BOTTLENECK LINKS, ESSENTIAL INTERMEDIARIES AND COMPETING PATHS OF DIFFUSION IN NETWORKS

MIHAI MANEA

Department of Economics, Stony Brook University, mihai.manea@stonybrook.edu

Abstract. We investigate how information goods are priced and diffused over links in a network. A new equivalence relation between nodes captures the effects of network architecture and locations of sellers on the division of profits and characterizes the topology of competing (and potentially overlapping) diffusion paths. Sellers indirectly appropriate profits over intermediation chains from buyers in their equivalence classes. Links within the same class constitute bottlenecks for information diffusion and confer monopoly power. Links bridging distinct classes are redundant for diffusion and generate competition among sellers. In dense networks, competition limits the scope of indirect appropriability, and intellectual property rights foster innovation.

1. Introduction

Information, knowledge and other replicable goods are often traded over links in a network. Digital goods (e.g., software, music and audiobooks) are copied and shared between friends. Farmers reproduce high quality plant and animal breeds and sell them in local markets (Boldrin and Levine 2008). Apprentices are willing to accept low wages or even pay fees to learn a trade and then set up their own businesses in which they, in turn, train apprentices in exchange for cheap labor (Frazer 2006). Insider tips about corporate events that impact financial markets are sometimes transmitted over four or more links in networks formed by family and friends, and tipsters are rewarded for insider information with gifts and jobs (Ahern 2017). Chains of intermediaries analyze, package and distribute financial and industry-specific information tailored for business solutions, accounting purposes or reporting (Sarvary 2011). This paper studies how the locations of the initial sources of a replicable good in a network shape diffusion paths and determine the profits that players at different positions in the network obtain from consuming and reselling the good.

We consider a market with a network structure in which some players are endowed with an identical information good. Information is an indivisible consumption good for which players have unit demand and heterogeneous values. We assume that there are no consumption externalities and that the good does not depreciate. Our model builds on work by Polanski

Date: May 4, 2021.
I thank Dilip Abreu, Kenneth Ahern, Nageeb Ali, Alessandro Bonatti, Aubrey Clark, Marc Claveria Mayol, Glenn Ellison, Andrea Galeotti, Ben Golub, Matt Jackson, Erik Madsen, Al Roth, Miklos Sarvary, Andrew Sinclair, Satoru Takahashi, Johan Walden, Andy Wei, Glen Weyl, and Alex Wolitzky for comments. This research has been funded by NSF grant 1260744.

(2007), who analyzed a market with a single seller and homogeneous values. Every player who has the good can replicate it at no cost and sell copies to his neighbors in the network. Players acquire the good for the opportunity to consume it and further resell copies. We refer to the players who own the good at a certain time as sellers and to the others as buyers. At every date, a buyer-seller pair linked in the network is randomly selected to bargain over the price of the good. We propose a Markovian solution concept under which the terms of trade for each matched buyer and seller are determined by Nash bargaining. The state of the market at each date, which determines the disagreement payoffs in every match, is given by the configuration of sellers in the network at that date.

While our analysis applies broadly to goods with the properties outlined above, it is convenient to frame concepts and results in terms of (indivisible) information diffusing through the network. The assumption that the same good is transmitted through the network need not be taken literally. Some players may alter the good and resell customized versions. For instance, Sarvary (2011) describes the "value chain" in the information industry on a data-information-knowledge continuum as intermediaries add context, patterns and causal links to raw data and sell different versions to multiple institutions.

Buyers serve as both consumers and intermediaries in the market. The intermediation role can be beneficial for sellers when buyers provide access to parts of the network that sellers cannot reach directly or when buyers enhance the good and make it useful for others. Each buyer derives a direct utility from consuming the good and can also earn profits from selling copies to other buyers. Sellers may indirectly extract profits from buyers via intermediation paths along which every player demands a share of the consumption value and the resale profits of the next buyer on the path. Liebowitz (1985) coined the term "indirect appropriability" for the idea that sellers can collect part of the profits gained by intermediaries who acquire the original good and resell copies.[1] Nevertheless, as the standard argument for intellectual property rights suggests, competition among sellers of the original good and buyers who resell copies drives prices down in secondary markets and eliminates opportunities for indirect appropriation. Our network formulation encapsulates the intermediation role of buyers who provide indispensable market access as well as competitive forces that restrict indirect appropriability.

The starting point of our analysis is the intuition that a seller s can appropriate profits directly or indirectly from a buyer b if and only if the following conditions hold: (1) there exists a unique path between s and b; and (2) any path from another seller to b is intermediated by s. When these conditions are met, every player along the path from s to b is the only potential source of information for the next buyer on the path. We prove that uniqueness of the path between a pair of nodes defines an equivalence relation over nodes in any given

[1] Liebowitz argues that indirect appropriability explains why the introduction of photocopiers in 1959 led publishers to increase price discrimination for individual and library journal subscriptions but has not harmed publisher profits.

network. This key equivalence relation facilitates a succinct statement of conditions (1) and (2) above. Specifically, seller s and buyer b satisfy the two conditions if and only if b belongs to the equivalence class of s in an auxiliary network that mainly differs from the original one in linking all pairs of sellers.

The partition of the set of nodes into equivalence classes derived from the auxiliary network reflects the effects of competition and the scope of indirect appropriability for every seller. This partition delivers a classification of links in terms of competitive and monopolistic functions, which also determines the consequences of removing a link for information diffusion and the distribution of profits in the network. Sellers extract profits only from buyers who belong to their blocks in the partition. Every block that does not include sellers contains one buyer—the dealer—who intermediates all diffusion paths between sellers and other nodes in the block. Links within blocks constitute bottlenecks for the diffusion of information. Removing such links disconnects the network and stops information from reaching some buyers. For this reason, bottleneck links confer monopoly power to sellers and generate positive externalities for all players. When trade takes place across a bottleneck link, the seller demands a fraction of the buyer's consumption value and resale profits, and the partition evolves to reflect the buyer's takeover of the submarket for which he provides essential intermediation. Links between blocks are redundant for diffusion. Removing any such link does not affect the ultimate spread of information. However, redundant links create competition and enable dealers to acquire information at zero price. Hence, sellers have incentives to sever redundant links.

Beside its economic relevance, the network partition we discover provides graph theoretic insights into the structure of competing diffusion paths. There is at most one seller in every block. Information invariably enters any block without sellers through the dealer of the block. Dealers can get information from multiple neighbors and always receive it via redundant links. Nodes along the unique path connecting a particular buyer to the seller or dealer in his block provide essential intermediation for conveying information to that buyer; information diffuses within blocks via bottleneck links. In particular, every non-dealer buyer can only obtain information from a single neighbor over a bottleneck link. Moreover, all diffusion paths that reach the same buyer via a given block must overlap within that block.

Our results indicate that sellers' profit blocks are small in networks that are sufficiently well-connected or clustered, as is the case for many large social and economic networks (Jackson 2008; Easley and Kleinberg 2010). In such networks, the possibility of reproducing the good and its competitive effects undermine the indirect appropriability argument. If the creation of the original good requires investments greater than the low profits sellers can earn in the network, then granting intellectual property rights to sellers may be socially optimal.[2] When replication and resale are not prohibited, sellers can still avoid the harm of competition by engineering certain features of the prototype in order to restrict trade.[3]

The interplay between competition and monopoly in this setting is reminiscent of the market forces emerging in the intermediation model of Manea (2018). In that model, a single non-replicable good is sequentially resold between linked intermediaries in a network until it reaches a consumer. At every point in the resale process, the price is determined by competition between buyers. Similarly, pricing in the present model is driven by competition between sellers, but this analogy is superficial as the option of selling multiple units to neighbors leads to distinct strategic considerations. Even in markets with no intermediaries where the seller is linked directly to several buyers, the seller may prefer to limit supply in order to enhance competition among buyers and charge higher prices. This observation offers a bargaining theory perspective on the price-quantity trade-off faced by a monopolist.[4] In his version of the model with a single initial seller, Polanski (2007) provides recursive equations for the evolution of payoffs as buyers acquire information. Those payoff equations capture transitions between "consecutive" market states and, as such, reflect local network effects, but do not elucidate how these effects aggregate to overall profits. We fill this gap by providing explicit payoff formulae that reflect the global network structure. Our network decomposition into equivalence classes identifies the effective market share of every seller. The novel graph theoretic concepts developed here offer a stark delineation between detrimental competition and beneficial intermediation in networks for information goods, and deliver a purely topological characterization of competing diffusion paths.

In a contemporaneous working paper, Ali et al. (2020) study markets for information goods in which every pair of players can trade. Their setting corresponds to a complete network.[5] Focusing on a complete network affords a characterization of the best and the worst equilibria for the information seller, and facilitates the design of a mechanism in which the seller sells tokens and delays the release of information until all but one buyer purchase a token. In the mechanism, buyers are effectively pre-paying for information, and the seller extracts the profits attainable when resale is prohibited.

---

[2]Nevertheless, small networks involving criminal activity such as the networks of inside traders mapped by Ahern (2017) are sparse, allowing sellers to indirectly appropriate significant profits.

[3]Roundup Ready seeds are genetically modified to be resistant to the herbicide Roundup but are at the same time designed to be sterile, so that farmers cannot reproduce and share them. Digital rights management schemes control how digital content can be accessed and shared. Sellers can also eliminate redundant links while authorizing bottleneck links by restricting resale markets via sublicensing agreements, or by selling encrypted versions of the good to buyers who generate competition on the primary market, while selling the production technology or blueprint to buyers who are essential for serving secondary markets.

[4]Abreu and Manea (2021) formalize a general class of "exclusion commitments" and characterize the optimal commitment for the seller.

[5]Ali et al. considered the case of incomplete networks in an early version of their paper, and the exposition of some results here benefited from a preview of their first draft.

In other related work, Novos and Waldman (1984), Besen and Kirby (1989), Bakos et al. (1999), and Varian (2000) investigate how producer profits and social welfare are affected by copying and sharing information goods. Boldrin and Levine (2002) argue that the creator of a good can earn profits in a market without copyright protection where users reproduce the good at a constant rate and rent out copies. Jovanovic and Wang (2020) study diffusion and pricing in an industry where firms can either innovate or randomly meet and imitate innovators, and the "idea" can be licensed or resold to imitators and consumers. Muto (1986) and Takeyama (1994) analyze the consequences of consumption externalities for the pricing of information goods in a complete network, and Polanski (2019) provides a treatment for general networks with a focus on link formation and an application to citation graphs. This paper contributes to the growing literature on intermediation and bargaining power in networks (Condorelli and Galeotti 2016; Manea 2016), which has focused on the trade of non-replicable goods so far.

The rest of this paper is organized as follows. Section 2 defines basic graph theory concepts necessary for the analysis. Section 3 introduces the information selling game and the solution concept. In Section 4, we develop the network decomposition into equivalence classes and the characterization of payoffs building on this decomposition. In Section 5, we discuss the independence of prices from the history of trades. Section 6 introduces the notions of bottleneck and redundant links and elucidates how they shape diffusion paths. In Section 7, we derive comparative statics for link removals, buyer values, and seller entry. Section 8 provides concluding remarks. Appendix A presents proofs omitted in the main text, whereas the Online Appendix explores the implications of our results for large random networks, and contrasts the present model with the intermediation game of Manea (2018).

2. Graph Theory Preliminaries

This section reviews standard graph theory notions needed for the analysis: undirected networks, links, paths, distance, connected components, cycles, trees, and forests. Readers familiar with these concepts are advised to proceed to the next section.

Let M be a finite set whose elements we call *nodes*. A *network* H linking the nodes in M is a subset of $M \times M \setminus \{(i,i)|i \in M\}$. The condition $(i,j) \in H$ is interpreted as the existence of a *link* between nodes i and j in network H. For brevity, we use the notation ij for the link (i,j). The network H is *undirected* if $ij \in H$ whenever $ji \in H$. All networks in our analysis are assumed to be undirected. If $ij \in H$, we say that i and j are *neighbors* in H. The *subnetwork* H' of H *induced* by a subset of nodes $M' \subseteq M$ is the network linking the nodes in M' formed by the set of links $H \cap (M' \times M')$.

A *path* connecting nodes i and j in network H is a sequence of distinct nodes $(i_0 = i, i_1, \ldots, i_{\bar{k}} = j)$ such that $i_k i_{k+1} \in H$ for all $k \in \{0, 1, \ldots, \bar{k}-1\}$. The *distance* between nodes i and j in H is the smallest length $\bar{k}$ of any path $(i_0 = i, i_1, \ldots, i_{\bar{k}} = j)$ connecting i and j in H (defined to be infinite if there is no path between the two nodes). A *connected* *component* of H is the subnetwork of H induced by any maximal (with respect to inclusion) set of nodes that are mutually connected by paths in H. It is known that the set of connected components of an undirected network partitions the sets of nodes and links. A network is *connected* if it has a single connected component. A *cycle* in H is a sequence of nodes $(i_0 = i, i_1, \ldots, i_{\bar{k}} = i)$ such that $i_k i_{k+1} \in H$ for all $k \in \{0, 1, \ldots, \bar{k} - 1\}$ with the property that the first $\bar{k}$ nodes are distinct. A connected network that does not contain any cycle is called a *tree*. A network without cycles is a *forest* (alternatively, a forest is a network whose connected components are all trees).

3. The Information Selling Game

A finite set of *players* N is linked by an undirected connected *network* G. Some of the players—the *initial sellers*—are endowed with an identical *information good*. Let $\underline{S} \subset N$ denote the non-empty set of initial sellers. We assume that information is a non-depreciating indivisible consumption good for which every player has unit demand. Sellers can replicate the good at zero cost and sell it sequentially to each of their neighbors in G.[6] Upon acquiring the good, player $i \in N$ enjoys a *consumption value* $v_i \geq 0$ and joins the set of sellers.[7] The market is open for an infinite number of discrete *dates* $t = 0, 1, \ldots$. Players do not discount future payoffs.

The *state of the market* at date t is described by the set of holders of the information good $S \supseteq \underline{S}$ at t. For a given state S, we refer to the players in S as *sellers* and to those in $N \setminus S$ as *buyers*. In state S, one randomly selected buyer-seller pair linked in G is presented with the opportunity to trade. Hence, the set of links across which trade is possible in state S is given by $\mathcal{L}(S) = \{bs \in G | b \in N \setminus S, s \in S\}$. Let $\mathcal{S}$ denote the set of seller configurations that may arise from $\underline{S}$ following a sequence of trades.[8] For every $S \in \mathcal{S} \setminus \{N\}$, a probability distribution $\pi(S)$ assumed to have full support on $\mathcal{L}(S)$ specifies the probability $\pi_{bs}(S)$ with which each link $bs \in \mathcal{L}(S)$ is selected for bargaining at any date when the seller configuration is S. If b and s agree to trade in state S at date t, then b pays the agreed price to s, consumes the good, and becomes a seller in the new state $S \cup b$ at $t + 1$.[9] The game ends when the market reaches the state N, in which all players have the good.

3.1. **The solution.**

We propose a cooperative solution concept with a Markov structure under which the expected payoff of every player and the probability of agreement for the selected link at each date t depend only on the set of sellers at date t. Let $u_i(S)$ denote the expected *payoff* of player $i \in N$ in state $S \in \mathcal{S}$. When the link $bs \in \mathcal{L}(S)$ is selected for

---

[6]The analysis extends to a model in which players have a common unit cost for producing copies of the good.

[7]There are no consumption externalities. Players i with $v_i = 0$ act exclusively as intermediaries in the market. Sellers in $\underline{S}$ are assumed to have consumed the good before date 0.

[8]Formally, $\mathcal{S}$ represents the collection of sets $S \supseteq \underline{S}$ with the property that every node in S is connected to a node in $\underline{S}$ by a path that contains only nodes in S.

[9]For notational convenience, we routinely write $X \cup y$ and $X \setminus y$ for the sets $X \cup \{y\}$ and $X \setminus \{y\}$, respectively.

bargaining at date t in state S, seller s and buyer b negotiate the price for the information good as follows. In the event of an agreement, the market transitions to state $S \cup b$ at $t+1$, and the price in the transaction between s and b is determined according to the *Nash bargaining solution* with weights $(p, 1-p)$, where $p \in (0,1)$ is an exogenous variable common to all seller-buyer interactions,[10] assuming that:

- the total surplus created by the agreement is $v_b + u_b(S \cup b) + u_s(S \cup b)$, which represents the sum of the consumption value of b and the continuation payoffs of b and s in the new state $S \cup b$;
- the threat points of b and s are given by their corresponding disagreement payoffs, $u_b(S)$ and $u_s(S)$.

Hence, the feasibility of an agreement between b and s in state S hinges on the *gains from trade*

$$w_{bs}(S) := v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S). \tag{1}$$

Specifically, the *probability* $\alpha_{bs}(S)$ of an *agreement* between b and s in state S must satisfy the following incentive constraints:

$$\forall bs \in \mathcal{L}(S): \ \alpha_{bs}(S) \begin{cases} = 1 & \text{if } w_{bs}(S) > 0 \\ \in [0,1] & \text{if } w_{bs}(S) = 0 \\ = 0 & \text{if } w_{bs}(S) < 0. \end{cases} \tag{2}$$

Conditional on s and b being matched to bargain in state S, their respective continuation payoffs are given by $u_s(S) + p\alpha_{bs}(S)w_{bs}(S)$ and $u_b(S) + (1-p)\alpha_{bs}(S)w_{bs}(S)$. In the event of an agreement between b and s in state S, the continuation payoff of player $i \in N \setminus \{b, s\}$ is given by $u_i(S \cup b)$, while in case of disagreement it remains $u_i(S)$. Hence, the expected payoffs for sellers $s \in S$ and buyers $b \in N \setminus S$ in state $S \in \mathcal{S} \setminus \{N\}$ solve the following equations:

$$\begin{aligned}
(3) \qquad \forall s \in S: \ u_s(S) &= \sum_{b': b's \in \mathcal{L}(S)} \pi_{b's}(S) \left(u_s(S) + p\alpha_{b's}(S) w_{b's}(S)\right) \\
&+ \sum_{b's' \in \mathcal{L}(S): s' \neq s} \pi_{b's'}(S) \left(\alpha_{b's'}(S) u_s(S \cup b') + (1 - \alpha_{b's'}(S)) u_s(S)\right) \\
(4) \ \forall b \in N \setminus S: \ u_b(S) &= \sum_{s': bs' \in \mathcal{L}(S)} \pi_{bs'}(S)(u_b(S) + (1-p)\alpha_{bs'}(S) w_{bs'}(S)) \\
&+ \sum_{b's' \in \mathcal{L}(S): b' \neq b} \pi_{b's'}(S) \left(\alpha_{b's'}(S) u_b(S \cup b') + (1 - \alpha_{b's'}(S)) u_b(S)\right).
\end{aligned}$$

[10] All results generalize to a setting in which for every state $S \in \mathcal{S}$ and pair $(s, b) \in S \times (N \setminus S)$, seller s and buyer b divide the gains from trade according to Nash bargaining with weights $(p(s,b), 1 - p(s,b))$.

If seller s and buyer b are matched to bargain and reach an agreement in state S, the implicit *price* $t_{bs}(S)$ at which s and b trade solves the equation $u_s(S \cup b) + t_{bs}(S) = u_s(S) + pw_{bs}(S)$. Hence, $t_{bs}(S) = u_s(S) - u_s(S \cup b) + pw_{bs}(S)$.

The equations above do not lead to any constraints on payoffs for states in which all agreement probabilities are 0. To avoid this degeneracy, we assume that trade takes place with positive probability for at least one link in every state, i.e.,

$$\forall S \in \mathcal{S} \setminus \{N\}, \exists bs \in \mathcal{L}(S) \text{ s.t. } \alpha_{bs}(S) > 0. \tag{5}$$

Naturally, continuation payoffs at the end of the game should be zero,

$$u_i(N) = 0, \forall i \in N. \tag{6}$$

For seller configurations S in which trade takes place with positive probability on a single link bs (i.e., $\alpha_{bs}(S) > 0$ and $\alpha_{b's'}(S) = 0$ for all $b's' \in \mathcal{L}(S) \setminus \{bs\}$), we need to impose an additional condition on the bargaining solution. In such situations, the payoff equation for seller s in state S boils down to

$$u_s(S) = u_s(S) + p\pi_{bs}(S)\alpha_{bs}(S)w_{bs}(S),$$

which is equivalent to $w_{bs}(S) = v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S) = 0$ (since $p\pi_{bs}(S)\alpha_{bs}(S) > 0$). The equation for $u_b(S)$ is equivalent to the same condition. The payoff equations for players $i \in N \setminus \{b, s\}$ do not provide any constraints on $u_s(S)$ and $u_b(S)$, as they reduce to

$$u_i(S) = (1 - \pi_{bs}(S)\alpha_{bs}(S))u_i(S) + \pi_{bs}(S)\alpha_{bs}(S)u_i(S \cup b),$$

which is equivalent to $u_i(S) = u_i(S \cup b)$. The indeterminacy of the bargaining solution for states S in which $\alpha_{bs}(S) > 0$ for a single link $bs \in \mathcal{L}(S)$ is a consequence of the assumption that threat points in the bilateral bargaining game between b and s are given by the solution itself in state S. When (b, s) is the only pair that trades in configuration S, it is more natural to assume that both players' threat points are 0 since the market is permanently shut down if b and s fail to reach an agreement in state S. Thus, we require that s and b split the gains $v_b + u_b(S \cup b) + u_s(S \cup b)$ from a potential agreement according to the Nash bargaining solution with weights $(p, 1-p)$ and disagreement payoffs of 0 for both players. Formally, we impose the following condition:

$$\{b's' \in \mathcal{L}(S) | \alpha_{b's'}(S) > 0\} = \{bs\} \Longrightarrow u_s(S) = p\left(v_b + u_b(S \cup b) + u_s(S \cup b)\right). \tag{7}$$

The formula for $u_s(S)$ in the condition above, along with the equation $v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S) = 0$, implies that $u_b(S) = (1-p)\left(v_b + u_b(S \cup b) + u_s(S \cup b)\right)$.

We are now prepared to define our solution concept. The profile (u, α) of expected payoffs $u = (u_i(S))_{i \in N, S \in \mathcal{S}}$ and agreement probabilities $\alpha = (\alpha_{bs}(S))_{bs \in \mathcal{L}(S), S \in \mathcal{S} \setminus \{N\}}$ constitutes a *bargaining solution* if it satisfies constraints (2)-(7) for every state $S \in \mathcal{S}$ (with the variables

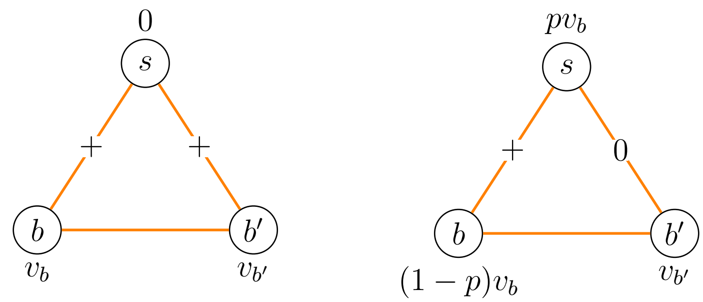

FIGURE 1. Multiple solutions.

$w_{bs}(S)$ derived from u via (1)). We say that the payoffs u are *consistent* with the agreement probabilities α if (u, α) constitutes a bargaining solution.

3.2. Uniqueness of the payoffs for a given a structure of agreements.

A contraction argument shows that the agreement probabilities α uniquely determine the payoffs u in every bargaining solution.

Proposition 1. *There exists at most one profile of expected payoffs that is consistent with a given profile of agreement probabilities.*

3.3. Multiplicity of the solution.

Polanski (2007) introduced a version of this model with a single initial seller ($|\underline{S}| = 1$) and symmetric consumption values ($v_i = 1$ for all $i \in N$), which we discuss in detail in Section 4.6. He shows that multiple solutions may coexist in his model, and this conclusion extends to our framework.[11] In light of Proposition 1, multiplicity necessarily stems from different profiles of agreement probabilties. The example from Figure 1 illustrates the multiplicity in a simple network with a single seller, player s, and two buyers, b and b'; the three players are linked with one other. In this example, after one of the buyers acquires the good from the seller, competitive forces imply that the other buyer obtains the good at zero price. Hence, Bertrand competition arises endogenously under our solution (see Appendix A for a proof of this claim). Based on this fact, we can construct several bargaining solutions in this network.[12]

The left panel of Figure 1 depicts a solution in which trading probabilities are positive over every link in all states. Under this solution, seller s suffers from a commitment problem (cf. Coase 1972) and does not make any profit. Each of the two buyers expects that s will eventually trade with the other buyer and can subsequently exploit the competition between s and the other buyer to acquire the good at zero price. Given these expectations, neither buyer is willing to pay a positive price for the good to seller s in the initial market. Payoffs are 0 for the seller, and v_b and $v_{b'}$ for b and b', respectively. All pairs of matched players are

[11]Polanski's solution concept allows some non-Markovian behavior that turns out to be inconsequential.
[12]The multiple solutions in this example are robust to the introduction of discounting and non-cooperative bargaining.

indifferent between trading and not trading in every market state (w takes value 0 for all links in every state), and thus the assumed structure of agreements is incentive compatible.

The right panel of Figure 1 illustrates a second solution, in which seller s "commits" not to trade with buyer b' in the initial market, but trade takes place with positive probability for all other matches and states.[13] After s trades with b, neither s nor b can extract any profit from b'. Since b' never acquires the good before b does given the assumed structure of agreements, bargaining between s and b proceeds as in a two-player network. Payoffs under this solution are pv_b for the seller, $(1-p)v_b$ for b, and $v_{b'}$ for b'. The agreement probabilities prescribed by the solution are incentive compatible. In particular, s and b' do not have incentives to trade in the initial market because $w_{b's}(\{s\}) = -pv_b < 0$.

The second solution does not amount to seller s permanently severing his link with buyer b' even though it stipulates that s cannot appropriate any fraction of the value of b'. Indeed, the Bertrand competition argument in Appendix A shows that even though the link $b's$ is not utilized for trade in the initial market state $\{s\}$ under this solution ($\alpha_{b's}(\{s\}) = 0$), it must be used with positive probability in state $\{b, s\}$ ($\alpha_{b's}(\{b, s\}) > 0$). Then, competition between seller s and buyer b drives down the price b' pays for the good to 0. By contrast, in the network obtained by removing the link $b's$, seller s does not have the option to compete with buyer b to sell to b', and the unique bargaining solution prescribes that b sells the good to b' at the bilateral monopoly price of $pv_{b'}$, and s demands a price of $pv_b + p^2 v_{b'}$ from b. Hence, seller s would earn higher profits if he could sever his link with buyer b'.

3.4. **Refinement of the solution.** To solve the multiplicity problem, we introduce a *refinement* of the bargaining solution similar to one proposed by Polanski (2007). We require that a bargaining solution (u, α) specifies a positive probability of agreement for every link in any configuration, i.e.,

$$\forall S \in \mathcal{S}, bs \in \mathcal{L}(S) : \alpha_{bs}(S) > 0. \quad (8)$$

Everywhere except Section 5, we restrict attention to bargaining solutions that satisfy this requirement and simply use the term *bargaining solution* to describe such profiles. We prove that a solution satisfying the refinement always exists and that the refinement selects unique payoffs u, which are consistent with any profile of agreement probabilities α (subject to (5)).

Note that under the bargaining solution illustrated in the right panel of Figure 1, buyer b acquires the good from seller s at price pv_b. However, in the "off-the-equilibrium-path" event that s trades with b' first, the market transitions to state $\{s, b'\}$, in which b obtains the good at price 0 from either s or b'. Hence, under the solution prescribing that s trade only with buyer b in the initial state, the price b pays depends on the history of trades (reflected in the market state).

[13]Another solution is obtained by interchanging the roles of b and b'.

By contrast, prices are history-independent under the bargaining solution illustrated in the left panel of Figure 1, which is selected by our refinement in the example. Indeed, any bargaining solution specifying that the seller trade with positive probability with either buyer in the initial market entails that each buyer obtains the good at price 0 in any state of the market. In Section 5, we show how this conclusion generalizes to arbitrary networks: the refinement selects the only bargaining solution payoffs that induce history-independent prices in trades over each link. Thus, under the refinement, bargaining between any buyer and seller does not require information about past trades.

4. Profits and a Network Decomposition

In our model, buyers act as both consumers and intermediaries. Upon acquiring the good, each buyer enjoys his consumption value and expects to collect a *resale value* that reflects the profits he earns by reselling the good to others. Thus, sellers may extract profits from buyers by means of direct links or indirect paths along which every player demands a fraction of the consumption and resale values of the next buyer on the path. This reasoning expresses the notion of *indirect appropriability* in a network setting (Liebowitz 1985; Johnson and Waldman 2005; Boldrin and Levine 2008).

The following definition, inspired by the example from Figure 1, is crucial for identifying which buyers a seller can indirectly appropriate profits from. Seller s is the *essential supplier* for buyer b in state S if the following conditions hold:

- there is a unique path $(s, b_1, \ldots, b_k = b)$ in G between s and b;
- any path in G from another seller in S to b contains s.

Under these conditions, seller s is the unique supplier of the good for all buyers on the path $(s, b_1, \ldots, b_k)$, and every player along the path is the only potential seller of the good for the next buyer on the path. Then, seller s exploits his *monopoly power* over buyer b_1 to get a fraction p of b_1's consumption and resale values. Likewise, b_1 acts as a monopolist for b_2 and demands a fraction p of b_2's consumption and resale values, and so on. These arguments suggest that s should obtain a share p^k of the consumption value v_b of buyer b.

Similarly, we say that buyer b is an *essential intermediary* for buyer b' in state S if the following conditions hold:

- there is a unique path $(b, b_1, \ldots, b_k = b')$ in G between b and b';
- every path in G from a node in S to b' passes through b.

These conditions imply that the good can reach b' only after b purchases it, and resale subsequently proceeds along the chain $(b, b_1, \ldots, b_k)$.

4.1. **The equivalence relation.** Since uniqueness of paths between pairs of nodes is central to the roles of essential suppliers and intermediaries, it is useful to study the properties of this relation. To this end, define a *binary relation* $\sim_H$ on the set of nodes of an arbitrary undirected network H as follows: $i \sim_H j$ if and only if nodes i and j are connected by a

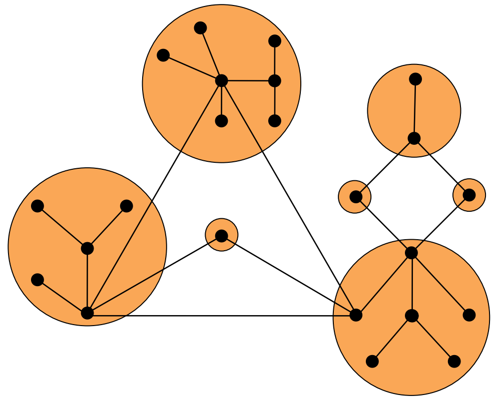

FIGURE 2. Equivalence classes in a network.

unique path in network H. We prove that $\sim_H$ constitutes an equivalence relation for every network H.

Lemma 1. *For every undirected network H, $\sim_H$ is an equivalence relation. Furthermore, if $i \sim_H j$, then all nodes on the unique path between nodes i and j in the network H belong to the same equivalence class under $\sim_H$.*

Figure 2 illustrates the partition of nodes in a network into equivalence classes of the binary relation. The set of nodes inside each circle constitutes an equivalence class. By Lemma 1, each equivalence class induces a tree in the underlying network.

Lemma 1 gives rise to an alternative interpretation of $\sim_H$. Let $\mathcal{F}(H)$ denote the network obtained from H by simultaneously removing every link that belongs to a cycle in H. Since $\mathcal{F}(H)$ has no cycles, it must be a forest. If $ij \in \mathcal{F}(H)$, then there is no cycle in H that contains the link ij, which means that the link constitutes the only path between i and j in H, so $i \sim_H j$. Since $\sim_H$ is an equivalence relation, every connected component of $\mathcal{F}(H)$ is included in the same equivalence class of $\sim_H$. If two nodes from different connected components of $\mathcal{F}(H)$ were in the same equivalence class of $\sim_H$, then Lemma 1 implies that all nodes along the unique path connecting them in H must be in the same equivalence class of $\sim_H$. However, in that case every link along the path represents the unique path in H between the two nodes, so the entire path must lie in $\mathcal{F}(H)$. This contradicts the assumption that the path connects different components of $\mathcal{F}(H)$. Therefore, the equivalence classes of $\sim_H$ are identical to the connected components of $\mathcal{F}(H)$. We refer to $\mathcal{F}(H)$ as the *forest derived by eliminating cycles* from H.[14]

4.2. **Characterization of essential suppliers and intermediaries.** The first condition required for s to serve as the essential supplier for b in state S can be restated as $b \sim_G s$. To articulate the second condition necessary for s to be the essential supplier for b in state

[14]A forest structure also drives market power and equilibrium prices in the model of segmented markets with interconnected exchanges analyzed by Malamud and Rostek (2021).

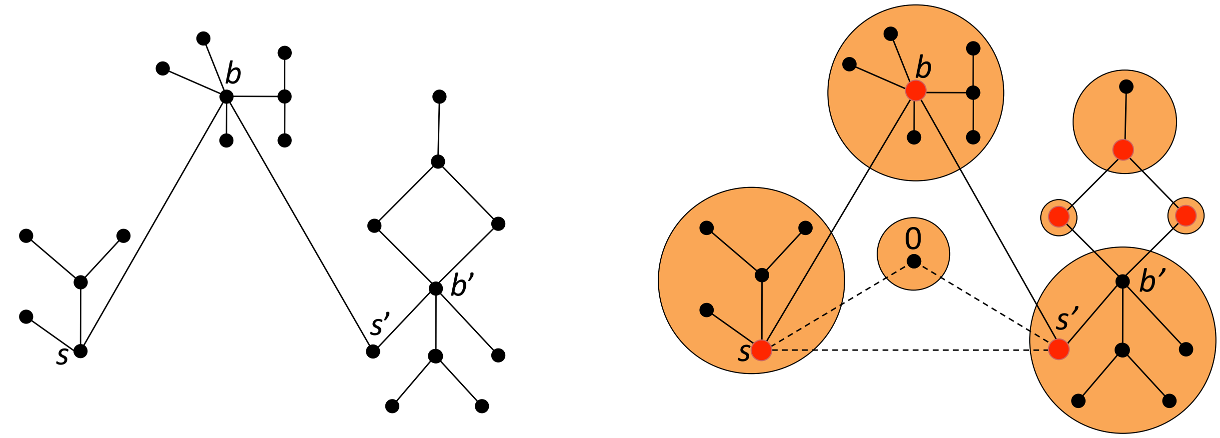

FIGURE 3. Equivalence classes for $G(\{s, s'\})$ and their dealers.

S, namely the requirement that any path in G from another seller in S to b contains s, we employ the binary relation $\sim$ for an auxiliary network. Consider the network $G(S)$ obtained by introducing a *dummy player* 0 and adding links between all pairs of nodes in the set $S \cup 0$. Let $C_i(S) = \{j \in N | j \sim_{G(S)} i\}$ denote the *equivalence class* of node i under $\sim_{G(S)}$ (or equivalence class of i in $G(S)$, for short) excluding the dummy player. The presence of the dummy player guarantees that no two nodes in S belong to the same equivalence class in $G(S)$ (its main purpose is to streamline notation and arguments for market states with two sellers). The right panel of Figure 3 shows how the network $G(S)$ and its equivalence classes are derived from the network G depicted in the left panel for the seller configuration $S = \{s, s'\}$.

We show that seller s is the essential supplier for buyer b in state S if and only if $b \sim_{G(S)} s$. Analogously, we find that buyer b is an essential intermediary for buyer b' in state S if and only if $b \sim_{G(S \cup b)} b'$. In other words, seller s is the essential supplier in state S for the set of buyers $C_s(S) \setminus s$, and buyer b is an essential intermediary in state S for the set of buyers $C_b(S \cup b) \setminus b$. Therefore, $C_s(S) \setminus s$ is the captive market of seller s, while $C_b(S \cup b) \setminus b$ represents the captive resale market of buyer b in state S.

Lemma 2. *Fix $S \in \mathcal{S}, s \in S$ and $b, b' \in N \setminus S$. Seller s is the essential supplier for buyer b in state S if and only if $b \sim_{G(S)} s$. Buyer b is an essential intermediary for buyer b' in state S if and only if $b \sim_{G(S \cup b)} b'$.*

4.3. **Dealers.** There may be equivalence classes in $G(S)$ that do not contain any seller. The next result shows that the good always "enters" each such class through the same node. Moreover, it characterizes this node as the only buyer in the class that can acquire the good from at least two of its neighbors. The following definition is useful in formalizing this statement: a neighbor j of node i in the network G is a *potential supplier* for buyer i in state S if there exists a path in G from a node in S to i that contains the link ij.

Lemma 3. *For every seller configuration $S \in \mathcal{S}$ and player $i \in N$, there exists a unique node $d(S, C_i(S))$ that is the first element of $C_i(S)$ along any path from S to $C_i(S)$ in the network G. Moreover, every buyer in $C_i(S) \setminus d(S, C_i(S))$ has a unique potential supplier in state S, and if $d(S, C_i(S)) \notin S$, then $d(S, C_i(S))$ has two or more potential suppliers in state S.*

Lemma 3 implies that in the seller configuration S, the players in $C_i(S) \setminus d(S, C_i(S))$ can only purchase the good via a sequence of trades that involves player $d(S, C_i(S))$ (re)selling the good. For this reason, we refer to $d(S, C_i(S))$ as the *dealer* of $C_i(S)$ in state S. Note that for $s \in S$, the definition naturally implies that seller s is the dealer for his equivalence class in $G(S)$, i.e., $d(S, C_s(S)) = s$. Recall that in this case, seller s is the essential supplier for the buyers in $C_s(S) \setminus s$ in state S. Likewise, every buyer b who is the dealer of his equivalence class $C_b(S)$ in state S is an essential intermediary for the buyers in $C_b(S) \setminus b$ in state S. In Figure 3, we indicate the dealer of each equivalence class in the network $G(\{s, s'\})$ by enlarging the corresponding node.

4.4. **Evolution of equivalence classes.** Clearly, if $i \sim_{G(S)} j$, then the assumption that G is connected implies that $i \sim_G j$, so i and j must belong to the same tree in the partition induced by $\sim_G$. In effect, $\sim_{G(S)}$ decomposes the equivalence classes under $\sim_G$ into smaller trees. In order to compute the bargaining solution payoffs, it is necessary to understand how equivalence classes evolve as trades take place. Consider a configuration of sellers $S \in \mathcal{S}$ and fix a seller $s \in S$ linked in G to a buyer $b \in N \setminus S$. We show that equivalence classes in $G(S)$ and $G(S \cup b)$ are identical, with one important exception: if b and s belong to the same equivalence class in $G(S)$, i.e., $b \sim_{G(S)} s$, then the equivalence class of s in $G(S)$ breaks up into two equivalence classes in $G(S \cup b)$ that separate b from s.

Proposition 2. *Fix $s \in S \in \mathcal{S}$ and $b \in N \setminus S$ such that $bs \in G$.*

(1) If $b \not\sim_{G(S)} s$, then $C_i(S \cup b) = C_i(S)$ for all $i \in N$.
(2) If instead $b \sim_{G(S)} s$, then $C_i(S \cup b) = C_i(S)$ for all $i \in N \setminus C_s(S)$, but $b \not\sim_{G(S \cup b)} s$ and $C_s(S \cup b) \cup C_b(S \cup b) = C_s(S)$.

4.5. **Main result: payoff formulae.** Suppose that seller s trades with buyer b in state S. Following the trade, buyer b realizes his consumption utility v_b and collects a resale value of $u_b(S \cup b)$ as a seller in the new state $S \cup b$, aggregating to total gains of $v_b + u_b(S \cup b)$.

Lemma 3 and Proposition 2 imply that if $s \not\sim_{G(S)} b$, then b is the dealer for his identical equivalence classes $C_b(S) = C_b(S \cup b)$ in the networks $G(S)$ and $G(S \cup b)$. By Lemma 3, dealer b can acquire the good via multiple intermediation paths from at least two potential suppliers in state S. The assumption that each buyer-seller matched pair trades with positive probability, which underlies our refinement of the bargaining solution, implies that b can delay trade until all players along the competing paths from sellers, including two of his potential suppliers, have the good. At that stage, *competition* between the potential suppliers drives the price that b pays for the good to 0. It follows that b eventually obtains the good for free, and his expected payoffs satisfy $u_b(S) = v_b + u_b(S \cup b)$.

If instead $s \sim_{G(S)} b$, then s is the dealer of $C_b(S)$ and b is not. Following the trade between b and s, the new seller b takes over the submarket $C_b(S \cup b) \subset C_s(S)$ for which he provides essential intermediation. Lemma 3 implies that s is the only potential supplier for b in state S. *Monopoly power* enables seller s to demand a fraction p of b's consumption and resale values. Hence, b acquires the good from s at a price of $p(v_b + u_b(S \cup b))$, and his expected payoff in state S is given by $u_b(S) = (1 - p)(v_b + u_b(S \cup b))$.

These intuitions pave the way to our main result, which leverages Lemma 2 to extend the formulae for buyer profits derived above to buyers b not directly linked to any seller in S. For $s \in S \in \mathcal{S}$, define

$$r_s(S) = \sum_{i \in C_s(S) \setminus s} p^{\delta(i,s)} v_i,$$

where $\delta(i, s)$ denotes the distance between nodes i and s in network G.[15]

Theorem 1. *The profile (u, α) constitutes a bargaining solution (under the refinement) if and only if for every $S \in \mathcal{S}$,*

$$\forall s \in S,\ u_s(S) = r_s(S)$$

$$\forall b \in N \setminus S,\ u_b(S) = \begin{cases} v_b + r_b(S \cup b) & \text{if } b = d(S, C_b(S)) \\ (1-p)\,(v_b + r_b(S \cup b)) & \text{if } b \neq d(S, C_b(S)) \end{cases} \tag{9}$$

and $\alpha_{bs}(S) \in (0, 1]$ for all $bs \in \mathcal{L}(S)$.

The proof shows that the unique bargaining solution payoffs u satisfy $v_b + u_b(S \cup b) + u_s(S \cup b) = u_b(S) + u_s(S)$ for all $bs \in \mathcal{L}(S)$ and $S \in \mathcal{S}$. Hence, players are indifferent between trading and not trading across every link, and the payoff profile u is consistent with any profile of agreement probabilities α.[16] The variables $r_s(S)$ and $r_b(S \cup b)$ in formulae (9) reflect the profits that seller s and buyer b indirectly appropriate from their captive markets $C_s(S) \setminus s$ and $C_b(S \cup b) \setminus b$, respectively, in state S. Lemma 2 implies the following restatement of Theorem 1. For any seller configuration S, seller s appropriates a fraction $p^{\delta(b,s)}$ of the consumption value v_b of each buyer b for whom s is the essential supplier in state S. Similarly, following any sequence of trades that conveys the good to buyer b, the resale value of buyer b aggregates a fraction $p^{\delta(b,b')}$ of the consumption value $v_{b'}$ of each buyer b' for whom b is an essential intermediary in state S. The price buyer b pays for the good is either 0 or a fraction p of his consumption and resale values in state S corresponding to whether b is a dealer for his equivalence class in $G(S)$ or not. Proposition 2 and Theorem 1 imply that

[15]The distances $\delta(i, s)$ appearing in the formula for $r_s(S)$ involve pairs (i, s) with $i \sim_{G(S)} s$, and hence $i \sim_G s$. In this case, $\delta(i, s)$ is simply the length of the unique path between i and s in G.

[16]This indifference in the frictionless model is the main obstacle in extending the analysis to the case of discounting. In a version of the model with discounting, it is difficult to prove that the indifference is always broken in favor of trade as postulated by our refinement.

prices decline along any trading path within an equivalence class and drop to 0 when the good is sold to a new class. Moreover, if b is a buyer in state S and $(d(S, C_b(S)), b_1, \ldots, b_k = b)$ denotes the unique path in G between the dealer for $C_b(S)$ and buyer b, the consumption value of buyer b is directly or *indirectly appropriated* by the players in the *intermediation chain* $(d(S, C_b(S)), b_1, \ldots, b_k)$ with corresponding shares $(p^k, (1-p)p^{k-1}, \ldots, (1-p)p, 1-p)$.

4.6. **The case with a single seller and Polanski's result.** Polanski (2007) provides a recursive system of payoff equations for a setting similar to the one studied here for markets with a single initial seller. For the special case with a single seller, $\underline{S} = \{s\}$, the equivalence relations $\sim_{G(\{s\})}$ and $\sim_G$ developed in our framework coincide (modulo the dummy player), so the formula for seller profits from Theorem 1 boils down to

$$u_s(\{s\}) = \sum_{i \sim_G s, i \neq s} p^{\delta(i,s)} v_i.$$

Therefore, in order to determine the profit of seller s, it is sufficient to consider equivalence classes under $\sim_G$. Examining equivalence classes in the auxiliary networks $G(S)$ is necessary only for computing buyers' payoffs and tracking the evolution of profits as other players acquire the good. The construction of the auxiliary network $G(S)$ has indeed been motivated by the intuition that the corresponding equivalence relation $\sim_{G(S)}$ succinctly describes competition among sellers in S.

Polanski's recursive equations capture transitions between "consecutive" market states by relating the payoff $u_i(S)$ to payoffs of the type $u_i(S \cup b)$. He finds that the terms of trade between a seller s and a buyer b depend on whether b belongs to a cycle that includes at least one seller. To extend his result to our setting with multiple initial sellers, we need to consider cycles in the network $G(S)$ rather than G for the payoff equations corresponding to state S. For $S \in \mathcal{S}, b \in N \setminus S$, define

$$c_b(S) = \begin{cases} 0 & \text{if there exists a cycle in } G(S) \text{ that contains } b \text{ and an element of } S \\ 1 & \text{otherwise.} \end{cases}$$

One can check that for $bs \in \mathcal{L}(S)$, we have $c_b(S) = 0$ if b is the dealer for $C_b(S)$ in state S, and $c_b(S) = 1$ otherwise. This observation leads to the following corollary of Theorem 1 and Proposition 2, which generalizes Polanski's result.

Corollary 1. *For any $s \in S \in \mathcal{S}$ and $b \in N \setminus S$ such that $bs \in G$, the bargaining solution payoffs satisfy*

$$\begin{aligned} u_s(S) &= u_s(S \cup b) + pc_b(S)(v_b + u_b(S \cup b)) \\ u_b(S) &= (1 - pc_b(S))(v_b + u_b(S \cup b)). \end{aligned}$$

For $s \in S \in \mathcal{S}$ and $b, b' \in N \setminus S$ such that $\mathcal{L}(S)$ does not contain any links of b or s, but contains a link of b', we have $u_s(S) = 0$ and $u_b(S) = u_b(S \cup b')$.

As Polanski points out, the identities from the corollary provide a computational procedure for evaluating the bargaining solution payoffs based on transitions between market states. These recursive payoff equations reflect local network effects. Our closed-form payoff formulae elucidate how the global network structure affects the division of gains from trade, and the decomposition of the network into equivalence classes delineates opportunities for indirect appropriability and provides a classification of links according to their monopolistic or competitive roles. We will see that this classification also translates into a taxonomy of links as either bottlenecks or redundant for information diffusion and network connectivity. We will demonstrate that the graph theoretic concepts introduced in this paper—the equivalence relation, the auxiliary network, dealers, essential intermediaries, and bottleneck and redundant links—characterize the overlapping structure of competing paths of diffusion (Theorem 2) and provide tractable tools for the derivation of comparative statics for buyer values (Corollary 2), seller entry (Proposition 4) and, most relevant for applications, link removals (Theorem 3). Another result that relies on this toolkit is that the refinement selects the unique payoff profile under which prices are history-independent (Proposition 3).[17]

5. Foundation for the Refinement

With the aim of providing a foundation for our refinement, we revert to using the terms *bargaining solution* for any profile (u, α) satisfying conditions (2)-(7) and *refinement of the bargaining solution* for profiles that additionally satisfy constraint (8). Fix a bargaining solution with payoff profile u. Recall that an agreement in state S between seller s and buyer b entails the price $t_{bs}(S) = u_s(S) - u_s(S \cup b) + pw_{bs}(S)$. We say that the *prices generated by u are history-independent* if for every $bs \in G$, we have $t_{bs}(S) = t_{bs}(S')$ for any pair of states $S, S' \in \mathcal{S}$ such that s is a seller and b is a buyer in both configurations S and S'. The interpretation of history-independence of prices is that the bargaining process for any buyer-seller link does not require information about prior trades.

In Section 3, we argued that prices under the bargaining solution ruled out by the refinement in the network from Figure 1 are not history-independent. The next result generalizes that conclusion: in every network, prices are history-independent only for the bargaining solution payoffs that survive the refinement.[18] Hence, our refinement selects the solutions that do not rely on the assumption that matched players observe the state of the market.

Proposition 3. *The refinement of the bargaining solution generates history-independent prices and selects the unique payoff profile for which prices are history-independent.*

[17] Polanski also notes history-independence of prices in his specification of the model and the solution, but does not argue that his refinement selects the unique payoffs with this property, which is technically more involved.

[18] The refinement has the additional property of inducing *seller-independent* prices, i.e., $t_{bs}(S) = t_{bs'}(S')$ for any pair of states $S, S' \in \mathcal{S}$ such that $s \in S,\ s' \in S'$, and $b \notin S \cup S'$.

6. The Anatomy of Diffusion Paths: Redundant and Bottleneck Links

Consider a link $ij \in G$ such that not both i and j are sellers in the initial state $\underline{S}$. We say that ij is a *redundant* link if i and j belong to distinct equivalence classes in the initial market, i.e., $i \not\sim_{G(\underline{S})} j$. For all $S \in \mathcal{S}$, we have that $G(\underline{S}) \subseteq G(S)$, so $i \not\sim_{G(\underline{S})} j$ implies $i \not\sim_{G(S)} j$. Thus, if ij is redundant, then i and j remain in distinct equivalence classes as the market evolves. We say that ij is a *bottleneck* link if it is not redundant, i.e., $i \sim_{G(\underline{S})} j$. Since equivalence classes induce trees in the network G and each trade breaks up at most one equivalence class into two distinct classes, Proposition 2 implies that the only pair of players linked in G that can be separated into different equivalence classes following a trade is the buyer-seller pair conducting the trade. Hence, if ij is a bottleneck link, then i and j are members of the same equivalence class in state $\underline{S}$ and continue to share an equivalence class until they trade with each other; that is, $i \sim_{G(S)} j$ for all $S \in \mathcal{S}$ such that $\{i, j\} \not\subseteq S$.

The evolution of equivalence classes as information diffuses (Proposition 2) can be restated in the language of redundant and bottleneck links as follows. Trading over a redundant link does not change the structure of equivalence classes, while trading over a bottleneck link breaks up the equivalence class containing the link into two classes that separate the buyer from the seller.

The partition of the network into equivalence classes and the ensuing concepts of redundant and bottleneck links lead to a systematic characterization of competing paths of diffusion in the network. By definition, each dealer buyer can receive the good only from neighbors outside his class. As links that span distinct equivalence classes are redundant, dealer buyers must acquire the good by means of redundant links. Moreover, Lemma 3 implies that dealer buyers have at least two potential suppliers.

By contrast, Lemma 3 shows that each non-dealer buyer has a single potential supplier, which is the neighbor of the buyer on the unique path connecting the buyer to the dealer of his equivalence class. This path is contained within the buyer's equivalence class and thus consists of bottleneck links. In particular, the non-dealer buyer acquires the good from his only potential supplier over a bottleneck link. Therefore, there is an implicit flow of trade over bottleneck links: diffusion within each equivalence class is described by a *directed tree rooted* at its dealer.

Consider now the collection of competing paths that deliver the good to a given buyer. Every path in this collection that "crosses" a certain equivalence class has to enter the class via its dealer. Logic similar to Lemma 3 shows that each such path must also exit the equivalence class through the same node. This implies that all paths conveying the good to the chosen buyer and intersecting a given equivalence class must cross the class only once and overlap within the class. The next result summarizes these observations.

Theorem 2. *A buyer is a dealer in state S if and only if he has two or more potential suppliers in state S. The good always reaches dealer buyers via redundant links and non-dealer buyers via bottleneck links. For any market state $S \in \mathcal{S}$ and buyer $b \in N \setminus S$, all paths in G that connect any seller in S to buyer b and intersect a given equivalence class $C_i(S)$ of $\sim_{G(S)}$ must enter $C_i(S)$ exactly once and overlap perfectly within $C_i(S)$.*

7. Comparative Statics

In this section, we present comparative statics results for buyer values, seller entry and link removals, and discuss the optimal removal of links for sellers.

7.1. Comparative statics for buyer values and seller entry.

Since $r_s(S)$ is increasing in v_b for all $s \in S$ and $b \in N \setminus S$, and equivalence classes are determined entirely by network topology, Theorem 1 has the following corollary.

Corollary 2. *For any $S \in \mathcal{S}$ and $b \in N \setminus S$, the payoffs of all players in state S are (weakly) increasing in v_b.*

Theorem 1 also delivers comparative statics with respect to the set of sellers. Suppose that new sellers enter the market and the initial state expands from $\underline{S}$ to $\underline{S}'$. Since $G(\underline{S})$ is a subnetwork of $G(\underline{S}')$, every pair of nodes related under $\sim_{G(\underline{S}')}$ is also related under $\sim_{G(\underline{S})}$. It follows that $C_i(\underline{S}') \subseteq C_i(\underline{S})$ for all $i \in N$. In particular, the captive market of every incumbent seller in $\underline{S}$ shrinks, and Theorem 1 implies that the profits of all these sellers decrease following the entry of the new sellers from $\underline{S}' \setminus \underline{S}$. By definition, dealer buyers (not in $\underline{S}'$) maintain dealer status. However, the captive markets of dealer buyers may shrink, resulting in lower profits. Finally, the set of buyers for whom non-dealer buyers (not in $\underline{S}'$) are essential intermediaries shrinks as well—for every buyer b, we have $C_b(\underline{S}' \cup b) \subseteq C_b(\underline{S} \cup b)$—but such buyers may become dealers in the new seller configuration $\underline{S}'$ due to competition created by the additional sellers. Such buyers acquire the good at zero price following the entry of new sellers, which may translate into higher payoffs. Whether the entry of the new sellers benefits non-dealer buyers depends on the trade-off between the lower acquisition price and the smaller captive market. The next result summarizes these findings.

Proposition 4. *Consider two initial market states $\underline{S} \subset \underline{S}'$. The payoffs of every seller in $\underline{S}$ and every buyer (not in $\underline{S}'$) who is a dealer in state $\underline{S}$ are weakly lower in market $\underline{S}'$ than in $\underline{S}$. The effect of the expansion of the set of sellers from $\underline{S}$ to $\underline{S}'$ for buyers who are not dealers in state $\underline{S}$ is ambiguous.*

7.2. Comparative statics for link removals.

We now investigate the effects of removing links from the network on information diffusion and intermediation profits. Fix a connected network G, a seller configuration $S \in \mathcal{S}$, and a link $ij \in G$ for which not both i and j belong to S (links between sellers are irrelevant in the game). Let G' denote the network obtained by removing link ij from G.[19]

Suppose first that ij is a bottleneck link. As argued in Section 6, the assumption that $\{i, j\} \not\subseteq S$ implies that $i \sim_{G(S)} j$. In particular, we have $i \sim_G j$, and hence deleting the link from G disconnects the network into two connected components. We prove that the sellers in S belong to the same connected component of the resulting network G' as the dealer $d(S, C_i(S))$ for the common equivalence class of i and j in $G(S)$. Hence, players in the other connected component of G' do not have access to any seller and obtain no profits. The removal of bottleneck link ij breaks up the equivalence class of i and j from $G(S)$ into two subclasses and does not affect the composition of other equivalence classes. Player $d(S, C_i(S))$ remains the dealer for his smaller equivalence class in G' but suffers a drop in profits. The loss of the link hurts both i and j: one of them becomes disconnected from sellers and gets zero payoff, while the other collects intermediation profits from a smaller equivalence class. Since the other equivalence classes of $\sim_{G(S)}$ contained in the connected component of node $d(S, C_i(S))$ in G' and their dealers are unaffected by the removal of link ij, players in those classes obtain the same payoffs in G and G'. For an illustration, consider the pair of nodes $b' \sim_{G(\{s,s'\})} s'$ linked in the network G from Figure 3. The removal of the link $b's'$ from G does not affect the payoffs of players in the equivalence classes of b and s in $G(\{s, s'\})$, but disconnects the buyers in other equivalence classes from the two sellers.

If instead ij is a redundant link, then we show that its removal from G does not prevent any player from acquiring the good.[20] The removal of the redundant link ij leads to a weak expansion in each player's equivalence class in state S. Theorem 1 implies that every seller's profit is weakly higher in G' than in G. Therefore, redundant links impose *negative externalities* on sellers. As the example from Section 3 demonstrates, a seller may benefit from severing one of his links. The set of players for whom each buyer serves as an essential intermediary also weakly expands. Lemma 2 and Theorem 1 imply that the payoffs of buyers who are not dealers in $G(S)$ weakly increase after link ij is deleted from G. The network from Figure 4 provides an example in which the profit of a non-dealer buyer strictly increases after deleting one of his redundant links. Indeed, if b deletes his redundant link with b'' in that network, then his equivalence class expands from $\{b, s\}$ to $\{b, b', b'', s\}$. Since b is not a dealer either before or after deleting the link bb'', the link deletion increases his payoff from $(1 - p)v_b$ to $(1 - p)(v_b + pv_{b'} + p^2 v_{b''})$.

However, dealer buyers may lose dealer status when a redundant link is deleted from the network. Such buyers exploit competition between sellers to obtain the good for free in the original network, but have to pay a fraction $1 - p$ of their consumption and resale values

[19]While G' may be disconnected, the results of previous sections apply to every connected component of G' that contains sellers, and we use this straightforward extension here.

[20]However, the ensuing network G' may be disconnected. For instance, in the network G with two sellers, s and s', linked to a single buyer, player b, we have $b \not\sim_{G(\{s,s'\})} s$. Removing the link bs from G disconnects the network but does not prevent b from acquiring the good from s'.

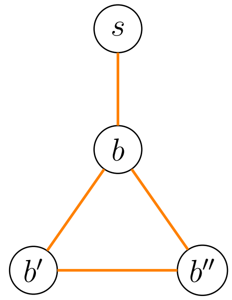

FIGURE 4. Non-dealer buyer b is better off if he severs his link with b''. Removing the redundant link bb'' also benefits dealer buyer b' if $v_{b'} < (1-p)v_{b''}$.

following the deletion of the redundant link, which may cause a decline in their overall profits. For example, consider the link between nodes b and s for which $b \not\sim_{G(\{s,s'\})} s$ in the network G from Figure 3. Removing the link bs from G leads to the merger of the equivalence classes of nodes b and s' from $G(\{s, s'\})$. After the link removal, buyer b is no longer a dealer, and seller s' is able to get a share $1-p$ of his consumption and resale values. Hence, the removal of link bs is beneficial for s' and detrimental for b. Removing redundant links can also have the opposite effect on dealer buyer payoffs. For instance, in the network from Figure 4, both buyers b' and b'' are dealers for singleton equivalence classes. Removing the redundant link bb'' leads to a network with a single equivalence class where buyer b' obtains payoff $(1-p)(v_{b'} + pv_{b''})$, which is greater than his payoff $v_{b'}$ in the original network if $(1-p)v_{b''} > v_{b'}$.

In the example from Figure 4, removing the link bb'' undermines the dealer status of buyer b', but expands his captive resale market at the same time. However, we show that removing one of a dealer's *own* links cannot simultaneously produce both these outcomes for the dealer. That is, if the removal of one of a dealer buyer's own links deprives him of dealer status, then that buyer's captive market cannot expand; in this case, the buyer's payoff unambiguously decreases. If instead the link removal does not affect the buyer's dealer role, then the buyer continues to secure the good at zero price after losing the link; as the buyer's captive market cannot shrink following the link removal, the buyer's payoff weakly increases. Thus, to predict the payoff consequences of losing a link for a dealer buyer, it is sufficient to determine whether the link loss changes that buyer's dealer status. A special category of redundant links is relevant for this phenomenon: a link ij is *pivotal* for dealer buyer i in state S if i has exactly two potential suppliers in state S, one of which is j (by Theorem 2, links between dealers and their potential suppliers are always redundant). We prove that removing the link ij from the network results in buyer i's loss of dealer status if and only if ij is a pivotal link for dealer i in state S. Therefore, if i is a dealer buyer in state S, then i is hurt by the loss of the link ij if and only if ij is a pivotal link for i in state S.

The following result, whose detailed proof is relegated to the Online Appendix given the extended sketch above, summarizes the comparative statics.

Theorem 3. *Consider a seller configuration $S \in \mathcal{S}$ in the connected network G and a link $ij \in G$ with $\{i, j\} \not\subseteq S$. Let G' be the network obtained by deleting the link ij from G.*

(1) If ij is a bottleneck link, then G' is a disconnected network formed by two connected components. Information does not reach the players in the connected component of G' that does not contain $d(S, C_i(S))$; thus, these players' payoffs drop to 0 when link ij is removed. The payoffs of players in $C_i(S)$ from the same connected component as $d(S, C_i(S))$ in G' weakly decrease after removing link ij. The payoffs of players i, j, and $d(S, C_i(S))$ strictly decrease following the link removal if $v_b > 0$ for all $b \in N \setminus \underline{S}$. The payoffs of all other players are identical in G and G'.

(2) If ij is a redundant link, then information diffuses to all players in G'. All sellers and the buyers who are not dealers in state S for network G weakly benefit from the removal of link ij. In general, the effect of removing the link on the payoffs of buyers who are dealers for their equivalence class in $G(S)$ is ambiguous. Nevertheless, if i is a dealer in state S and $v_b > 0$ for all $b \in N \setminus \underline{S}$, then the loss of link ij strictly reduces i's payoff if and only if ij is a pivotal link for i in state S.

The result above considers the effects of removing a single redundant link from the network. If instead we remove all redundant links from G at the same time, which leads to the forest $\mathcal{F}(G(\underline{S}))$, then the profits of sellers do not change. However, the simultaneous removal of redundant links blocks the spread of information to buyers whose equivalence class under $\sim_{G(\underline{S})}$ does not contain sellers and reduces these buyers' payoffs to 0.

The classification of links emerging from Theorems 2 and 3 leads to the following conclusions regarding seller profits and information transmission. Bottleneck links confer monopoly power to sellers. The deletion of a bottleneck link disconnects the network, blocks the spread of information, and hurts sellers. Redundant links create competition among sellers. The deletion of a redundant link does not prevent the diffusion of information and benefits sellers.

7.3. **Optimal link removals.** Consider now a situation with a single seller s, who can prohibit trade on a subset of links by engineering features of the good as explained in footnote 3. Theorem 1 implies that seller s would optimally allow trade only over the links of a tree T, which is a subnetwork of G that maximizes the expression

$$\sum_{i \in N \setminus s} p^{\delta^T(i,s)} v_i,$$

where $\delta^T(i, s)$ represents the distance between nodes i and s in tree T. Note that any restructuring of a tree whereby a given buyer b who originally receives the good from a node b' severs his link with b' and creates a new link with a node closer to s is beneficial for the seller. In particular, the *star* network, in which the seller is linked to all buyers and there are no links between buyers, maximizes seller profit among all networks.

We can similarly characterize the subnetwork of G that maximizes the joint profits of a group of competing sellers $\underline{S}$. In this case, the optimal subset of trading links is described by a partition $(N_s)_{s\in\underline{S}}$ of the set of nodes N and a collection of associated trees $(T_s)_{s\in\underline{S}}$ such that $s \in N_s$ and T_s consists of a subset of the links in G between pairs of nodes in N_s for all $s \in \underline{S}$. The partition should maximize the expression

$$\sum_{s\in\underline{S}} \sum_{i\in N_s\setminus s} p^{\delta^{T_s}(i,s)} v_i.$$

Sellers prefer to remove all links from G not belonging to the forest $\cup_{s\in\underline{S}} T_s$. This means that sellers divide the market into a set of non-overlapping trees from which they indirectly appropriate profits and commit to not competing with one another for any buyer.

8. Conclusion

We studied a model in which players consume, replicate, and resell copies of a good in a network. In the model, buyers may intermediate trade and indirectly transfer profits from far-away buyers to sellers as the good is sequentially resold over the links of the network. However, buyers who acquire copies of the good may also create competition for sellers of the original good, and this limits opportunities for indirect profit appropriation. Our network formulation thus captures the antithesis between two central concepts in the research on copying and intellectual property—indirect appropriability versus competition. We discovered a key equivalence relation that describes the roles of essential suppliers and intermediaries for the diffusion of the good. Sellers collect profits from buyers for whom they are essential suppliers, while buyers make profits by conveying the good to other buyers for whom they provide essential intermediation. Equivalence classes of the relation delineate the captive markets of every seller and buyer in the network.

The price a buyer pays for the good is either zero or a fixed fraction of his consumption and resale values corresponding to whether the buyer is able to exploit competition among multiple neighbors supplying the good or is subject to a monopoly in which a single neighbor provides access to the good. Links that induce competition among sellers are redundant for the diffusion of the good through the network and generate negative externalities for sellers, while links that enable monopolies constitute bottlenecks for diffusion and produce positive externalities for all players. Redundant links bridge distinct equivalence classes, while bottleneck links are enclosed in the same equivalence class. The network partition into equivalence classes reveals rich structural properties of competing paths of diffusion. Our analysis shows that in networks that are well-connected or clustered, competition obstructs indirect appropriability. In such situations, granting intellectual property rights fosters the creation of information goods.

In order to obtain theoretical results for general networks, we have made a number of simplifying assumptions, among which we enumerate: the network structure and buyer values are exogenous and commonly known; players do not discount payoffs; the original good and its copies are perfect substitutes; the solution concept is cooperative and favors trade; sales contracts are bilateral and cannot specify restrictions on replication and resale. In future work, it would be useful to extend the analysis to markets in which some of these modeling assumptions are unrealistic. Nevertheless, the graph theoretic byproducts of this research—including the concepts of equivalence classes, essential suppliers and intermediaries, dealers, and bottleneck and redundant links—do not hinge on the particular model specification and are likely to play an important role in other models of diffusion in networks.

Appendix A. Proofs

Proof of Proposition 1. We proceed by contradiction. Suppose that (u, α) and (u', α) constitute two bargaining solutions with distinct payoffs u and u' but identical agreement probabilities α. Let $S \in \mathcal{S}$ be a set of maximal cardinality for which there exists $i \in N$ such that $u_i(S) \neq u'_i(S)$. By definition, $S \neq N$, $\mathcal{L}(S) \neq \emptyset$, and

$$u_i(S \cup b) = u'_i(S \cup b), \forall i \in N, b \in N \setminus S \text{ (such that } bs \in \mathcal{L}(S) \text{ for some } s \in S). \tag{10}$$

Then the payoff equations for the solutions (u, α) and (u', α) lead to

$$\begin{aligned} u_s(S) &= \left(\sum_{b':b's \in \mathcal{L}(S)} \pi_{b's}(S)(1 - p\alpha_{b's}(S)) + \sum_{b's' \in \mathcal{L}(S):s' \neq s} \pi_{b's'}(S)(1 - \alpha_{b's'}(S)) \right) u_s(S) \\ &+ p \sum_{b':b's \in \mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S) \left(v_{b'} + u_{b'}(S \cup b') + u_s(S \cup b') - u_{b'}(S)\right) \\ &+ \sum_{b's' \in \mathcal{L}(S):s' \neq s} \pi_{b's'}(S)\alpha_{b's'}(S) u_s(S \cup b') \end{aligned} \tag{11}$$

$$\begin{aligned} u'_s(S) &= \left(\sum_{b':b's \in \mathcal{L}(S)} \pi_{b's}(S)(1 - p\alpha_{b's}(S)) + \sum_{b's' \in \mathcal{L}(S):s' \neq s} \pi_{b's'}(S)(1 - \alpha_{b's'}(S)) \right) u'_s(S) \\ &+ p \sum_{b':b's \in \mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S) \left(v_{b'} + u_{b'}(S \cup b') + u_s(S \cup b') - u'_{b'}(S)\right) \\ &+ \sum_{b's' \in \mathcal{L}(S):s' \neq s} \pi_{b's'}(S)\alpha_{b's'}(S) u_s(S \cup b'). \end{aligned} \tag{12}$$

Let $\Delta_1 = \max_{s \in S} |u_s(S) - u'_s(S)|$ and $\Delta_2 = \max_{b \in N \setminus S} |u_b(S) - u'_b(S)|$. We prove that $\Delta_1 = \Delta_2 = 0$, which contradicts the assumption that $u_i(S) \neq u'_i(S)$ for some $i \in N$.

Fix $s \in S$ such that $|u_s(S) - u'_s(S)| = \Delta_1$. Let X denote the probability that the matched pair does not reach agreement under α in a period with seller configuration S, Y_s the probability that seller s reaches an agreement in such a period, and Z_s the sum of terms that do not involve the variables $(u_i(S))_{i\in N}$ in (11). Mathematically,

$$X = \sum_{b's'\in\mathcal{L}(S)} \pi_{b's'}(S)(1-\alpha_{b's'}(S)), \quad Y_s = \sum_{b':b's\in\mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S),$$

$$Z_s = p\sum_{b':b's\in\mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S)\left(v_{b'} + u_{b'}(S\cup b') + u_s(S\cup b')\right) + \sum_{b's'\in\mathcal{L}(S):s'\neq s} \pi_{b's'}(S)\alpha_{b's'}(S)u_s(S\cup b').$$

We have

$$1 - X - (1-p)Y_s = \sum_{b's'\in\mathcal{L}(S):s'\neq s} \pi_{b's'}(S)\alpha_{b's'}(S) + p\sum_{b':b's\in\mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S) > 0$$

because $p > 0$, $\pi(S)$ places positive probability on every link in $\mathcal{L}(S) \neq \emptyset$, and condition (5) requires that the probability of agreement under α is positive for at least one link in state S. Collecting the variables $u_s(S)$ in (11) and $u'_s(S)$ in (12), we obtain

$$\begin{aligned} u_s(S)\left(1 - X - (1-p)Y_s\right) &= Z_s - p\sum_{b':b's\in\mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S)u_{b'}(S) \\ u'_s(S)\left(1 - X - (1-p)Y_s\right) &= Z_s - p\sum_{b':b's\in\mathcal{L}(S)} \pi_{b's}(S)\alpha_{b's}(S)u'_{b'}(S), \end{aligned}$$

or equivalently

$$\begin{aligned} u_s(S) &= \frac{Z_s}{1 - X - (1-p)Y_s} - p\sum_{b':b's\in\mathcal{L}(S)} \frac{\pi_{b's}(S)\alpha_{b's}(S)}{1 - X - (1-p)Y_s} u_{b'}(S) \\ u'_s(S) &= \frac{Z_s}{1 - X - (1-p)Y_s} - p\sum_{b':b's\in\mathcal{L}(S)} \frac{\pi_{b's}(S)\alpha_{b's}(S)}{1 - X - (1-p)Y_s} u'_{b'}(S). \end{aligned}$$

The triangle inequality implies that

$$\begin{aligned} \Delta_1 = |u_s(S) - u'_s(S)| &\leq p\sum_{b':b's\in\mathcal{L}(S)} \frac{\pi_{b's}(S)\alpha_{b's}(S)}{1 - X - (1-p)Y_s} |u_{b'}(S) - u'_{b'}(S)| \\ &\leq p\sum_{b':b's\in\mathcal{L}(S)} \frac{\pi_{b's}(S)\alpha_{b's}(S)}{1 - X - (1-p)Y_s} \Delta_2 = \frac{pY_s}{1 - X - (1-p)Y_s}\Delta_2. \end{aligned}$$

We can define buyer-side variables b and Y_b analogous to the seller-side ones s and Y_s, respectively, and derive the inequality

$$\Delta_2 \leq \frac{(1-p)Y_b}{1 - X - pY_b}\Delta_1. \tag{13}$$

It follows that

$$\Delta_1 \leq \frac{pY_s}{1 - X - (1-p)Y_s}\Delta_2 \leq \frac{pY_s}{1 - X - (1-p)Y_s} \times \frac{(1-p)Y_b}{1 - X - pY_b}\Delta_1,$$

which implies that

$$\left(1 - \frac{pY_s}{1 - X - (1-p)Y_s} \times \frac{(1-p)Y_b}{1 - X - pY_b}\right)\Delta_1 \leq 0. \tag{14}$$

If $\Delta_1 = 0$, then (13) implies that $\Delta_2 = 0$ and hence $u_i(S) = u'_i(S)$ for all $i \in N$—a contradiction. Therefore, $\Delta_1 > 0$, which along with (14) leads to

$$\frac{pY_s}{1 - X - (1-p)Y_s} \times \frac{(1-p)Y_b}{1 - X - pY_b} \geq 1. \tag{15}$$

As $1 - X - Y_s \geq 0$, we have $pY_s/(1 - X - (1-p)Y_s) \leq 1$, with equality if and only if $1 - X - Y_s = 0$, which means that the total probability of an agreement that does not involve player s under the profile $\alpha(S)$ is 0. Similarly, $(1-p)Y_b/(1 - X - pY_b) \leq 1$, with equality if and only if $\alpha(S)$ places positive probability only on links in $\mathcal{L}(S)$ that involve node b. Thus, (15) holds if and only if $\alpha(S)$ places positive probability only on the link bs. Then, constraint (7) in the definition of bargaining solutions implies that the payoffs u and u' in state S must satisfy

$$\begin{aligned}
u_s(S) &= p\left(v_b + u_b(S \cup b) + u_s(S \cup b)\right) \\
u_b(S) &= (1-p)\left(v_b + u_b(S \cup b) + u_s(S \cup b)\right) \\
u'_s(S) &= p\left(v_b + u'_b(S \cup b) + u'_s(S \cup b)\right) \\
u'_b(S) &= (1-p)\left(v_b + u'_b(S \cup b) + u'_s(S \cup b)\right).
\end{aligned}$$

Condition (10) leads to $u_s(S) = u'_s(S)$ and $u_b(S) = u'_b(S)$. The choice of s and b implies that $\Delta_1 = \Delta_2 = 0$, so $u_i(S) = u'_i(S)$ for all $i \in N$, which contradicts the definition of S. □

Proof of claim regarding Bertrand competition for the network in Figure 1. Suppose that s trades with b first, and consider the ensuing seller configuration $S = \{b, s\}$. We show that $u_{b'}(S) = v_{b'}$. By (6), we have that $u_b(N) = u_{b'}(N) = u_s(N) = 0$. The payoff of buyer b' in state S solves equation (4):

$$u_{b'}(S) = u_{b'}(S) + (1-p)\left(\pi_{bb'}(S)\alpha_{bb'}(S)w_{bb'}(S) + \pi_{b's}(S)\alpha_{b's}(S)w_{b's}(S)\right).$$

Since $1 - p > 0, \pi_{bb'}(S) > 0, \pi_{b's}(S) > 0$ and the incentive constraints (2) imply that $\alpha_{bb'}(S)w_{bb'}(S) \geq 0$ and $\alpha_{b's}(S)w_{b's}(S) \geq 0$, it must be that $\alpha_{bb'}(S)w_{bb'}(S) = \alpha_{b's}(S)w_{b's}(S) = 0$. The payoff equations (3) for players b and s in state S, along with $u_b(N) = u_s(N) = 0$, imply that

$$\begin{aligned}
u_b(S) &= \pi_{bb'}(S)u_b(S) + \pi_{b's}(S)(1 - \alpha_{b's}(S))u_b(S) \\
u_s(S) &= \pi_{b's}(S)u_s(S) + \pi_{bb'}(S)(1 - \alpha_{bb'}(S))u_s(S),
\end{aligned}$$

which reduce to $u_b(S)\pi_{b's}(S)\alpha_{b's}(S) = u_s(S)\pi_{bb'}(S)\alpha_{bb'}(S) = 0$. Since $\pi_{b's}(S) > 0$ and $\pi_{bb'}(S) > 0$, we have $u_b(S)\alpha_{b's}(S) = u_s(S)\alpha_{bb'}(S) = 0$. Condition (5) requires that $\alpha_{bb'}(S) > 0$ or $\alpha_{b's}(S) > 0$. Without loss of generality, assume that $\alpha_{bb'}(S) > 0$. In that case, we have $u_s(S) = 0$. If $\alpha_{b's}(S) = 0$, then constraint (7) leads to $u_b(S) = pv_{b'}$ and $u_{b'}(S) = (1-p)v_{b'}$. We obtain $w_{b's}(S) = v_{b'} - u_{b'}(S) - u_s(S) = pv_{b'} > 0$. Then (2) implies that $\alpha_{b's}(S) = 1$, a contradiction with the assumption that $\alpha_{b's}(S) = 0$. Hence, we also have that $\alpha_{b's}(S) > 0$, which leads to $u_b(S) = 0$. The conditions $\alpha_{bb'}(S) > 0$ and $\alpha_{bb'}(S)w_{bb'}(S) = 0$ imply that $w_{bb'}(S) = v_{b'} - u_b(S) - u_{b'}(S) = v_{b'} - u_{b'}(S) = 0$, so $u_{b'}(S) = v_{b'}$, as claimed. □

Proof of Lemma 1. Let $\delta(i, j)$ denote the distance between nodes i and j in network H. Suppose, by contradiction, that $\sim_H$ is not an equivalence relation. Pick a triple (x, y, z) with $x \sim_H y, y \sim_H z, x \not\sim_H z$ that minimizes the expression $\delta(x, y) + \delta(y, z)$. If there were any common node $t \neq y$ on the unique paths from x to y and y to z, respectively, then $x \sim_H t$ and $t \sim_H z$ and $\delta(x, t) + \delta(t, z) < \delta(x, y) + \delta(y, z)$. Hence, (x, t, z) would contradict the minimality of the counterexample (x, y, z). Thus, y is the only common node of the paths from x to y and y to z. This implies the existence of a path P from x to z obtained by appending the path from x to y to the one from y to z.

Since $x \not\sim_H z$, there exists an alternative path Q between x and z that excludes at least one of the links ij in P. Without loss of generality, assume that ij belongs to the path between x and y. Let $\tilde{H}$ denote the network obtained by removing link ij from H. It must be that y and z belong to the same connected component of $\tilde{H}$, as the path connecting them in H overlaps only at node y with the path between x and y in H and is thus contained in $\tilde{H}$. Since ij does not belong to Q, nodes x and z also belong to the same connected component in $\tilde{H}$. Thus, x and y must lie in the same connected component of $\tilde{H}$, which means that there exists a path between x and y in $\tilde{H}$. By definition, this path lies in H and excludes link ij, contradicting the fact that ij belongs to the unique path between x and y in H.

The second part of the lemma follows from the observation that if node k belongs to the unique path connecting nodes i to j in H, then the subpath of this path between i and k is the only path connecting i to k in H. □

Proof of Lemma 2. To prove the first statement, assume first that $b \sim_{G(S)} s$. Then, there exists a unique path P between b and s in $G(S)$. Since G is a connected subnetwork of $G(S)$, P must also be the unique path between b and s in G. This shows that b and s satisfy the first condition required for s to be the essential supplier for b in state S. Lemma 3 (whose proof does not rely on the current result) implies that s is the dealer of $C_b(S)$ in state S and any path between a node in S and b passes through s, which is the second necessary condition for s to be the essential supplier for b in state S. We have established that the relationship $b \sim_{G(S)} s$ implies that s is the essential supplier for b in state S.

Suppose next that $b \not\sim_{G(S)} s$. Then, there exist two distinct paths between b and s in $G(S)$. If neither of these paths contains a node from $S \cup 0$ different from s, then both paths are contained in G, which means that the first necessary condition for s to serve as the essential supplier for b in state S is violated. If one of the paths contains a node from $S \cup 0$ different from s, then the node with this property that is closest to b along that path must be an element of $S \setminus s$ (node 0 is linked only to nodes in S) and the subpath connecting that node to b not contain s. Then b and s do not satisfy the second condition required for s to be the essential supplier for b in state S. Therefore, if $b \not\sim_{G(S)} s$, then s is not the essential supplier for b in state S.

The second statement of the result follows from the first part and the observation that b is an essential intermediary for b' in state S if and only if b is the essential supplier for b' in state $S \cup b$. □

Proof of Lemma 3. Fix a seller configuration S, a seller $s \in S$, and a player $i \in N$. Since G is a connected network, it contains at least one path connecting s to i (if $s = i$, this is the degenerate path formed by the single node i and no links). Let x be the first element of $C_i(S)$ along the path, and let P denote the subpath between s and x (if $s \in C_i(S)$, then $x = s$ and P is the degenerate path consisting solely of node s). We argue that x is the first point of intersection with $C_i(S)$ of any other path in G from a node in S to a node in $C_i(S)$.

We proceed by contradiction. If the claim is not true, then there exists a path Q in G that connects a node $s' \in S$ to a node $y \neq x$ in $C_i(S)$ and contains no other node from $C_i(S)$. If there are nodes that belong to both P and Q, let z be the common node that is the smallest number of links away from x along P. Since by construction x is the only node from $C_i(S)$ contained in P and similarly $y \neq x$ is the only node from $C_i(S)$ contained in Q, we have that $z \notin C_i(S)$. Then we can form a path from x to y in G by following P from x to z and subsequently Q from z to y. As $x \sim_{G(S)} i \sim_{G(S)} y$, the resulting path must be the unique path connecting x and y in $G(S)$. By Lemma 1, any node along this path, including z, must belong to $C_i(S)$—a contradiction.

If P and Q do not have any nodes in common, then $s \neq s'$ and we can construct a path between x and y in $G(S)$ by appending the sequence of links from x to s in P with the link $ss' \in G(S)$ and subsequently the links between s' and y in Q. Since $x, y \in C_i(S)$, Lemma 1 implies that all the nodes along this path, including s and s', must belong to $C_i(S)$. Then $s \sim_{G(S)} s'$, which is impossible for $s \neq s' \in S$.

We are left to prove the second part of the lemma. Suppose that $x \notin S$. If x has a single potential supplier y in state S, then every path from S to x contains the link xy. This implies that the only path between x and y in G is the link xy (otherwise, we could construct a path from S to x that does not include the link xy). It follows that y should be an essential intermediary (or supplier) for x in state S, so $x \sim_{G(S \cup y)} y$ by Lemma 2. However, $x \sim_{G(S \cup y)} y$ implies that $x \sim_{G(S)} y$ and $y \in C_x(S)$. Since y is a potential supplier for x in state S, there exists a path in G from a node in S to $y \in C_x(S)$ that does not contain node x. This contradicts the finding that x should be the first point of intersection with $C_x(S)$ of any path in G from a node in S to a node in $C_x(S)$. The contradiction proves that x must have at least two potential suppliers in state S.

Consider now a buyer $b \in C_x(S) \setminus x$. Any path from a node in S to b must include node x. Since $b \sim_{G(S)} x$, there exists a unique path from x to b in G. Therefore, any path from S

to b must contain the unique path between x and b. The node preceding b on this path is the unique potential supplier for b. $\square$

Proof of Proposition 2. We first prove that if $b \not\sim_{G(S)} s$, then $C_i(S \cup b) = C_i(S)$ for all $i \in N$. Since $G(S) \subset G(S \cup b)$, it must be that $C_i(S \cup b) \subseteq C_i(S)$. For a proof by contradiction, suppose that there exists $i \in N$ for which $C_i(S \cup b) \neq C_i(S)$, so that we can find $j \in C_i(S)$ with $j \notin C_i(S \cup b)$. The condition $j \in C_i(S)$ implies the existence of a unique path P between i and j in $G(S)$, which contains only nodes from $C_i(S)$. Since $j \notin C_i(S \cup b)$, there must be a path P' distinct from P between i and j in $G(S \cup b)$. As P is the unique path between i and j in $G(S)$, P' must contain some links from the set $G(S \cup b) \setminus G(S) \subset \{bs' | s' \in S \cup 0\}$. All such links include b, so P' involves either two links $bs', bs'' \in G(S \cup b) \setminus G(S)$ or a single such link $bs' \in G(S \cup b) \setminus G(S)$. We consider each of these cases in turn.

If P' contains two links bs' and bs'' with $s', s'' \in S \cup 0$, we can replace them with the link $s's'' \in G(S)$ to obtain another path P'' connecting i to j in $G(S)$. Since P is the unique such path, it must be that P'' is identical to P. Hence P contains s' and s'', which means that $s' \sim_{G(S)} s''$. Since all nodes in $S \cup 0$ are mutually linked, $s' \sim_{G(S)} s''$ is only possible if S contains a single seller, so $S = \{s\}$ and $\{s', s''\} = \{0, s\}$. However, node $0 \in \{s', s''\}$ cannot belong to P since $i, j \neq 0$ and 0 has a single link in $G(\{s\})$, namely the link with s.

Suppose instead that P' contains a single link $bs' \in G(S \cup b) \setminus G(S)$. If s does not belong to P', then we can replace the link bs' with the pair of links $bs, ss' \in G(S)$ to obtain a path P'' connecting i to j in $G(S)$. It must be that P'' coincides with P. By an argument similar to the one above, we need $s \sim_{G(S)} s' = 0$ and $S = \{s\}$. We reach a contradiction using the fact that $i, j \neq 0$ and node 0 has a single link in $G(\{s\})$. Thus, s must belong to P'. Note that $s \neq s'$ since $bs \in G(S)$, while $bs' \notin G(S)$. We construct a path P'' by replacing the portion of P' between s and s' with the link $ss' \in G(S)$. If P'' is contained in $G(S)$, we obtain a contradiction as before. Therefore, P'' must include the link bs'. We can now replace the links bs' and ss' in P'' with the link $bs \in G(S)$ to obtain another path P'''. Since P''' connects i to j using only links in $G(S)$, it must be that P''' is identical to P. Then $P''' = P$ contains the link bs, which implies that $b \sim_{G(S)} s$—a contradiction with our initial assumption.

We now turn to the case $b \sim_{G(S)} s$. The proof that $C_i(S \cup b) = C_i(S)$ for all $i \in N \setminus C_s(S)$ follows exactly the same steps as in the case $b \not\sim_{G(S)} s$ except for the final contradiction, which is reached by noting that since the path $P''' = P$ from i to j in $G(S)$ contains the link bs and $i \sim_{G(S)} j$ by assumption, we have $i \sim_{G(S)} s$ or, equivalently, $i \in C_s(S)$.

We are left to prove that if $b \sim_{G(S)} s$, then $b \not\sim_{G(S \cup b)} s$ and $C_s(S \cup b) \cup C_b(S \cup b) = C_s(S)$. Since b and s are directly linked in $G(S) \subset G(S \cup b)$ and are also connected by the path $(b, 0, s)$ in $G(S \cup b)$, we have $b \not\sim_{G(S \cup b)} s$. Hence $C_s(S \cup b) \cap C_b(S \cup b) = \emptyset$. Clearly, $C_s(S \cup b) \subset C_s(S)$ and $C_b(S \cup b) \subset C_s(S)$. To establish that $C_s(S \cup b) \cup C_b(S \cup b) = C_s(S)$, we need to show that for every $i \in C_s(S)$, either $i \in C_b(S \cup b)$ or $i \in C_s(S \cup b)$. Fix $i \in C_s(S)$. Then $b, s \in C_s(S)$

implies that $G(S)$ contains a unique path P from i to b and similarly a unique path Q from i to s. If node s does not belong to P, then we can augment P by adding the link bs to obtain a path from i to s in $G(S)$. This path must coincide with Q, and hence Q contains the link bs. Similarly, if b does not belong to Q, then P should contain the link bs.

Suppose that Q contains the link bs. We set out to prove that $i \in C_b(S \cup b)$. If this is not the case, there is a path P' distinct from P connecting i to b in $G(S \cup b)$. This path must contain a link $bs' \in G(S \cup b) \setminus G(S)$ with $s' \in S \cup 0$. If node s belongs to P', then the subpath of P' from i to s excludes b. Hence, this subpath lies in $G(S)$ and has to be identical to the unique path Q from i to s in $G(S)$. However, Q contains node b by assumption, which means that P' passes through b twice, a contradiction. This reasoning proves that s does not belong to P'. If we replace the link bs' in P' with the link $ss' \in G(S)$, we obtain a path Q' that lies in $G(S)$ and connects i to s. It follows that Q' coincides with Q. Since Q' does not contain b, neither should Q, a contradiction with the hypothesis that Q includes the link bs.

Finally, assume that P contains the link bs. Suppose, by contradiction, that $i \notin C_s(S \cup b)$. Then, there exists a path Q' that connects i to s in $G(S \cup b)$ and includes node b with links in $G(S \cup b) \setminus G(S)$. We construct a path Q'' by replacing the subpath between b and s in Q' with the link bs. If Q'' lies entirely within $G(S)$, then $Q'' = Q$ and b is the neighbor of s in Q. However, in that case the subpath of Q from i to b must be identical to P, so it contains the link bs by assumption. Hence the link bs appears on the path Q twice, a contradiction which implies that Q'' includes a link $bs' \in G(S \cup b) \setminus G(S)$ with $s' \in S \cup 0$. If we modify Q'' by replacing its links bs and bs' with the link $ss' \in G(S)$ we obtain a path Q''' in $G(S)$ that connects i to s. It must be that $Q''' = Q$, which leads to the conclusion that $s \sim_{G(S)} s' = 0$ and $S = \{s\}$ as above, contradicting the fact that node $s' = 0$ has a single link in $G(\{s\})$ and appears on the path Q from $i \neq 0$ to $s \neq 0$. □

Proof of Theorem 1. We establish that the payoffs u defined by equation (9) along with any profile of agreement probabilities α such that $\alpha_{bs}(S) > 0$ for all $bs \in \mathcal{L}(S)$ and $S \in \mathcal{S}$ constitute a bargaining solution. Proposition 1 then implies that u represents the payoff profile in all bargaining solutions that satisfy the refinement.

The following properties of the payoffs u for $S \in \mathcal{S}$ are central to the proof:

(a) $u_s(S) = u_s(S \cup b')$ whenever $b's' \in \mathcal{L}(S)$ and $s \neq s' \in S$;
(b) $u_b(S) = u_b(S \cup b')$ whenever $b's' \in \mathcal{L}(S)$ and $b' \neq b \notin S$;
(c) $v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S) = 0$ for all $bs \in \mathcal{L}(S)$;
(d) if $\mathcal{L}(S) = \{bs\}$, then $u_s(S) = p(v_b + u_b(S \cup b) + u_s(S \cup b))$.

For claim (a), we need to show that if $b's' \in \mathcal{L}(S)$ and $s \neq s' \in S$, then $u_s(S) = u_s(S \cup b')$. As $s \in S$, this is equivalent to $r_s(S) = r_s(S \cup b')$. To prove this identity, it is sufficient to show that $C_s(S) = C_s(S \cup b')$. By Proposition 2, adding b' to S following his agreement with s' can only affect the equivalence class of $\sim_{G(S)}$ that contains s'. Then $s \neq s' \in S$

and $s \not\sim_{G(S)} s'$ imply that the equivalence class of s is identical under $\sim_{G(S)}$ and $\sim_{G(S\cup b)}$, so $C_s(S) = C_s(S \cup b')$, as desired.

For claim (b), we must show that $u_b(S) = u_b(S \cup b')$ for $b's' \in \mathcal{L}(S)$ with $b' \neq b \notin S$. We first argue that $C_b(S \cup b) = C_b(S \cup b \cup b')$, which implies that $r_b(S \cup b) = r_b(S \cup b \cup b')$. Since both b and s' are sellers in state $S \cup b$, we have $b \not\sim_{G(S\cup b)} s'$. Then an agreement between s' and b' in state $S \cup b$, which leads to state $S \cup b \cup b'$, cannot affect the equivalence class of b, so $C_b(S \cup b) = C_b(S \cup b \cup b')$, as desired. Given the definition of u_b, establishing that $u_b(S) = u_b(S \cup b')$ reduces to showing that either $d(S, C_b(S)) = d(S \cup b', C_b(S \cup b')) = b$ or $d(S, C_b(S)) \neq b \neq d(S \cup b', C_b(S \cup b'))$. We proceed by considering two possible cases separately: $b \not\sim_{G(S)} s'$ and $b \sim_{G(S)} s'$.

If $b \not\sim_{G(S)} s'$, then the equivalence class of b remains unchanged when b' joins S, so $C_b(S) = C_b(S \cup b')$. Hence, $d(S \cup b', C_b(S \cup b')) = d(S, C_b(S))$ because Lemma 3 implies that $d(S \cup b', C_b(S \cup b'))$ represents the only node in $C_b(S \cup b') = C_b(S)$ that belongs to all paths from $S \cup b'$ to $C_b(S \cup b')$ in G, and there exists a path from s' to $C_b(S)$ in G whose only intersection with $C_b(S)$ is $d(S, C_b(S))$. Since $r_b(S \cup b) = r_b(S \cup b' \cup b)$ and $d(S, C_b(S)) = d(S \cup b', C_b(S \cup b'))$, the definition of u_b implies that $u_b(S) = u_b(S \cup b')$.

If instead $b \sim_{G(S)} s'$, then either $b \sim_{G(S\cup b')} s'$ or $b \sim_{G(S\cup b')} b'$. In the former case, $d(S \cup b', C_b(S \cup b')) = s'$, while in the latter $d(S \cup b', C_b(S \cup b')) = b'$ since both s and b' are sellers in the new configuration $S \cup b'$. As $b \notin \{b', s'\}$, we have $d(S, C_b(S)) \neq b \neq d(S \cup b', C_b(S \cup b'))$ in either case. Since $r_b(S \cup b) = r_b(S \cup b' \cup b)$ and $d(S, C_b(S)) \neq b \neq d(S \cup b', C_b(S \cup b'))$, the definition of u_b implies that $u_b(S) = u_b(S \cup b') = (1 - p)\,(v_b + r_b(S \cup b))$.

To prove claim (c), consider first a link $bs \in \mathcal{L}(S)$ with $b \not\sim_{G(S)} s$. An agreement between b and s leaves all equivalence classes unchanged, i.e., $\sim_{G(S)}$ and $\sim_{G(S\cup b)}$ represent the same equivalence relation. In particular, $C_b(S \cup b) = C_b(S)$ and $C_s(S \cup b) = C_s(S)$. Hence $u_s(S \cup b) = r_s(S \cup b) = r_s(S) = u_s(S)$. Moreover, since s is linked to b, it must be that $d(S, C_b(S)) = b$, which means that $u_b(S) = v_b + r_b(S \cup b)$. Since b is a seller in the configuration $S \cup b$, we have by definition that $u_b(S \cup b) = r_b(S \cup b)$. It follows that

$$v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S) = v_b + r_b(S \cup b) + r_s(S) - (v_b + r_b(S \cup b)) - r_s(S) = 0.$$

Assume next that $bs \in \mathcal{L}(S)$ with $b \sim_{G(S)} s$. Then, an agreement between b and s splits s's equivalence class into two classes, $C_s(S) = C_s(S \cup b) \cup C_b(S \cup b)$. Since b and s are sellers in the configuration $S \cup b$, we have by definition that

$$\begin{aligned}
u_s(S) &= r_s(S) = \sum_{i \in C_s(S)\setminus s} p^{\delta(i,s)} v_i \\
u_s(S \cup b) &= r_s(S \cup b) = \sum_{i \in C_s(S\cup b)\setminus s} p^{\delta(i,s)} v_i \\
u_b(S \cup b) &= r_b(S \cup b) = \sum_{i \in C_b(S\cup b)\setminus b} p^{\delta(i,b)} v_i.
\end{aligned}$$

By Lemma 2, b is an essential intermediary and s is the essential supplier in state S for the buyers in $C_b(S \cup b) \setminus b$. Hence, for all $i \in C_b(S \cup b) \setminus b$, the link bs belongs to the unique path connecting s to i and $\delta(i, s) = \delta(i, b) + 1$. Since $C_s(S) \setminus s = (C_s(S \cup b) \setminus s) \cup b \cup (C_b(S \cup b) \setminus b)$, $\delta(b, s) = 1$, and $\delta(i, s) = \delta(i, b) + 1$ for $i \in C_b(S \cup b) \setminus b$, the formula for $u_s(S)$ can be rewritten as follows:

$$u_s(S) = \sum_{i \in C_s(S \cup b) \setminus s} p^{\delta(i,s)} v_i + p v_b + \sum_{i \in C_b(S \cup b) \setminus b} p^{\delta(i,s)} v_i = r_s(S \cup b) + p v_b + \sum_{i \in C_b(S \cup b) \setminus b} p^{\delta(i,b)+1} v_i$$
$$= r_s(S \cup b) + p v_b + p \sum_{i \in C_b(S \cup b) \setminus b} p^{\delta(i,b)} v_i = r_s(S \cup b) + p(v_b + r_b(S \cup b)).$$

As $s \in C_b(S)$, we have $d(S, C_b(S)) = s$, and hence by definition,

$$u_b(S) = (1 - p)(v_b + r_b(S \cup b)).$$

The equalities above imply that

$$v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S)$$
$$= v_b + r_b(S \cup b) + r_s(S \cup b) - (1 - p)(v_b + r_b(S \cup b)) - (r_s(S \cup b) + p(v_b + r_b(S \cup b))) = 0.$$

For a proof of claim (d), suppose that $\mathcal{L}(S) = \{bs\}$. Then, seller s has no neighbor left to sell to when all players in $S \cup b$ have the good. Hence, $C_s(S \cup b) = \{s\}$ and $u_s(S \cup b) = r_s(S \cup b) = 0$. Since $\mathcal{L}(S) = \{bs\}$, we have $b \sim_{G(S)} s$, which implies that $C_b(S \cup b) = C_s(S) \setminus C_s(S \cup b) = C_s(S) \setminus s$. As $\delta(i, s) = 1 + \delta(i, b)$ for all $i \in N \setminus S$, it follows that

$$u_s(S) = r_s(S) = \sum_{i \in C_s(S) \setminus s} p^{\delta(i,s)} v_i = \sum_{i \in C_b(S \cup b)} p^{\delta(i,s)} v_i$$
$$= p v_b + \sum_{i \in C_b(S \cup b) \setminus b} p^{1+\delta(i,b)} v_i = p v_b + p \sum_{i \in C_b(S \cup b) \setminus b} p^{\delta(i,b)} v_i = p(v_b + u_b(S \cup b)).$$

Then, $u_s(S \cup b) = 0$ leads to $u_s(S) = p\,(v_b + u_b(S \cup b) + u_s(S \cup b))$, as asserted.

Consider now a profile (u, α) satisfying the hypotheses of the theorem. To prove that (u, α) is a bargaining solution, fix a state $S \in \mathcal{S}$. Claim (c) implies that $w_{bs}(S) = 0$ for all $bs \in \mathcal{L}(S)$. Hence, (u, α) satisfies the incentive constraints (2). Claims (a), (b) and (c) imply that the profile (u, α) solves the payoff equations (3) and (4). If $S \neq N$, then the set $\mathcal{L}(S)$ is nonempty because the network G is assumed to be connected. Thus, the agreement profile α meets the requirement (5) since it assigns positive probability of agreement for every link in $\mathcal{L}(S)$. By construction, the payoffs u satisfy condition (6). Finally, to verify that (u, α) has property (7), suppose that $\alpha_{bs}(S) > 0$ for a single link $bs \in \mathcal{L}(S)$. As α specifies a positive probability of agreement for any trading link in every state, it must be that $\mathcal{L}(S) = \{bs\}$. Claim (d) then implies (7). We have shown that (u, α) satisfies conditions (2)-(7) for every state $S \in \mathcal{S}$ and thus constitutes a bargaining solution. The proof is concluded as outlined in the preamble. $\square$

Proof of Proposition 3. We first show that the refinement of the bargaining solution generates history-independent prices. Let u^* be the payoffs under the refinement with associated gains from trade and prices denoted by w^* and t^*, respectively. Step (c) in the proof of Theorem 1 shows that for all $S \in \mathcal{S}$ and $bs \in \mathcal{L}(S)$, we have $w^*_{bs}(S) = 0$ and thus $t^*_{bs}(S) = u^*_s(S) - u^*_s(S \cup b)$. To establish history-independence of prices under u^*, it is sufficient to argue that $t^*_{bs}(S) = t^*_{bs}(S \cup b')$ for any $b' \in N \setminus (S \cup b)$ such that $S \cup b' \in \mathcal{S}$. Fix b, b', s, S with the properties listed above. We have to check that the payoffs selected by the refinement solve the equation $u^*_s(S) - u^*_s(S \cup b) = u^*_s(S \cup b') - u^*_s(S \cup \{b, b'\})$, or equivalently, that $r_s(S) - r_s(S \cup b) = r_s(S \cup b') - r_s(S \cup \{b, b'\})$. Given the formula for r, the latter equation is equivalent to

$$\sum_{i \in C_s(S) \setminus C_s(S \cup b)} p^{\delta(i,s)} v_i = \sum_{i \in C_s(S \cup b') \setminus C_s(S \cup \{b,b'\})} p^{\delta(i,s)} v_i.$$

Therefore, it is sufficient to prove that

$$C_s(S) \setminus C_s(S \cup b) = C_s(S \cup b') \setminus C_s(S \cup \{b, b'\}). \tag{16}$$

If $b' \not\sim_{G(S)} s$, then Proposition 2 implies that $C_s(S) = C_s(S \cup b')$. Moreover, $b' \not\sim_{G(S \cup b)} s$ and Proposition 2 also leads to the conclusion that $C_s(S \cup b) = C_s(S \cup \{b, b'\})$. Hence, (16) holds in this case.

If $b \not\sim_{G(S)} s$, then Proposition 2 implies that $C_s(S) = C_s(S \cup b)$, so $C_s(S) \setminus C_s(S \cup b) = \emptyset$. Moreover, $b \not\sim_{G(S \cup b')} s$ and Proposition 2 also leads to $C_s(S \cup b') = C_s(S \cup \{b, b'\})$, which means that $C_s(S \cup b') \setminus C_s(S \cup \{b, b'\}) = \emptyset$. Hence, (16) holds in this case as well.

We are left with the case $b \sim_{G(S)} s \sim_{G(S)} b'$. Since $S \cup b' \in \mathcal{S}$, it must be that b' is linked to a node in S. By Lemma 2, the relationship $b' \sim_{G(S)} s$ implies that s is the essential supplier for b' in state S and thus belongs to any path from a node in S to b', including any link connecting b' to S. It follows that $b's \in G$. Since $b \sim_{G(S)} s$, Proposition 2 implies that $C_s(S) \setminus C_s(S \cup b) = C_b(S \cup b)$. Note that $b \not\sim_{G(S \cup b')} b'$ because b and b' are connected by the paths (b, s, b') and $(b, s, 0, b')$ in $G(S \cup b')$. Applying Proposition 2 again, we have $C_s(S) = C_s(S \cup b') \cup C_{b'}(S \cup b')$. As $b \in C_s(S)$ but $b \notin C_{b'}(S \cup b')$, we infer that $b \in C_s(S \cup b')$ and thus $b \sim_{G(S \cup b')} s$. Proposition 2 leads to $C_s(S \cup b') \setminus C_s(S \cup \{b, b'\}) = C_b(S \cup \{b, b'\})$. Then (16) follows from the fact that $C_b(S \cup b) = C_b(S \cup \{b, b'\})$, which is a consequence of step (b) in the proof of Theorem 1.

We next prove that every bargaining solution with history-independent prices must generate the payoffs selected by the refinement. Fix a bargaining solution (u, α) under which prices are history-independent. We need to show that $u(S) = u^*(S)$ for every $S \in \mathcal{S}$. The proof of this claim proceeds by induction on $|N \setminus S|$. For the base case $|N \setminus S| = 0$, we have that $S = N$, and the claim follows trivially from assumption (6).

For the inductive step, fix $S \subset N$ and assume that the induction hypothesis holds for every set in $\mathcal{S}$ of greater cardinality than S. In particular, $u(S \cup b) = u^*(S \cup b)$ for every $b \in N \setminus S$ that is linked to a node in S. Since G is connected and $S \subset N$, there exists at least one node $b \in N \setminus S$ such that $S \cup b \in \mathcal{S}$. We consider two cases, depending on whether there exists only one such node or there are multiple ones.

First, assume that there exists only one $b \in N \setminus S$ such that $S \cup b \in \mathcal{S}$. Then, all links in $\mathcal{L}(S)$ contain node b. In this case, the payoff equations along with condition (5) imply that $u_{b'}(S) = u_{b'}(S \cup b)$ and $u^*_{b'}(S \cup b) = u^*_{b'}(S)$ for all $b' \in N \setminus (S \cup b)$. Since $u_{b'}(S \cup b) = u^*_{b'}(S \cup b)$ by the induction hypothesis, it follows that $u_{b'}(S) = u^*_{b'}(S)$ for all $b' \in N \setminus (S \cup b)$. Furthermore, $u_s(S) = u^*_s(S) = 0$ for all sellers s not linked to b in G.

The payoff equation for buyer b in state S leads to

$$u_b(S) = \sum_{s: bs \in \mathcal{L}(S)} \pi_{bs}(S)(u_b(S) + (1-p)\alpha_{bs}(S) w_{bs}(S)).$$

Since $\sum_{s: bs \in \mathcal{L}(S)} \pi_{bs}(S) = 1$ and $\pi_{bs}(S) > 0$ and $\alpha_{bs}(S) w_{bs}(S) \geq 0$ for $bs \in \mathcal{L}(S)$, it must be that $\alpha_{bs}(S) w_{bs}(S) = 0$ for all s such that $bs \in \mathcal{L}(S)$.

The payoff equation for any seller s in state S linked to node b in network G reduces to

$$\begin{aligned} u_s(S) &= \pi_{bs}(S)\left(u_s(S) + p\alpha_{bs}(S) w_{bs}(S)\right) \\ &+ \sum_{s' \neq s: bs' \in \mathcal{L}(S)} \pi_{bs'}(S)\left(\alpha_{bs'}(S) u_s(S \cup b) + (1 - \alpha_{bs'}(S)) u_s(S)\right) \\ &= \pi_{bs}(S) u_s(S) + \sum_{s' \neq s: bs' \in \mathcal{L}(S)} \pi_{bs'}(S)(1 - \alpha_{bs'}(S)) u_s(S), \end{aligned}$$

where we took into account that $\alpha_{bs}(S) w_{bs}(S) = 0$ and $u_s(S \cup b) = 0$ (s is not linked to any buyer in state $S \cup b$). It follows that

$$u_s(S) \sum_{s' \neq s: bs' \in \mathcal{L}(S)} \pi_{bs'}(S) \alpha_{bs'}(S) = 0,$$

which is possible only if either $u_s(S) = 0$ or $\alpha_{bs'}(S) = 0$ for all $s' \neq s$ such that $bs' \in \mathcal{L}(S)$.

Suppose first that

$$\exists s \in S \text{ s.t. } bs \in \mathcal{L}(S) \text{ and } \alpha_{bs'}(S) = 0, \forall s' \neq s \text{ with } bs' \in \mathcal{L}(S). \tag{17}$$

Then, constraint (5) implies that there exists exactly one s satisfying this condition and $\alpha_{bs}(S) > 0$. Assumption (7) leads to $u_s(S) = p(v_b + u_b(S \cup b) + u_s(S \cup b)) = p(v_b + u_b(S \cup b))$ and $u_b(S) = (1-p)(v_b + u_b(S \cup b) + u_s(S \cup b)) = (1-p)(v_b + u_b(S \cup b))$.

If $\mathcal{L}(S) = \{bs\}$, then we also have that $u^*_s(S) = p(v_b + u^*_b(S \cup b) + u^*_s(S \cup b))$ and $u^*_b(S) = (1-p)(v_b + u^*_b(S \cup b) + u^*_s(S \cup b))$, which along with the induction hypothesis implies that $u_s(S) = u^*_s(S)$ and $u_b(S) = u^*_b(S)$.

We are left to consider the case $|\mathcal{L}(S)| \geq 2$. In this case, there exists $s' \in S \setminus s$ such that $bs' \in \mathcal{L}(S)$ and $\alpha_{bs'}(S) = 0$. As argued above, $\alpha_{bs}(S) > 0$ implies that $u_{s'}(S) = 0$. Hence, $u_{s'}(S) = u_{s'}(S \cup b) = 0$. Since $\alpha_{bs'}(S) = 0$, we have $w_{bs'}(S) \leq 0$, and thus $v_b + u_b(S \cup b) + u_{s'}(S \cup b) - u_b(S) - u_{s'}(S) = v_b + u_b(S \cup b) - u_b(S) \leq 0$. Then, we have

$v_b + u_b(S \cup b) \leq u_b(S) = (1-p)(v_b + u_b(S \cup b))$, which contradicts the conditions $p > 0$, $v_b > 0$, and $u_b(S \cup b) \geq 0$. We demonstrated that (17) implies that $|\mathcal{L}(S)| = 1$ and $u(S) = u^*(S)$.

Suppose next that statement (17) is false. Then, it must be that $u_s(S) = 0$ for all $s \in S$, $|\mathcal{L}(S)| \geq 2$, and b is a dealer in state S, while each seller forms a singleton equivalence class in $G(S)$. It follows that $u_s(S) = 0 = u_s^*(S)$ for all $s \in S$. There exists $s \in S$ with $bs \in \mathcal{L}(S)$ such that $\alpha_{bs}(S) > 0$, which implies that $w_{bs}(S) = 0$. For such an s, we have $u_b(S) = v_b + u_b(S \cup b) + u_s(S \cup b) - u_s(S) = v_b + u_b^*(S \cup b) = u_b^*(S)$. The second equality relies on $u_b(S \cup b) = u_b^*(S \cup b)$ (induction hypothesis) and $u_s(S) = u_s(S \cup b) = 0$, while the third follows from the dealer status of buyer b in state S. We have shown that the negation of (17) implies that $u(S) = u^*(S)$, which completes the proof of the inductive step for the case in which $S \cup b \in \mathcal{S}$ for a single $b \in N \setminus S$.

Finally, consider the case in which there exist $b \neq b' \in N \setminus S$ with the property that $S \cup b, S \cup b' \in \mathcal{S}$. For such pairs (b, b'), the induction hypothesis implies that $t_{bs}(S \cup b') = t_{bs}^*(S \cup b')$ whenever $bs \in \mathcal{L}(S)$. History independence of prices under u and u^* requires that $t_{bs}(S) = t_{bs}(S \cup b')$ and $t_{bs}^*(S) = t_{bs}^*(S \cup b')$, and hence, $t_{bs}(S) = t_{bs}^*(S)$ for $bs \in \mathcal{L}(S)$. We have shown that in this case, $t_{bs}(S) = t_{bs}^*(S)$ for every link $bs \in \mathcal{L}(S)$.

Fix $s \in S$. The payoff equation for seller s in state S can be rewritten as follows:

$$\begin{aligned} u_s(S) &= \sum_{b: bs \in \mathcal{L}(S)} \pi_{bs}(S) \left((1 - \alpha_{bs}(S)) u_s(S) + \alpha_{bs}(S)(u_s(S \cup b) + t_{bs}(S))\right) \\ &+ \sum_{bs' \in \mathcal{L}(S): s' \neq s} \pi_{bs'}(S) \left(\alpha_{bs'}(S) u_s(S \cup b) + (1 - \alpha_{bs'}(S)) u_s(S)\right). \end{aligned}$$

Since $t_{bs}(S) = t_{bs}^*(S)$ and $u_s(S \cup b) = u_s^*(S \cup b)$ in the equation above, we have

$$\begin{aligned} u_s(S) &= \sum_{b: bs \in \mathcal{L}(S)} \pi_{bs}(S) \left((1 - \alpha_{bs}(S)) u_s(S) + \alpha_{bs}(S)(u_s^*(S \cup b) + t_{bs}^*(S))\right) \\ &+ \sum_{bs' \in \mathcal{L}(S): s' \neq s} \pi_{bs'}(S) \left(\alpha_{bs'}(S) u_s^*(S \cup b) + (1 - \alpha_{bs'}(S)) u_s(S)\right). \end{aligned}$$

Recall that the payoffs u^* are consistent with any profile of agreement probabilities, including α. Therefore, we also have that

$$\begin{aligned} u_s^*(S) &= \sum_{b: bs \in \mathcal{L}(S)} \pi_{bs}(S) \left((1 - \alpha_{bs}(S)) u_s^*(S) + \alpha_{bs}(S)(u_s^*(S \cup b) + t_{bs}^*(S))\right) \\ &+ \sum_{bs' \in \mathcal{L}(S): s' \neq s} \pi_{bs'}(S) \left(\alpha_{bs'}(S) u_s^*(S \cup b) + (1 - \alpha_{bs'}(S)) u_s^*(S)\right). \end{aligned}$$

Subtracting the two equalities above and rearranging terms, we obtain

$$(u_s(S) - u_s^*(S)) \sum_{bs' \in \mathcal{L}(S)} \pi_{bs'}(S) \alpha_{bs'}(S) = 0.$$

Condition (5) implies that the summation in the equation above is positive, so it must be that $u_s(S) = u_s^*(S)$.

We have argued that $u_s(S) = u_s^*(S)$ for all $s \in S$. A similar logic proves that $u_b(S) = u_b^*(S)$ for all $b \in N \setminus S$ and completes the proof of the inductive step for the case under consideration. □

Proof of Theorem 2. The first two statements of the result have been proven in Section 6. To prove the third statement, consider a seller configuration $S \in \mathcal{S}$ and a buyer $b \in N \setminus S$. All paths in G connecting any seller in S to buyer b that intersect some equivalence class $C_i(S)$ must enter $C_i(S)$ via its dealer $d(S, C_i(S))$ and thus can cross $C_i(S)$ only once. If two paths in this collection exit $C_i(S)$ through nodes $x \neq y \in C_i(S)$, then we obtain the contradiction that $x \not\sim_{G(S)} y$ by "pasting" the subpaths from x to b and from b to y and eliminating potential overlap as in the proof of Lemma 3. Therefore, every path that connects a node in S to buyer b in G and intersects $C_i(S)$ must enter $C_i(S)$ via node $d(S, C_i(S))$ and exit through the same node x. Hence, all such paths must overlap in $C_i(S)$ with the unique path between $d(S, C_i(S))$ and x in G. □

References

1. Abreu, D. and M. Manea (2021) Bargaining with Exclusionary Commitments, mimeo.
2. Ahern, K.R. (2017) Information Networks: Evidence from Illegal Insider Trading Tips, *J. Fin. Econ.*, **125**, 26-47.
3. Ali, N.S., A. Chen-Zion, and E. Lillethun (2020), Reselling Information, mimeo.
4. Bakos, Y., E. Brynjolfsson, and D. Lichtman (1999) Shared Information Goods, *J. Law Econ.*, **42**, 117-156.
5. Besen, S.M. and S.N. Kirby (1989) Private Copying, Appropriability, and Optimal Copying Royalties, *J. Law Econ.*, **32**, 255-280.
6. Boldrin, M. and D.K. Levine (2002) The Case against Intellectual Property, *Amer. Econ. Rev.*, **92**, 209-212.
7. Boldrin, M. and D.K. Levine (2008) *Against Intellectual Monopoly*, Cambridge University Press.
8. Coase, R. (1972) Durability and Monopoly, *J. Law Econ.*, **5**, 143-149.
9. Condorelli, D. and A. Galeotti (2016) Strategic Models of Intermediation Networks, Ch. 27 in *The Oxford Handbook on the Economics of Networks* (eds: Y. Bramoulle, A. Galeotti, and B. Rogers), Oxford University Press.
10. Easley, D. and J. Kleinberg (2010) *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*, Cambridge University Press.
11. Frazer, G. (2006) Learning the master's trade: Apprenticeship and human capital in Ghana, *J. Dev. Econ.*, **81**, 259-298.
12. Jackson, M.O. (2008) *Social and Economic Networks*, Princeton University Press.
13. Johnson, J.P. and M. Waldman (2005) The Limits of Indirect Appropriability in Markets for Copiable Goods, *Rev. Econ. Res. Copyright*, **2**, 19-37.
14. Jovanovic, B. and Z. Wang (2020) Idea Diffusion and Property Rights, mimeo.
15. Liebowitz, S.J. (1985) Copying and Indirect Appropriability: Photocopying of Journals, *J. Pol. Econ.*, **93**, 945-957.
16. Malamud, S. and M. Rostek (2021) Bilateral Market Power in Fragmented Markets, mimeo.

17. Manea, M. (2016) Models of Bilateral Trade in Networks, Ch. 26 in *The Oxford Handbook on the Economics of Networks* (eds: Y. Bramoulle, A. Galeotti, and B. Rogers), Oxford University Press.
18. Manea, M. (2018) Intermediation and Resale in Networks, *J. Pol. Econ.*, **126**, 1250-1301.
19. Muto, S. (1986) An Information Good Market with Symmetric Externalities, *Econometrica*, 295-312.
20. Novos, I.E. and M. Waldman (1984) The Effects of Increased Copyright Protection: An Analytic Approach, *J. Pol. Econ.*, **92**, 236-46.
21. Polanski, A. (2007) A Decentralized Model of Information Pricing in Networks, *J. Econ. Theory*, **136**, 497-512.
22. Polanski, A. (2019) Communication Networks, Externalities and the Price of Information, *RAND J Econ*, **50**, 481-502.
23. Sarvary, M. (2011) *Gurus and Oracles: The Marketing of Information*, MIT Press.
24. Takeyama, L.N. (1994) The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities, *J. Ind. Econ.*, 155-166.
25. Varian, H.R. (2000) Buying, Sharing and Renting Information Goods, *J. Ind. Econ*, **48**, 473-488.