1 the unique solution to $V_1(\mu_1) = V_0(\mu_1)$.

Continuing inductively, given the triple (p_k, V_k, μ_k) , we define the triple $(p_{k+1}, V_{k+1}, \mu_{k+1})$ as follows. First, p_{k+1} makes the high-value buyer indifferent between accepting and rejecting the offer, i.e., $\delta v_H - p_{k+1} = \delta(\delta v_H - p_k)$. Second, for all $\hat{\mu} \in [\mu_k, 1]$, $V_{k+1}(\hat{\mu})$ is the seller's expected payoff when her current belief is $\hat{\mu}$, her next period belief is μ_k , and bargaining continues for an additional $k + 1$ periods, i.e.,

$$V_{k+1}(\hat{\mu}) = \frac{\hat{\mu} - \mu_k}{1 - \mu_k} p_{k+1} + \frac{1 - \hat{\mu}}{1 - \mu_k} \delta V_k(\mu_k).$$

Third, the threshold μ_{k+1} makes the seller indifferent between bargaining for an additional $k + 1$ periods and k periods, i.e., μ_{k+1} is the unique solution to $V_{k+1}(\mu_{k+1}) = V_k(\mu_{k+1})$.¹⁰ Note that we have $\mu_{k+1} > \mu_k$, and $V_{k+1}(\hat{\mu}) - V_k(\hat{\mu}) \geq 0$ for all $\hat{\mu} \geq \mu_{k+1}$.

Finally, we determine T as the solution to $\mu_{T+1} > \mu \geq \mu_T$. The following proposition summarises the equilibrium construction (with a slight abuse of notation, we let $\mu_T = \mu$).

Proposition 2 (Hart (1989)). *The bilateral bargaining game has a (generically) unique equilibrium, where (i) bargaining ends on or before period T with probability one, (ii) at each period t , the seller offers p_{T-t} , (iii) at each period $t \leq T - 1$, the high-value buyer accepts the offer p_{T-t} with probability $(\mu_{T-t} - \mu_{T-t-1}) / ((1 - \mu_{T-t-1})\mu_{T-t})$, and (iv) the low-value buyer accepts the offer p_0 with probability one. As $\delta \rightarrow 1$, the seller's expected payoff converges to R_1 .*

3.1.2. Layers and equilibrium construction. We construct the following sets, called the *layers*, recursively. First, we define $L_0 = \{i \in N : N_i = \emptyset\}$; then, for $k \geq 1$, we define $L_k = \{i \in N \setminus \cup_{\ell=0}^{k-1} L_\ell : N_i \subseteq \cup_{\ell=0}^{k-1} L_\ell\}$. Since G is a multi-layer network,

10. Note that $V_{k+1}(1) - V_k(1) = p_{k+1} - p_k > 0$, $V_{k+1}(\mu_k) - V_k(\mu_k) = (\delta - 1)V_k(\mu_k) < 0$. Thus, the existence and uniqueness of a solution follows from the linearity of $\hat{\mu} \mapsto (V_{k+1}(\hat{\mu}) - V_k(\hat{\mu}))$, which is increasing.

there exists $k^* < \infty$ such that $\{L_0, \dots, L_{k^*}\}$ is a partition of N and $L_{k^*} = \{s^0\}$.¹¹ This partition creates an ordering of traders based on their resale opportunities. Traders in L_0 have no resale opportunities (we refer to them as end-traders); traders in L_k can resell only to subsets of traders belonging to $\cup_{\ell=0}^{k-1} L_\ell$. In Figure 1, we have that $k^* = 3$, and $L_0 = \{3, 5, 6, 7\}$, $L_1 = \{4\}$, $L_2 = \{1, 2\}$ and $L_3 = \{s^0\}$.

Suppose that trader $i_t \in L_k$ acquires the object at some history $h^t \oplus (i_t, p_t)$. If $k = 0$, trader i_t consumes and we set $R_{i_t} = v_L$. If $k = 1$, in view of Proposition 1, high-value trader i_t consumes at the beginning of period $t + 1$. Low-value trader i_t may bargain with her trading partners $b \in N_{i_t} \subseteq L_0$. Since high-value trader i_t consumes, it is common belief that trader i_t has a low-value if she makes an offer. Hence, trader i_t and her trading partners $b \in N_{i_t}$ face a multi-lateral bargaining game, with $R_b = v_L$ for all $b \in N_{i_t}$. We assume that trader i_t and her trading partners $b \in N_{i_t}$ follow the equilibrium strategies of Lemma 1, regardless of the history $h^t \oplus (i_t, p_t)$ at which trader i_t has acquired the object. We let R_{i_t} be the continuation payoff of trader i_t if she acquires the object. If $k > 1$, having defined R_i for all $i \in \cup_{k' < k} L_{k'}$, we complete the construction of the equilibrium by proceeding recursively in the layer structure, appealing to the equilibrium strategies of Lemma 2 between trader i_t and her trading partners $N_{i_t} \subseteq \cup_{k' < k} L_{k'}$.

To recap, in equilibrium, the current owner of the object makes a sequence of offers at δv_H to all her trading partners but the one with the highest resale value. If all these offers are rejected, the owner engages in bilateral bargaining with her trading partner having the highest resale value. (If there are multiple trading partners with the highest resale value, choose one arbitrarily.) Thus, a trader receives either no offers, or offers at her discounted resale value, which we call *resale offers*, or offers that makes the high-value trader indifferent between accepting and rejecting, which we call *consumption offers*.¹² As we shall see later, these are general features of what we call *regular equilibria* in Section 4. In the rest of the paper, we call *dealers* the traders who, along the equilibrium path, receive a resale offer with positive probability.

11. See Lemma 1, p. 513, of Renou and Tomala (2012) for a proof.

12. Strictly speaking, an offer is a pair (i, p) . With a slight abuse of terminology, we refer to the price associated with an offer as the offer itself.

3.2. Dealers, payoffs and price dynamic

The next proposition characterises the set of dealers, their payoffs, and the price dynamic, as traders becomes infinitely patient.

For a directed path (i_1, \dots, i_m) from trader i_1 to trader i_m , we let $\mathfrak{s}((i_1, \dots, i_m)) := \sum_k |N_{i_k}|$ be the sum of outgoing edges and $l((i_1, \dots, i_m))$ the length of the path (i.e., $m - 1$). Note that $\mathfrak{s}((i_1, \dots, i_m))$ is the number of traders who are potentially offered the good if trader i_1 owns the good, should the trading route be (i_1, \dots, i_m) . Let $\kappa((i_1, \dots, i_m)) := \mathfrak{s}((i_1, \dots, i_m)) - l((i_1, \dots, i_m))$. It is worth noting that $n - 1 \geq \kappa((i_1, \dots, i_m)) \geq 0$ for all paths; the lower bound obtains when the network is a line, whereas the upper bound obtains when the network is a star, with i_1 as the center.¹³ In Figure 1, $\mathfrak{s}((s^0, 1, 4)) = 7$, $\mathfrak{s}((s^0, 2, 4)) = 6$, and $l((s^0, 1, 4)) = l((s^0, 2, 4)) = 2$.

Proposition 3. *Assume that there exists a unique path (i_1^*, \dots, i_m^*) that maximises κ among all paths (i_1, \dots, i_m) from the seller to traders in layer L_1 (i.e., $(i_1, i_m) \in \{s^0\} \times L_1$). As $\delta \rightarrow 1$, the sequence of equilibria we have constructed satisfies the following at the limit:*

- (i) *The sequence of dealers converges to (i_2^*, \dots, i_m^*) .*
- (ii) *The resale offer to dealer i_k^* , $k \geq 2$, converges to $v_H - (1 - \mu)^{\kappa((i_k^*, \dots, i_m^*))}(v_H - v_L)$.*
- (iii) *The payoff of the initial seller converges to $v_H - (1 - \mu)^{\kappa((i_1^*, \dots, i_m^*))}(v_H - v_L)$, and the payoff of high-value dealer i_k^* , $k \geq 2$, converges to*

$$(1 - \mu)^{\left(\sum_{\ell < k} |N_{i_\ell^*}\right) - 1} (1 - \mu)^{\kappa((i_k^*, \dots, i_m^*))} (v_H - v_L) = (1 - \mu)^{\kappa((i_1^*, \dots, i_m^*)) + k - 2} (v_H - v_L).$$

As an illustration of Proposition 3, the unique maximising path is $(s^0, 1, 4)$ in Figure 1. The resale offer to trader 1 converges to $v_H - (1 - \mu)^4(v_H - v_L)$, which is higher than trader 4's resale offer of $v_H - (1 - \mu)^3[v_H - v_L]$. High-value trader 1's expected payoff is $(1 - \mu)^5[v_H - v_L]$, which is higher than high-value trader 4's expected payoff of $(1 - \mu)^6[v_H - v_L]$.

A number of insights emerge from the above characterisation. As we show in Section 4, these insights are not limited to the equilibrium we have constructed, nor are they to multi-layer networks.

13. The length of the empty path is normalised to zero.

First, the path (i_1^*, \dots, i_m^*) maximises the number of traders who potentially receive a consumption offer. Along that path, the initial seller and the dealers are able to price-discriminate most. Note, however, that the path (i_1^*, \dots, i_m^*) is not necessarily one of the paths that maximise the likelihood of a high-value trader consuming the good and, thus, the expected surplus. For an illustration, consider the multi-layer network in Figure 2. We have that $L_0 = \{3, 4, 7\}$, $L_1 = \{1, 6\}$, $L_2 = \{5\}$, $L_3 = \{2\}$, $L_4 = \{s^0\}$, the unique maximising path is $(s^0, 1)$, and $\kappa((s^0, 1)) = 3$, whereas $\kappa((s^0, 2, 5, 6)) = 2$. Therefore, for sufficiently high discount factor, there exists an equilibrium in which trader 1 is the dealer. Since $\mathfrak{s}((s^0, 1)) = 4$, the ex-ante expected total surplus generated converges to $v_H - (v_H - v_L)(1 - \mu)^4$, as δ goes to one. However, the ex-ante total surplus is higher if the object flows from s^0 to 7 along $(s^0, 2, 5, 6)$ with a consumption offer to 1; it is $v_H - (v_H - v_L)(1 - \mu)^5$. Although along the path $(s^0, 2, 5, 6)$, there is a higher number of traders that can receive an offer relative to $(s^0, 1)$, the dealers 5 and 6 have a single resale opportunity each. This results in a low resale value for trader 2. The alternative path $(s^0, 1)$ is short, but each dealer has two resale opportunities, leading to a higher resale value for trader 1. We can, indeed, rank networks based on the expected surplus they generate at the equilibrium described in Proposition 3. Let $(i_1^*, \dots, i_m^*)_G$ be the path that maximises κ among all paths (i_1, \dots, i_m) from the seller to traders in layer L_1 in network G . Then, as δ goes to one, the equilibrium expected surplus in G is higher than in G' whenever $\mathfrak{s}((i_1^*, \dots, i_m^*)_G) > \mathfrak{s}((i_1^*, \dots, i_m^*)_{G'})$. For example, fixing the number of nodes, both the star and line networks maximise the equilibrium expected surplus.¹⁴

Second, resale offers are declining as the object flows from the initial seller to end-traders through the path (i_1^*, \dots, i_m^*) . Since $\kappa((i_k^*, \dots, i_m^*)) = |N_{i_k^*}| - 1 + \kappa((i_{k+1}^*, \dots, i_m^*))$, we have that $R_{i_k^*} = (1 - (1 - \mu)^{|N_{i_k^*}| - 1})v_H + (1 - \mu)^{|N_{i_k^*}| - 1}R_{i_{k+1}^*} \geq R_{i_{k+1}^*}$ (see Equation (1) for a generalisation to arbitrary networks). Resale offers reflect the expected surplus that future trade generates; these surpluses decrease as the object flows from the initial seller to end-traders because later dealers have fewer resale opportunities. Moreover, the sequence of offers is of the form $(v_H, \dots, v_H, R_{i_2^*}, v_H, \dots, v_H, R_{i_k^*}, v_H, \dots)$ and is, in

14. Manea (2014) studies a model of bilateral bargaining on acyclic networks. He assumes complete information and that offers are made either by the buyer or the seller; the proposer is randomly selected. His characterisation relies on a decomposition of the network in a sequence of layers. In his model, traders in the same layer face a similar downstream competition. This contrasts with our model where traders in the same layer face similar resale opportunities. Like us, Manea shows how strategic intermediation in networks can lead to inefficient outcomes.




FIGURE 2
Inefficiency.

general, non-monotonic. As mentioned in the introduction, this is reminiscent of fire-sale and hot-potato trade dynamics. If dealer i_k^* fails to sell the good at price v_H to one of her trading partners, she ends up selling the good at price $R_{i_{k+1}^*}$ to dealer i_{k+1}^* , thus making a loss. Yet, the sequence of prices at which transactions take place is decreasing, with the possible exception of the last one. Indeed, as soon as a consumption offer is accepted, the game ends next period and, consequently, transaction prices must be resale offers, with the possible exception of the last one.

Third, Proposition 3 states that the initial seller as well as high-value dealers expect a positive profit. No other traders make a positive profit. To see this, note that if trader i receives no offers, her payoff is zero. Next, assume that trader i receives consumption offers only, and focus on the last consumption offer that trader i receives (this offer exists because the game ends in finite time). Since, by definition, consumption offers make high-value traders indifferent between accepting and rejecting them, the last consumption offer to trader i must come at a price of v_H .¹⁵ This implies that all previous offers (if any) must be at a price of v_H , as well. Trader i 's payoff is, therefore, zero.¹⁶ Turning to dealers, all low-value dealers obtain zero profit by Proposition 1.

Fourth, the payoff of high-value dealers is declining with their distance from the initial seller. The payoff of a high-value dealer is determined by the probability of receiving

15. More precisely, it comes at a price of δv_H , which converges to v_H as δ converges to 1.

16. It is worth noting that, in addition to resale offers, dealers receive consumption offers too, when they engage in bilateral bargaining.

her first resale offer as well as its price. Since resale offers decline along the equilibrium path, conditional on receiving a resale offer, earlier high-value dealers acquire the object at a higher price than later high-value dealers (i.e., $R_{i_k^*} > R_{i_{k'}^*}$ for $k < k'$). However, earlier high-value dealers have a higher probability of receiving the resale offer than latter high-value dealers (i.e., $(1 - \mu)^{\left(\sum_{\ell < k} |N_{i_\ell^*}\right)^{-1}} > (1 - \mu)^{\left(\sum_{\ell < k'} |N_{i_\ell^*}\right)^{-1}}$ for $k < k'$). The second effect dominates the former. Intuitively, the decline in resale offers offsets the expected demand of the traders that receive consumption offers, but it does not incorporate the possibility that dealers consume the object. As a consequence, the difference in resale values between later and earlier high-value dealers does not compensate for the decrease in the probability of obtaining the offer.

Finally, the payoff of the initial seller increases with a change in the network that increases $\kappa((i_1^*, \dots, i_m^*))$. It is the minimum in the line, where $\kappa((i_1^*, \dots, i_m^*)) = 1$, and maximum in the star, where $\kappa((i_1^*, \dots, i_m^*)) = n - 1$.

4. EQUILIBRIA IN GENERAL TRADING NETWORKS

The equilibrium we have constructed for the multi-layer network trading game satisfies two key properties: (i) the skimming property and (ii) that every equilibrium offer above the discounted resale value of a trader makes the high-value trader indifferent between accepting and rejecting it. We now show that *all* network trading games have an equilibrium satisfying these two properties. Most importantly, these two properties imply the results on pricing and payoffs that we have already emphasised in the previous section.

Definition 1. *A wMPBE is regular if it satisfies the following conditions:*

- a. *Skimming: For all periods t , for all histories $h^t \oplus (i_t, p_t)$, if the low-value trader i_t accepts the offer p_t with positive probability, then the high-value trader i_t accepts p_t with probability one.*
- b. *Indifference: For all periods t , for all histories h^t , let $p_{i_t}^*(h^t)$ be the highest (supremum) price that trader i_t accepts with probability one over all histories $\{h^t \oplus (i_t, p) : p \in \mathbb{R}_+\}$. If $p_{i_t}^*(h^t)$ is higher than trader i_t 's discounted resale value at h^t , then the high-value trader i_t is indifferent between accepting and rejecting $p_{i_t}^*(h^t)$.*

Proposition 4. *A regular equilibrium exists for every network trading game.*

As discussed in Section 3, the existence of an equilibrium is proved by induction on beliefs in classical bilateral bargaining problems. This logic may not extend to network trading games. Indeed, with multiple buyers and resale, the acceptance strategy of a buyer as well as the offers a seller makes depend on the profile of beliefs of *all* traders, a multi-dimensional object. So, instead of doing an induction on the profile of beliefs, we do an induction argument on time, a one-dimensional object. First, we prove that each network trading game with a finite horizon has a regular equilibrium. Then, using the fact bargaining ends in finite time in all regular equilibria (see Proposition 5), we complete the proof by an argument analogous to that of Chatterjee and Samuelson (1988). Despite its simple logic, the proof is fraught with technical difficulties and relegated to online Appendix A.

We now discuss our definition of regularity. The *skimming property* ensures that, for each buyer, the common belief that she has high-value decreases as she rejects offers. This, in turn, is key to proving that the game ends in finite time, a result that we use throughout. As already mentioned, the skimming property holds in classical bilateral bargaining problems. Unfortunately, we have not been able to prove that it also holds in network trading games. The main difficulty is that resale values change over time as the profile of beliefs changes, which implies that the valuation a low-value trader has for the object changes over time, possibly non-monotonically. This sharply contrasts with the classical model, where the valuation is fully persistent (given by the consumption value).

The *indifference condition* allows us to prove that, in every regular equilibrium, offers are at a price that makes either the low-value or the high-value buyer indifferent between accepting and rejecting. The indifference property is satisfied in classical bilateral bargaining problems. It is guaranteed by the fact that the buyer's equilibrium payoff correspondence is monotonic in her belief. However, when a seller faces multiple buyers and buyers can resell, a buyer's continuation payoff depends on the entire profile of beliefs, and we cannot guarantee that her equilibrium payoff correspondence is monotonic in her own belief. As a consequence, we cannot rule out equilibria in which the high-value buyer accepts up to a price that gives her a strictly positive payoff, but rejects with probability one any slightly higher price. The indifference condition allows us to tie consumption offers with continuation payoffs and, thus, with future offers. Without this property, it

is doubtful whether one could say anything meaningful about the price dynamic.

To sum up, we have neither been able to prove that all wMPBE are regular nor that there exist non-regular wMPBE. We focus on regular equilibria as it allows us to obtain sharp results thanks to the following proposition.

Proposition 5. *In a regular equilibrium:*

- (1) *There exists $T^* < \infty$ such that consumption takes place before round T^* in any on-path terminal history.*
- (2) *Each offer made to trader i is either a resale offer at price δR_i , or a consumption offer.*

4.1. Price dynamics and payoffs

Regular equilibria retain most of the features of the equilibrium we have constructed for multi-layer networks. A high-value trader who acquires the object consumes it. A low-value trader who acquires the object makes a sequence of consumption offers to her trading partners and, in the event that all these offers are rejected, makes a resale offer, which is accepted (unless, at some point, the discounted resale value of the object is lower than her own consumption value). Hence, in equilibrium, any on-path terminal history of the game can be summarised by a list of consumption offers and resale offers, ending with consumption.

Fix a regular equilibrium. For an arbitrary on-path terminal history, let p_ℓ^s indicate the ℓ -th consumption offer that the s -th seller makes, r_s^{s-1} indicate the resale offer that the $(s-1)$ -th seller makes to the s -th seller, and $(p_1^s, p_2^s, \dots, r_{s_1}^s, \dots, r_s^{s-1}, p_1^s, \dots, p_k^s, r_{s+1}^s, p_1^{s+1}, \dots)$ be the sequence of consumption and resale offers. If sellers do not randomise, the resale offer that seller $s-1$ makes to s is given by

$$\begin{aligned} r_s^{s-1} &= \delta R_s \\ &= \delta \left[\alpha_1 p_1^s + \delta(1 - \alpha_1) \alpha_2 p_2^s + \dots + \prod_{\ell=1}^k (1 - \alpha_\ell) \delta^{k-1} r_{s+1}^s \right], \end{aligned} \tag{4.1}$$

where α_ℓ indicates the probability that the consumption offer p_ℓ^s is accepted, and k the number of consumption offers seller s makes.¹⁷ In words, the discounted resale value of a

¹⁷ In general, a seller could be indifferent between different offers, in which case there may be equilibria at which the seller randomises. In these equilibria, the sequence of consumption and resale

trader who acquires the object equals, in equilibrium, the expected discounted sale price. A similar formula already appeared in Section 3 (see Proposition 3 and its discussion). The following proposition generalizes Proposition 3 to arbitrary trading networks.

Proposition 6. *In a regular equilibrium, the following conditions hold:*

- (1) *Along all terminal histories, discounted resale offers are decreasing in time – i.e., if r_t is the resale offer at t and $r_{t'}$ is the resale offer at t' , with $t' > t$, then $\delta^t r_t \geq \delta^{t'} r_{t'}$.*
- (2) *Along all terminal histories, any consumption offer is greater than the first ensuing resale offer – i.e., if p_t is the consumption offer at t and $r_{t'}$ is the resale offer at $t' > t$, and every offer from $t + 1$ to $t' - 1$ is a consumption offer, then $p_t > r_{t'}$.*
- (3) *The initial seller and high-value dealers have a positive expected payoff; all other traders have a zero payoff. Moreover, if all offers to trader j are always preceded by a resale offer to dealer i , then the expected payoff of high-value dealer i is strictly greater than the expected payoff of trader j .*

As illustrated in Section 3, in multi-layer networks, resale offers decline over time primarily because later dealers have fewer resale opportunities. This is not necessarily true in arbitrary networks, where many dealers may have the same resale opportunities (e.g., if the network is a wheel). To understand the first part of Proposition 6, consider two consecutive resale offers, the first at period t and the second at period $t' > t$. The trader who receives the resale offer at t' knows that all consumption offers from period $t + 1$ to period $t' - 1$ have been rejected. When a trader rejects a consumption offer, all other traders update downward their belief that she has a high consumption value; thus, the trader who receives the resale offer at period t' is more pessimistic than the trader who receives the resale offer at period t . Hence, in our bargaining game where sellers make all the offers, resale offers, evaluated at a given fixed date, decline over time.

Turning to the second part of Proposition 6, the result that consumption offers are above the ensuing resale offers reflects the ability of sellers to use their local bargaining power to demand a high price from some of their trading partners, before passing the object to another dealer.

offers, along the equilibrium path, follows a stochastic process determined by the sellers' equilibrium strategies. Equation (4.1) must hold for every possible sequence of offers.

Finally, the third part of Proposition 6 follows from the very same logic as in Section 3. Yet, we note that a resale offer is never at a price higher than $\delta^2 v_H$, since the resale value of every trader is bounded from above by δv_H . Thus, although we do not have an explicit formula for the payoff of high-value dealers, a high-value dealer must obtain at least $\delta v_H(1 - \delta) > 0$, if she accepts a resale offer.

Proposition 6 has two immediate implications. First, prices demanded over time are non-monotonic, a fact we already pointed out in Section 3. This follows by combining two observations: (i) equilibrium discounted resale values decline over time (part 1 of Proposition 6) and (ii) in equilibrium, the discounted resale value of a trader equals the expected discounted price at which she will sell the object (Equation 4.1).

Second, there is a clear relationship between the location of traders on the trading network and their payoffs, as shown in the next corollary. We say that trader j is *essential* for trader i if j belongs to every (directed) path from the initial seller to trader i . Trader i is an *end-trader* if $ij \notin E$ for all j (e.g., in a multi-layer network traders in L_0 are end-traders).

Corollary 1. *In a regular equilibrium:*

- 1) *Every end-trader has a zero payoff.*
- 2) *If trader j is essential for trader i , then high-value trader j obtains a higher expected payoff than high-value trader i .*

The corollary points out the importance of the location of a trader in a trading network. It emphasises that traders who are essential in connecting other traders to the initial owner obtain a payoff advantage. Among the essential traders, the ones who are located more upstream in the network (i.e., closer to the seller) enjoy higher surplus.

4.2. *Trading cycles*

In multi-layer networks, a trader does not have the opportunity to buy back the object, so that the object cannot cycle. This is no longer true in general trading networks. We illustrate the possibility of a cycle with the help of the following example. Consider the




FIGURE 3

An example of trading cycle.

network in Figure 3, and assume that $1 > \mu_3 \geq \mu_4 > 0$ and $\mu_2 = \mu_1 = 0$, i.e., traders 1 and 2 have valuation v_L with probability one.¹⁸

If trader 2 acquires the good, she can obtain a payoff $\mu_4 \delta v_H + (1 - \mu_4) \delta v_L$ by offering δv_H to trader 4 and consuming upon rejection. Thus, trader s^0 can obtain at least $\delta(\mu_4 \delta v_H + (1 - \mu_4) \delta v_L)$, extracting the entire surplus from trader 2. However, trader s^0 can do better by first offering the object to trader 1, with the implicit promise to buy it back later in case trader 1 does not sell the object to trader 3.

More formally, for δ sufficiently high, there exists a regular equilibrium with the following equilibrium path. Initially, trader s^0 makes a resale offer to trader 1 at price

$$\delta^2 \mu_3 v_H + \delta^5 (1 - \mu_3) (\mu_4 v_H + (1 - \mu_4) v_L).$$

Trader 1 accepts the offer and makes a consumption offer to trader 3 at price δv_H . If the offer is rejected, trader 1 resales the object to trader s^0 at price $\delta^3 (\mu_4 v_H + (1 - \mu_4) v_L)$. Trader s^0 accepts and makes a resale offer to trader 2 at price $\delta^2 (\mu_4 v_H + (1 - \mu_4) v_L)$. Trader 2 buys and, in turn, makes a consumption offer to trader 4 at price δv_H . Finally, if the offer is rejected, trader 2 consumes the object.

The intuition as to why cycling occurs is that, for sufficiently large discount factors, the consumption value of a low-value seller is lower than the resale value of her trading partners who can resell to traders, still believed to have a high-value with positive

18. The same equilibrium outcome obtains if μ_1 and μ_2 are positive, but arbitrarily small. A full specification of the equilibrium strategy of the example is in on-line Appendix B.

probability. Hence, the object tends to move toward traders whose value is still uncertain. We thus conjecture that, in a strongly connected network, as traders become fully patient, the object will indeed cycle until it becomes a common belief that all traders in the network have low value – at which point the current owner consumes.

5. CONCLUDING DISCUSSION

Our model is an abstraction of many decentralised markets in which trade proceeds from a producer to a final customer, traveling through many possible layers of intermediation. We have shown that: (i) the sequence of prices at which the object is exchanged is not constant over time, (ii) the seller does not extract the entire surplus; dealers extract some of the surplus too, (iii) the allocation is not necessarily efficient, and (iv) the object may cycle. The assumptions that we employ are all relatively standard in the literature, and were chosen to build the simplest possible model that incorporates resale and asymmetric information. Nonetheless, a few of them deserve further discussion.

To start with, we assume incomplete information. This assumption is obviously natural and essential for our results. Indeed, assume complete information, i.e., $\mu_i \in \{0, 1\}$ for all $i \in N$. It is not difficult to see that the following properties hold in equilibrium: (i) the sequence of prices at which the object is exchanged is constant over time, (ii) the seller extracts all the entire surplus, (iii) the outcome is allocative efficient, and (iv) the object never cycles in the network, e.g., a trader receives at most one offer. These properties differ sharply from the one we have highlighted.

We assume that actions are perfectly observable. This assumption is at odds with the idea that information is dispersed across traders, but relaxing it would entail considerable technical difficulties. Suppose that only the buyer who received the offer could observe it. Then, as the game proceeds, two traders may end up with different beliefs about the value of a third trader. For instance, if seller s has made an offer to buyer b_1 , she may have a belief about the value of buyer b_1 different than the one held by buyer b_2 , who may not be certain whether buyer b_1 has received an offer yet. Higher-order beliefs would then start to play a non-trivial role in the analysis.

The assumption that each trader has only two possible values assures that bargaining takes place under one-sided asymmetric information. This is necessary to progress in the analysis, given the additional intricacies that the resale opportunities introduce, see

Zheng (2002). However, a minimal feasible generalisation of our information structure consists in assuming that traders are heterogeneous with regard to their low value (i.e., $v_i \in \{v_L^i, v_H\}$ with $v_L^i < v_H$ for all $i \in N$). Note that, in our model, the willingness to pay of a low-value buyer is $\max\{\delta v_L, \delta R_b\}$ and since resale values are generally different across traders, our current analysis already deals with such heterogeneity.

For the sake of tractability, we also assume that the seller makes all the offers: allowing buyers to make offers alters their signalling possibilities. We acknowledge that some of our results hinge on this assumption.¹⁹ Nonetheless, we expect our results to be robust in giving a certain amount of bargaining power to buyers. In fact, the fundamental force toward price decline that we have uncovered (i.e., the decrease in everybody's resale value as consumption offers are being made and rerejected) would still be present with alternating-offer bargaining.

Finally, in our model, any two traders have the same information about the valuation of any third trader. An alternative assumption would be that each trader knows her valuation and, in addition, the valuation of her closest trading partners, but remains uncertain about the valuation of all other traders. For multi-layer networks, we can slightly modify the construction derived in Section 3.1.1 to characterise an equilibrium in this environment. To see this, consider Figure 1 and assume that $\delta v_H > v_L$. Low-value trader 4 can ask and get δv_H if one of her trading partner has high value; otherwise, she consumes, and gets v_L . Suppose that low-value trader 1 has the object. Again, if either trader 3 or 4 has a high value, then trader 1 obtains δv_H . If they both have a low value, then trader 1 faces a bargaining game where trader 3 has valuation δv_L with probability one, whereas trader 4 has valuation δv_L with probability $(1 - \mu)^3$ and has valuation δv_H otherwise. By deriving the equilibrium in this subgame in the standard way, we can compute how much the initial seller s expects low-value trader 1 to be willing to pay in the event that all of trader 1's trading partners have a low value. The method to construct the equilibrium then follows the same logic that we developed in Section 3.1.1.

19. Suppose the game has alternating-offer and an initial seller s is facing a single buyer b . Let $v_L = 0$. If s is sufficiently pessimistic about the value of b , then she will bargain exclusively with the low-value b , and will only extract a share of the available gains from trade (e.g., $\delta R_b / (1 + \delta)$ assuming Rubinstein's bargaining game). Hence, for a sufficiently high discount factor, s 's resale value will be lower than the discounted resale value of b . Also note that the payoff of s could be lower than that of b . See Ausubel et al. (2002) for the analysis of the buyer-seller game.

APPENDIX

A. MULTI-LAYER NETWORKS

Proposition A characterises an equilibrium for the auxiliary game introduced in Section 3. Lemma 1 and Lemma 2 follow directly from Proposition A. For a given strategy profile and belief system, let h^ℓ be an history in which $\mu_1 \leq \mu$, there are $\ell \geq 0$ buyers $i \neq 1$ with $\mu_i = \mu$ and the remaining buyers with $\mu_i = 0$. The set of all such histories is denoted by \widehat{H} . As we shall show below, in the equilibrium that we characterise, all histories are of type h^ℓ , with $\ell \geq 0$. We also denote by $\Pi_{1s}(\hat{\mu}, \delta)$ and $\Pi_{s1}(\hat{\mu}, \delta)$ the payoff of the buyer and seller, respectively, in the unique equilibrium of the bilateral bargaining game with prior $\hat{\mu}$ summarised in Proposition 2.

Proposition A. Consider the auxiliary network trading game. The following is an equilibrium. At every history h^ℓ with $\ell \geq 0$:

1. Buyers $i \neq 1$ and low-value buyer 1 accept all offers at a price $p \leq \delta v_i$, and reject all other offers.
2. Suppose $\delta R_1 \leq v_L$. High-value buyer 1 accepts $p \leq \delta v_H$. If $\ell = 0$, the seller consumes if $v_L \geq \delta[\mu_1 v_H + (1 - \mu_1)v_L]$, otherwise asks δv_H to buyer 1. If $\ell > 0$, the seller consumes if $v_L \geq \delta[\mu v_H + (1 - \mu)v_L]$, otherwise asks δv_H to the highest index buyer $i \neq 1$ with $\mu_i = \mu$.
3. Suppose $\delta R_1 > v_L$ and $R_1 \geq \mu v_H + (1 - \mu)\delta R_1$. High-value buyer 1 accepts $p \leq \delta v_H(1 - \delta) + \delta^2 R_1$. The seller offers δR_1 to buyer 1.
4. Suppose $\delta R_1 > v_L$ and $R_1 < \mu v_H + (1 - \mu)\delta R_1$. High-value buyer 1 accepts every price $p \leq \delta v_H - \delta^{\ell+1}(1 - \mu)^\ell \Pi_{1s}(0, \delta)$, rejects every price $p \geq \delta v_H - \delta^{\ell+1}(1 - \mu)^\ell \Pi_{1s}(\mu_1, \delta)$, and, every other price is accepted with probability $\lambda(p)$, which is a left-continuous function such that $\delta v_H - p = \delta^{\ell+1}(1 - \mu)^\ell \Pi_{1s}(\mu'_1, \delta)$, where $\mu'_1 = (1 - \lambda(p))\mu_1 / (1 - \lambda(p)\mu_1)$. If $\ell = 0$, the seller makes an offer to buyer 1 following the strategy defined in Hart (1989). If $\ell > 0$, the seller asks δv_H to the highest index player $i \neq 1$ with $\mu_i = \mu$.
5. Beliefs are updated according to Bayes' rule whenever possible. Otherwise, we assume that beliefs are passive.

Proof of Proposition A. We first note that given the strategy profile specified in part 1-4, and given that beliefs are updated according to part 5, all equilibrium histories

belong to \hat{H} . Furthermore, starting from $h \in \hat{H}$, if a buyer deviates by accepting an offer, the game ends. If, starting from $h \in \hat{H}$, a buyer deviates by rejecting an offer, given the above strategy we will enter a new history $h' \in \hat{H}$. Finally, if starting from $h \in \hat{H}$ the seller deviates, part 5 implies that we remain in \hat{H} . Hence, we can focus, without loss of generality, on the set \hat{H} . We now prove each part of the proposition maintaining the assumption that we are in history $h^\ell \in \hat{H}$ with $\ell \geq 0$.

Part 1. It is immediate, given seller's strategy defined in parts 2-4, to verify the optimality of the strategy of buyers $i \neq 1$ and the optimality of low-value buyer 1's strategy, both defined in part 1.

Part 2. It is also immediate to check the optimality of high-value buyer 1's strategy in part 2. We now establish the optimality of seller's behaviour. Since $\delta R_1 \leq v_L$ and since high-value buyers accept every offer up to δv_H , the seller either asks δv_H or consumes at h^ℓ , for $\ell \geq 0$. Next, when $\ell = 0$, making an offer at δv_H to buyer 1 is optimal if and only if $v_L < \delta[v_H\mu_1 + (1 - \mu_1)v_L]$. Suppose we are at h^ℓ with $\ell > 0$. If the seller consumes she gets v_L . If it is optimal to make an offer δv_H , since $\mu_1 \leq \mu$, the seller must make an offer to buyer $i \neq 1$ with $\mu_i = \mu$. Moreover, if this offer is rejected, then we enter $h^{\ell-1}$. If $\ell - 1 > 0$, because of stationarity, then it will still be optimal to make an offer at δv_H to a buyer $i \neq 1$, with $\mu_i = \mu$. So, if the seller does not consume, her expected payoff is

$$P(\ell) := v_H\delta\mu \sum_{j=0}^{\ell-1} (1-\mu)^j \delta^j + (1-\mu)^\ell \delta^\ell \max\{v_L, \delta[v_H\mu_1 + (1-\mu_1)v_L]\}.$$

Note that $P(1) = \delta[\mu v_H + (1-\mu) \max\{v_L, \delta[v_H\mu_1 + (1-\mu_1)v_L]\}]$ and that $P(\ell) = \delta[\mu v_H + (1-\mu)P(\ell-1)]$ for all $\ell \geq 1$. Furthermore, if $v_L \geq \delta[\mu v_H + (1-\mu)v_L]$, then, because $\mu_1 \leq \mu$, we have that $v_L \geq P(1)$ and that $P(\ell) \leq \delta v_L$. Hence, it is a best reply for the seller to consume. If, on the other hand, $v_L < \delta[\mu v_H + (1-\mu)v_L]$, then $P(\ell) > \delta v_L$ and it is optimal for the seller to ask δv_H to the highest index i with $\mu_i = \mu$.

Part 3. The proof of part 3 follows the same logic of the proof of part 2, and therefore is omitted.

Part 4. Consider that high-value buyer 1 receives an offer. If she accepts, she gets $\delta v_H - p$. If buyer 1 rejects it, we enter history $h^\ell \in \hat{H}$ with $\mu'_1 \in [0, \mu_1]$ and the seller asks, in sequence starting with the highest index, δv_H to all buyers $i \neq 1$ with $\mu_i = \mu$. From part 1, only high-value buyers accept, and, upon accepting they consume. If all these offers are rejected, we enter history h^0 , and the equilibrium payoff of 1 is $\Pi_{1s}(\mu'_1, \delta)$. Hence, the continuation payoff upon rejection for buyer 1 is $(1-\mu)^\ell \delta^{\ell+1} \Pi_{1s}(\mu'_1, \delta)$. Given buyer 1's

strategy, if high-value buyer 1 rejects $p \leq \delta v_H - \delta^{\ell+1}(1-\mu)^\ell \Pi_{1s}(0, \delta)$, then, by part 5, $\mu'_1 = 0$ and therefore it holds that accepting such an offer p is a best reply. Analogously, if player 1 rejects $p \geq \delta v_H - \delta^{\ell+1}(1-\mu)^\ell \Pi_{1s}(\mu_1, \delta)$, then, by part 5, $\mu'_1 = \mu_1$, and therefore it holds that rejecting such an offer p is a best reply. For all other offers, Hart (1989) assures that there exists a left-continuous function in p , $\lambda(p)$, such that if player 1 accepts p with probability $\lambda(p)$, player 1 is indifferent between accepting and rejecting.²⁰

We now show that the prescribed strategy is optimal for the seller. We focus on the case where $\ell = 1$ and call 2 the player $i \neq 1$, with belief μ . The seller expected payoff by following the strategy is

$$V^s := \delta[v_H \mu + (1-\mu)\Pi_{s1}(\mu, \delta)]. \quad (\text{A.1})$$

The seller does not wish to deviate by asking a lower price to 2. So, suppose the seller makes an offer to 1 at p , and, if rejected, follows the equilibrium strategy. Note that we use the one-shot deviation principle, hence we look at deviations followed by equilibrium play. We call $p_{\mu_2}(\mu')$ the highest price that, if rejected by buyer 1, induces belief $\mu' \leq \mu$ about player 1, when the prior about buyer 2 is μ_2 . We say that μ' is implementable from μ if there exists such a price $p_\mu(\mu')$. The fact that the acceptance strategy of high-value buyer 1 is left continuous implies that there is a finite set of μ' that are implementable. The payoff of the seller from a deviation that asks the price $p_\mu(\mu')$ is

$$V_d^s := p_\mu(\mu') \frac{\mu - \mu'}{1 - \mu'} + \frac{1 - \mu}{1 - \mu'} \delta^2 [v_H \mu + (1 - \mu) \Pi_{s1}(\mu', \delta)], \quad (\text{A.2})$$

where $\frac{\mu - \mu'}{1 - \mu'}$ is the unconditional probability of buyer 1 accepting the price $p_\mu(\mu')$.

Next, given high-value buyer 1's strategy, the set of beliefs μ' that are implementable from μ does not depend on $\mu_2 \in \{0, \mu\}$. Hence, since in the history in which $\mu_2 = 0$ and $\mu_1 = \mu$, we have postulated that the seller and buyer plays according to the equilibrium of Hart (1989), the seller's pricing must be optimal, which implies that

$$\Pi_{s1}(\mu, \delta) \geq p_0(\mu') \frac{\mu - \mu'}{1 - \mu'} + \frac{1 - \mu}{1 - \mu'} \delta \Pi_{s1}(\mu', \delta). \quad (\text{A.3})$$

20. This follows from the result in Hart (1989) that the high-value buyer 1's equilibrium payoff correspondence as μ_1 varies has a closed graph, it is weakly decreasing and single valued except at a finite number of points.

This allows us to relate the seller's payoff by following the strategy, expression A.1, to her deviating payoff, expression A.2, as follows:

$$\begin{aligned} V^s &= \delta[v_H\mu + (1-\mu)\Pi_{s1}(\mu, \delta)] \geq \delta v_H\mu + (1-\mu)\delta \left[p_0(\mu') \frac{\mu - \mu'}{1 - \mu'} + \frac{1-\mu}{1-\mu'} \delta \Pi_{s1}(\mu', \delta) \right] \\ &= \mu\delta v_H \left[1 - \frac{1-\mu}{1-\mu'} \delta \right] + (1-\mu)\delta p_0(\mu') \frac{\mu - \mu'}{1 - \mu'} - p_\mu(\mu') \frac{\mu - \mu'}{1 - \mu'} + V_d^s \end{aligned}$$

where the first inequality is obtained by replacing $\Pi_{s1}(\mu, \delta)$ with its lower bound defined in inequality (A.3), and the second equality follows by using equation (A.2) that relates $\Pi_{s1}(\mu', \delta)$ with V_d^s . To show that the strategy of the seller is a best reply, i.e., $V^s \geq V_d^s$, it is, then, sufficient to show that

$$V^s - V_d^s = \mu\delta v_H \left[1 - \frac{1-\mu}{1-\mu'} \delta \right] + (1-\mu)\delta p_0(\mu') \frac{\mu - \mu'}{1 - \mu'} - p_\mu(\mu') \frac{\mu - \mu'}{1 - \mu'} \geq 0 \quad (\text{A.4})$$

Finally, note that $p_\mu(\mu')$ makes high-value buyer 1 indifferent between accepting $p_\mu(\mu')$ or rejecting it. The indifferent condition is therefore $\delta v_H - p_\mu(\mu') = (1-\mu)\delta^2 \Pi_{1s}(\mu', \delta)$. Analogously, $p_0(\mu')$ must solve $\delta v_H - p_0(\mu') = \delta \Pi_{1s}(\mu', \delta)$. Substituting the expressions for $p_\mu(\mu')$ and $p_0(\mu')$, and after some algebra, condition (A.4) becomes $\delta v_H(1-\delta)\mu'(1-\mu)/(1-\mu') \geq 0$, which is always satisfied. ■

Proof of Proposition 3. For a seller i , let $b^* = \arg \max_{b \in N_i} R_b$. We first show that there exists $\bar{\delta} < 1$ such that for all $\delta \in [\bar{\delta}, 1)$ there is an equilibrium in which the resale values are

$$R_i = \begin{cases} v_L, & \text{if } i \in L_0, \\ \delta v_H \mu \sum_{k=1}^{|N_i|} (1-\mu)^{k-1} \delta^{k-1} + (1-\mu)^{|N_i|} \delta^{|N_i|} v_L, & \text{if } i \in L_1, \\ \delta v_H \mu \sum_{k=1}^{|N_i|-1} (1-\mu)^{k-1} \delta^{k-1} + (1-\mu)^{|N_i|-1} \delta^{|N_i|-1} \Pi_{ib^*} & \text{if } i \in L_k, k \geq 2. \end{cases} \quad (\text{A.5})$$

First select $i \in L_1$ and consider the auxiliary game in which $i \in L_1$ is the seller. Note that, for each $j \in N_i$, $R_j = v_L$, and therefore part 2 of Proposition A applies. Note also that, when δ is sufficiently high, $v_L < \delta[\mu v_H + (1-\mu)v_L]$. Hence, i makes a sequence of offers at δv_H to all her buyers, each of these offers is accepted with probability μ , and if all the offers are rejected, i consumes. This gives an expected payoff to i which equals expression A.5. We can use the same logic to construct the resale value of $i \in L_k$, for all $k > 1$.

From Hart (1989) we know that $\Pi_{ib^*}(\mu, \delta)$ is continuous in δ and, as δ goes to one, the expected payoff of seller i bargaining with buyer b^* converges to the resale value of

b^* , i.e., $\lim_{\delta \rightarrow 1} \Pi_{ib^*}(\mu, \delta) = R_{b^*}$. It follows that R_i is continuous in δ for each $i \in L_k$, $k = 0 \dots k^*$ and that, as δ goes to one, resales converges to

$$R_i = \begin{cases} v_L, & \text{if } i \in L_0 \\ v_H - (1 - \mu)^{|N_i|} [v_h - v_L], & \text{if } i \in L_1 \\ v_H - (1 - \mu)^{|N_i|-1} [v_H - \max_j R_j] & \text{if } i \in L_k, k \geq 2 \end{cases} \quad (\text{A.6})$$

When $\delta = 1$, one can recursively construct the equilibrium path by selecting, for each $i \in L_k$, $k = 1 \dots k^*$, the trading partner with the highest resale, who must be on the path (i_0^*, \dots, i_m^*) that maximizes κ . Since it is unique (by assumption) and since resale values are continuous in δ , (i_0^*, \dots, i_m^*) is indeed the trading path for sufficiently high δ . The rest of the proposition follows by simple algebra. \blacksquare

B. OTHER PROOFS

Proof of Proposition 1.

First part. We first note that, at each round t , the total t -period payoff (i.e., sum of payoffs over all players) is v_i if there is consumption by player i and 0, otherwise. Consequently, regardless of the realisation of types, the maximal discounted sum of total payoffs from period $t+1$ onwards is $\delta^{t+1}v_H$ (i.e., when consumption takes place at period $t+1$).

Next, assume, in contradiction with the proposition, that there exists an equilibrium where player i owns the object at period t , his type is v_H , and she does not consume. To be an equilibrium, it must be that player i 's equilibrium payoff from period t onwards is at least $\delta^t v_H$, since she can consume the object at period t . However, given that the total payoff at period t is 0 (since player i does not consume the object) and the maximal discounted sum of total payoffs from period $t+1$ onwards is $\delta^{t+1}v_H$, it must be that at least one type of some player expects a strictly negative equilibrium payoff from period t onwards. This is impossible, as each type of each player can guarantee herself a payoff of at least 0 by rejecting all offers.

Second part. In what follows, for a given equilibrium, $V_i^H(h)$ (or $V_i^L(h)$) represents the continuation payoff of player i with value v_H (or v_L) at history h . Furthermore, we denote by $\mu_i(h)$ the (common) probability that i has high-value at (public) history h ; $\mu(h) = (\mu_1(h), \dots, \mu_n(h))$ is the (common) profile of beliefs.

We show that in a wMPBE, at any history h where s is a seller, s never makes an offer (i, p) , to some $i \in N_s$, such that $p < V_i^L(h \oplus (i, p) \oplus a)$. The following lemma is instrumental in proving this result.

Lemma 3. *In a wMPBE, if trader i is not the current owner at h and $\mu_i(h) = 1$, then $V_i^L(h) = V_i^H(h) = 0$.*

Proof of Lemma 3. We show that no seller will ask to i a price strictly below δv_H starting from h . Take a seller starting from h and suppose that she has a high-value. In light of the first part of Proposition 1, the seller can obtain δv_H by consuming in the next round and therefore she will make no offer below δv_H (unless such offer was refused with probability one, which can't happen in a wMPBE).

Next, we consider low-value sellers. Let \underline{p} be the infimum price that is ever offered to i from h onward at both on-path and off-path histories. We show that $\underline{p} = \delta v_H$.

By way of contradiction, suppose that $\underline{p} < \delta v_H$. By definition of \underline{p} , there must exist an history h' such that i is offered $p' \in [\underline{p}, \delta v_H - \delta(\delta v_H - \underline{p})]$. Note that $\delta v_H - \delta(\delta v_H - \underline{p})$ is the price that makes high-value player i indifferent between accepting that price and consume tomorrow, or rejecting and accept \underline{p} tomorrow and then consume. By construction of $\delta v_H - \delta(\delta v_H - \underline{p})$, it is a strict best response of high-value player i to accept the offer at p' . Therefore the seller that makes offer (i, p') obtains an expected payoff of p' because she anticipates that the offer is accepted with probability one, and that i consumes the object; recall that $\mu_i(h) = 1$ and therefore $\mu_i(h') = 1$ because degenerate beliefs are never updated. But note that the seller can strictly increase her profits by making offer $(i, p' + \varepsilon)$, with $p' + \varepsilon < \delta(\delta v_H - \underline{p})$; in fact, player i has also a strict incentive to accept this offer. ■

For a player $i \in N$ we define $H(i)$ as the set of histories, including off-path histories, such that: (a) player i receives an offer at $h \in H(i)$ and (b) no prior (observable) deviation has taken place at the consumption stage. It is important to note that at any history $h \in H(i)$, the seller that makes the offer to i , say s , has $\mu_s(h) = 0$. This holds by combining: part (b) in the definition of $H(i)$, high-value owners are always expected to consume (from part 1 of Proposition 1), and degenerate beliefs are never updated.

Let $p(h)$ be the price asked in h , and let $\bar{u} = \sup_{h \in H(i)} \{V_i^L(h \oplus (i, p(h)) \oplus a) - p(h)\}$ ■ be the maximal payoff attainable by the low-value player i from accepting an offer. We

show that $\bar{u} = 0$. By way of contradiction, assume that $\bar{u} > 0$. Note that the maximal payoff attainable by i by refusing some offer is at most $\delta\bar{u}$. Second, observe that by definition of \bar{u} as a supremum, there exists $\bar{h} \in H(i)$ and offer $(i, p(\bar{h}))$ such that $\delta V_i^L(\bar{h} \oplus (i, p(\bar{h})) \oplus a) - p(\bar{h}) \in (\delta\bar{u}, \bar{u})$.

We next argue that $V_i^L(\bar{h} \oplus (i, p) \oplus a) = V_i^L(\bar{h} \oplus (i, p(\bar{h})) \oplus a)$ for any p . This follows because the continuation payoff of the low-value i in $\bar{h} \oplus (i, p) \oplus a$ only depends on the beliefs over traders different from i . In fact, if she consumes, her payoff is v_L . If she does not consume, the belief of the other traders with regard to i becomes degenerate (because of part 1 Proposition 1) and, therefore, it stays constant in any future on-path history. Furthermore, since $\bar{h} \in H(i)$, we have that $\mu_s(\bar{h}) = 0$, and therefore this belief is not affected by the seller posting a different price.

Hence, there exists $p' \in (p(\bar{h}), V_i^L(\bar{h} \oplus (i, p(\bar{h})) \oplus a) - \delta\bar{u})$, which is accepted with probability one by the low-value player i , as it provides a payoff larger than $\delta\bar{u}$. It follows that $p(\bar{h})$ and p' are also accepted by the high-value player i with probability one, as refusal would imply $\mu_i(\bar{h} \oplus (i, p(\bar{h})) \oplus na) = 1$ and provide zero payoff to i in light of Lemma 3.

But then the seller can strictly increase her payoff at \bar{h} if she makes offer (i, p') , instead of the postulated optimal offer $(i, p(\bar{h}))$. This is the case because both p' and $p(\bar{h})$ are accepted with probability one by player i , and the continuation payoff of the seller, following acceptance of $(i, p(\bar{h}))$ or following acceptance of (i, p') must be the same. In fact, in a wMPBE the continuation payoff of the seller following an offer to i which is accepted can only depend on i (i.e., the new seller) and on the state of beliefs at the moment in which i acquires the object. We have already argued that the belief about the seller cannot change following the two different offers because $\mu_s(\bar{h}) = 0$. The belief about i is identical in both cases because i accepts with probability one both offers. ■

Proof of Proposition 5

First part. The proof of this part follows closely the proof for an analogous statement in Fudenberg et al. (1985), Lemma 3. For a related proof in the case of a seller bargaining with multiple buyers, see Lemma 2 in De Fraja and Muthoo (2000).

By way of contradiction, assume that there is a set H^∞ of infinite public histories that has strictly positive probability in the support of the equilibrium.

First, note that if at some generic beginning-of-period history h we have $\mu_i(h) = 0$

$\forall i \in N$, then the game ends at h with probability one. Consumption takes place at h since the maximum attainable continuation payoff for any type in any equilibrium in which there is no consumption at h is equal to δv_L .

Then, for any $h \in H^\infty$, define $\{h^0(h), h^1(h), \dots\}$ as the beginning-of-period sub-histories of h (i.e., $h^0(h) = \emptyset$ and $h^t(h)$ is an history which includes all events that took place in h from period 0 up until and including the acceptance decision in round $t - 1$). The skimming property implies that $\mu_i(h^t(h))$ is weakly decreasing in t . By the argument above, and remembering that degenerate beliefs never change, for any $h \in H^\infty$, there exists (maximal) non-empty $Z(h) \subset N$ such that $\forall i \in Z(h)$ and $t \geq 1$, we have $\mu_i(h^t(h)) > 0$.

Fix $h \in H^\infty$ and $i \in Z(h)$, then $\forall (\epsilon, k)$ with $k \geq 1$ and $\epsilon > 0$ there exists $t_i^h(\epsilon, k)$ such that $\forall t > t_i^h(\epsilon, k)$, $\mu_i(h^{t+k}(h)) - \mu_i(h^t(h)) < \epsilon$. If this was not the case, there would exist a time period T_i^* such that $\mu_i(h^{T_i^*}) = 0$ (for details see De Fraja and Muthoo (2000), page 863). Then, let $t^*(\epsilon, k) = \sup_{h \in H^\infty} \sup_{i \in Z(h)} \{t_i^h(\epsilon, k), \tilde{t}^h\}$, where \tilde{t}^h is the period such that, for all $t \geq \tilde{t}^h$ we have $\mu_j(h^t(h)) = 0$ for all $j \neq Z(h)$. If the supremum is not finite, we select $t^*(\epsilon, k)$ arbitrarily large, in such a way that the effect on expected payoffs of histories that exceed the bound is made negligible. Without loss of generality assume that there are only infinite histories lasting beyond $t^*(\epsilon, k)$.

Next, let $Z = \cup_{h \in H^\infty} Z(h)$ and suppose that all offers to $i \in Z$ between $t > t^*(\epsilon, k)$ and $t + k$ along h come at prices above the resale value of the low-value. (Otherwise by the skimming property high-value player i would accept with probability one and consume, contrary to the stated assumption). Note that then the probability with which the high-value i accepts an offer between k and $t + k$ along $h \in H^\infty$ is less than or equal to $\mu_i(h^{t+k}(h)) - \mu_i(h^t(h))$. This follows from Bayesian updating given that the offers must come at price that the low-value refuses with probability one.

We are now ready to conclude the proof by contradiction. We now argue that for any $h \in H^\infty$ there exists a period at which the owner of the good consumes (i.e., consumption almost surely takes place, hence the set H^∞ has zero probability). In fact, suppose that at some node $h^{t^*(\epsilon, k)}(h)$ in history $h \in H^\infty$, the seller's continuation payoff is bounded from above by $\delta v_H(1 - (1 - \epsilon)^z) + \delta^k v_H$ where z is the numbers of traders in Z . Then for all ϵ small and k large enough this continuation would be strictly lower than v_L which is the payoff from consuming. (Given that there is discounting and equilibrium per-period payoffs are bounded, this follows even in case $t^*(\epsilon, k)$ has been selected arbitrarily large).

A contradiction to the hypothesis that the game continues beyond $t^*(\widehat{\epsilon}, \widehat{k})(h)$.

To conclude the proof we now verify that the payoff of the seller is bounded above by $\delta v_H(1 - (1 - \epsilon)^z) + \delta^k v_H$. To see this observe that: (i) any offer that high-value traders accept with some probability, between between $t^*(\epsilon, k)$ and $t^*(\epsilon, k) + k$, regardless of the history, comes at a price at most equal to δv_H ; (ii) the probability that such offers are accepted is at most ϵ for each $i \in Z$ between $t^*(\epsilon, k)$ and $t^*(\epsilon, k) + k$, regardless of the history; (iii) offers to a low-value traders, if any, come to traders with low-value with probability one and at a price equal to the continuation payoff of such traders, which can only make the bound less tighter.

Second part. First recall that, along the path, a seller never makes an offer that is rejected with probability one. Next, note that part 2 of Proposition 1 implies that if the seller makes an offer to i , then she will ask either her discounted resale value, or a strictly higher price. If she asks the discounted resale value, then, by definition, low value trader i is indifferent between accepting and rejecting the offer. If the seller asks a price strictly higher than i 's discounted resale, then the low-value trader i rejects it with probability one. Since the seller never makes offers that are rejected with probability one, the high-value trader i must accept the offer with strictly positive probability. If the high-value trader accepts with probability one, then the indifference property of regular equilibria guarantees the result. If the high-value buyer mixes between acceptance and rejection, then the result is true in light of the indifference condition for mixing. ■

Proof of Proposition 6. Part 1. Part (1) follows from part (2) and equation (4.1) from the main text.

Part 2. The proof is by induction. We discuss the base case and then move to the induction step. We focus on histories that terminate with consumption in finite time.

As our induction *base case*, we show that there exists t^* such that the statement is true for any offer made from round t^* onward. By part 1 of Proposition 5 there exists a T^* that is the maximum round in which consumption takes place. Let $t^* + 1 < T^*$ be the latest round in which a resale offer is made in equilibrium. Let r^* be the price asked in the unique resale offer made in that round. (The argument is analogous if (i) the seller randomises among different resale offers in $t^* + 1$ or (ii) there are multiple on-path histories on which a resale offer is made in the $t^* + 1$ period.) We show that, along the equilibrium path, for all consumption offers the price p^* asked in round t^* must

be strictly greater than r^* . Assuming that p^* is followed by r^* with probability one (the conclusion holds a fortiori otherwise), there are two possibilities. First, the offer in t^* goes to a player j which is different than the player that receives offer at price r^* ; in this case, $p^* = \delta v_H$ because player j is indifferent between accepting and rejecting and there are no further resale offers after $t^* + 1$; therefore $p^* = \delta v_H > r^* = \delta R^*$, because $v_H > R^*$. Second, the offer in t^* goes to the same player that receives the offer at price r^* ; in this case, since p^* is a consumption offer, by part 2 of Proposition 5 the indifference condition for the high-value dictates that $\delta v_H - p^* = \delta(\delta v_H - r^*)$. To see this point observe that, by the skimming property, an high-value player must accept a resale offer. We conclude that $p^* = \delta v_H(1 - \delta) + \delta r^* > r^*$, given that $\delta v_H > r^* = \delta R^*$.

The *induction step*, consists in showing that the statement is valid for any consumption offer (i, p) in round t made by some seller s after some history h^t , given that it is true for any consumption offer made from round $t + 1$ onward. Note that since (i, p) is a consumption offer, consumption by the high-value follows acceptance. Henceforth, to simplify the exposition, we focus on the case in which no other consumption offer is made to i following a rejection of offer (i, p) and consumption offers are accepted with probability one, but with extra notational burden the proof extends to this case as well.

The following notation is used in the rest of the proof. Starting with the rejection of offer (i, p) made by seller s and, restricting attention to the equilibrium path, consider all the different sequences of offers leading to a resale offer being made by s . In those cases in which there are on-path terminal histories where s makes more than one resale offer, then we construct a single sequence, truncating it at the first resale offer that s makes. The set of all such sequences is denoted Σ . For a given $\sigma \in \Sigma$, call $r(\sigma)$ the price quoted in the last offer of σ (which is a resale offer), $t(\sigma)$ the length of the sequence, and call $1 - \beta(\sigma)$ the probability that the resale offer $r(\sigma)$ occurs in equilibrium, evaluated conditional on refusal of offer (i, p) .²¹

Analogously, we denote by Σ_i the set of sequences of offers starting with the rejection of offer (i, p) and terminating with a resale offer to i . In those cases in which there are

21. To fix ideas, suppose that after (i, p) is refused, s consumes with probability x_1 and otherwise makes an offer (j, p) ; that such offer is accepted with probability x_2 and otherwise refused; that in case of refusal s makes offer (j', p') with probability x_3 and offer (j'', p'') with the remaining probability, and that both of these offers are resale offers, both accepted with probability one. Then $\Sigma = \{\sigma_1, \sigma_2\}$ where $\sigma_1 = \{(j, p), (j', p')\}$ and $\sigma_2 = \{(j, p), (j'', p'')\}$, $t(\sigma_1) = t(\sigma_2) = 2$, and $1 - \beta(\sigma_1) = (1 - x_1)(1 - x_2)x_3$, $1 - \beta(\sigma_2) = (1 - x_1)(1 - x_2)(1 - x_3)$.

on-path terminal histories where i receives more than one resale offer, then we construct the set Σ_i truncating each sequence at the first resale offer to i . For any given $\sigma \in \Sigma_i$ we call $r^i(\sigma)$ the price of the resale offer to i that ends the sequence, we call $t(\sigma)$ the length of the sequence, and we call $(1 - \beta(\sigma))$ the probability that the offer that ends the sequence occurs according to the equilibrium, evaluated after the refusal of (i, p) . Observe that both Σ and Σ_i are either empty or well defined, given that the game ends in finite time.

We can begin the analysis, by noting that if Σ is empty then the statement is trivially true. Similarly, if Σ_i is empty, then trader i , upon rejecting (i, p) , does not receive any further offer and obtains a continuation of 0. The indifference condition with respect to consumption offer (i, p) implies, then, that $p = \delta v_H$. Since $\delta v_H > r(\sigma)$ for all $\sigma \in \Sigma$, we obtain that $p > r(\sigma)$ for all $\sigma \in \Sigma$.

For the case in which Σ_i is non empty, the induction argument is concluded by means of the two following lemmas. Lemma 4 provides a lower bound of the price p associated to consumption offer (i, p) . This lower bound is obtained by using the indifference property of regular equilibria. Lemma 5 provides an upper bound of the first resale offer following refusal of (i, p) . This uses our induction hypothesis.

Lemma 4. *There exists $\hat{\sigma} \in \Sigma_i$ such that $p > \delta v_H \beta(\hat{\sigma}) + (1 - \beta(\hat{\sigma}))r^i(\hat{\sigma})$*

Proof of Lemma 4. In a regular equilibrium, the high-value trader i is indifferent between accepting and rejecting the consumption offer (i, p) . Therefore, recalling that the high-value i must accept any resale offer with probability one and that his payoff is zero whenever she does not receive a resale offer (as in that case he only obtains offers at δv_H), we have

$$\delta v_H - p = \int_{\Sigma_i} (1 - \beta(\sigma)) \delta^{t(\sigma)} (\delta v_H - r^i(\sigma)) d\sigma.$$

We conclude that there exists $\hat{\sigma} \in \Sigma_i$ such that:

$$\delta v_H - p \leq \delta(1 - \beta(\hat{\sigma})) \delta^{t(\hat{\sigma})} (\delta v_H - r^i(\hat{\sigma})). \quad (\text{B.1})$$

and therefore

$$\begin{aligned} p &\geq \delta v_H - \delta(1 - \beta(\hat{\sigma})) \delta^{t(\hat{\sigma})} (\delta v_H - r^i(\hat{\sigma})) \\ &> \delta v_H - (1 - \beta(\hat{\sigma})) (\delta v_H - r^i(\hat{\sigma})). \end{aligned}$$

where the first inequality is a rewriting of inequality (B.1), and the second inequality follows by replacing $\delta^{t(\hat{\sigma})} < 1$ with 1. ■

Lemma 5. *For all $\sigma \in \Sigma$, we have $r(\sigma) \leq \delta v_H \beta(\hat{\sigma}) + (1 - \beta(\hat{\sigma}))r^i(\hat{\sigma})$, where $\hat{\sigma} \in \Sigma_i$ is derived from Lemma 1.*

Proof of Lemma 5. For any $\sigma \in \Sigma$, let $V_s^L(\sigma)$ be the continuation payoff of the seller following rejection of (i, p) , computed assuming that when each of her offer is refused she makes offers according to σ .²²

It is immediate to observe that $V_s^L(\sigma) = V_s^L(\sigma')$ for any $\sigma, \sigma' \in \Sigma$ (recalling that Σ only contains sequences of on-path offers). In fact, the low-value seller s expects no further surplus in the continuation game ensuing after any of her offer is accepted and both $r(\sigma)$ and $r(\sigma')$ are accepted with probability one. Hence, if, by contradiction, $V_s^L(\sigma) > V_s^L(\sigma')$, the seller would strictly prefer to implement the sequence of offers σ , and therefore σ' could not be part of the equilibrium.

Next, by the induction hypothesis, every price that s offers, following the rejection of offer (i, p) until her resale offer $r(\sigma)$, is a consumption offer and is strictly larger than $r(\sigma)$. Therefore, $V_s^L(\sigma) \geq r(\sigma)$.

We can then conclude that for all $\sigma \in \Sigma$ we have

$$r(\sigma) \leq V_s^L(\sigma) = V_s^L(\hat{\sigma}) \leq \delta v_H \beta(\hat{\sigma}) + (1 - \beta(\hat{\sigma}))r^i(\hat{\sigma}),$$

where the last inequality follows because we have computed the upper bound of $V_s^L(\hat{\sigma})$ by assuming that if acceptance of an offer (or consumption by s) takes place along $\hat{\sigma}$ then the highest possible and undiscounted surplus is collected by s . ■

The two lemmas imply that $p > r(\sigma)$ for every $\sigma \in \Sigma$ and conclude the induction step and the proof of part (2) of the Proposition. ■

Part 3. We maintain the definition of Σ_i which has been introduced above, with the proviso that the set is now defined at the empty history (beginning of the game) with an initial seller denoted s with $s = s^0$. For $\sigma \in \Sigma_i$, the definitions of $r^i(\sigma)$, $(1 - \beta(\sigma))$ and $t(\sigma)$ are as above.

22. To pursue the example in the previous footnote, $V_s^L(\sigma_1) = x_1 v_s + (1 - x_1)x_2 p + (1 - x_1)(1 - x_2)p'$ and $V_s^L(\sigma_2) = x_1 v_s + (1 - x_1)x_2 p + (1 - x_1)(1 - x_2)p''$.

In addition, we need the following notation. For each $\sigma \in \Sigma_i$, define the set $\Gamma_j(\sigma)$ to include all the sequences of offers that start after the resale offer $(i, r^i(\sigma))$ and that end with a resale offer to player j (again, focus on the first resale offer to j). Note that $\Gamma_j(\sigma)$ may be empty for some $\sigma \in \Sigma_i$. For any $\gamma \in \Gamma_j(\sigma)$ call $r_\sigma^j(\gamma)$ the resale offer received by j at the end of the sequence γ , let $t_\sigma(\gamma)$ be the length of the sequence and $1 - \beta_\sigma(\gamma)$ be the probability, evaluated after the acceptance of $r^i(\sigma)$, that the game reaches offer $r_\sigma^j(\gamma)$.

We can now write the interim utility of the high-value i as

$$V_i^H(\emptyset) = \int_{\Sigma_i} (1 - \beta(\sigma)) \delta^{t(\sigma)} (\delta v_H - r^i(\sigma)) d\sigma. \quad (\text{B.2})$$

Similarly, noting that all offers received by j must originate through some sequence $\sigma \in \Sigma_i$, we can write:

$$V_j^H(\emptyset) = \int_{\Sigma_i} (1 - \beta(\sigma)) \delta^{t(\sigma)} \left(\int_{\Gamma(\sigma)} (1 - \beta_\sigma(\gamma)) \delta^{t_\sigma(\gamma)} (\delta v_H - r_\sigma^j(\gamma)) d\gamma \right) d\sigma. \quad (\text{B.3})$$

Note that we are implicitly ignoring consumption offers that i and j may receive along the various sequences of offers in virtue of the indifference property.

Next, we observe that, for any $\sigma \in \Sigma_i$ and $\gamma(\sigma) \in \Gamma_j(\sigma)$:

$$r^i(\sigma) < \beta_\sigma(\gamma) \delta v_H + (1 - \beta_\sigma(\gamma)) \delta^{t_\sigma(\gamma)} r_\sigma^j(\gamma). \quad (\text{B.4})$$

This inequality follows by exploiting two facts. First, every seller intermediating the object between the initial resale offer and the final resale offer must be indifferent between the different sequences of offers that she may implement by randomisation (see the discussion in the proof of proposition 6). Second, the right hand side of the above expression is (more than) the total expected surplus that can be generated along the sequence $\gamma(\sigma)$. The strictness of the inequality is guaranteed by the fact that when i receives her first resale offer she will accept and consume, which happens with positive probability. (The proposition will hold a fortiori in the case in which the high-value i has accepted a previous consumption offer with probability one.)

For every $\sigma \in \Sigma_i$ and any $\gamma \in \Gamma_j(\sigma)$, using inequality (B.4) we can write:

$$\begin{aligned} \delta v_H - r^i(\sigma) &> (1 - \beta_\sigma(\gamma)) (\delta v_H - \delta^{t_\sigma(\gamma)} r_\sigma^j(\gamma)) \\ &> (1 - \beta_\sigma(\gamma)) \delta^{t_\sigma(\gamma)} (\delta v_H - r_\sigma^j(\gamma)), \end{aligned}$$

where the first inequality follows from subtracting δv_H to both sides of inequality (B.4), and the second by multiplying δv_H by the factor $\delta^{t_\sigma(\gamma)} < 1$.

Using expression (B.2) for $V_i^H(\emptyset)$, expression (B.3) for $V_j^H(\emptyset)$ and $\delta v_H - r^i(\sigma)$ as an upper bound, we obtain the desired conclusion

$$\begin{aligned} V_j^H(\emptyset) &= \int_{\Sigma_i} (1 - \beta(\sigma)) \delta^{t(\sigma)} \left(\int_{\Gamma(\sigma)} (1 - \beta_\sigma(\gamma)) \delta^{t_\sigma(\gamma)} (\delta v_H - r_\sigma^j(\gamma)) d\gamma \right) d\sigma \\ &< \int_{\Sigma_i} (1 - \beta(\sigma)) \delta^{t(\sigma)} \left(\int_{\Gamma(\sigma)} (\delta v_H - r^i(\sigma)) d\gamma \right) d\sigma \leq V_i^H(\emptyset). \end{aligned}$$

■

Acknowledgements. We thank an Editor and four Referees for their comments. We thank Yeon-Koo Che, Sanjeev Goyal, Philippe Jehiel, David Levine, Mihai Manea, Claudio Mezzetti, Alessandro Pavan, Brian Rogers, Mark Sattthertwaite, Balazs Szentes, Rakesh Vohra and Asher Wolinsky for detailed comments. We also thank seminar audiences in the many departments and conferences where the paper has been presented. Daniele Condorelli and Andrea Galeotti thank the ESRC for financial support via grant RES-000-22-3902. Andrea Galeotti is grateful to the European Research Council for support through ERC-starting Grant Award (Number 283454) and to the Leverhulme Trust for support through the Philip Leverhulme Prize. Ludovic Renou acknowledges the financial support of the Agence Nationale pour la Recherche, ANR Cigne (ANR-15-CE38-0007-01).

Supplementary Data

Supplementary materials are available at *Review of Economic Studies* online.

REFERENCES

- Aker, J. (2010). Information from markets near and far: mobile phones and agricultural markets in Niger. *American Economic Journal: Applied Economics* 2(3), 46–59.
- Allen, F. and A. Babus (2009). Networks in Finance. In P. Kleindorfer and J. Wind (Eds.), *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World*, pp. 367. Wharton School Publishing.
- Allen, T. (2014). Information frictions in trade. *Econometrica* 82(6), 2041–83.
- Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002). Bargaining with incomplete information. In R. Aumann and S. Hart (Eds.), *Handbook of Game Theory with Economic Applications*, Volume 3, Chapter 50, pp. 1897–1945. Elsevier.
- Ausubel, L. M. and R. J. Deneckere (1989). Reputation in bargaining and durable goods monopoly. *Econometrica* 57(3), pp. 511–531.
- Blume, L., D. Easley, J. Kleinberg, and E. Tardos (2009). Trading networks with price-setting agents. *Games and Economic Behavior* 67(1), 36 – 50. Special Section of Games and Economic Behavior Dedicated to the 8th ACM Conference on Electronic Commerce.
- Calzolari, G. and A. Pavan (2006). Monopoly with resale. *RAND Journal of Economics* 37(2), 362–375.
- Chatterjee, K. and L. Samuelson (1988). Bargaining under two-sided incomplete information: The unrestricted offers case. *Operation Research* 36(4), 605–618.
- Choi, S., A. Galeotti, and S. Goyal (2016). Trading networks: Theory and experiment. *Journal of the European Economic Association*, Forthcoming.
- Condorelli, D. and A. Galeotti (2015). Strategic intermediation in networks. *Forthcoming, The Oxford Handbook on the Economics of Networks*. Eds: Y. Bramoulle, A. Galeotti, and B. Rogers..

- De Fraja, G. and A. Muthoo (2000, November). Equilibrium partner switching in a bargaining model with asymmetric information. *International Economic Review* 41(4), 849–69.
- Duffie, D., N. Garleanu, and L. H. Pedersen (2005, November). Over-the-counter markets. *Econometrica* 73(6), 1815–1847.
- Fafchamps, M. and B. Minten (1999). Relationships and traders in madagascar. *The Journal of Development Studies* 35(6), 1–35.
- Fafchamps, M. and B. Minten (2001, January). Property rights in a flea market economy. *Economic Development and Cultural Change* 49(2), 229–67.
- Fudenberg, D., D. K. Levine, and J. Tirole (1985). *Bargaining with Incomplete Information*, Chapter Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information, pp. 73–98. Cambridge, UK and New York: Cambridge University Press.
- Fudenberg, D., D. K. Levine, and J. Tirole (1987). Incomplete-information bargaining with outside opportunities. *Quarterly Journal of Economics* 52(1), 37–50.
- Fudenberg, D. and J. Tirole (1983). Sequential bargaining with incomplete information. *The Review of Economic Studies* 50(2), 221–247.
- Fudenberg, D. and J. Tirole (1991). *Game Theory*. MIT Press.
- Gale, D. M. and S. Kariv (2007, May). Financial networks. *American Economic Review* 97(2), 99–103.
- Gale, D. M. and S. Kariv (2009, August). Trading in networks: A normal form game experiment. *American Economic Journal: Microeconomics* 1(2), 114–32.
- Garratt, R. and T. Tröger (2006). Speculation in standard auctions with resale. *Econometrica* 74(3), 753–769.
- Gofman, M. (2011). A network-based analysis of over-the-counter markets. *Mimeo*.
- Golosov, M., G. Lorenzoni, and A. Tsyvinski (2014). Decentralized trading with private information. *Econometrica* 82(3), 1055–1091.
- Gul, F., H. Sonnenschein, and R. Wilson (1986, June). Foundations of dynamic monopoly and the coase conjecture. *Journal of Economic Theory* 39(1), 155–190.
- Hafalir, I. and V. Krishna (2008, March). Asymmetric auctions with resale. *American Economic Review* 98(1), 87–112.
- Haile, P. A. (2003, January). Auctions with private uncertainty and resale opportunities. *Journal of Economic Theory* 108(1), 72–110.
- Hart, O. (1989). Bargaining and strikes. *Quarterly Journal of Economics* 104(1), 25–44.
- Jehiel, P. and B. Moldovanu (1999, October). Resale markets and the assignment of property rights. *Review of Economic Studies* 66(4), 971–91.
- Jensen, R. (2007, August). The digital provide: information (technology), market performance, and welfare in the south indian fisheries sector. *Quarterly Journal of Economics* 122(3), 879–924.
- Kotowski, M. and M. Leister (2014). Trading networks and equilibrium intermediation. *Working Paper*.
- Li, D. and N. Schurhoff (2014). Dealer networks. *CEPR Discussion Paper No. DP10237*. SSRN: 2023201.
- Lyons, R. K. (1997). A simultaneous trade model of foreign exchange hot potato. *Journal of International Economics* 42(3–4).
- Manea, M. (2014). Intermediation in networks. *Working Paper*.
- Manea, M. (2015). Models of bilateral trade in networks. *Forthcoming, The Oxford Handbook on the Economics of Networks*. Eds: Y. Bramoulle, A. Galeotti, and B. Rogers.

- Nava, F. (2015). Efficiency in decentralized oligopoly markets. *Journal of Economic Theory* 157, 315–348.
- Renou, L. and T. Tomala (2012). Mechanism design and communication networks. *Theoretical Economics* 7(3), 489–533.
- Rubinstein, A. and A. Wolinsky (1985, September). Equilibrium in a market with sequential bargaining. *Econometrica* 53(5), 1133–50.
- Satterthwaite, M. and A. Shneyerov (2007). Dynamic matching, two-sided incomplete information, and participation costs: Existence and convergence to perfect competition. *Econometrica* 75(1), 155–200.
- Siedlarek, J.-P. (2015). Intermediation in networks. *Working Paper*.
- Svensson, J. and D. Yanagizawa (2009). Getting prices right: the impact of the market information service in uganda. *Journal of the European Economic Association* 7(2-3), 435–445.
- Zheng, C. Z. (2002, November). Optimal auction with resale. *Econometrica* 70(6), 2197–2224.