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

MIHAI MANEA

Department of Economics, MIT, manea@mit.edu

Abstract. We investigate how information goods are priced and diffused over links in a network. Buyers have idiosyncratic consumption values for information and, after acquiring it, can replicate it and resell copies to uninformed neighbors. A partition of the network captures the effects of network architecture and locations of information sellers on player profits and the structure of competing diffusion paths. Sellers indirectly appropriate profits over intermediation chains from buyers in their block of the partition. Links within blocks are critical for connecting the network and constitute bottlenecks for information diffusion. Links bridging distinct blocks are redundant for diffusion and impose negative externalities on sellers. Information enters each block not containing a seller via a single node—the dealer of the block. Dealers can receive information over redundant links from multiple neighbors and benefit from competitive pricing. Every non-dealer buyer can acquire information from a single neighbor via a bottleneck link and is subject to a monopoly. In dense networks, competition limits the scope of indirect appropriability, and intellectual property rights foster innovation.

Keywords: networks, diffusion, indirect appropriability, captive markets, intermediation, competition, bottlenecks, redundant links, information goods, copying, intellectual property.

1. Introduction

Information is often traded over links in a network. Digital goods (e.g., software, movies, music, and computer games) are replicated and sold in local markets or shared among friends in exchange for favors. Insider trading tips about corporate events that impact financial markets are sometimes transmitted over four or more links of trust in networks formed by family, friends, and coworkers, and tipsters are rewarded for insider information with goodwill, cash, gifts, other insider tips, and jobs [1]. Professional expertise and agricultural know-how (e.g., use of new tools and development of better plant or animal breeds) are also acquired through personal contacts [6]. Technological innovation spreads via partnerships among firms in an industry, while trade secrets are acquired by poaching employees from

Date: September 7, 2017.
I thank Kenneth Ahern, Nageeb Ali, Aubrey Clark, Glenn Ellison, Andrea Galeotti, Ben Golub, Matt Jackson, Andy Wei, Glen Weyl, and Alex Wolitzky for comments. This research is funded by NSF Grant 1260744.

rival companies [23]. News and event invitations are communicated through social networks. Similarly to information flowing in a network, book publishing rights are sublicensed via sequences of agreements demarcating markets defined by geographic regions and languages along with the corresponding media (print, audio, or electronic). This paper studies how the locations of initial sources of information in a network shape diffusion paths and determine the profits that players at different positions in the network obtain from consuming and reselling information.

We consider a market with a network structure in which some players are endowed with an identical information good. Information is an indivisible, non-depreciating, and non-rival consumption good for which players have unit demand and heterogeneous values. Every player who has the good can replicate it for free and sell it to his neighbors in the network. Each player who acquires the good enjoys his consumption value and gets the opportunity to replicate and resell it subsequently to any neighbor.[1] We refer to players who own the good at a certain time as sellers and to the other players 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 every matched buyer-seller pair are determined according to the Nash bargaining solution under the assumption that the seller has bargaining power p and the buyer has bargaining power $1-p$. The state of the market at each date, which determines the disagreement payoffs for every match, is described by the configuration of sellers in the network at that date. While our analysis applies broadly to goods with the properties outlined here, it will often be convenient to frame concepts and results in terms of (indivisible) information diffusing through the network. Polanski (2007) studied a related model. We discuss his contribution after an exposition of our results.

Buyers serve as both consumers and intermediaries in the market. Upon acquiring the good, each buyer enjoys a consumption value as well as a resale value that reflects the profits he earns by selling the good to other buyers. In every state of the market, sellers may indirectly extract profits from buyers via intermediation paths along which every player demands a fraction p of the consumption and resale values of the next buyer on the path. Liebowitz (1985) coined the term "indirect appropriability" for the idea that sellers can capture part of the profits gained by intermediaries who acquire the original good and resell copies to buyers. 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. Boldrin and Levine (2008) describe the concept of indirect appropriability in the context of agricultural innovation: "competing agricultural innovators captured a substantial share of the value of all future profits accruing to

[1]Alternatively, ownership of the good may indicate membership in a club. In that interpretation of the model, existing club members can invite others to join the club.

subsequent users of the new plant or animal." However, several studies [2, 14, 18, 19, 21, 27] point out that competition among sellers of the original good and buyers who resell copies may drive prices down in secondary markets and eliminate opportunities for indirect appropriation of profits. Our network formulation encapsulates both the intermediation role of buyers who provide essential access to parts of the market not directly accessible to sellers as well as competitive forces that restrict indirect appropriability.

The main contribution of this paper is a network partition that reflects the effects of competition and the scope of indirect appropriability for every seller in the network. Sellers extract profits only from buyers who belong to their block in the partition. 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 places across a bottleneck link, the seller demands a fraction p of the buyer's consumption and resale values, and the partition into profit blocks evolves to reflect the buyer's takeover of the submarket for which he provides essential intermediation in the block. 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 have negative externalities for sellers. Some sellers have incentives to sever redundant links. Information is sold at zero price over redundant links.

Our analysis indicates 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 documented in empirical research [9, 12]. In such networks, the possibility of reproducing the good and its competitive effects severely 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 sellers intellectual property rights may be socially optimal. When replication and resale are not prohibited by law, sellers can attempt to eliminate the negative effects of competition by engineering certain features of the prototype in order to restrict trade. For instance, digital rights management schemes ranging from passwords and activation keys to built-in software and hardware incompatibilities control how digital goods can be accessed, copied, shared, or converted to other formats and effectively remove redundant links from the network.

Besides 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 any block. Information invariably enters any block without sellers through the same node—the dealer of the block. Dealers have multiple paths of access to information and always receive information via redundant links. The nodes lying along the unique path between a particular buyer and the seller or dealer of 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 any given block must overlap within that block.

Due to a multiplicity problem, we focus on a refinement of the bargaining solution whereby trade takes place with positive probability and is incentive compatible for each matched buyer-seller pair in all market states. We prove that the refinement selects unique payoffs for all players. We provide the following theoretical foundation for our refinement: it generates the unique bargaining solution payoffs with the property that the price each buyer pays to acquire the good is independent of the history of trades. Thus, prices under the refinement are robust with respect to the amount of information players receive about the state of the market.

The starting point of our analysis is the intuition that each seller can extract profits directly or indirectly only from buyers for whom he is the exclusive supplier of the good. Formally, a seller is an *exclusive supplier* for a buyer if the following conditions hold: (1) there exists a unique path between the seller and the buyer; and (2) any path from another seller to the buyer is intermediated by the particular seller. These conditions imply that every player along the path is the only potential source of information for the next buyer on the path. Then, players intermediating trade along the path sequentially take advantage of their *monopoly power* over the rest of the path.

We introduce a binary relation over nodes in a network that succinctly captures conditions (1) and (2) above. For an arbitrary undirected network, two nodes are related if they are connected by a unique path in the network. We establish that this binary relation is an equivalence relation for every network and that each one of its equivalence classes induces a tree in the underlying network. Condition (1) requires that the particular buyer-seller pair belongs to the same equivalence class of the binary relation in the original network. To express condition (2), we consider an auxiliary network derived from the original one by adding a dummy player and linking all sellers in the market with one another and with the dummy player. The construction of the auxiliary network ensures that no two sellers are in the same equivalence class of the corresponding binary relation. We show that a buyer and a seller satisfy conditions (1) and (2) in a given market if and only if they belong to the same equivalence class in the auxiliary network. This finding leads to a formula for seller profits in every market state. The formula reveals that each seller indirectly appropriates a fraction of the consumption value of every buyer from his equivalence class in the auxiliary network that declines exponentially, as a power of the parameter p, with the distance to the buyer and does not earn profits from other buyers. Therefore, a seller's equivalence class represents his *captive submarket*.

The main result of the paper extrapolates the formula for seller profits to characterize the payoffs of all players in every state of the market. An important step in determining the division of the gains from trade identifies the nodes that provide critical access to information in equivalence classes that do not include sellers. We show that every such class contains one buyer—the *dealer*—who intermediates all diffusion paths between sellers and other members of the class. The assumption that each buyer-seller matched pair trades with positive probability, which underlies the refinement of the solution, implies that all players along the competing diffusion paths that lead to a particular dealer—including two neighbors of the dealer—eventually acquire the good. The dealer can then exploit the competition between his neighbors to purchase the good at zero price. Buyers not serving as dealers for their equivalence classes do not benefit from competitive pricing. Indeed, every non-dealer buyer can purchase the good from a single neighbor and has to pay a fraction p of his consumption and resale values in that transaction. After any sequence of trades conveying the good to a buyer, the buyer's resale value is given by the formula for seller profits in the ensuing market.

The network decomposition reveals that every trade is governed by either competitive or monopolistic forces. Trades across equivalence classes entail *competition* between sellers. The expansion of the set of sellers generated by such trades does not affect the composition of equivalence classes in the auxiliary network. By contrast, trades within the same equivalence class involve *monopolies*. Such trades split the common equivalence class of the buyer and the seller into two classes, reflecting the buyer's takeover of the seller's share of the market for which the buyer provides *essential intermediation*.

The classification of links in terms of competitive and monopolistic functions also proves useful in understanding the role each link plays in information transmission as well as the effects of removing a link on the distribution of profits in the network. We find that links between nodes in the same equivalence class constitute *bottlenecks* for the diffusion of information. The removal of a link contained in an equivalence class disconnects the network into two connected components and blocks the spread of information to buyers from the component that does not include the dealer of that class. All players from the equivalence class containing the link suffer from its removal. The deletion of the link does not affect the payoffs of players from other equivalence classes that remain connected to sellers. Hence, bottleneck links provide *positive externalities* for all players.

Any link bridging different equivalence classes is *redundant* for information diffusion. Deleting such a link from the network does not prevent any buyer from acquiring the good. However, the removal of a redundant link may lead to the merger of some equivalence classes. All sellers and all buyers who are not dealers in the original network benefit from the removal of a redundant link as their respective captive markets grow. In particular, redundant links impose *negative externalities* on sellers. Dealer buyers can lose dealer status and suffer a drop in profits following the removal of redundant links. Nonetheless, the removal of a redundant link may also generate profit boosting expansions in the equivalence classes of some dealers.

The interplay between competition and monopoly in the present setting is reminiscent of the market forces emerging in the non-cooperative intermediation game of Manea (forthcoming). In that game, a non-replicable good is sequentially resold between linked intermediaries in a network until a player consumes it. In equilibrium, at every point in the resale process, the owner of the good experiences either a competitive situation in which he obtains second-price auction profits from his neighbors or a bilateral monopoly scenario in which he is held up by the neighbor with the highest resale value. In the former case, the seller is able to take advantage of competition among buyers, while in the latter the seller is effectively subject to a monopsony. The competitive environment is transposed in the context of the present model, with buyers and sellers switching roles: dealer buyers exploit competition among sellers, while non-dealers are monopolized by their sole suppliers.

Our model builds on work by Polanski (2007). In Polanski's framework, a single seller is initially endowed with the information good, and all buyers derive the same utility from consuming the good. Polanski provides recursive equations that describe the evolution of payoffs as buyers acquire information. His 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. Polanski finds that the price a buyer pays for the good depends on whether he belongs to a cycle that includes a seller in the prevailing market. Our analysis complements Polanski's result by explicitly computing payoffs in terms of global network structure in addition to extending the model to competing sellers and asymmetric consumption values. The network decomposition we discover and the payoff formulae stemming from it reflect the effective market share of every seller and the revenues sellers collect directly or indirectly from any buyer. Another novel contribution of our research is the classification of links into competitive and monopolistic roles as well as the exact correspondence between these roles and network connectivity, information diffusion, and prices. Our version of the solution refinement and its foundation are also new.

In a contemporaneous paper, Ali et al. (2016) study a market for an information good in which every seller can trade with every buyer. Their setting corresponds to a complete network in our model.[2] Focusing on a complete network affords a characterization of the best and the worst equilibria for sellers of information and permits the investigation of complementary economic issues such as costly innovation and information acquisition, optimal patent policies, and first-mover advantage deriving from delays in imitation.

[2] 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, 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 among members of clubs. Boldrin and Levine (2002) argue that the creator of a good can earn substantial profits in a market without copyright protection where users reproduce the good at a constant rate and rent copies at competitive prices over time. Muto (1986), Takeyama (1994), and Polanski (2017) analyze the consequences of consumption externalities for the pricing and diffusion of information goods. Varian (2005) reviews economic issues related to copying and copyright law, while Novos and Waldman (2013) discuss the evolution of piracy of digital goods. The paper contributes to the growing literature on intermediation and bargaining power in networks (see Condorelli and Galeotti (2016) and Manea (2016) for recent surveys), which thus far has focused on the trade of non-replicable goods.

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 resale 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. Section 5 introduces the concepts of bottleneck and redundant links and describes how they shape diffusion paths. Section 6 characterizes the roles of bottleneck and redundant links for network connectivity and profit distribution. In Section 7, we present foundations for the refinement of the solution. Section 8 develops results showing that competition limits the scope of indirect appropriability in dense networks, and Section 9 provides concluding remarks. Proofs omitted in the main text are available in the Appendix.

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 Resale 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 sellers. We assume that information is a homogeneous, non-depreciating, and non-rival consumption good and that every player has unit demand for the good. Sellers can replicate and sell the good sequentially to each of their neighbors in G without any production or transaction cost.[3] Upon acquiring the good, player $i \in N$ enjoys a *consumption value* of $v_i > 0$ and joins the set of sellers.[4] The market is open at an infinite number of discrete *dates* $t = 0, 1, \ldots$. Players are infinitely patient.

The *state* of the market (or *configuration of sellers*) 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 players in S as *sellers* and to those in $N \setminus S$ as *buyers*. In state S, a 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. That is, $\mathcal{S}$ denotes 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. For every $S \in \mathcal{S} \setminus \{N\}$, a probability distribution $\pi(S)$ assumed to have full support over $\mathcal{L}(S)$ 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 $t + 1$.[5] The game ends when the market reaches state N in which all players have the good.

To describe "equilibrium" outcomes in this market, we propose a cooperative solution concept with a Markov structure. We assume that payoffs, trading probabilities, and prices at each date t depend only on the prevailing seller configuration at t. Let $u_i(S)$ denote the

[3] The analysis extends to a model in which players have a common unit cost for making copies of the good.
[4] Players in $\underline{S}$ are assumed to have consumed the good before the beginning of the game.
[5] For notational convenience, we routinely write $X \cup y$ for the set $X \cup \{y\}$.

payoff of player $i \in N$ in state $S \in \mathcal{S}$. When the link $bs \in \mathcal{L}(S)$ is selected for 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 b and s is determined according to the *Nash bargaining solution*, assuming that:

- s has bargaining power p and b has bargaining power $1-p$, where $p \in (0,1)$ is an exogenous variable common to all buyer-seller interactions;[6]
- 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 values 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 trade 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 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:

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

[6]All results generalize to a model in which for every state $S \in \mathcal{S}$ and pair $(s, b) \in S \times (N \setminus S)$, seller s has bargaining power $p(s, b)$ and buyer b has bargaining power $1 - p(s, b)$ in state S (in this extension of the model, bargaining power depends on the buyer-seller pair but not directly on the market state).

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, $t_{bs}(S) = u_s(S) - u_s(S \cup b) + pw_{bs}(S)$.

Note that 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 respective bargaining powers p and $1 - p$ and disagreement payoffs of 0 for both players. Formally, we impose the following condition:

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

Note that 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 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 conditions (2)-(7) for every state $S \in \mathcal{S}$ (with the variables

$w_{bs}(S)$ 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.

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

Theorem 1. *At most one payoff profile is consistent with any specific profile of agreement probabilities.*

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 $i \in N$). He shows that multiple bargaining solutions may coexist in his model, and this conclusion extends to our framework.[7] Indeed, different payoffs may be consistent with different profiles of agreement probabilities. 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.[8] Based on this fact, we can construct several bargaining solutions in this network.[9]

The first panel of Figure 1 depicts a solution in which trading probabilities are positive on every link in all states. Under this solution, seller s suffers from a commitment problem and does not make any profit. Each of the two buyers expects that s will eventually trade with the other buyer and can then 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 indifferent between trading

[7] His solution concept allows some non-Markovian behavior that turns out to be inconsequential.

[8] For a proof, 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):

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

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

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

which reduce to $u_{b'}(S)\pi_{bs}(S)\alpha_{bs}(S) = u_s(S)\pi_{bb'}(S)\alpha_{bb'}(S) = 0$. Since $\pi_{bs}(S) > 0$ and $\pi_{bb'}(S) > 0$, we have $u_{b'}(S)\alpha_{bs}(S) = u_s(S)\alpha_{bb'}(S) = 0$. Condition (5) requires that $\alpha_{bb'}(S) > 0$ or $\alpha_{bs}(S) > 0$. Without loss of generality, assume that $\alpha_{bb'}(S) > 0$. In that case, we have $u_s(S) = 0$. If $\alpha_{bs}(S) = 0$, then condition (7) leads to $u_{b'}(S) = pv_b$ and $u_b(S) = (1-p)v_b$. We obtain $w_{bs}(S) = v_b - u_b(S) - u_s(S) = pv_b > 0$. Then (2) implies that $\alpha_{bs}(S) = 1$, a contradiction with the assumption that $\alpha_{bs}(S) = 0$. Hence, we also have that $\alpha_{bs}(S) > 0$, which leads to $u_{b'}(S) = 0$. Then, $\alpha_{bb'}(S) > 0$ 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.

[9] One can show that the multiple payoffs supported by the solution in this example are robust to the introduction of discounting and non-cooperative bargaining.

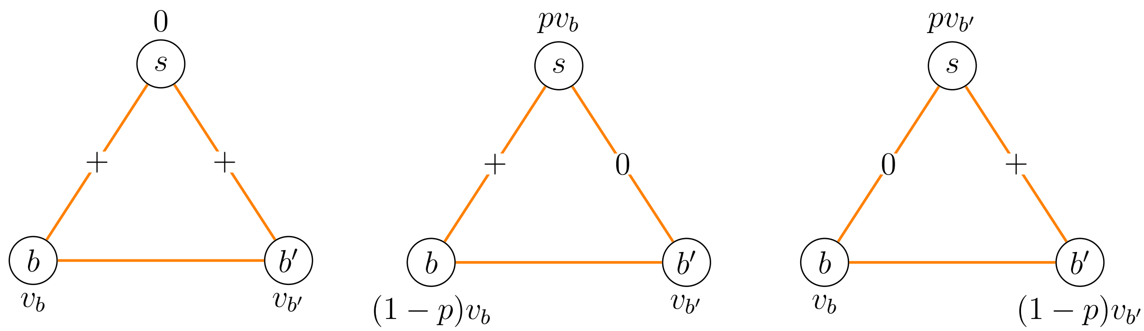

FIGURE 1. Multiplicity example

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

The second panel of Figure 1 represents a second solution, in which seller s commits to not trading with buyer b' in the initial market, but trade takes place with positive probability for all other matches and states. After s trades with b, neither s nor b can extract any profit from b'. Given that b' never acquires the good before b does, 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 $w_{b's}(\{s\}) = -pv_b < 0$. Another solution, illustrated in the third panel of the figure, is obtained by interchanging the roles of b and b' in the solution just described.

In the example above, if the seller severs one of his links, then the bargaining solution in the ensuing network is unique. For instance, if s removes his link with b', then his payoff under the solution for the new network is $pv_b + p^2 v_{b'}$. Hence, seller s has incentives to sever one of his links to avoid the zero prices of the first solution for the original network.

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.,

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

In Sections 4 and 6, we restrict attention to bargaining solutions that satisfy this requirement and simply use the term *bargaining solution* to describe such profiles. We prove that the refinement generates unique bargaining solution payoffs u, which are consistent with any profile of agreement probabilities α that satisfies (8) and do not depend on the matching technology π.

Note that under the bargaining solution illustrated in the second 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' in the initial state, the market transitions to state $\{s, b'\}$, in which competitive forces embedded in the definition of the solution drive the price that b pays to either s or b' for the good to 0 (see footnote 8). Hence, under the solution prescribing that s trade only with buyer b in the initial state, prices depend on the history of trades (reflected in the market state).

By contrast, prices are history-independent under the bargaining solution illustrated in the first 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 zero price in any state of the market. In Section 7, 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 as well as a *resale value* that reflects the profits he can gain from reselling the good to other buyers. 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; Waldman 2014).

The first step in developing payoff formulae for the bargaining solution selected by the refinement identifies the buyers from whom each seller can extract positive profits directly or indirectly. The following definition, inspired by the example from Figure 1, plays a key role in this problem. We say that seller s is the *exclusive 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 obtains 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';

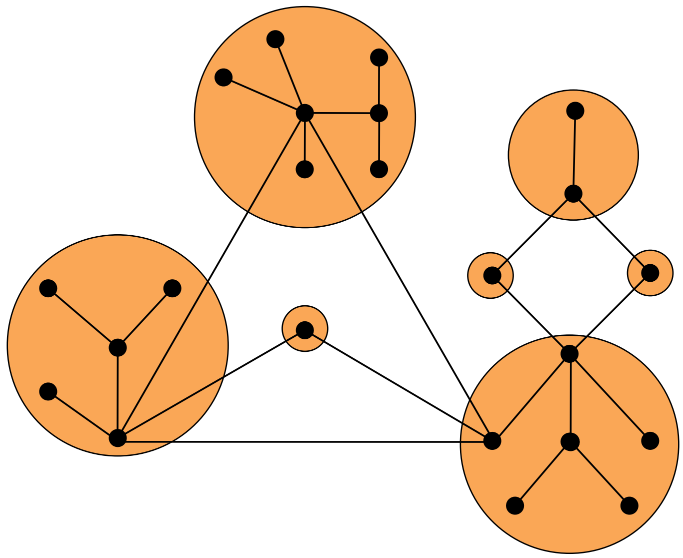

FIGURE 2. Equivalence classes in a network

- 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 then resale proceeds along the chain $(b, b_1, \ldots, b_k)$. We will show that after acquiring the good, buyer b extracts a fraction p^k of the consumption value $v_{b'}$ of buyer b'.

We can formally express the roles of both exclusive suppliers and essential intermediaries using the following concept. Define the *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 unique path in network H. We first show 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 network H belong to the same equivalence class of $\sim_H$.*

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 by Lemma 1, 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

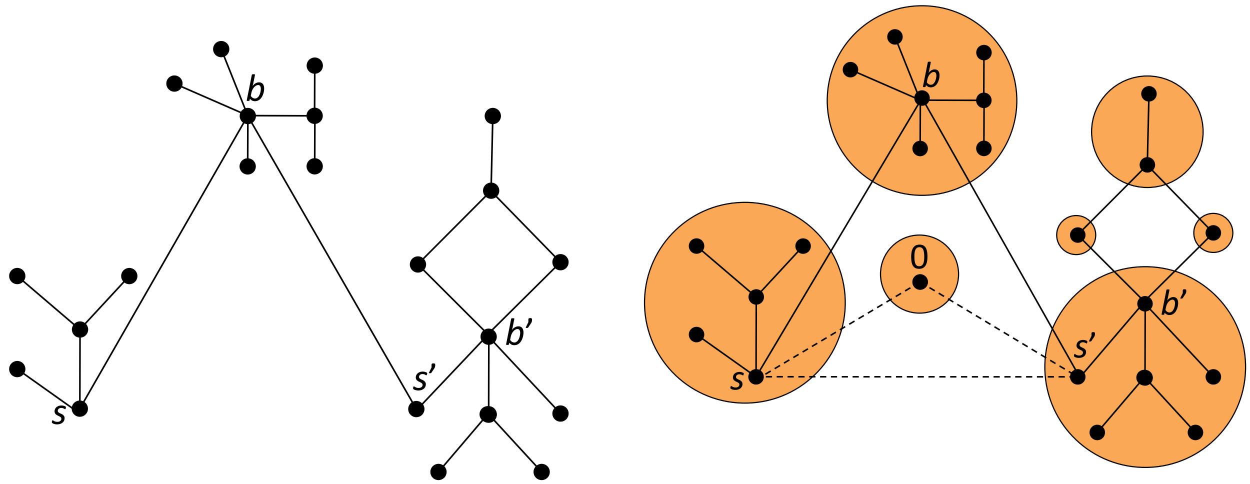

FIGURE 3. Equivalence classes for $G(\{s, s'\})$

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.

Note that the first condition required for s to serve as the exclusive 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 exclusive supplier for b in state 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 exclusive 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 exclusive 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.

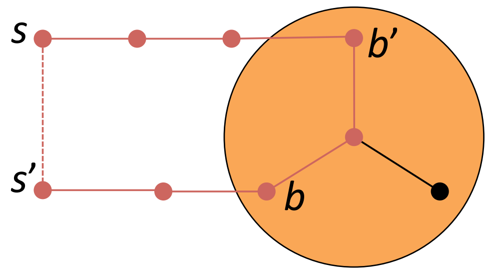

FIGURE 4. Intuition for the existence of dealers

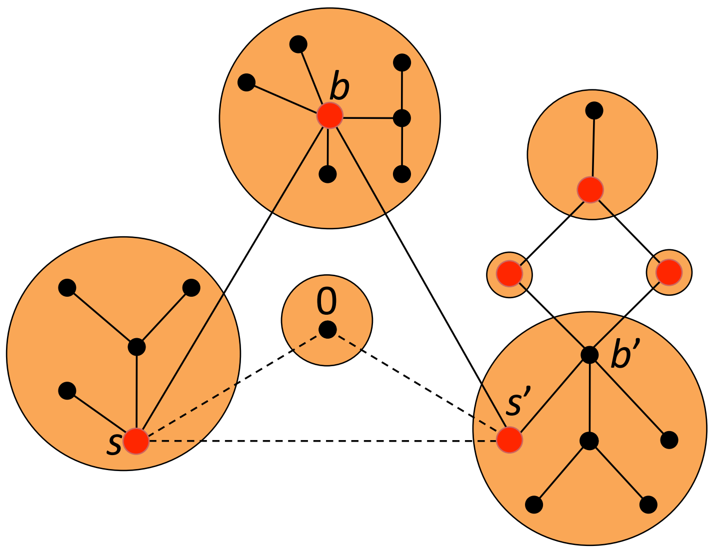

FIGURE 5. Dealers for equivalence classes in $G(\{s, s'\})$

Lemma 2. *Fix a seller configuration $S \in \mathcal{S}$. Seller s is the exclusive 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 $b \sim_{G(S \cup b)} b'$.*

There may be equivalence classes in $G(S)$ that do not contain any seller. The next result shows that the good always "enters" such classes through the same node.

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 network G.*

Figure 4 provides intuition for this result. Assume that buyers b and b' belong to an equivalence class in $G(S)$ that does not contain a seller. If the good can enter $C_b(S)$ via both b and b', then the construction of $G(S)$ implies the existence of a path between b and b' containing nodes outside $C_b(S)$ (if the paths from sellers to b and b' overlap, then we need to consider the "last" node where the two paths intersect), which contradicts Lemma 1.

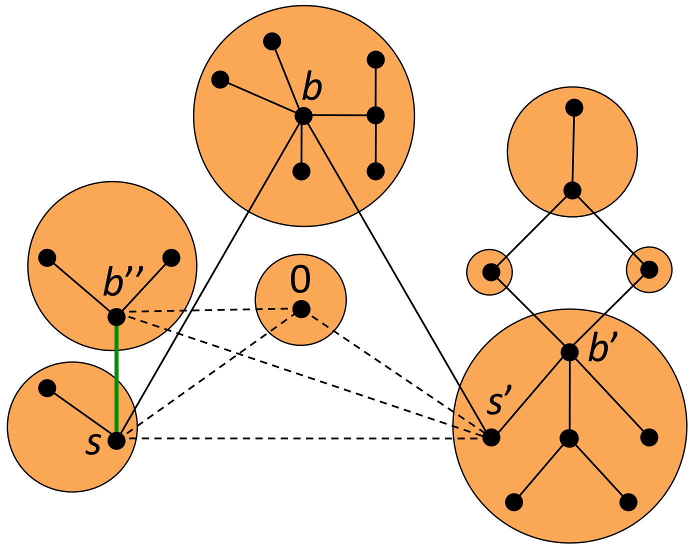

FIGURE 6. Equivalence classes in the network $G(\{s, s', b''\})$ emerging after buyer b'' acquires the good from seller $s \sim_{G(\{s,s'\})} b''$

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* for $C_i(S)$ 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 exclusive 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 5, we indicate the dealer of each equivalence class in the network $G(\{s, s'\})$ from Figure 3 by enlarging the corresponding node.

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. Figure 6 illustrates the evolution of equivalence classes following a trade between seller s and buyer b'', who are members of the same equivalence class in the network $G(\{s, s'\})$ depicted in Figure 3.

Theorem 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 $C_s(S \cup b) \cup C_b(S \cup b) = C_s(S)$.

Lemma 3 and Theorem 2 show that if buyer b purchases the good from seller $s \not\sim_{G(S)} b$ in state S, then b is the dealer for his identical equivalence classes $C_b(S) = C_b(S \cup b)$ in the networks $G(S)$ and $G(S \cup b)$.[10] When $s \not\sim_{G(S)} b$, dealer b can acquire the good via multiple intermediation paths that involve at least two of his neighbors.[11] 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 neighbors, have the good. At that stage, *competitive* forces drive the price that b pays for the good to zero. Thus, b eventually obtains the good at no cost and becomes the exclusive supplier for players in $C_b(S)$. Buyer b obtains a resale value of $u_b(S \cup b)$ following any sequence of trades that delivers the good to him.

However, if buyer b purchases the good from seller $s \sim_{G(S)} b$, then s is the exclusive supplier for b in state S. *Monopoly power* enables seller s to demand a fraction p of b's consumption and resale values. 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. Then, buyer b's resale value reflects the profits $u_b(S \cup b)$ he can extract as a seller from the buyers in $C_b(S \cup b)$. Hence, b acquires the good from s at a price of $p(v_b + u_b(S \cup b))$. These intuitions pave the way to our main result. 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.[12]

Theorem 3. *The profile (u, α) constitutes a bargaining solution 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 $bs \in \mathcal{L}(S)$.

[10]It is worth noting that if $i, j \notin S$, then $ij \in G$ and $i \not\sim_{G(S)} j$ do not imply that i and j are the dealers of their respective equivalence classes in $G(S)$. For example, in the network G from Figure 7 in Section 6, the link bb' connects distinct equivalence classes in $G(\{s\})$, but seller s is the dealer for b's equivalence class $\{b, s\}$ in $G(\{s\})$.

[11]For a proof, note that if b had a single neighbor i providing access to sellers in S, then $i \sim_{G(S)} b$ and i would appear on a path connecting b to a seller, contradicting b's dealer status for $C_b(S)$.

[12]The distances $\delta(i, s)$ appearing in the formula above 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.

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. Recall that Lemma 2 shows that $C_s(S) \setminus s$ is the set of buyers for whom seller s is the exclusive supplier in state S, while $C_b(S \cup b) \setminus b$ is the set of buyers for whom buyer b is an essential intermediary in state S. Then, Theorem 3 can be restated as follows. 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 exclusive 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 zero 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. Prices decline along any trading path within an equivalence class and drop to zero when the good is sold to a new class.

Consider now a buyer b in state S and let $(d(S, C_b(S)), b_1, \ldots, b_k = b)$ denote the unique path in G between the dealer for $C_b(S)$ in state S and buyer b. By definition, the players on the path form the set of essential intermediaries (and the exclusive supplier) for buyer b in state S, and buyer b can acquire the good only after it is resold along the path. The arguments above show that 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 $(p^k, (1-p)p^{k-1}, \ldots, (1-p)p, 1-p)$.

The proof of Theorem 3 shows that the unique bargaining solution payoffs u are consistent with any profile of agreement probabilities α that satisfies (8), so players are indifferent between trading and not trading across every link. Formally, the unique payoffs u have the property that $v_b + u_b(S \cup b) + u_s(S \cup b) = u_b(S) + u_s(S)$ for all $bs \in \mathcal{L}(S)$ and $S \in \mathcal{S}$. An economic interpretation of this formula suggested by Polanski (2007) is that every pair of players who trade with each other captures all the gains created by the trade. The proof of Theorem 3 also demonstrates that $u_i(S) = u_i(S \cup b)$ for all $S \in \mathcal{S}$, $bs \in \mathcal{L}(S)$, and $i \in N \setminus \{b, s\}$. Hence, each trade affects only the payoffs of the two players involved in the trade and leaves the payoffs of the other players unchanged. Relatedly, in Section 7, we show that the prices induce by the refinement of the bargaining solution are independent of the market state. Another remarkable property of the bargaining solution, also noted by Polanski, is that the payoffs do not depend on the matching technology π.

Since $r_s(S)$ is increasing in v_b for all $s \in S$ and $b \in N \setminus S$, Theorem 3 has the following corollary.

Corollary 1. *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.*

Polanski (2007) provides a recursive system of payoff equations for the setting with a single seller and identical consumption values. His payoff equations capture transitions between "consecutive" market states by relating the payoff $u_i(S)$ to payoffs of the type $u_i(S \cup b)$. Polanski 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. In order 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

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

Equivalently, $c_b(S) = 0$ if b has two paths in G with no common interior nodes to (possibly identical) sellers. One can check that for $bs \in \mathcal{L}(S)$, we have $c_b(S) = 0$ if $b \not\sim_{G(S)} s$ and $c_b(S) = 1$ if $b \sim_{G(S)} s$. Note that if $bs \in \mathcal{L}(S)$, then b is the dealer for $C_b(S)$ in state S if and only if $b \not\sim_{G(S)} s$. Hence, 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. These observations lead to the following corollary of Theorems 2 and 3, which generalizes Polanski's result.

Corollary 2. *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 explicit formulae for the payoffs elucidate how the global network structure affects the division of gains from trade among players. The decomposition of the network into equivalence classes delineates opportunities for indirect appropriability and provides a taxonomy of links according to their monopolistic or competitive roles. It also captures the contribution of every individual link to information diffusion, network connectivity, and intermediation profits, as the next two sections demonstrate.

5. 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, Theorem 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 uncovered by Theorem 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, each dealer buyer can purchase the good from multiple neighbors (see footnote 11). Since dealers are essential intermediaries for buyers in their equivalence classes, each non-dealer buyer obtains the good after it is acquired by the dealer of his class and is then resold along the unique path between the dealer and the buyer, which is contained in the class. In particular, each non-dealer buyer can buy the good from a single neighbor over a bottleneck link.

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 4. *The good always reaches dealer buyers via redundant links and non-dealer buyers via bottleneck links. Each dealer buyer may acquire the good from multiple neighbors, while non-dealer buyers can acquire the good from only one neighbor. 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 $C_i(S)$.*

6. Network Comparative Statics

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 for the game). Let G' denote the network obtained by removing the link ij from G. 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 in what follows.

Suppose first that ij is a bottleneck link. As argued in the previous section, 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 the 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' and suffers a drop in profits following the deletion of the link ij. The loss of the link hurts both i and j: one of the two players becomes disconnected from sellers and gets zero payoff, while the other extracts intermediation profits from a smaller equivalence class upon acquiring the good. 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 the link ij, Theorem 3 implies that 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 from other equivalence classes from the two sellers.

If ij is a redundant link instead, then we show that its removal from G does not prevent any player from acquiring the good.[13] The removal of the redundant link ij leads to a weak expansion in each player's equivalence class in state S. Theorem 3 implies that every seller's payoff 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 from whom buyers serve as essential intermediaries and extract resale profits also weakly expands. Theorem 3 implies that the payoffs of buyers who are not dealers in $G(S)$ weakly increase after the link ij is deleted from G. The network from Figure 7 provides an example in which the profit of a non-dealer buyer

[13]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'.

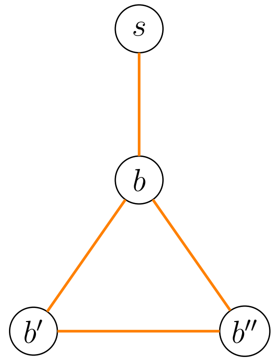

FIGURE 7. Non-dealer buyer b is better off if he severs his link with b''. Removing the redundant link bb'' also benefits dealer buyer b' if $v_{b'} < (1-p)v_{b''}$.

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'', Theorem 3 implies that the deletion of the link 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 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 removal of the link, buyer b is not a dealer anymore and seller s' is able to get a share $1-p$ of his consumption and resale values. Hence, the removal of the link bs is beneficial for s' and detrimental for b. The removal of redundant links can also have the opposite effect on dealer buyer payoffs. For instance, in the network from Figure 7, 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 a payoff of $(1-p)(v_{b'} + pv_{b''})$, which is greater than his payoff $v_{b'}$ in the original network if $v_{b'} < (1-p)v_{b''}$.

The following result, whose detailed proof can be found in the Appendix, summarizes the comparative statics.

Theorem 5. *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 the 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 removal of the link. 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. The effect of removing the link on the payoffs of buyers who are dealers for their equivalence class in $G(S)$ is ambiguous.*

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 2-5 leads to the following conclusions regarding seller profits and information transmission. Bottleneck links confer monopoly power to sellers. The deletion of bottleneck links disconnects the network, blocks the spread of information, and hurts sellers. Redundant links create competition among sellers. The deletion of redundant links does not prevent the diffusion of information and benefits sellers.

Consider now a situation with a single seller s, who can prohibit trade on a subset of links. Digital rights management tools—such as passwords, product keys, limited install activations, encryptions coupled with specific software, hardware, or world regions, digital watermarks, and streaming content—illustrate ways in which sellers can implement trade restrictions.[14] Theorem 3 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 the tree T. Note that any restructuring of a tree whereby a given buyer b who originally receives the good from 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.

7. Foundation for the Refinement

This section provides a foundation for the refinement of the bargaining solution. For this purpose, we revert to using the terms *bargaining solution* for any profile (u, α) satisfying

[14] Concrete examples can be found at https://en.wikipedia.org/wiki/Digital_rights_management.

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.

At the end of Section 3, we argued that prices under the two classes of bargaining solutions 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.[15] Hence, our refinement selects the solutions that do not rely on the assumption that matched players observe the state of the market.

Theorem 6. *The refinement of the bargaining solution selects the only payoff profile for which prices are history-independent.*

8. Limits to Indirect Appropriability

We conclude our analysis with a set of results suggesting that protection of intellectual property is necessary for providing sellers with incentives to create the good in many markets. We show that in sufficiently "dense" networks, the effects of competition between sellers of the original good and buyers of copies are extreme and eliminate indirect appropriability. If creating the prototype requires large investments, sellers do not have incentives to produce it even when production is welfare enhancing. Then, prohibiting the reproduction of the good is socially optimal.

Theorem 3 implies that buyer b receives the good for free from seller s in state S if and only if $b \not\to_{G(S)} s$. Since $b \not\to_{G(S)} s$ whenever $b \not\to_G s$, seller s obtains zero profit from trading with buyer b if $b \not\to_G s$. The latter condition is equivalent to the fact that removing the link bs from network G does not disconnect the network. If b is linked to any other neighbor of s in G, then the network obtained by removing the link bs from G is connected, so seller s must trade with buyer b at zero price. Hence, if G is sufficiently "clustered," in the sense that neighbors of s tend to be neighbors with each other,[16] then s is unable to extract any profits from his neighbors. Furthermore, if seller s has at least two links in G and the network obtained by removing node s (and its links) from G is connected, then the network

[15]The refinement of the solution 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 $b \notin S \cup S'$.

[16]This principle, known as *triadic closure*, has been popularized by the work of Granovetter (1973).

obtained by removing any link of s from G is also connected, so s obtains zero total profit in state S. Another immediate observation is that if there exists a cycle in G that contains all nodes—conventionally called a *Hamiltonian cycle*—then for any $S \in \mathcal{S}$, all equivalence classes of $\sim_{G(S)}$ are singletons, and Theorem 3 implies that no seller makes profits in state S. Intuitively, the previous two statements suggest that sellers are unable to generate any profits if G is "sufficiently connected." We established the following result.

Theorem 7. *Fix a seller configuration $S \in \mathcal{S}$ in the network G.*

(1) If every neighbor of seller $s \in S$ in G is linked in G to at least one other neighbor of s, then s makes no profit in state S.
(2) If seller $s \in S$ has at least two links in G and the network obtained by removing node s from G is connected, then s makes no profit in state S.
(3) If G contains a Hamiltonian cycle, then no seller earns any profit in state S.

One can asymptotically estimate probabilities related to connectivity in the context of large random networks. We focus on the well-known random graph model of Erdos and Renyi (1959),[17] for which the relevant asymptotic results are readily available. Our exposition of theorems here relies on the monograph of Bollobas (2001). A (Erdos-Renyi) *random graph* with parameters (n, q) is defined by the probability distribution over networks with a fixed set of n nodes in which each link is present independently with probability q or alternatively by the random variable $\mathbf{G}_{n,q}$ that has this distribution. In what follows, let ω be any function of n such that $\omega(n) \to \infty$ as $n \to \infty$.

Define $q^C(n) = (\log n + \omega(n))/n$. Theorem 7.3 in Bollobas (2001) implies that if $q_n \geq q^C(n)$ for all n, then the probability that the random graph $\mathbf{G}_{n,q_n}$ is connected converges to 1 as $n \to \infty$.[18] A rough interpretation of this result is that random networks with n nodes and an average degree slightly greater than $\log n$ are asymptotically connected for large n. Fix a seller s who belongs to $\mathbf{G}_{n,q_n}$ with $q_n \geq q^C(n)$ for all n. Given the link independence assumption embedded in the definition of random graphs, the network $\mathbf{G}'_{n-1,q_n}$ obtained by removing node s from $\mathbf{G}_{n,q_n}$ is a random graph with parameters $(n-1, q_n)$. Applying the result above for the sequence $(\mathbf{G}'_{n-1,q_n})_{n\geq 2}$ (with a simple adjustment in the corresponding function ω), we conclude that $\mathbf{G}'_{n-1,q_n}$ is connected with limit probability 1 as $n \to \infty$. Since $q_n \geq q^C(n)$ for all n, the probability that s has at least two links in $\mathbf{G}_{n,q_n}$ converges to 1 as $n \to \infty$. The second part of Theorem 7 then implies that seller s gets zero profit in $\mathbf{G}_{n,q_n}$ with limit probability 1 as $n \to \infty$.

[17]This first article of Erdos and Renyi on the topic considered a variation of the model presented here, but follow-up work developed parallel results for the two versions of the model.

[18]It is remarkable that, as Bollobas explains, the threshold function q^C is sharp in the following sense: if alternatively $q_n \leq (\log n - \omega(n))/n$ for all n, then $\mathbf{G}_{n,q_n}$ has an isolated node, and is thus not connected, with limit probability 1 as $n \to \infty$.

Similarly, Theorem 8.9 from Bollobas (2001) states that if $q_n \geq q^H(n) =: (\log n + \log\log n + \omega(n))/n$ for all n, then the probability that the random graph $\mathbf{G}_{n,q_n}$ contains a Hamiltonian cycle converges to 1 as $n \to \infty$.[19] Thus, a relatively small increase in the average degree of $\mathbf{G}_{n,q_n}$ by the amount $\log\log n$ over the threshold $\log n$ needed for $\mathbf{G}_{n,q_n}$ to be asymptotically connected generates a clear instance of connectedness—the existence of a Hamiltonian cycle. Based on the third part of Theorem 7, we conclude that if $q_n \geq q^H(n)$ for all n, then all sellers make zero profits in $\mathbf{G}_{n,q_n}$ with limit probability 1 as $n \to \infty$. The next result summarizes our findings related to random graphs.

Theorem 8. *Consider a sequence of random networks* $(\mathbf{G}_{n,q_n})_{n\geq 1}$ *and a function* ω *such that* $\lim_{n\to\infty} \omega(n) = \infty$.

(1) If $q_n \geq (\log n + \omega(n))/n$ *for all* $n \geq 1$*, then any particular seller who belongs to all networks in the sequence earns 0 profit in* $\mathbf{G}_{n,q_n}$ *with limit probability 1 as* $n \to \infty$.

(2) If $q_n \geq (\log n + \log\log n + \omega(n))/n$ *for all* $n \geq 1$*, then all sellers in* $\mathbf{G}_{n,q_n}$ *obtain 0 profits with limit probability 1 as* $n \to \infty$.

Note that both parts of Theorem 8 apply when the number of sellers changes arbitrarily with the size of the network. Versions of this result in which only the network of buyers is random and sellers are linked to several buyers can be derived using the same ideas.

The negative effects of competition on seller profits may be more pronounced in large networks observed in applications than the Erdos-Renyi model suggests. Empirical research provides extensive evidence that social and economic networks are highly clustered.[20] For such networks, the first part of Theorem 7 implies that it is difficult for sellers to earn high profits when reproduction and resale are allowed. While the refinement of the bargaining solution favors trade and generates extreme competition, clustering represents an obstacle to indirect appropriability even for solutions that do not survive the refinement. Indeed, under all solutions, the existence of a link between a pair of a seller's neighbors implies that the seller cannot extract any profits from one of the two neighbors (see footnote 8), a point echoed by Ali et al. (2016).

It should be acknowledged, nevertheless, that the studies about high connectivity and clustering primarily concern large networks. Small networks involving criminal activity, such as the networks of insider traders mapped by Ahern (forthcoming), may be sparse and generate less competition between sellers. In sparse networks with short paths, the indirect

[19]By analogy with the remark from footnote 18, Bollobas argues the threshold function q^H is sharp: if $q_n \leq (\log n + \log\log n - \omega(n))/n$ for all n instead, then the probability that at least one node has fewer than two neighbors in $\mathbf{G}_{n,q_n}$, and hence $\mathbf{G}_{n,q_n}$ does not contain any Hamiltonian cycle, converges to 1 as $n \to \infty$.

[20]See the books of Jackson (2008) and Easley and Kleinberg (2010) for references. In social networks, clustering captures the idea that individuals who have common friends are more likely to be friends with each other. Another expression of this phenomenon, highlighted by the random graph model of Jackson and Rogers (2007), is that individuals are typically friends with the friends of their friends.

appropriability argument of Liebowitz (1985) and Boldrin and Levine (2008) remains valid, and sellers may extract a fraction of the social value of the good which covers their production costs. In general, the efficiency of patents and copyright depends on network density and diameter.

9. 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 found that a key equivalence relation derived from a suitably modified network formally describes the roles of exclusive suppliers and essential intermediaries for the diffusion of the good. Sellers obtain profits from buyers for whom they represent exclusive suppliers, while buyers make profits by conveying the good to buyers for whom they provide essential intermediation. Equivalence classes of the relation delineate the captive submarkets 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 delivers a complete description of the anatomy of competing paths of diffusion. Our analysis reveals that in networks that are fairly 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; there is no bound on the number of copies players can produce; 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. The graph-theoretic byproducts of this research—including the concepts of equivalence classes, exclusive suppliers, essential intermediaries, dealers, and bottleneck and redundant links—are relevant beyond the model studied here and are likely to play an important role in other models of diffusion in networks.

Appendix: Proofs

Proof of Theorem 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,

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

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

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

Then 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. $\square$

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. $\square$

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 exclusive 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 exclusive supplier for b in state S. We have established that the relationship $b \sim_{G(S)} s$ implies that s is the exclusive 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 exclusive 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 exclusive supplier for b in state S. Therefore, if $b \not\sim_{G(S)} s$, then s is not the exclusive 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 exclusive supplier for b' in state $S \cup b$. $\square$

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 then 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 then 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 $s \neq s' \in S$. $\square$

Proof of Theorem 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 then obtain 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'. Then, 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 $s \neq 0$. $\square$

Proof of Theorem 3. 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. Then, Theorem 1 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 Theorem 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 $C_b(S) = C_b(S \cup b')$. Then, $d(S \cup b', C_b(S \cup b')) = d(S, C_b(S))$ 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 $b \not\sim_{G(S)} s$. Then 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 an exclusive 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:

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

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

$$\begin{aligned} &v_b + u_b(S \cup b) + u_s(S \cup b) - u_b(S) - u_s(S) \\ &= v_b + r_b(S \cup b) + r_s(S \cup b) - (1 - p)(v_b + r_b(S \cup b)) - (r_s(S \cup b) + p(v_b + r_b(S \cup b))) = 0, \end{aligned}$$

as desired.

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

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

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 Theorem 4. The first two statements of the result have been proven in Section 5. 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. Then, all such paths must overlap in $C_i(S)$ with the unique path between $d(S, C_i(S))$ and x in G. $\square$

Proof of Theorem 5. For a general network H, let $H \setminus ij$ denote the network obtained by deleting the link ij from H (which is identical to H if $ij \notin H$). Fix a connected network G with $ij \in G$ and let $G' = G \setminus ij$. When the network G' is not connected, the proof relies on applications of earlier results to the connected components of G'. For every seller configuration $S \in \mathcal{S}$, let $G'(S)$ denote the network derived from G' in the same fashion $G(S)$ is derived from G. Note that $G'(S) = G(S)$ if $i, j \in S$ and $G'(S) = G(S) \setminus ij$ otherwise. We use the notation $C'_k(S)$ for the equivalence class of k under $\sim_{G'(S)}$, $u'_k(S)$ for the payoff of player k in network G' in state S, and $\delta'(k,l)$ for the distance between nodes k and l in network G'. Fix a seller configuration $S \in \mathcal{S}$ and assume that $\{i,j\} \not\subseteq S$, so $G'(S) = G(S) \setminus ij$.

Suppose that ij is a bottleneck link. As argued in Section 5, the condition $\{i,j\} \not\subseteq S$ implies that $i \sim_{G(S)} j$. Then, the link ij represents the unique path between i and j in $G(S)$, which implies that it is also the unique path connecting i and j in G. Since $G' = G \setminus ij$ and $G'(S) = G(S) \setminus ij$, both G' and $G'(S)$ are disconnected. Each of G' and $G'(S)$ must have exactly two connected components, which separate i from j, because G and $G(S)$ are connected. Furthermore, the partition of (non-dummy) players into the two components is identical for the two networks. Since $G'(S)$ is disconnected and all sellers in S are linked with one another in $G'(S)$, information does not reach all players in G'.

The relation $i \sim_{G(S)} j$ implies that there is no cycle in $G(S)$ that contains link ij. Then every link that is part of a cycle in $G(S)$ is also part of a cycle in $G'(S)$. It follows that the forests derived by eliminating cycles from $G(S)$ and $G'(S)$ satisfy $\mathcal{F}(G'(S)) = \mathcal{F}(G(S)) \setminus ij$. The removal of link ij from the forest $\mathcal{F}(G(S))$ breaks up the connected component of $\mathcal{F}(G(S))$ containing i and j into two components and does not affect other components. Therefore, $C_i(S) = C'_i(S) \cup C'_j(S)$ with $C'_i(S) \cap C'_j(S) = \emptyset$ and $C'_k(S) = C_k(S)$ for all $k \not\sim_{G(S)} i$. Theorem 3 implies that sellers outside $C_i(S)$ obtain the same profits in G and G'. If $d(S, C_i(S))$ is a seller, Theorem 3, along with $C'_{d(S,C_i(S))}(S) \subset C_i(S)$, implies that $d(S, C_i(S))$'s profit is strictly lower in G' than in G.

To investigate the effects of ij's removal from G on information diffusion and buyer payoffs, suppose without loss of generality that $d(S, C_i(S)) \in C'_i(S)$ (it is possible that $d(S, C_i(S)) = i$). Then, i and $d(S, C_i(S))$ are in the same connected component of G', which is different from j's component. There is a path in G from a seller in S to $d(S, C_i(S))$ that does not contain any other node from $C_i(S)$ and, in particular, does not contain the link ij. Hence, $d(S, C_i(S))$ is in the same connected component as a seller in $G'(S)$. Since sellers are linked to one another in $G'(S)$, all nodes in S must be in the same connected component of $G'(S)$ as i and $d(S, C_i(S))$. This implies that the good cannot reach the players in j's connected component in G' (this component is a superset of $C'_j(S)$; it can be a strict superset formed by the union of $C'_j(S)$ and some of the sets $C_k(S)$ with $k \not\sim_{G(S)} i$). Hence, players in j's connected component in G' obtain zero payoffs in G'.

Consider now a buyer b from i's connected component in G'. As $j \notin S \cup b$, we have $G'(S \cup b) = G(S \cup b) \setminus ij$. Since the links in $G'(S \cup b) \setminus G'(S)$ connect only nodes in the set $S \cup \{b, 0\}$, which is disjoint from j's connected component in $G'(S)$, it must be that $G'(S \cup b)$ and $G'(S)$ have identical connected components. Thus, i and j are in distinct components of $G'(S \cup b)$, which means that the link ij constitutes the only path in $G(S \cup b)$ between i and j and hence $i \sim_{G(S \cup b)} j$. Arguments analogous to those above then show that $\mathcal{F}(G'(S \cup b)) = \mathcal{F}(G(S \cup b)) \setminus ij$. If $b \notin C_i(S)$, then $b \notin C_i(S \cup b)$, which implies that

$C_b'(S \cup b) = C_b(S \cup b)$. If $b \in C_i(S)$, then we have that $C_b'(S \cup b) \subseteq C_b(S \cup b)$, with strict inclusion if $b \in \{i, d(S, C_i(S))\}$. Note that b is a dealer for $C_b'(S)$ in state S if and only if b is a dealer for $C_b(S)$ in state S. Theorem 3 then implies that all buyers in i's connected component in G' that do not belong to $C_i(S)$ obtain the same payoffs in G and G', while buyers in $C_i'(S)$ have weakly lower payoffs in G' than in G, with i and $d(S, C_i(S))$ having strictly lower payoffs in G'.

Suppose next that ij is a redundant link, i.e., $i \not\sim_{G(\underline{S})} j$. Then, we also have that $i \not\sim_{G(S)} j$, so there exists a path between i and j in $G(S)$ that does not involve link ij. Since $G'(S) = G(S) \setminus ij$ and $G(S)$ is connected, the path is contained in $G'(S)$, and $G'(S)$ is also connected. This means that every buyer is connected to a seller by a path in G', so information reaches all buyers eventually. The removal of link ij leads to a weak expansion in each player's equivalence class in $G(S)$. For a proof, fix a player $k \in N$. Since $i \not\sim_{G(S)} j$, it cannot be that both i and j belong to $C_k(S)$. By Lemma 1, every pair of nodes in $C_k(S)$ is connected by a unique path in $G(S)$, which necessarily contains only nodes in $C_k(S)$ and thus excludes link ij. As $G'(S) = G(S) \setminus ij$, every pair of nodes in $C_k(S)$ is connected by a unique path in $G'(S)$ as well. Hence, all nodes in $C_k(S)$ are in the same equivalence class of $\sim_{G'(S)}$, i.e., $C_k(S) \subseteq C_k'(S)$. Theorem 3 implies that every seller's payoff is weakly higher in G' than in G. Indeed, for all $s \in S$, $C_s(S) \subseteq C_s'(S)$ implies that

$$u_s(S) = \sum_{k \in C_s(S) \setminus s} p^{\delta(k,s)} v_k \leq \sum_{k \in C_s'(S) \setminus s} p^{\delta'(k,s)} v_k = u_s'(S).$$

The inequality above relies on the fact that $\delta(k, s) = \delta'(k, s)$ for all $k \in C_s(S)$. This follows from the observation that there is a single path in G between s and any node $k \in C_s(S)$, which does not include the link ij and hence constitutes the unique path between s and k in G'. We have established that the payoffs of all sellers weakly increase when the redundant link ij is removed from G.

Similarly, for every buyer b, we have $i \not\sim_{G(S \cup b)} j$, so $C_b(S \cup b) \subseteq C_b'(S \cup b)$. Suppose that b is not the dealer for $C_b(S)$ in state S. Then, b is not the dealer for $C_b'(S)$ in state S either. For a proof by contradiction, assume that there is a path in G' from a seller in S to b that does not contain any node from $C_b'(S)$ except for b. The path also lies in G because $G' = G \setminus ij$. Since $C_b(S) \subseteq C_b'(S)$, the path does not contain any node from $C_b(S)$ other than b. Then b should be the dealer for $C_b(S)$ in state S, a contradiction. Theorem 3, along with the condition $C_b(S \cup b) \subseteq C_b'(S \cup b)$ and the equality $\delta(b, k) = \delta'(b, k)$ for $k \in C_b(S \cup b)$, implies that

$$u_b(S) = (1-p)(v_b + \sum_{k \in C_b(S \cup b) \setminus b} p^{\delta(b,k)} v_k) \leq (1-p)(v_b + \sum_{k \in C_b'(S \cup b) \setminus b} p^{\delta'(b,k)} v_k) = u_b'(S).$$

This proves that non-dealer buyers weakly benefit from the removal of the redundant link ij from G.

We demonstrated that the removal of a redundant link has ambiguous payoff consequences for dealer buyers before the statement of Theorem 5. □

Proof of Theorem 6. 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 3 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 Theorem 2 implies that $C_s(S) = C_s(S \cup b')$. Moreover, $b' \not\sim_{G(S \cup b)} s$ and Theorem 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 Theorem 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 Theorem 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 an exclusive 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 $b \sim_{G(S)} s$, Theorem 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 Theorem 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 $b \sim_{G(S \cup b')} s$. Theorem 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 3.

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 $bs \in \mathcal{L}(S)$.

Then, 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, condition (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 have 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. □

References

1. Ahern, K.R. (forthcoming) Information Networks: Evidence from Illegal Insider Trading Tips, *J. Fin. Econ.*
2. Ali, N.S., A. Chen-Zion, and E. Lillethun (2016), Reselling Information, mimeo.
3. Bakos, Y., E. Brynjolfsson, and D. Lichtman (1999) Shared Information Goods, *J. Law Econ.*, **42**, 117-156.
4. Besen, S.M. and S.N. Kirby (1989) Private Copying, Appropriability, and Optimal Copying Royalties, *J. Law Econ.*, **32**, 255-280.
5. Boldrin, M. and D.K. Levine (2002) The Case against Intellectual Property, *Amer. Econ. Rev.*, **92**, 209-212.
6. Boldrin, M. and D.K. Levine (2008) *Against Intellectual Monopoly*, Cambridge University Press.
7. Bollobas, B. (2001) *Random Graphs*, Cambridge Studies in Advanced Mathematics, **73**, Cambridge University Press.
8. Condorelli, D. and A. Galeotti (2016) Strategic Models of Intermediation Networks, Chapter 27 of *The Oxford Handbook on the Economics of Networks* (eds: Y. Bramoulle, A. Galeotti, and B. Rogers), Oxford University Press.
9. Easley, D. and J. Kleinberg (2010) *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*, Cambridge University Press.
10. Erdos, P. and A. Renyi (1959) On Random Graphs I, *Publicationes Mathematicae*, **6**, 290-297.
11. Granovetter, M. (1973) The Strength of Weak Ties, *Amer. J. Soc.*, **78**, 1360-1380.
12. Jackson, M.O. (2008) *Social and Economic Networks*, Princeton University Press.
13. Jackson, M.O. and B.W. Rogers (2007) Meeting Strangers and Friends of Friends: How Random Are Social Networks?, *Amer. Econ. Rev.*, **97**, 890-915.
14. Johnson, J.P. and M. Waldman (2005) The Limits of Indirect Appropriability in Markets for Copiable Goods, *Rev. Econ. Res. Copyright*, **2**, 19-37.

15. Liebowitz, S.J. (1985) Copying and Indirect Appropriability: Photocopying of Journals, *J. Pol. Econ.*, **93**, 945-957.
16. Manea, M. (2016) Models of Bilateral Trade in Networks, Chapter 26 of *The Oxford Handbook on the Economics of Networks* (eds: Y. Bramoulle, A. Galeotti, and B. Rogers), Oxford University Press.
17. Manea, M. (forthcoming) Intermediation and Resale in Networks, *J. Pol. Econ.*
18. Muto, S. (1986) An Information Good Market with Symmetric Externalities, *Econometrica*, 295-312.
19. Novos, I.E. and M. Waldman (1984) The Effects of Increased Copyright Protection: An Analytic Approach, *J. Pol. Econ.*, **92**, 236-46.
20. Novos, I.E. and M. Waldman (2013) Piracy of Intellectual Property: Past, Present, and Future, *Rev. Econ. Res. Copyright*, **10**, 1-26.
21. Polanski, A. (2007) A Decentralized Model of Information Pricing in Networks, *J. Econ. Theory*, **136**, 497-512.
22. Polanski, A. (2017) Communication Networks, Externalities and the Price of Information, mimeo.
23. Romer, P.M. (1990) Are Nonconvexities Important for Understanding Growth?, *Amer. Econ. Rev.*, **80**, 97-103.
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.
26. Varian, H.R. (2005) Copying and Copyright, *J. Econ. Perspect.*, **19**, 121-138.
27. Waldman, M. (2014) What Limits Indirect Appropriability?, in R. Watt (ed.), *Handbook on the Economics of Copyright: A Guide for Students and Teachers*, Edward Elgar Publishing, 26-47.