ISSN 0956-8549

Arbitrage Networks

By

Rohit Rahi
Jean-Pierre Zigrand

FINANCIAL MARKETS GROUP DISCUSSION PAPER NO. 928

June 2025

Arbitrage Networks*

by

Rohit Rahi
and
Jean-Pierre Zigrand
Department of Finance
and Financial Markets Group,
London School of Economics,
Houghton Street, London WC2A 2AE

April 18, 2023

Abstract

This paper studies the general equilibrium implications of arbitrage trades in segmented financial markets. Arbitrageurs choose a category of trades to specialize in. This results in an equilibrium network in which the various market segments are linked by arbitrageurs. Arbitrageurs exert externalities on each other depending on their position in the network. Due to these externalities, the complete network architecture, in which all links are feasible, is in general suboptimal for arbitrageurs; it is dominated by a hub-spoke architecture. The hub acts as a repository of liquidity, facilitating trades with minimal price impact. For an arbitrary architecture, as the mass of arbitrageurs grows, equilibrium prices converge to those of the frictionless economy with no segmentation. On the other hand, even if the architecture is complete, equilibrium networks may not be complete or even connected, regardless of the mass of arbitrageurs.

Journal of Economic Literature classification numbers: D53, D85.

Keywords: Arbitrage, segmented markets, networks, externalities.

*This paper has benefited from comments by Piero Gottardi, Giulia Iori, Tarun Ramadorai, and especially Sanjeev Goyal and Frank Page. We also thank seminar participants at several universities for helpful discussions.

1 Introduction

In the standard frictionless model of financial markets, all assets are traded on a centralized exchange, with a Walrasian auctioneer determining a price vector that clears all markets simultaneously. Clearly actual markets do not correspond to such an idealization. Assets are traded in a variety of markets, separated by both geography and function, such as stock exchanges, options and futures exchanges, alternative trading venues (such as Multilateral Trading Facilities, Electronic Communication Networks and Dark Pools), as well as over-the-counter, i.e. in direct and private arms-length transactions bypassing organized markets. Moreover, these market segments or trading platforms, which we shall refer to collectively as *exchanges* in this paper, attract distinct clienteles. Very few retail clients trade on more than one exchange or across all asset classes, let alone all asset classes across all trading venues simultaneously.

Market segmentation opens up the possibility for agents, such as market makers, mutual and hedge funds, and proprietary trading desks of investment banks, to profit by intermediating and facilitating trade across the various exchanges. A very large percentage of asset trading is attributable to such traders, whom we shall refer to simply as *arbitrageurs*. Our main object of interest is the resulting network, and the externalities that arbitrageurs exert on each other due to their position in the network.

In actual trading networks, arbitrageurs operate on only a few exchanges at best. For instance, pairs trading is a fashionable component of equity long-short hedge funds. What is more, even if these traders do operate between a number of exchanges at the institutional level, the various trading desks do not seem to coordinate in general. Anecdotal evidence puts this down to informational and other frictions (for instance, Agnes (2000) cites local "market feel" as the reason for a concerted strategy among global swaps banks to decentralize non-US swaps books to their natural markets), to the fact that each desk is allocated a capital limit and operates roughly as a stand-alone profit center, as well as to the fact that compensations almost exclusively depend on a desk's own P&L and lead to a natural rivalry among dealers within the same institution (refer for instance to Drobny (2000)).

In the present paper, we allow arbitrageurs to link only two exchanges, but let them choose which ones. Formally, we consider a two-period economy in which assets are traded at date 0 and uncertainty is resolved at date 1. Trading occurs on multiple exchanges. We assume that the aggregate excess demand function of the local investors on each exchange is affine in prices. This is the case, for instance, if the CAPM holds locally on all exchanges. Arbitrageurs, each of whom can choose a pair of exchanges on which to trade, seek to exploit price differentials across exchanges. Arbitrageurs behave competitively and face an arbitraging cost that bounds their trades.

There is a given network architecture, which specifies which pairs of exchanges can be arbitraged, i.e. which links between exchanges are admissible. For an arbitrary network architecture, we characterize the endogenous distribution of arbitrageur activity across admissible links, and the resulting asset prices that clear markets on each exchange. Comparing different architectures allows us to understand the externalities generated by arbitrageur trading.

Generally speaking, the complete architecture, in which all links are admissible, is suboptimal for arbitrageurs; a hub-spoke architecture leads to higher profits. This is because of externalities in the provision of liquidity. Imagine the equilibrium state-price vectors on the various exchanges as points, or "nodes", in Euclidean space, with node k corresponding to exchange k. Controlling for arbitraging costs, the profit from arbitraging a pair of exchanges, k and ℓ, is proportional to the distance (appropriately defined) between nodes k and ℓ. As more arbitrageurs set up shop on the link (k, ℓ), they pull these nodes closer together, affecting the distance between them and other nodes $\ell' \neq k, \ell$. This is an externality, which may be positive or negative, on arbitrageurs active on links (k, ℓ') and (ℓ, ℓ'). If all trades are channeled through a node (the "hub") that lies towards the center of all the nodes, arbitrageurs on one side of the hub exert a positive externality on arbitrageurs on the other side. This allows mispricings to be exploited with as little market impact as possible, with the hub serving as a repository of liquidity. However, if arbitrageurs can operate on any link, this cannot be sustained in equilibrium. Each arbitrageur has an incentive to deviate and arbitrage two exchanges on opposite sides of the hub, since there is a larger mispricing on such a link. Instead of contributing to liquidity, the deviating agent uses up liquidity at both ends. All other agents act similarly, leading to a Prisoner's Dilemma style suboptimal outcome.

As the mass of arbitrageurs goes to infinity, state prices on all exchanges converge to the state prices of the frictionless integrated economy. Thus, in the limiting case, arbitrageurs connect markets and carry out trades in the aggregate that achieve exactly the transfers of state-contingent consumption that a global auctioneer would have performed. This is true for an arbitrary network architecture as long as it is connected (i.e. it is possible to connect any pair of exchanges via a sequence of links), despite the inefficiencies arising from network externalities, and from the fact that each arbitrageur is allowed to operate on only one link.

On the other hand, even if all links are admissible, equilibrium networks need not be complete, or indeed connected, regardless of the mass of arbitrageurs. Even though asymptotically equilibrium prices are those that would obtain in a frictionless economy, network structure does not become irrelevant.

Related Literature. This paper traces its antecedents to Zigrand (2004, 2006), where the general segmented markets framework is introduced, and in particular to Rahi and Zigrand (2009), which specializes this framework to the CAPM setting. In these papers, each arbitrageur is simultaneously active on all exchanges, so network effects are absent.

The usefulness of a general segmentation setup has been recognized for a long time, in theory[1] and in applied work as documented for instance in the success of

[1]Allais (1967) argued for a more realistic "economy of markets" in lieu of a "market economy."

the market segmentation hypothesis (Culbertson (1957)) and the preferred habitat hypothesis (Modigliani and Sutch (1966)) in fixed income analytics. For example, banks and building societies concentrate a large part of their activity at the short end of the interest rate term structure, both for asset-liability and for regulatory reasons, while pension funds and insurance companies operate at the long end. There is a growing empirical literature in support of more general segmentation in financial markets (see Rahi and Zigrand (2009) for a discussion of this literature). A stylized fact that emerges from this research is that assets in different market segments are priced by distinct groups of investors.

In order to focus on network structure, we assume in this paper that asset markets are complete on every exchange. In actual fact, a considerable number of securities are issued by agents whom we call arbitrageurs. In Rahi and Zigrand (2013) we allow arbitrageurs to design the securities that they trade. Our characterization in the present paper of equilibrium networks with complete markets in fact also holds when markets are incomplete, but with asset payoffs that are determined endogenously.

In this paper each exchange is characterized by an affine demand function. We do not explicitly model the behavior of the local investors or clientele. This can easily be done, and indeed must be done, in order to carry out a complete welfare analysis. We leave this task to a companion paper.

In the literature on networks in economics (see, for example, Goyal (2007), Jackson (2008) and Fique et al. (2013)), a typical network formation game consists of nodes corresponding to players and links that these players form with other players. Our approach to network formation is quite different. For us, the nodes are not players at all. Instead, agents choose to link an arbitrary pair of nodes. Moreover, the relative location of the nodes in Euclidean space (a property that is irrelevant in most of the networks literature) is endogenously determined, and indeed is the key attribute which characterizes the network. In that sense, the networks we study are novel. In the present paper, we limit ourselves to competitive arbitrageurs, but the theory can be extended to agents who behave strategically.

The paper is organized as follows. In the next section we formalize the notion of a network architecture in the context of which arbitrageur activity takes place: each arbitrageur chooses an admissible link and how much to buy or sell on the two exchanges on which he is active. In Section 3, we solve for market-clearing prices for an arbitrary distribution of arbitrageurs across links. We endogenize this distribution in Section 4 and outline some basic features of the equilibrium. In Section 5 we provide some benchmark results for equilibrium networks with a large mass of arbitrageurs. In Section 6 we characterize equilibrium networks as polytopes in Euclidean space. We then proceed to explicitly solve for equilibrium networks

In his Nobel speech he says: "... I was led to discard the Walrasian general model of the market economy, characterized at any time, whether there be equilibrium or not, by a single price system, the same for all the operators, - a completely unrealistic hypothesis, - and to establish the theory of economic evolution and general equilibrium, of maximum efficiency, and of the foundations of economic calculus, on entirely new bases resting on ... a new model, the model of the economy of markets (in the plural)."

for economies in which autarky state prices are distributed symmetrically. Section 7 considers complete architectures, and Section 8 hub-spoke architectures, followed by a comparison in Section 9. Proofs are in the appendix.

2 The Setup

There are two dates, $t = 0, 1$. Uncertainty, realized at date 1, is parametrized by S states of the world, with typical state s. Assets, which are in zero net supply, are traded on several exchanges, with complete markets on each exchange. The set of exchanges is $\mathcal{K} := \{0, \ldots, K\}$, $K \geq 1$.

We denote the state-price deflator (or pricing kernel) on exchange $k \in \mathcal{K}$ by $\hat{p}^k$, a vector in $\mathbb{R}^S$ ($\hat{p}^k_s$ is the price of the s'th Arrow security divided by the probability of state s).[2] Date 0 consumption serves as the numeraire. The aggregate excess demand for date 1 state-contingent consumption on exchange k is $\theta^k(\hat{p}^k)$. We assume that this demand function is affine, taking the following form:

$$\theta^k(\hat{p}^k) = \frac{1}{\beta^k}(p^k - \hat{p}^k). \tag{1}$$

The exogenously given vector p^k is exchange k's autarky state-price deflator (net demand on k is zero when $\hat{p}^k = p^k$). The coefficient β^k captures the price impact of an additional unit of state-contingent consumption supplied to exchange k; thus $1/\beta^k$ is the "depth" of exchange k. We assume that $p^k \geq 0$, and $\beta^k > 0$, for all $k \in \mathcal{K}$, and that the p^k's are not all the same.

One way to generate demand functions of the form given in (1) is to assume that, associated with each exchange, there is a group of competitive investors with quadratic preferences (as in Rahi and Zigrand (2009)). For the purposes of the present paper, however, it does not matter how these demand functions are microfounded.

While exchanges do not interact with each other directly, there is a mass (or measure) N of competitive arbitrageurs who can exploit price differentials across them in the manner described below. Arbitrageurs have no endowments, and care only about date 0 consumption.

Given the set of exchanges $\mathcal{K}$, we specify a set $\mathcal{A}$ of *links* (k, ℓ), i.e. $\mathcal{A} \subset \{(k, \ell) : k, \ell \in \mathcal{K}, k \neq \ell\}$. We will use the abbreviated notation $k\ell$ instead of (k, ℓ). To avoid notational ambiguity, links $k\ell$ and ℓk are taken to be the same link. Each arbitrageur chooses to arbitrage one of the links in $\mathcal{A}$. Let $N^{k\ell}$ denote the mass of arbitrageurs on link $k\ell \in \mathcal{A}$. We have $\sum_{k\ell \in \mathcal{A}} N^{k\ell} = N$. For notational convenience we define $N^{k\ell}$ to be zero if $k\ell \notin \mathcal{A}$.

Formally, $\mathcal{G} := (\mathcal{K}, \mathcal{A})$ is a *graph*, with *nodes* $\mathcal{K}$ and *links* $\mathcal{A}$.[3] We say that ℓ is a *neighbor* of k if $k\ell \in \mathcal{A}$. The graph is *complete* if every node is a neighbor of every

[2]While we do not impose nonnegativity in the definition of a state-price deflator, the $\hat{p}^k$'s will be nonnegative in equilibrium.

[3]A standard reference on graph theory is Diestel (2005). We employ the terms "node" and

other node; if not, it is *incomplete*. If $\mathcal{A} = \{k\ell : \ell \in \mathcal{K}, \ell \neq k\}$, we say that $\mathcal{G}$ is a *hub-spoke* graph with node k as the hub and the other nodes as the spokes. There is a *path* connecting k and ℓ if there is a sequence of distinct nodes $\{k_1, \ldots, k_I\}$ in $\mathcal{K}$ such that $k_1 = k$, $k_I = \ell$ and $\{k_1k_2, k_2k_3, \ldots, k_{I-1}k_I\} \subset \mathcal{A}$. We say that $\mathcal{G}$ is *connected* if there is a path connecting any pair of nodes $k, \ell \in \mathcal{K}$; if not, it is *disconnected*. A maximal connected subgraph of $\mathcal{G}$ is called a *component* of $\mathcal{G}$. It is convenient to denote a component of $\mathcal{G}$ simply by its nodes $\mathcal{C} \subset \mathcal{K}$, where it is understood that the links between these nodes are those that are inherited from $\mathcal{A}$.

While we have introduced the above (standard) terminology for the graph $\mathcal{G}$, it applies of course to any other graph that we consider in the paper (typically a subgraph of $\mathcal{G}$). $\mathcal{G}$ itself will be referred to as a *network architecture* (or, simply, architecture), with $\mathcal{A}$ being the set of *admissible links*. While $\mathcal{G}$ is not necessarily complete, we assume that it is connected. This is without loss of generality as each component of $\mathcal{G}$ can be analyzed as a separate economy.

Each arbitrageur chooses a link on which to trade and how much to supply to the two exchanges connected by this link. It is convenient to think of this competitive interaction as occurring in two stages, with arbitrageurs choosing links at the *network formation stage*, and subsequently choosing their supplies at the *trading stage*. We solve for equilibrium by backwards induction, solving first for market-clearing prices at the trading stage for an arbitrary distribution of arbitrageurs, and then solving for an equilibrium distribution of arbitrageurs at the network formation stage.

Given an architecture $\mathcal{G}$, let $\mathfrak{N} := \{N^{k\ell}\}_{k\ell \in \mathcal{A}}$ be an arbitrary distribution of arbitrageurs across links that are admissible in that architecture. A *network* is a tuple $(\mathcal{G}, \mathfrak{N})$. We say that an admissible link $k\ell$ is *active* if $N^{k\ell} > 0$; likewise, a node is active if it is an endpoint of an active link. Let $\mathcal{A}^*(\mathfrak{N}) \subset \mathcal{A}$ be the set of active links for arbitrary $\mathfrak{N}$.

3 The Trading Stage

We begin by studying optimal arbitrageur supplies and market-clearing prices for $(\mathcal{G}, \mathfrak{N})$, where $\mathfrak{N}$ is an arbitrary distribution of arbitrageurs. Let $y^k_{k\ell} \in \mathbb{R}^S$ be the supply of state-contingent consumption on exchange k of a typical arbitrageur active on link $k\ell \in \mathcal{A}^*(\mathfrak{N})$. The arbitrageur faces a no-default constraint at date 1, namely $y^k_{k\ell} + y^\ell_{k\ell} \leq 0$. Since he cares only for date 0 consumption, this constraint will necessarily hold with equality. We assume that he faces an arbitraging cost, which for tractability is quadratic and takes the form $\frac{1}{2}\delta^{k\ell} {y^k_{k\ell}}^\top \Pi y^k_{k\ell}$, where $\delta^{k\ell}$ is a positive cost parameter, and Π is the $S \times S$ diagonal matrix whose s'th diagonal element is the probability of state s.[4] Taking prices as given, the arbitrageur maximizes (net)

"link" instead of "vertex" and "edge", reserving the latter terminology for its standard usage in the theory of polytopes, which we make extensive use of later.

[4]If we think of $y^k_{k\ell}$ as a random variable, the arbitraging cost is proportional to the size of the induced state-contingent consumption $E(y^k_{k\ell})^2$.

profits from his arbitrage trades, i.e. he solves

$$\max_{y^k_{k\ell}\in\mathbb{R}^S} {y^k_{k\ell}}^\top \Pi(\hat{p}^k - \hat{p}^\ell) - \frac{1}{2}\delta^{k\ell}\, {y^k_{k\ell}}^\top \Pi y^k_{k\ell}.$$

The following is immediate:

Lemma 3.1 (Arbitrageur supplies) *Given $(\mathcal{G}, \mathfrak{N})$, the optimal supply of an arbitrageur on link $k\ell$ in $\mathcal{A}^*(\mathfrak{N})$ is given by $y^k_{k\ell} = -y^\ell_{k\ell} = \frac{1}{\delta^{k\ell}}(\hat{p}^k - \hat{p}^\ell)$.*

Arbitrageurs on link $k\ell$ supply consumption in state s to exchange k when the price that agents on exchange k are willing to pay for a unit of state s consumption exceeds the price at which the arbitrageurs can procure that unit on exchange ℓ. Arbitrageur trades are scaled by the cost parameter $\delta^{k\ell}$. For convenience, we define $y^k_{k\ell}$ to be zero if $k\ell \notin \mathcal{A}^*(\mathfrak{N})$.

We can solve for market-clearing state-price deflators $\hat{p}^k$, $k \in \mathcal{K}$, as follows. Let $\alpha^{k\ell} := N^{k\ell}/\delta^{k\ell}$ for $k\ell \in \mathcal{A}^*(\mathfrak{N})$, and zero otherwise. Also define y^k to be the aggregate supply on exchange k of all arbitrageurs in the economy. From Lemma 3.1,

$$y^k = \sum_{\ell\in\mathcal{K}} N^{k\ell} y^k_{k\ell} = \sum_{\ell\in\mathcal{K}} \alpha^{k\ell}(\hat{p}^k - \hat{p}^\ell). \tag{2}$$

Setting y^k equal to the aggregate demand on exchange k, θ^k (given by (1)), we have

$$\hat{p}^k = p^k - \beta^k y^k. \tag{3}$$

Let $\mathbb{C}(\mathfrak{N})$ be the set of components of the graph $(\mathcal{K}, \mathcal{A}^*(\mathfrak{N}))$, with typical element $\mathcal{C}$. Then $\{\hat{p}^k\}_{k\in\mathcal{K}}$ is a solution to the following system of equations:

$$\hat{p}^k + \beta^k \sum_{\ell\in\mathcal{C}} \alpha^{k\ell}(\hat{p}^k - \hat{p}^\ell) = p^k, \qquad k \in \mathcal{C},\, \mathcal{C} \in \mathbb{C}(\mathfrak{N}). \tag{4}$$

Notice that the equations corresponding to any one component are independent of those corresponding to the other components. Hence we can solve separately for market-clearing prices $\{\hat{p}^k\}_{k\in\mathcal{C}}$ for each component $\mathcal{C}$.

Lemma 3.2 (Market-clearing prices) *Given $(\mathcal{G}, \mathfrak{N})$, there exists a unique profile of market-clearing state-price deflators $\{\hat{p}^k\}_{k\in\mathcal{K}}$. For any $\mathcal{C} \in \mathbb{C}(\mathfrak{N})$ and $k \in \mathcal{C}$, we have $\hat{p}^k = \sum_{\ell\in\mathcal{C}} \eta^{k\ell} p^\ell$, for positive weights $\{\eta^{k\ell}\}_{\ell\in\mathcal{C}}$ that satisfy (a) $\sum_{\ell\in\mathcal{C}} \eta^{k\ell} = 1$; (b) $\beta^\ell \eta^{k\ell} = \beta^k \eta^{\ell k}$, $\forall \ell \in \mathcal{C}$; and (c) $\eta^{kk} > \eta^{\ell k}$, $\forall \ell \in \mathcal{C}, \ell \neq k$.*

Thus the market-clearing outcome at the trading stage is unique and symmetric, with all arbitrageurs on a given link supplying the same amount, and a unique market-clearing state-price deflator for each exchange.

Note that $\hat{p}^k \geq 0$ due to our assumption that $p^k \geq 0$. The market-clearing state-price deflator on exchange k is a convex combination of the autarky state-price deflators of all exchanges to which it is linked directly or indirectly, with the weights depending on the depths, arbitraging costs and distribution of arbitrageur activity for these exchanges. Consider, for example, a neighbor ℓ of k. The higher is $N^{k\ell}$, and the lower is $\delta^{k\ell}$, the greater is the arbitrageur-mediated transfer of state-contingent consumption between k and ℓ. This activity reduces the mispricing $\hat{p}^k - \hat{p}^\ell$, and increases the influence of p^ℓ on $\hat{p}^k$. Furthermore, given that yet other arbitrageurs transfer resources between ℓ' and ℓ, where ℓ' is a neighbor of ℓ but not of k, the state prices of ℓ' also find their way into $\hat{p}^k$. And likewise for the neighbors of ℓ', and so on.

Moreover, if p_s^k, the autarky valuation of state s on exchange k, rises exogenously, then $\hat{p}_s^\ell$ increases on all exchanges ℓ in the component to which k belongs, with the largest increase occurring on k itself. While a local shock affects local state prices the most, it affects the state prices of all exchanges directly or indirectly connected to it.

For a vector v in $\mathbb{R}^S$, the $L^2(\Pi)$-norm of v is defined as follows: $\|v\|_2 := (v^\top \Pi v)^{\frac{1}{2}}$. The following is immediate from Lemma 3.1:[5]

Lemma 3.3 (Arbitrageur profits) *Given $(\mathcal{G}, \mathfrak{N})$, the profit of an arbitrageur on $k\ell$ in $\mathcal{A}^*(\mathfrak{N})$ is given by $\varphi^{k\ell} := \frac{1}{2\delta^{k\ell}}\|\hat{p}^k - \hat{p}^\ell\|_2^2$.*

For an inactive link $k\ell \in \mathcal{A}$, $\varphi^{k\ell}$ has the interpretation of the "potential" profit on $k\ell$.

4 Equilibrium Networks: Preamble

We have shown that, for any given distribution of arbitrageurs $\mathfrak{N}$, there exists a unique market-clearing price vector for each exchange. We are now in a position to analyze the network formation stage in which $\mathfrak{N}$ is determined. Whenever we wish to emphasize that an arbitrageur distribution is an equilibrium distribution, we write it as $\mathfrak{N}(N)$, i.e. as a function of the mass of arbitrageurs N. All the variables introduced earlier, such as the set of active links, prices and profits, depend on $\mathfrak{N}(N)$. To save on notation, we write $\mathcal{A}^*(N)$ instead of $\mathcal{A}^*(\mathfrak{N}(N))$. Similarly (unless it is clear from the context), we write the equilibrium state-price deflator on exchange k as $\hat{p}^k(N)$; notice that $\hat{p}^k(0) = p^k$. Given the graph $\mathcal{G}^*(N) := (\mathcal{K}, \mathcal{A}^*(N)) \subset \mathcal{G}$, we refer to $(\mathcal{G}^*(N), \mathfrak{N}(N))$ as an *equilibrium network*.

Since each arbitrageur is atomistic, and hence has no impact on prices, in an equilibrium network profits must be equal on all active links, with (weakly) lower potential profits on all inactive admissible links. Formally:

Lemma 4.1 *$\mathfrak{N}$ is an equilibrium distribution of arbitrageurs if and only if it satisfies the following condition:*

(N) *There is a $\Phi(\mathfrak{N}) > 0$ such that $\varphi^{k\ell}(\mathfrak{N}) \leq \Phi(\mathfrak{N})$, for all $k\ell \in \mathcal{A}$, and $\varphi^{k\ell}(\mathfrak{N}) = \Phi(\mathfrak{N})$, for all $k\ell \in \mathcal{A}^*(\mathfrak{N})$.*

[5] If we view $\hat{p}^k$ as a random variable with realization $\hat{p}_s^k$ in state s, then $\|\hat{p}^k - \hat{p}^\ell\|_2^2$ can be written as $E(\hat{p}^k - \hat{p}^\ell)^2$, the mean-square distance between $\hat{p}^k$ and $\hat{p}^\ell$.

We show in the appendix that there exists an $\mathfrak{N}$ satisfying condition **(N)**, so that:

Proposition 4.1 (Existence) *There exists an equilibrium network.*

However, in general, there is not a unique equilibrium network, as the following example shows. Here, and elsewhere, when we speak of the arbitraging cost for link $k\ell$, we refer to the parameter $\delta^{k\ell}$.

Example 1 (Multiplicity of equilibrium networks) Suppose there are four exchanges, with $p^0 = p^1$ and $p^2 = p^3$. The architecture is complete, arbitraging costs are the same for every link, and all the exchanges have the same depth. It is easy to verify that each of the following is an equilibrium arbitrageur distribution (where we have specified only the active links): (i) $N^{02} = N^{13} = N/2$; (ii) $N^{03} = N^{12} = N/2$; and (iii) $N^{02} = N^{03} = N^{12} = N^{13} = N/4$. In all three cases, $\hat{p}^0 = \hat{p}^1$ and $\hat{p}^2 = \hat{p}^3$. $\diamondsuit$

We turn now to an analysis of the properties of equilibrium networks. We seek to answer two sets of questions. The first relates to features of an equilibrium network, for a given architecture. Which links attract the most arbitrageur activity? What connectivity properties emerge in equilibrium? Are all admissible links active? If not, is an equilibrium network still connected? The second set of questions pertain to the "comparative statics" of equilibrium networks with respect to the architecture. How are arbitrageur profits affected by the architecture? What connectivity properties do desirable architectures possess?

Most of these questions boil down to a combinatorial problem which, as one might expect, leads to very few clear-cut general results since so many tradeoffs must be balanced, such as the various depths and arbitraging costs, and the various initial mispricings across admissible links, taking into account that the prices on each exchange depend on all flows across the network, no matter how "remote", i.e. no matter how many links away. In particular, one should not expect to obtain general results of the sort "every equilibrium network is hub-spoke", as have been derived for instance in Bala and Goyal (2000), for in our paper nodes are exchanges with heterogeneous intrinsic characteristics. Indeed, any given connectivity structure can be perturbed by varying these underlying parameters. For instance, consider an equilibrium network with a particular connectivity structure $\mathcal{A}^*(N)$. Pick an autarky state-price deflator, say p^0, and move it in $\mathbb{R}^S$ space further away from the other p^k's.[6] At some stage the resulting equilibrium network becomes hub-spoke with 0 as the hub. On the other hand, if an equilibrium network has 0 as the hub, this can be perturbed away by raising the costs of arbitraging with 0.

Since "scale" can always make or break any particular connectivity structure, the most interesting effects are what one might call "network effects", keeping scale fixed (for example, taking the β^k's and $\delta^{k\ell}$'s to be equal, and/or imposing symmetry

[6] When we provide Euclidean geometric intuition, we view the space $L^2(\Pi)$ with norm $\|v\|_2 := (v^\top \Pi v)^{\frac{1}{2}}$ as the Euclidean space $\mathbb{R}^S$ with norm $\|w\| := (w^\top w)^{\frac{1}{2}}$ via the isomorphism $v \mapsto w := \Pi^{1/2} v$. For notational simplicity, we will not make this transformation explicit.

restrictions on the p^k's). Before we study network effects in more detail, the next example illustrates pure scale effects arising from the relative location of the p^k's.

Example 2 (Pure scale effects) Consider the class of networks in which only a single link is admissible, and hence there are no externalities across links. Suppose arbitraging costs are the same for every link, and all exchanges have the same depth. Arbitrageur profits on an admissible link $k\ell$ can be calculated directly from (4): $\varphi^{k\ell} = \frac{\delta}{2[\delta+2\beta N]^2} \cdot \|p^k - p^\ell\|_2^2$. Therefore, the optimal single-link architecture for arbitrageurs is the one that maximizes the autarky gains from trade, $\|p^k - p^\ell\|_2^2$. $\diamondsuit$

Assuming equal depths and arbitraging costs, it is tempting to conjecture that, for an arbitrary architecture, $N^{k\ell} > N^{k'\ell'}$ whenever autarky gains from trade are higher on $k\ell$ than on $k'\ell'$. However, this is not true in general because of network externalities. The following example illustrates.

Example 3 (Network externalities in a hub-spoke architecture) Consider a hub-spoke architecture with four exchanges. Exchange 0 is the hub, and exchanges $1, 2$ and 3 are the spokes. Arbitraging costs are the same for all admissible links, the exchanges are equally deep, and the autarky state-price deflators lie on a straight line segment with p^1 at one extremity, $p^2 = p^3$ at the other extremity, and p^0 half way in between (see Figure 1). Even though $\|p^k - p^0\|_2$ is the same for all $k \neq 0$,

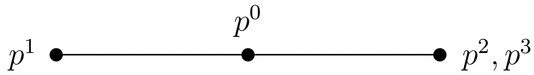

Figure 1: **Network externalities**

there is no equilibrium with $N^{0k} = N/3$, $k = 1, 2, 3$. If this were the case, there would be mass $2N/3$ of arbitrageurs pulling $\hat{p}^0$ towards $p^2 = p^3$ and only $N/3$ pulling it towards p^1. With an equal mass of arbitrageurs pulling $\hat{p}^k$, $k = 1, 2, 3$, towards the middle, profits on link 01 would be higher than on 02 or 03 (this can be verified via an explicit calculation). In equilibrium, we must have $N^{01} > N^{02} = N^{03}$. It is clear that this is robust to small perturbations of the p^k's. In particular, there will be more arbitrageurs on 01 even if $\|p^1 - p^0\|_2$ is (slightly) lower than $\|p^2 - p^0\|_2$ and $\|p^3 - p^0\|_2$. $\diamondsuit$

5 Equilibrium Networks with Many Arbitrageurs

Equilibrium networks with a large mass of arbitrageurs provide a useful benchmark for our analysis. We show that, as the mass of arbitrageurs N goes to infinity, equilibrium state prices on every exchange converge to the equilibrium state prices of the entire integrated economy. This is true for an arbitrary architecture (provided it is connected, which we have assumed throughout).

For any subset $\mathcal{K}'$ of $\mathcal{K}$, let $p^*_{\mathcal{K}'}$ be the (unique) state-price deflator for which excess demand for state-contingent consumption, aggregated across the exchanges in $\mathcal{K}'$, is zero. From (1), we see that $p^*_{\mathcal{K}'} = \sum_{k \in \mathcal{K}'} \lambda^k_{\mathcal{K}'} p^k$, where $\lambda^k_{\mathcal{K}'} := \frac{(\beta^k)^{-1}}{\sum_{j \in \mathcal{K}'} (\beta^j)^{-1}}$. We interpret $p^*_{\mathcal{K}'}$ as the equilibrium state-price deflator for the integrated subeconomy corresponding to $\mathcal{K}'$, with no arbitrageurs. It is equal to the average willingness to pay in this subeconomy, with the willingness to pay on each exchange weighted by its relative depth. Let $p^* := p^*_{\mathcal{K}}$ be the global equilibrium state-price deflator,[7] with $\lambda^k := \lambda^k_{\mathcal{K}}$. Thus $p^* = \sum_{k \in \mathcal{K}} \lambda^k p^k$. In fact, multiplying equation (4) by λ^k and summing over $k \in \mathcal{K}$, we get $p^* = \sum_{k \in \mathcal{K}} \lambda^k \hat{p}^k$, for an arbitrary distribution of arbitrageurs $\mathfrak{N}$. In particular, for an equilibrium distribution of arbitrageurs $\mathfrak{N}(N)$, we have:

$$p^* = \sum_{k \in \mathcal{K}} \lambda^k \hat{p}^k(N), \qquad \forall N \geq 0. \tag{5}$$

We write $\hat{p}^k(\infty)$ for $\lim_{N \to \infty} \hat{p}^k(N)$, and similarly for other variables.

Proposition 5.1 (Convergence) *State prices on all exchanges converge to the state prices of the integrated economy, i.e.* $\hat{p}^k(\infty) = p^*$*, for all* $k \in \mathcal{K}$*. Moreover, if* $\mathcal{G}^*(\infty)$ *exists, then* $p^*_{\mathcal{C}} = p^*$*, for all components* $\mathcal{C}$ *of* $\mathcal{G}^*(\infty)$*.*

As the mass of arbitrageurs increases without bound, all mispricings across exchanges vanish, even though no single arbitrageur ties all the markets together. From Lemma 3.1, individual arbitrageur trades also vanish as N goes to infinity. Convergence need not be monotone, however, as we shall see in Example 6 below.

Notice that price differentials between a pair of exchanges go to zero even if there is no active link between them. Indeed, there may be admissible links that remain inactive for all N:

Example 4 (Incomplete equilibrium network) Consider the complete architecture with three equally deep exchanges, and the same arbitraging costs for all links. The autarky state-price deflators for exchanges 1 and 2 are the same (see Figure 2). There is a unique equilibrium network, with the following arbitrageur distri-

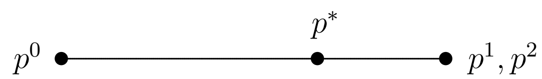

Figure 2: **Incomplete equilibrium network**

bution: $N^{01} = N^{02} = N/2$, and $N^{12} = 0$. Notice that this network is hub-spoke with exchange 0 as the hub. Link 12 is admissible but inactive for all N, since $\hat{p}^1(N) = \hat{p}^2(N)$.

[7] If the demand function on each exchange is generated by competitive investors, as in Rahi and Zigrand (2009), then p^* is the Walrasian state-price deflator of the integrated economy.

Now suppose we perturb this economy so that p^1 and p^2 are close but not equal. Initially, as N increases, there will be arbitrageur activity only on the longer of the two links, 01 and 02. Once profits are equalized across these two links, they will remain equal in length as they shrink to p^*. It can be shown via an explicit computation that the 12 link remains less profitable for all N; as N increases, profits fall on 01 and 02, but so does the potential profit on 12 as $\hat{p}^1$ and $\hat{p}^2$ get pulled even closer together. Thus, once again, the equilibrium network is hub-spoke, with the 12 link remaining inactive even asymptotically. $\Diamond$

This is a robust example of a complete architecture for which the equilibrium network is incomplete for all N. In fact, the equilibrium network may even be disconnected. For $v, w \in \mathbb{R}^S$, we denote by $L[v, w]$ the line segment joining the points v and w.

Example 5 (Disconnected equilibrium network) Suppose there are four equally deep exchanges with nodes $\{p^k\}$ as in Figure 3, the architecture is complete, and arbitraging costs are the same for every link. The integrated economy state-price

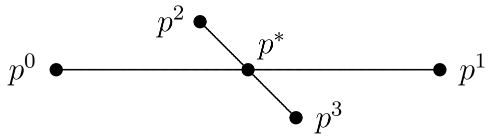

Figure 3: **Disconnected equilibrium network**

deflator p^* is at the common midpoint of the two line segments $L[p^0, p^1]$ and $L[p^2, p^3]$. For small N, all arbitrageurs are on the longer link 01. As N grows, $\hat{p}^0(N)$ and $\hat{p}^1(N)$ converge along $L[p^0, p^1]$ until $N = \bar{N}$ for which $\|\hat{p}^0(\bar{N}) - \hat{p}^1(\bar{N})\|_2 = \|p^2 - p^3\|_2$. For $N > \bar{N}$, two active links appear, 01 and 23, and the network consists of two disjoint subnetworks. As N increases without bound, all four nodes converge to p^*. Even as profits on 01 and 23 converge to zero, they are higher than potential profits on any of the other links. $\Diamond$

This is an example in which the nodes $\{p^k\}$ are symmetrical with respect to p^*. We provide a precise characterization of this class of disconnected equilibrium networks later (Proposition 7.1, part (iii)).

Notice that while the asymptotic incompleteness property of Example 4 is robust to perturbations of the p^k's, this is not the case for the asymptotic disconnectedness property of Example 5; the latter depends on the integrated economy state-price deflator being the same for the two components.

6 The Geometry of Equilibrium Networks

We now proceed to derive some general results on the properties of equilibrium networks. As we argued in Section 4, such results require some sort of symmetry.

Our first task then is to describe the geometry of networks, wherein a notion of symmetry can be formalized.

Recall that the *convex hull* of a Euclidean subset A, denoted by $\text{conv}(A)$, is the set of all convex combinations of points in A. The convex hull of a finite number of points in Euclidean space is called a *polytope*.[8] Geometrically we can view the nodes of a network as points in $\mathbb{R}^S$. The nodes are given by $\{\hat{p}^k(N)\}_{k\in\mathcal{K}}$ (we refer to both $k \in \mathcal{K}$ and $\hat{p}^k \in \mathbb{R}^S$ as the node corresponding to exchange k). Let $\hat{\mathcal{P}}(N) := \text{conv}(\{\hat{p}^k(N)\})$ and $\mathcal{P} := \hat{\mathcal{P}}(0) = \text{conv}(\{p^k\})$. Thus $\mathcal{P}$ and $\hat{\mathcal{P}}(N)$ are polytopes with $\hat{\mathcal{P}}(N) \subset \mathcal{P}$ for all N (from Lemma 3.2), and $\hat{\mathcal{P}}(\infty) = \{p^*\}$ (from Proposition 5.1). If nodes k and ℓ are connected by an active link, the squared distance between them is $2\delta^{k\ell}$ times the equilibrium profit of an arbitrageur (this follows from Lemma 3.3).

We will need some basic definitions and facts about polytopes. Consider a polytope P. The *circumsphere* of P, if it exists, is the sphere that circumscribes P, i.e. whose surface contains all the vertices of P. If the circumsphere exists, its center is called the *circumcenter* of P; it is the point from which all the vertices are equidistant. A polytope P is *centrally symmetric* about 0 if $P = -P$. P is centrally symmetric about a point x if $P = x + P'$, and P' is centrally symmetric about 0; x is called the *center of symmetry* of P. Centrally symmetric polytopes have an even number of vertices: each vertex is symmetric with respect to another vertex. The line segment joining such a pair of vertices is called an *axis* of P. We say that a point is the *center* of P if it is either the circumcenter or the center of symmetry of P (it is both if and only if P is centrally symmetric with equal axes).

We refer to a d-dimensional polytope as a d-polytope. A *simplex* can be defined inductively as follows: a 1-simplex is a line segment, a 2-simplex is a triangle, a 3-simplex is a pyramid bounded by triangles, and so on in higher dimensions. A *regular polytope* is a generalization of a regular polygon, i.e. a polygon that is equilateral and equiangular. Roughly speaking, a polytope is regular if it has congruent faces and angles. All regular polytopes are centrally symmetric, except the simplices of dimension greater than or equal to 2 and the odd polygons (i.e. 2-polytopes with an odd number of vertices). For formal definitions of the dimension of a polytope, of a simplex, and of a regular polytope, see Appendix B.

We call a polytope *symmetric* if it is either regular or centrally symmetric. Of course, it can be both, e.g. the cube. Examples of non-regular polytopes that are centrally symmetric are the parallelogram, or any prism based upon an even regular polygon. Symmetric polytopes provide us with a clear intuition as well as tractable closed-form solutions. Imposing symmetry amounts to assuming that state prices are distributed evenly in Euclidean space.

We refer to nodes $\hat{p}^k$ that are not vertices of $\hat{\mathcal{P}}$ as *internal* nodes (note that the internality of a given node depends on N, and also that an internal node may not be in the interior of $\hat{\mathcal{P}}$).

Proposition 6.1 (Internal nodes) *Suppose that $\mathcal{A}$ is either complete or hub-spoke,*

[8] Only convex polytopes are considered in this paper.

and $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$. Then an equilibrium network never has an active internal node unless $\mathcal{A}$ is hub-spoke with the internal node being the hub.

In particular, if the architecture is complete, an equilibrium network can never be hub-spoke with an internal hub. It can, however, be hub-spoke with a vertex hub, as in Example 4.

In the next section we characterize equilibrium networks when the architecture is complete. In Section 8, we consider equilibrium networks when the architecture is hub-spoke. We say that a network is symmetric if the corresponding polytope is symmetric. We obtain explicit characterizations under the assumption of symmetry, but these results serve to guide our intuition in the general non-symmetric case as well.

We refer to an architecture which is hub-spoke with node k as the hub as the h_k-architecture. We denote by Φ^c and Φ^{h_k} the profit of an arbitrageur (which is the same on all active links) associated with the complete architecture and the h_k-architecture respectively.

7 Equilibrium Networks: Complete Architecture

The main geometric intuition is clearest when the architecture is complete, and arbitraging costs are the same for every link. Consider the polytope $\hat{\mathcal{P}}(N)$. As N increases from zero, arbitrageurs locate on its vertices. They never trade with internal nodes. $\hat{\mathcal{P}}(N)$ is a subset of $\mathcal{P}$. If $N^{k\ell}(N) > 0$, then $\|\hat{p}^k(N) - \hat{p}^\ell(N)\|_2 = [2\delta\Phi^c(N)]^{\frac{1}{2}} = \max_{x,y\in\hat{\mathcal{P}}(N)} \|x-y\|_2$. In other words, an active link arises only between those pairs of vertices that are furthest apart. All such links generate the same profits, i.e. the linked vertices are equally far apart. Not all vertices are necessarily active. But as N increases, the equilibrium polytope contracts along the links that are active, while remaining "centered" around p^*, which is equal to $\sum \lambda^k \hat{p}^k(N)$ for all N. For N large enough, this implies that an inactive vertex becomes active as the length of the active links contracts to the length of the longest link emanating from the hitherto inactive vertex. At the same time, as the polytope contracts, an internal node (typically) becomes a vertex for some N large enough and, at some yet higher N, becomes an active vertex. This pattern continues until the polytope converges to the point p^*.

Under the assumption of symmetry, this convergence is very regular and well-behaved:[9]

Proposition 7.1 (Symmetric networks with complete architecture) *Suppose $\mathcal{A}$ is complete, $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$, and $\beta^k = \beta$ for all $k \in \mathcal{K}$. Suppose further that there is an $\bar{N} \geq 0$ such that $\hat{\mathcal{P}}(\bar{N})$ is symmetric with vertex set $\{\hat{p}^k(\bar{N})\}_{k\in\mathcal{K}}$.[10] Then there is an equilibrium network that can be characterized as follows: for $N \geq \bar{N}$, p^* is the center of $\hat{\mathcal{P}}(N)$, and*

[9] In this result, as well as in later results where we invoke symmetry, we conjecture that this equilibrium network computed in closed form is in fact unique.

[10] In particular, this means that the vertices $\{\hat{p}^k(\bar{N})\}_{k\in\mathcal{K}}$ are distinct.

i. If $\hat{\mathcal{P}}(\bar{N})$ is a regular simplex, so is $\hat{\mathcal{P}}(N)$. The equilibrium network is complete.

ii. If $\hat{\mathcal{P}}(\bar{N})$ is a regular odd polygon, so is $\hat{\mathcal{P}}(N)$. The equilibrium network is a cycle: the neighbors of node k are the vertices of the segment opposite to k.

iii. If $\hat{\mathcal{P}}(\bar{N})$ is centrally symmetric, so is $\hat{\mathcal{P}}(N)$. The equilibrium network is not connected for $K > 1$: the active links correspond to the axes of maximal length. Moreover, there is an $\bar{\bar{N}} \geq \bar{N}$ such that, for all $N \geq \bar{\bar{N}}$, $\hat{\mathcal{P}}(N)$ is centrally symmetric with equal axes.

In all three cases, $N^{k\ell}$ is strictly increasing in N, for $k\ell \in \mathcal{A}^*(N)$, $\Phi^c(N)$ is strictly decreasing in N, and $\hat{p}^k(N)$ converges monotonically to p^* along the line segment $L[\hat{p}^k(\bar{N}), p^*]$, for all $k \in \mathcal{K}$. Furthermore, in cases (i) and (ii), $\mathcal{A}^*(N) = \mathcal{A}^*(\bar{N})$, for $N \geq \bar{N}$, while in case (iii), $\mathcal{A}^*(N) = \mathcal{A}^*(\bar{\bar{N}})$, for $N \geq \bar{\bar{N}}$.[11] Finally, if $\bar{N} = \bar{\bar{N}} = 0$, there is an equal mass of arbitrageurs on each active link.

Note that the cases (i)–(iii) in the proposition cover all possible symmetric (i.e. regular or centrally symmetric) polytopes. We can specialize the proposition to the case where $\bar{N} = 0$, so that $\mathcal{P}$ is symmetric with vertex set $\{p^k\}_{k\in\mathcal{K}}$. Then $\hat{\mathcal{P}}(N)$ is a smaller symmetric polytope within the autarky polytope $\mathcal{P}$, and contracts evenly to p^* as N goes to infinity, with each state-price deflator $\hat{p}^k$ converging on a straight line segment towards p^*. Equilibrium profits converge monotonically to zero.

In case (i), arbitrageurs spread equally across all the edges of the simplex. The same is true in case (ii) if the polygon has three vertices (and is therefore a simplex). If the polygon has five or more vertices,[12] the equilibrium network is connected but not complete: it is a *cycle*, i.e. the $K+1$ nodes can be ordered as $\{k_1, \dots, k_{K+1}\}$ such that $\mathcal{A}^*(N) = \{k_1k_2, k_2k_3, \dots, k_Kk_{K+1}, k_{K+1}k_1\}$. In a cycle, each node has precisely two neighbors. The cycle should obviously not be visualized as the polygon itself, since the neighbors of k are not the nodes adjacent to it in the polytope but the ones that are maximally distant from k. In case (iii), arbitrageurs gravitate to the links that correspond to the longest axes of the polytope. As N increases, these axes become shorter until there is activity on all the axes. From this point on, the equilibrium polytope is centrally symmetric with equal axes. Example 5 is an illustration of this: the autarky polytope is a parallelogram which converges to a rectangle, after which the rectangle shrinks uniformly to its center. While in the case of the simplex every edge corresponds to an active link, no edge is active in the other cases.

For arbitrary polytopes, convergence of state-price deflators to p^* need not be along a linear trajectory, either globally or piecewise. However, even if $\mathcal{P}$ is not symmetric, $\hat{\mathcal{P}}(\bar{N})$ may be for some $\bar{N}$. Interestingly, convergence is linear from that $\bar{N}$ onwards, with active links as described in the proposition.

[11] We define $\mathcal{A}^*(0) := \lim_{N\to 0} \mathcal{A}^*(N)$.

[12] Note that a regular even polygon is centrally symmetric, and hence covered by case (iii) of the proposition.

If $\hat{\mathcal{P}}(\bar{N})$ is symmetric but has some internal nodes, the proposition still applies for $N \geq \bar{N}$ as long as the nodes that are internal for $\hat{\mathcal{P}}(\bar{N})$ are also internal for $\hat{\mathcal{P}}(N)$. Let $N_{\max}$ be the maximum N for which this is the case and let $\bar{\mathcal{K}}$ be the vertex set of $\hat{\mathcal{P}}(\bar{N})$. Then we have linear convergence of $\hat{p}^k$ to $p^*_{\bar{\mathcal{K}}}$, for $N \in [\bar{N}, N_{\max}]$, for all $k \in \bar{\mathcal{K}}$.

Figure 4 depicts the case where the autarky polytope is a simplex with node 0 at the center. Node 0 is internal, and hence inactive, for all N. Proposition

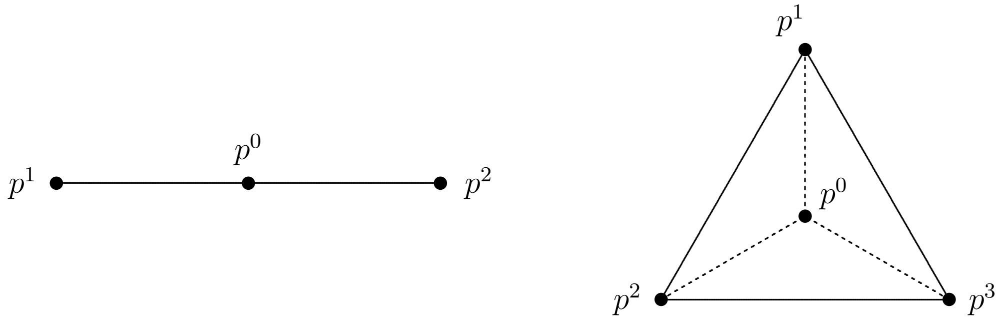

(a) 1-simplex

(b) 2-simplex

Figure 4: **Convergence in the complete architecture** ($p^0 = p^*$)

7.1 applies to the polytope generated by $\{p^k\}_{k \neq 0}$. In the 1-simplex (which is also centrally symmetric), $\hat{p}^1$ and $\hat{p}^2$ converge along the segment to the center $p^0 = p^*$, while in the 2-simplex the equilibrium deflators on the vertices converge linearly to the center along the dotted lines.

8 Equilibrium Networks: Hub-Spoke Architecture

We now analyze the geometry of hub-spoke networks. The following proposition provides a characterization of symmetric networks when the hub is at the center of the polytope.

Proposition 8.1 (Symmetric networks with central hub) *Consider the h_0-architecture with $K \geq 2$. Suppose $\delta^{0k} = \delta$ and $\beta^k = \beta$, for all $k \neq 0$. Suppose further that $\mathcal{P}$ is symmetric with vertex set $\{p^k\}_{k \neq 0}$, and $p^0 = p^*$. Then there is an equilibrium network that can be characterized as follows: for all N, p^* is the center of $\hat{\mathcal{P}}(N)$, and*

i. If $\mathcal{P}$ is a regular simplex or a regular odd polygon, so is $\hat{\mathcal{P}}(N)$, and $N^{0k} = N/K$ for all $k \neq 0$.

ii. If $\mathcal{P}$ is centrally symmetric, so is $\hat{\mathcal{P}}(N)$. There is an $\bar{N}$ such that, for all $N \geq \bar{N}$, $\hat{\mathcal{P}}(N)$ is centrally symmetric with equal axes, and all spokes are active. If $\mathcal{P}$ is centrally symmetric with equal axes, then $N^{0k} = N/K$ for all $k \neq 0$.

In both cases, $\Phi^{h_0}(N)$ is strictly decreasing in N, $\hat{p}^0(N) = p^0 = p^*$, and $\hat{p}^k(N)$ converges monotonically to p^* along $L[p^k, p^*]$, for all $k \neq 0$.

Due to the symmetry of the nodes with respect to the central hub, the state-price deflator on the hub remains unchanged at p^* as N increases. As in the case of the complete architecture, the equilibrium polytope contracts to its center p^*, with the state-price deflators on the spokes following a linear trajectory towards it. For example, in the case of the simplices shown in Figure 4, convergence is along the same trajectory as in the complete architecture. However, as we shall see in the next section, convergence is slower as N increases, and arbitrageurs earn higher profits for any given N.

The analysis of hub-spoke networks is less straightforward if state prices on the hub depend on the mass of arbitrageurs. We can, however, obtain an explicit characterization in the case where the autarky polytope is a simplex and the hub is one of its vertices. When we consider vertex hubs we adopt the convention of choosing exchange 1 as the hub (this is to facilitate comparison with a central hub).

Proposition 8.2 (Simplex networks with vertex hub) *Consider the h_1-architecture with $K \geq 2$. Suppose $\delta^{1k} = \delta$ for all $k \neq 1$, and $\beta^k = \beta$ for all $k \in \mathcal{K}$. Suppose further that $\mathcal{P}$ is a regular simplex with vertex set $\{p^k\}_{k\in\mathcal{K}}$. Then, p^* is the center of $\mathcal{P}$, and there is an equilibrium network with the following properties: $N^{1k} = N/K$ for all $k \neq 1$, $\Phi^{h_1}(N)$ is strictly decreasing in N, $\hat{p}^1(N)$ converges monotonically to p^* along $L[p^1, p^*]$, and $\hat{p}^k(N)$ converges monotonically to $\hat{p}^1(N)$ along $L[p^k, \hat{p}^1(N)]$, for $k \neq 1$.*

Under the symmetry assumptions of the proposition, the equilibrium state-price deflator on the hub is pulled evenly "from all sides" and follows a linear trajectory towards p^*. Equilibrium state-price deflators on other exchanges converge linearly to the equilibrium deflator on the hub, and hence converge to p^* along an arc. This is illustrated in Figure 5(b), where the dotted lines indicate the path along which the equilibrium deflators travel as N increases. Here p^0 is at the center of the simplex,

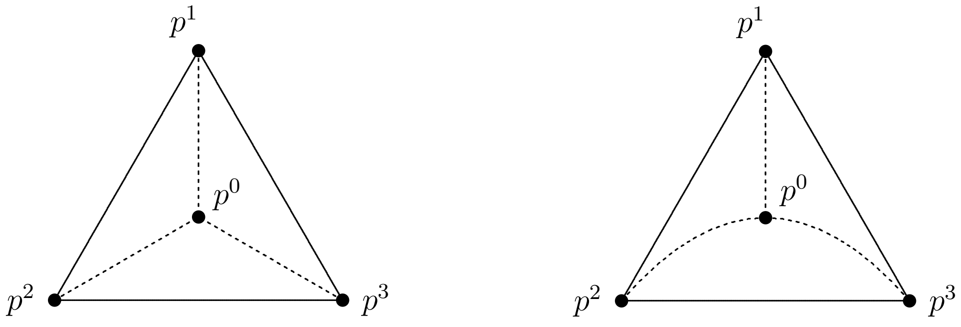

(a) Complete architecture or 0-hub

(b) 1-hub

Figure 5: **Convergence in a 2-simplex ($p^0 = p^*$)**

so that Proposition 8.2 applies to the simplex with vertices $\{p^k\}_{k\neq 0}$. For comparison, Figure 5(a) shows the case where exchange 0 is the hub. As we remarked above, this is the same pattern of convergence as in the complete architecture (Figure 4(b)).

More generally, convergence need not be monotonic, as we show in the following example:

Example 6 (Non-monotonic convergence) Consider the h_0-architecture with four equally deep exchanges.[13] Arbitraging costs are the same for all admissible links. Autarky state-price deflators lie on a straight line in $\mathbb{R}^S$, as depicted in Figure 6. The

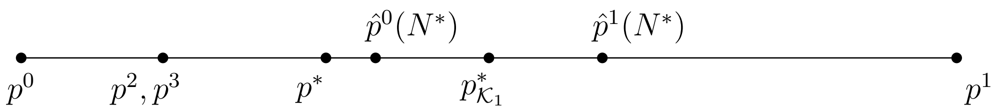

Figure 6: **Non-monotonic convergence of $\hat{p}^0(N)$ to p^***

nodes p^0 and p^1 are the endpoints of the segment, while $p^2 = p^3$ is located between p^0 and $p^*_{\mathcal{K}_1}$, $\mathcal{K}_1 = \{0, 1\}$. Note that $p^*_{\mathcal{K}_1} = \frac{1}{2}(p^0 + p^1)$. Since the gains from trade between 0 and 1 are the largest, for small enough N we have $N^{01} = N$, with $\hat{p}^0(N)$ and $\hat{p}^1(N)$ converging linearly to $p^*_{\mathcal{K}_1}$. This will be true until $N = N^*$ which satisfies $\hat{p}^0(N^*) - p^2 = \hat{p}^1(N^*) - \hat{p}^0(N^*)$. Due to (5), we have $p^*_{\mathcal{K}_1} = \frac{1}{2}[(\hat{p}^0(N^*) + \hat{p}^1(N^*)]$, so that $\hat{p}^1(N^*) - \hat{p}^0(N^*) = 2[p^*_{\mathcal{K}_1} - \hat{p}^0(N^*)]$. Therefore, $\hat{p}^0(N^*) = \frac{1}{3}p^2 + \frac{2}{3}p^*_{\mathcal{K}_1}$. For $N > N^*$, all admissible links are active and $\hat{p}^0(N)$ converges linearly to p^*. Now $p^* = \frac{1}{4}(p^0 + p^1 + p^2 + p^3) = \frac{1}{2}(p^2 + p^*_{\mathcal{K}_1})$, i.e. p^* lies between p^0 and $\hat{p}^0(N^*)$. So for $N > N^*$, $\hat{p}^0(N)$ reverts back in the direction of p^0 towards its limit p^*. $\diamondsuit$

9 Comparing Network Architectures

In this section we study comparative statics with respect to the network architecture. We first consider the case where the autarky polytope is symmetric, and compare the complete architecture to the hub-spoke architecture in which the hub is at the center of the polytope.

Proposition 9.1 (Complete architecture vs central hub) *Suppose $K \geq 2$, $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$, and $\beta^k = \beta$ for all $k \neq 0$. Suppose further that $\mathcal{P}$ is symmetric with vertex set $\{p^k\}_{k\neq 0}$, and $p^0 = p^*$. Then, in the complete architecture, node 0 is inactive for all N. We have $\Phi^{h_0}(N) > \Phi^c(N)$, provided $N > \frac{\delta}{\beta}|\mathcal{A}^*|$, where $|\mathcal{A}^*|$ is the number of active links in the complete architecture.*[14]

[13] Examples of non-monotonic convergence can also be constructed for the complete architecture. We choose to present an example with a hub-spoke architecture because it is simpler.

[14] From Proposition 7.1 it follows that $\mathcal{A}^*$ does not depend on N.

Thus, if all exchanges other than 0 have equal depth, and arbitraging costs are the same for all admissible links, the complete architecture is dominated by the central hub architecture as long as the mass of arbitrageurs is sufficiently large. The central hub architecture leads to higher arbitrageur payoffs even if the depth of the hub is lower than that of the spokes (and regardless of how low that depth is).

For the sake of intuition, let us first focus on the case where $\mathcal{P}$ is centrally symmetric (the simplest case is that of a line segment, as in Figure 4(a)). Then the vertices come in pairs (k, ℓ_k) that are symmetrical with respect to p^*. Equilibrium profits are higher with a central hub, as long as N is sufficiently large. The reason is the positive externality that arbitrageurs on link $0k$ exert on arbitrageurs on the symmetrical link $0\ell_k$. Arbitrageurs on $0k$ pull $\hat{p}^0$ towards p^k, thereby increasing $\|\hat{p}^0 - \hat{p}^{\ell_k}\|_2$. Arbitrageurs on $0\ell_k$ pull $\hat{p}^0$ in the opposite direction, towards p^{ℓ_k}. Due to symmetry, the net impact on state prices on the hub is in fact zero, i.e. $\hat{p}^0 = p^0$. The aggregate supply of arbitrageurs on the hub is also zero; any state-contingent consumption that is supplied to the hub by arbitrageurs on one spoke is absorbed by arbitrageurs on the symmetrical spoke. Thus the hub acts as a liquidity repository, channeling trades in such a manner that the two groups of arbitrageurs on each pair of spokes complement each other. This network-induced complementarity is sufficient to compensate for the fact that the autarky gains from trade between the center and any one of the extremes are considerably lower than the gains from trade between two symmetrically located extremes, ignoring the central exchange altogether.

When the architecture is complete, the market mechanism fails to achieve this outcome due to a Prisoner's Dilemma. If arbitrageurs could agree to designate the central exchange as the hub, and trade only through it, they would be able to minimize price impact and increase profits. But, given the opportunity, each arbitrageur would rather arbitrage one of the longer links corresponding to an axis of the polytope. The result is a suboptimal arrangement with all arbitrageurs on the axes.

Notice that the common measure of liquidity as depth does not do justice to exchange 0. If exchange 0 is the hub, it will attract a lot of trade with zero equilibrium price impact, irrespective of β^0. On the other hand, in the complete architecture, there is no trade with exchange 0. An important determinant of the liquidity of an exchange is the network architecture, and the position of the exchange in the network.

For the case where $\mathcal{P}$ is symmetric but not centrally symmetric, essentially the same intuition applies. Once again, consider the hub-spoke architecture, with the central exchange 0 as the hub. For each vertex k there is a facet that is opposite to it. The vertices of this facet pull $\hat{p}^0$ away from p^k. Due to symmetry, the net impact on $\hat{p}^0$ is zero. This is illustrated in Figure 5(a).

In fact, if $\mathcal{P}$ is a regular simplex, a hub-spoke architecture delivers higher profits for arbitrageurs than the complete architecture even if the hub is a vertex of the simplex. We have the following result:[15]

[15] The condition $K \geq 3$ is only needed here to ensure that the simplex is not simply a line segment, for in that case the complete and vertex hub architectures give rise to the same equilibrium network.

Proposition 9.2 (Complete architecture vs central/vertex hub) *Suppose $K \geq 3$, $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$, and $\beta^k = \beta$ for all $k \in \mathcal{K}$. Suppose further that $\mathcal{P}$ is a regular simplex with vertex set $\{p^k\}_{k \neq 0}$, and $p^0 = p^*$. Then, in the complete and h_1-architectures, node 0 is inactive for all N. We have $\Phi^{h_0}(N) > \Phi^{h_1}(N) > \Phi^c(N)$, provided $N \geq \frac{\delta}{\beta} K$.*

Since node 0 is always inactive in the complete and h_1-architectures, the ranking of profits in these architectures also holds if $\mathcal{P}$ has no internal nodes (replacing K by $K+1$ in the proposition).

The reason why a vertex hub leads to higher profits is similar to the one we gave for the central hub. Consider vertices k and ℓ, neither of which is the hub. As N increases, $\hat{p}^k$ and $\hat{p}^\ell$ converge to p^* along an arc, while $\hat{p}^1$ moves more quickly[16] towards the middle between $\hat{p}^k$ and $\hat{p}^\ell$ (see Figure 5(b)). The effect of this is again to induce a complementarity, with arbitrageurs on opposite sides of the hub providing liquidity to each other to some extent (but less than in the central hub case). If trade on $k\ell$ was now allowed, every arbitrageur would want to deviate to that link. Again a Prisoner's Dilemma type of result obtains. In particular, in Figure 5(b), adding the link 23, and thus moving from the 1-hub architecture to the complete architecture, makes all arbitrageurs worse off.

Thus, for both the central and vertex hub-spoke architectures, the restrictions implicit in the architecture coordinate arbitrageur actions by pooling liquidity and by preventing Prisoner's Dilemma type deviations. The state prices on the hub are in equilibrium (though not necessarily in autarky) to some extent "in-between" the other state prices. Arbitrageurs on one side of the hub generate positive externalities for arbitrageurs on the other side.

One way to interpret Propositions 9.1 and 9.2 is that profits are lower in the complete architecture because convergence to the economy-wide state-price deflator p^* is faster as we increase N. Let $d^{k,c}(N)$ and $d^{k,h_\ell}(N)$ denote $\|\hat{p}^k(N) - p^*\|_2^2$ for the complete and h_ℓ-architectures respectively. The following result complements Proposition 9.1:

Proposition 9.3 (Speed of convergence: complete architecture vs central hub) *Suppose $K \geq 2$, $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$, and $\beta^k = \beta$ for all $k \neq 0$. Suppose further that $\mathcal{P}$ is symmetric with vertex set $\{p^k\}_{k \neq 0}$, and $p^0 = p^*$. Then, for all $N > 0$,*

$$\begin{array}{lllll} d^{0,c}(N) & = & d^{0,h_0}(N) & = & 0, \\ d^{k,c}(N) & < & d^{k,h_0}(N), & & k \neq 0. \end{array}$$

Also, $d^{k,c}(N)$ and $d^{k,h_0}(N)$ are strictly decreasing in N for $k \neq 0$.

For example, in Figure 5(a), while the pattern of convergence is the same for both architectures, it is faster in the case of the complete architecture. Our final result provides a similar interpretation for Proposition 9.2:

[16] Proposition 9.4 below implies that $\hat{p}^1(N)$ is closer to p^* than $\hat{p}^k(N)$, for any vertex $k \neq 1$.

Proposition 9.4 (Speed of convergence: complete vs central/vertex hub) *Suppose $K \geq 3$, $\delta^{k\ell} = \delta$ for all $k\ell \in \mathcal{A}$, and $\beta^k = \beta$ for all $k \in \mathcal{K}$. Suppose further that $\mathcal{P}$ is a regular simplex with vertex set $\{p^k\}_{k \neq 0}$, and $p^0 = p^*$ Then, for all $N > 0$,*

$$
\begin{array}{lclcll}
d^{0,c}(N) & = & d^{0,h_1}(N) & = & d^{0,h_0}(N) & = \ 0, \\
d^{1,h_1}(N) & < & d^{1,c}(N) & < & d^{1,h_0}(N), & \\
d^{k,c}(N) & < & d^{k,h_1}(N) & < & d^{k,h_0}(N), & \qquad k \notin \{0,1\}.
\end{array}
$$

Also, the distance of each node k, $k \neq 0$, from p^* is strictly decreasing in N for all three architectures.

Like Proposition 9.2, this result holds for the complete and h_1-architectures even if $\mathcal{P}$ has no internal nodes. In Figure 5, the speed of convergence of $\hat{p}^2$ and $\hat{p}^3$ towards the center is fastest when the architecture is complete, and slowest when the center is the hub. On the other hand, $\hat{p}^1$ converges to the center at the fastest rate when node 1 is the hub.

While the results in this section require the assumption of symmetry, the intuition is clearly more general. Profits are higher when the architecture induces arbitrageurs to exert positive externalities on each other. In such an architecture, trades are channeled through a hub (or possibly several hubs). A good candidate for a hub is a node that lies towards the center of all the nodes.

10 Conclusion

Network structure is irrelevant in frictionless financial markets. Actual markets are segmented, however, and a natural question that arises is how the various market segments are linked to each other. In this paper we obtain a tractable framework by assuming that demand functions are affine.

Equilibrium networks display several subtle and interesting features. Even if all links are admissible, an equilibrium network may not be complete or even connected. Prices nevertheless converge (not necessarily monotonically) to those that would obtain in a frictionless economy, as the mass of intermediaries grows without bound. In particular, price differentials between a pair of nodes converge to zero even if there is no active link between these nodes.

While the connectivity properties of equilibrium networks can be completely arbitrary in general, symmetric networks have an elegant and regular structure. We focus on complete and hub-spoke network architectures, and in particular on the nature of network externalities induced by these architectures. Hub-spoke architectures generally lead to higher payoffs for intermediaries by pooling liquidity.

We explicitly characterize equilibrium networks only in the symmetric case. But the intuitions are more general. Centrally located nodes can serve as a repository of liquidity, even if their depth is low. Arbitrageurs on one "side" of a centrally located node exert positive externalities on arbitrageurs on the other "side". These externalities are of the Prisoner's Dilemma type.

A number of interesting questions remain. We address some of these, in particular financial innovation and welfare in arbitrage networks, in ongoing work. The question of how local shocks are propagated through the economy, via the endogenous linkages created by intermediaries, is a topic for future research; this is one area in which network structure is likely to play a key role. Also of interest are various network-theoretic questions, such as the properties of other architectures that have been studied in the networks literature (e.g. interlinked stars and core-periphery networks), and the relative importance of different links for agents' payoffs.

Appendices

A Proofs

Proof of Lemma 3.2 Fix a $\mathcal{C} \in \mathbb{C}(\mathfrak{N})$. For $k \in \mathcal{C}$, let $\alpha^k := \sum_{\ell\in\mathcal{K}} \alpha^{k\ell} = \sum_{\ell\in\mathcal{C},\ell\neq k} \alpha^{k\ell}$. Then we can write the equation system (4) for the given $\mathcal{C}$ as follows:

$$(1 + \beta^k\alpha^k)\hat{p}^k - \beta^k \sum_{\ell\in\mathcal{C},\ell\neq k} \alpha^{k\ell}\hat{p}^\ell = p^k, \qquad k \in \mathcal{C}. \tag{6}$$

Define the matrix $M = [m_{k\ell}]_{k,\ell\in\mathcal{C}}$ by $m_{kk} := 1 + \beta^k\alpha^k$ and $m_{k\ell} := -\beta^k\alpha^{k\ell}$ for $\ell \neq k$. Since $\sum_{\ell\neq k} |m_{k\ell}| = \beta^k\alpha^k < |m_{kk}|$, M is strictly row diagonally dominant. Let $\mathbf{M} := M \otimes I_{S\times S}$. Then, letting $\hat{p} := \{\hat{p}^k\}_{k\in\mathcal{C}}$ and $p := \{p^k\}_{k\in\mathcal{C}}$, the equation system (6) can be written as $\mathbf{M}\hat{p} = p$. $\mathbf{M}$ inherits the property of strict row diagonal dominance from M.

We appeal to the theory of M-matrices; see Berman and Plemmons (1979), henceforth BP (an M-matrix is a square matrix of the form $sI - A$, where $A \geq 0$, and $s \geq \mathrm{rad}(A)$, the spectral radius of A). By Theorem 6.2.3 in BP (in particular, condition M_{35} on page 137), both $\mathbf{M}$ and M are nonsingular M-matrices. Hence, there exists a unique $\hat{p}$ solving $\mathbf{M}\hat{p} = p$, namely $\hat{p} = \mathbf{M}^{-1}p$. Now fix a $k \in \mathcal{C}$. Since $\mathbf{M}^{-1} = M^{-1} \otimes I_{S\times S}$, we can write $\hat{p}^k = \sum_{\ell\in\mathcal{C}} \eta^{k\ell}p^\ell$, where $\eta^{k\ell}$ is element (k, ℓ) of M^{-1}. The matrix M is irreducible, also called indecomposable (indeed it is irreducible if and only if $\mathcal{C}$ is connected; see Theorem 2.2.7 in BP). Hence $M^{-1} \gg 0$ by Theorem 6.2.7 in BP, i.e. $\eta^{k\ell} > 0$, all $\ell \in \mathcal{C}$. Let $\mathbf{1} := (1 \dots 1)^\top$. Since $M\mathbf{1} = \mathbf{1}$ we also have $M^{-1}\mathbf{1} = \mathbf{1}$, i.e. $\sum_{\ell\in\mathcal{C}} \eta^{k\ell} = 1$. Let B be the diagonal matrix with typical diagonal entry β^ℓ, $\ell \in \mathcal{C}$. Notice that $B^{-1}M$ is symmetric. Therefore, $M^{-1}B$ is symmetric as well, i.e. $\beta^\ell\eta^{k\ell} = \beta^k\eta^{\ell k}$, $\ell \in \mathcal{C}$. Finally, from Theorem 2.5.12 in Horn and Johnson (1991), M^{-1} is strictly diagonally dominant of its column entries: $\eta^{kk} > \eta^{\ell k}$, $\forall \ell \in \mathcal{C}, \ell \neq k$. $\square$

Proof of Proposition 4.1 Let $n^{k\ell} := N^{k\ell}/N$, and $d := |\mathcal{A}|$. Then $\boldsymbol{n} := \{n^{k\ell}\}_{k\ell\in\mathcal{A}}$ lies in the standard simplex $\Delta_{d-1} := \{x \in \mathbb{R}^d_+ : \sum_{i=1}^d x_i = 1\}$. For given N, we can regard $\varphi^{k\ell}$ as a function of $\boldsymbol{n}$. Also let $\Phi(\boldsymbol{n}) := \sum_{k\ell\in\mathcal{A}} n^{k\ell}\varphi^{k\ell}(\boldsymbol{n})$, the average profit, and define $f(\boldsymbol{n}) := [\varphi^{k\ell}(\boldsymbol{n}) - \Phi(\boldsymbol{n})]_{k\ell\in\mathcal{A}}$. Then $f : \Delta_{d-1} \to \mathbb{R}^d$ is continuous, and $f(\boldsymbol{n}) \cdot \boldsymbol{n} = 0$, for all $\boldsymbol{n}$. By Theorem 8.3 in Border (1989), there is an $\boldsymbol{n}^*$ such that $f(\boldsymbol{n}^*) \leq 0$. Moreover, since $f(\boldsymbol{n}^*) \cdot \boldsymbol{n}^* = 0$, we must have $n^{*k\ell} f^{k\ell}(\boldsymbol{n}^*) = 0$, for all $k\ell \in \mathcal{A}$. Therefore, the arbitrageur distribution $N\boldsymbol{n}^*$ satisfies condition (**N**). The result then follows from Lemma 4.1. $\square$

Proof of Proposition 5.1 The equilibrium profit of an arbitrageur is $\Phi(N) = \frac{1}{2\delta^{k\ell}}\|\hat{p}^k(N) - \hat{p}^\ell(N)\|_2^2$, $k\ell \in \mathcal{A}^*(N)$. We claim that $\Phi(\infty) = 0$. If not, there exists a constant $\underline{\kappa} > 0$ such that $\Phi(N) > \underline{\kappa}$ for arbitrarily large N. Now, as N goes to infinity, $N^{k\ell}$ is unbounded, for some $k\ell \in \mathcal{A}$. For such a link $k\ell$, total arbitrageur supply on k by arbitrageurs on $k\ell$, given by $N^{k\ell} y^k_{k\ell} = \frac{N^{k\ell}}{2\delta^{k\ell}}[\hat{p}^k(N) - \hat{p}^\ell(N)]$, is unbounded: for any constant $\bar{\kappa}$, howsoever large, there is an N and a state s for which this supply is greater than $\bar{\kappa}$ in absolute value. Suppose that the supply is in fact positive (if it is negative, then we can consider instead the state-s supply by arbitrageurs on $k\ell$ to ℓ). Since $\hat{p}^k \geq 0$, (3) implies that y^k is bounded above. Due to (2), $N^{kk_1} y^k_{kk_1}$ must be unboundedly negative in state s, for some k_1. But then the state-s supply on exchange k_1 by arbitrageurs active on kk_1 is unboundedly positive, which places k_1 in the same situation as k was. Eventually the unboundedly large supply in state s must end up on some exchange k_m. If $k_m \neq k$, the condition that $\hat{p}^{k_m} \geq 0$ will be violated due to (3). If $k_m = k$, consider the following sequence of inequalities that must hold: $\hat{p}^k_s < \hat{p}^{k_1}_s < \hat{p}^{k_2}_s < \ldots < \hat{p}^{k_m}_s = \hat{p}^k_s$, a contradiction. Basically, arbitrageurs cannot be trading unboundedly large amounts without running afoul of the fact that equilibrium prices are nonnegative on every exchange.

Since $\Phi(N) \to 0$, and $\Phi(N) \geq \frac{1}{2\delta^{k\ell}}\|\hat{p}^k(N) - \hat{p}^\ell(N)\|_2^2$, for all $k\ell \in \mathcal{A}$, we must have $\hat{p}^k(N) - \hat{p}^\ell(N)$ converging to zero, for all $k\ell \in \mathcal{A}$. We claim that this is in fact true for all $k, \ell \in \mathcal{K}$. For arbitrary k and ℓ, there is a path connecting them, since $\mathcal{G}$ is connected, i.e. there is a sequence of distinct vertices $\{k_1, \ldots, k_I\}$ in $\mathcal{K}$ such that $k_1 = k$, $k_I = \ell$ and $(k_i, k_{i+1}) \in \mathcal{A}$ for all $i = 1, \ldots, I-1$. By the triangle inequality, $\|\hat{p}^k(N) - \hat{p}^\ell(N)\|_2^2 \leq \sum_{i=1}^{I-1} \|\hat{p}^i(N) - \hat{p}^{i+1}(N)\|_2^2$. Since each of the terms in the sum converges to zero, $\hat{p}^k(N) - \hat{p}^\ell(N)$ converges to zero as well.

Using (5), and the triangle inequality, we have, for all $k \in \mathcal{K}$,

$$\|\hat{p}^k(N) - p^*\|_2 = \left\|\sum_{j\in\mathcal{K}} \lambda^j \left[\hat{p}^k(N) - \hat{p}^j(N)\right]\right\|_2 \leq \sum_{j\in\mathcal{K}} \lambda^j \|\hat{p}^k(N) - \hat{p}^j(N)\|_2.$$

Since each term in the last sum converges to zero, $\hat{p}^k(N)$ converges to p^*, for all $k \in \mathcal{K}$.

Now suppose $\mathcal{G}^*(\infty)$ exists, and consider a component $\mathcal{C}$ of $\mathcal{G}^*(\infty)$. Define $\bar{N}$ large enough so that for all $N > \bar{N}$, $\mathcal{G}^*(N) = \mathcal{G}^*(\infty)$. Multiplying the equation system (4), corresponding to the component $\mathcal{C}$, by $\lambda^k_{\mathcal{C}}$, and summing over $k \in \mathcal{C}$, we have $\sum_{k\in\mathcal{C}} \lambda^k_{\mathcal{C}} \hat{p}^k(N) = p^*_{\mathcal{C}}$, for all $N > \bar{N}$. Taking limits as N goes to infinity, we get $p^*_{\mathcal{C}} = p^*$. $\square$

Proof of Proposition 6.1 Consider first the complete architecture and suppose

$\hat{p}^j$ is internal. Then $\hat{p}^j = \sum_{k\in\mathcal{K}^*} \nu^k \hat{p}^k$, where $\mathcal{K}^*$ is a subset of the vertex set of $\hat{\mathcal{P}}$, $|\mathcal{K}^*| \geq 2$, and the weights $\{\nu^k\}_{k\in\mathcal{K}^*}$ are strictly positive and sum to one. For any $\ell \in \mathcal{K}$, we have

$$\begin{aligned}\|\hat{p}^j - \hat{p}^\ell\|_2^2 = \Big\| \sum_{k\in\mathcal{K}^*} \nu^k (\hat{p}^k - \hat{p}^\ell) \Big\|_2^2 &\leq \Big[\sum_{k\in\mathcal{K}^*} \nu^k \|\hat{p}^k - \hat{p}^\ell\|_2 \Big]^2 &(7)\\ &< \sum_{k\in\mathcal{K}^*} \nu^k \|\hat{p}^k - \hat{p}^\ell\|_2^2 &(8)\\ &\leq \max_{k\in\mathcal{K}^*} \|\hat{p}^k - \hat{p}^\ell\|_2^2,\end{aligned}$$

where (7) follows from the triangle inequality and (8) from Jensen's inequality. Hence $\varphi^{j\ell} < \varphi^{k\ell}$, for some $k \in \mathcal{K}^*$, so that $j\ell$ is inactive.

For a hub-spoke architecture the same argument goes through, taking ℓ to be the hub. □

Proof of Proposition 7.1 Since $\beta^k = \beta$, all k, (5) implies that

$$p^* = \frac{1}{K+1} \sum_{k\in\mathcal{K}} \hat{p}^k(N), \qquad \forall N. \tag{9}$$

Therefore, p^* is the center of $\hat{\mathcal{P}}(N)$ as long as this polytope is symmetric.

Let $\{\bar{N}^{k\ell}\}$ be the equilibrium arbitrageur distribution for $N = \bar{N}$. We first consider cases (i) and (ii). We conjecture that, for $N \geq \bar{N}$, there is a strictly decreasing function $x(N)$, with $x(\bar{N}) = 1$ and $x(\infty) = 0$, such that

$$\hat{p}^k(N) = x(N)\hat{p}^k(\bar{N}) + [1 - x(N)]p^*, \tag{10}$$

$$N^{k\ell}(N) = \frac{1}{x(N)} \left[\bar{N}^{k\ell} + \frac{Nx(N) - \bar{N}}{|\mathcal{A}^*(\bar{N})|} \right], \qquad k\ell \in \mathcal{A}^*(N), \tag{11}$$

$$\mathcal{A}^*(N) = \mathcal{A}^*(\bar{N}), \tag{12}$$

$$\Phi^c(N) = [x(N)]^2 \Phi^c(\bar{N}). \tag{13}$$

It follows from (11) and (12) that $\sum_{k\ell\in\mathcal{A}^*(N)} N^{k\ell}(N) = N$. Also, for $N \geq \bar{N}$, we have $\mathbb{C}(N) = \mathbb{C}(\bar{N})$, where $\mathbb{C}(N)$ is the set of components of $\mathcal{G}^*(N)$.

We see from (4) that, for $N \geq \bar{N}$, equilibrium prices solve the following system of equations:

$$\Delta_k := \hat{p}^k(N) - p^k + \frac{\beta}{\delta} \sum_{\ell\in\mathcal{C}} N^{k\ell}[\hat{p}^k(N) - \hat{p}^\ell(N)] = 0, \qquad k \in \mathcal{C},\, \mathcal{C} \in \mathbb{C}(\bar{N}). \tag{14}$$

Using (10) and (11):

$$\begin{aligned}\Delta_k = x(N)\hat{p}^k(\bar{N}) + [1 - x(N)]p^* - p^k + \frac{\beta}{\delta} \sum_{\ell\in\mathcal{C}} \bar{N}^{k\ell}[\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})]\\ + \frac{\beta}{\delta} \cdot \frac{Nx(N) - \bar{N}}{|\mathcal{A}^*(\bar{N})|} \sum_{\ell\in\mathcal{C}} [\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})].\end{aligned}$$

Using (14) evaluated at $N = \bar{N}$, the above equation simplifies to:

$$\Delta_k = [1 - x(N)][p^* - \hat{p}^k(\bar{N})] + \frac{\beta}{\delta} \cdot \frac{Nx(N) - \bar{N}}{|\mathcal{A}^*(\bar{N})|} \sum_{\ell \in \mathcal{C}} [\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})]. \tag{15}$$

If $\hat{\mathcal{P}}(\bar{N})$ is a regular simplex, we conjecture that all links are active, so that there is a single component $\mathcal{C}$, equal to the set $\mathcal{K}$ of all the nodes, and

$$|\mathcal{A}^*(\bar{N})| = \frac{1}{2} K(K+1). \tag{16}$$

Using (9),

$$\sum_{\ell \in \mathcal{C}} [\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})] = (K+1)[\hat{p}^k(\bar{N}) - p^*].$$

Therefore, from (15), $\Delta_k = 0$, for all k, if and only if

$$x(N) = \frac{\delta K + 2\beta \bar{N}}{\delta K + 2\beta N}. \tag{17}$$

It follows from (10) that

$$\|\hat{p}^k(N) - \hat{p}^\ell(N)\|_2^2 = [x(N)]^2 \|\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})\|_2^2, \tag{18}$$

which is independent of $k\ell$ due to the regularity of the simplex $\hat{\mathcal{P}}(\bar{N})$. Therefore, profits are equalized across all links, with $\Phi^c(N) = [x(N)]^2 \Phi^c(\bar{N})$, and the network is complete, for all $N \geq \bar{N}$. Indeed, as N increases, all the edges of $\hat{\mathcal{P}}(N)$ contract uniformly, so that $\hat{\mathcal{P}}(N)$ is a smaller K-simplex within $\hat{\mathcal{P}}(\bar{N})$, with the same center p^*.

Now suppose $\hat{\mathcal{P}}(\bar{N})$ is a regular odd polygon. Clearly, the nodes that have maximal distance from k are the vertices of the segment that is opposite to k. Let us denote these by ℓ_k and m_k. Then we have, for $k \in \mathcal{C}$,

$$\begin{aligned}\sum_{\ell \in \mathcal{C}} [\hat{p}^k(\bar{N}) - \hat{p}^\ell(\bar{N})] &= 2 \left[\hat{p}^k(\bar{N}) - \frac{1}{2} [\hat{p}^{\ell_k}(\bar{N}) + \hat{p}^{m_k}(\bar{N})] \right] \\ &= 2 \left[\hat{p}^k(\bar{N}) - p^* \right] \left[1 + \cos\left(\frac{\pi}{K+1} \right) \right],\end{aligned}$$

where the second equality follows from a simple trigonometric calculation (see Coxeter (1963), Fig. 1.1A). Also,

$$|\mathcal{A}^*(\bar{N})| = K + 1. \tag{19}$$

Therefore, from (15), $\Delta_k = 0$, for all k, if and only if

$$x(N) = \frac{\delta(K+1) + 2\beta \bar{N} \left[1 + \cos\left(\frac{\pi}{K+1} \right) \right]}{\delta(K+1) + 2\beta N \left[1 + \cos\left(\frac{\pi}{K+1} \right) \right]}. \tag{20}$$

Just as in the case of the simplex, equation (18) holds. The sides of the polygon shrink uniformly as N increases.

Finally, we turn to case (iii) in which $\hat{\mathcal{P}}(\bar{N})$ is centrally symmetric. As in cases (i) and (ii), we conjecture that equations (10)–(13) hold, but with $(\bar{N}, \bar{N}^{k\ell})$ replaced by $(\bar{\bar{N}}, \bar{\bar{N}}^{k\ell})$. Then, clearly, equation (14) applies and hence so does (15), once we replace $\bar{N}$ by $\bar{\bar{N}}$.

Consider a link $k\ell_k$ corresponding to an axis of $\hat{\mathcal{P}}(\bar{N})$ that is maximal in length (there can be many such axes). Clearly there is no activity on any of the shorter axes. There is also no activity on any link $m\ell$ that does not correspond to an axis, since

$$\begin{aligned}
\|\hat{p}^k(\bar{N}) - \hat{p}^{\ell_k}(\bar{N})\|_2 &= \|\hat{p}^k(\bar{N}) - p^*\|_2 + \|p^{\ell_k}(\bar{N}) - p^*\|_2 && (21)\\
&\geq \|\hat{p}^m(\bar{N}) - p^*\|_2 + \|\hat{p}^\ell(\bar{N}) - p^*\|_2 && (22)\\
&> \|[\hat{p}^m(\bar{N}) - p^*] + [p^* - \hat{p}^\ell(\bar{N})]\|_2 && (23)\\
&= \|\hat{p}^m(\bar{N}) - \hat{p}^\ell(\bar{N})\|_2.
\end{aligned}$$

Equation (21) follows from the symmetry of k and ℓ_k with respect to p^*, (22) holds as an equality if and only if both m and ℓ correspond to one end of an axis of maximal length, and (23) follows from the triangle inequality, which is a strict inequality because because m and ℓ are not the end points of an axis (and hence $\hat{p}^m(\bar{N}) - p^*$ is not proportional to $\hat{p}^\ell(\bar{N}) - p^*$). As N increases, the axes of maximal length shrink uniformly towards the center of symmetry p^* until the length of all axes is equalized at some $N = \bar{\bar{N}}$.

At $\bar{\bar{N}}$, the active links are $k\ell_k$, all $k \in \mathcal{K}$. Thus

$$|\mathcal{A}^*(\bar{\bar{N}})| = \frac{1}{2}(K+1). \tag{24}$$

Since $\hat{p}^{\ell_k}(\bar{\bar{N}}) = -\hat{p}^k(\bar{\bar{N}}) + 2p^*$, we have, for $k \in \mathcal{C}$,

$$\sum_{\ell \in \mathcal{C}} [\hat{p}^k(\bar{\bar{N}}) - \hat{p}^\ell(\bar{\bar{N}})] = \hat{p}^k(\bar{\bar{N}}) - \hat{p}^{\ell_k}(\bar{\bar{N}}) = 2(\hat{p}^k(\bar{\bar{N}}) - p^*).$$

Therefore, equation (15) holds for all k, with $\bar{N}$ replaced by $\bar{\bar{N}}$, if and only if

$$x(N) = \frac{\delta(K+1) + 4\beta\bar{\bar{N}}}{\delta(K+1) + 4\beta N}. \tag{25}$$

Also, from (10), with $\bar{N}$ replaced by $\bar{\bar{N}}$, we have

$$\|\hat{p}^k(N) - \hat{p}^{\ell_k}(N)\|_2^2 = [x(N)]^2 \|\hat{p}^k(\bar{\bar{N}}) - \hat{p}^{\ell_k}(\bar{\bar{N}})\|_2^2,$$

which is independent of k since the axes of $\hat{\mathcal{P}}(\bar{\bar{N}})$ are equal. Therefore, profits are equalized across all such links. Clearly, $\hat{\mathcal{P}}(N)$ contracts uniformly as N increases beyond $\bar{\bar{N}}$. So the same pairs of nodes remain symmetric and maximally distant.

In all the three cases, it can be verified that $N^{k\ell}(N)$ is strictly increasing in N, and if $\bar{N} = \bar{\bar{N}} = 0$, $N^{k\ell}$ is the same for all active links. $\square$

Proof of Proposition 8.1 Using (4), the general form of equilibrium prices in the h_0-architecture, for an arbitrary arbitrageur distribution $\mathfrak{N}$, is as follows:

$$\hat{p}^0 = \sum_{k\in\mathcal{K}} \gamma^k p^k,$$

where

$$\gamma^0 := \frac{1}{1+\beta^0\sum_j \frac{\alpha^{0j}}{1+\beta^j\alpha^{0j}}}; \qquad \gamma^k := \frac{\frac{\beta^0\alpha^{0k}}{1+\beta^k\alpha^{0k}}}{1+\beta^0\sum_j \frac{\alpha^{0j}}{1+\beta^j\alpha^{0j}}}, \quad k\neq 0,$$

and

$$\hat{p}^k = (1+\beta^k\alpha^{0k})^{-1}(p^k + \beta^k\alpha^{0k}\hat{p}^0), \qquad k \neq 0. \tag{26}$$

From (26), the profit of an arbitrageur on link $0k$ is

$$\varphi^{0k} = \frac{\delta^{0k}}{2[\delta^{0k}+\beta^k N^{0k}]^2}\cdot \|p^k - \hat{p}^0\|_2^2, \qquad k\neq 0. \tag{27}$$

For the proof of this proposition as well as that of Proposition 8.2, it is useful to specialize these formulas to the case where $\beta^k = \beta$, $\delta^{0k} = \delta$, and $N^{0k} = N/K$, for all $k \neq 0$. It will turn out that, in the applications we have in mind, the specified arbitrageur distribution is in fact an equilibrium arbitrageur distribution. Anticipating this, we write state-price deflators and other variables as functions of N, in keeping with our convention that this notation indicates that these variables are associated with an equilibrium arbitrageur distribution $\mathfrak{N}(N)$. Direct substitution into the above formulas yields:

$$\gamma^0(N) = \frac{\delta K + \beta N}{\delta K + \beta N + \beta^0 NK}; \qquad \gamma^k(N) = \frac{\beta^0 N}{\delta K + \beta N + \beta^0 NK}, \quad k\neq 0,$$

so that

$$\hat{p}^0(N) = \gamma^0(N)p^0 + \sum_{k\neq 0}\gamma^k(N)p^k = \gamma^0(N)p^0 + [1-\gamma^0(N)]\left(\frac{1}{K}\sum_{k\neq 0}p^k\right). \tag{28}$$

Also, since $\beta^k = \beta$ for $k \neq 0$,

$$p^* = \sum_{k\in\mathcal{K}}\lambda^k p^k = (\beta + \beta^0 K)^{-1}\left[\beta p^0 + \beta^0\sum_{k\neq 0}p^k\right]. \tag{29}$$

Solving for $\sum_{k\neq 0} p^k$, and substituting into (28), we obtain:

$$\hat{p}^0(N) = y(N)p^0 + [1-y(N)]p^*, \tag{30}$$

where

$$y(N) = \frac{\delta K}{\delta K + \beta N + \beta^0 NK}. \tag{31}$$

From (26), we have:

$$\hat{p}^k(N) = z(N)p^k + [1 - z(N)]\hat{p}^0(N), \qquad k \neq 0, \tag{32}$$

where

$$z(N) = \frac{\delta K}{\delta K + \beta N}. \tag{33}$$

Finally, from (27),

$$\varphi^{0k} = \frac{[z(N)]^2}{2\delta}\|p^k - \hat{p}^0(N)\|_2^2, \qquad k \neq 0. \tag{34}$$

We now proceed with the proof of the proposition. Since $p^0 = p^*$, we see from (29) that $p^* = \frac{1}{K}\sum_{k\neq 0} p^k$. Therefore, p^* is the center of $\mathcal{P}$. Suppose first that p^* is the circumcenter of $\mathcal{P}$ (this covers case (i) of the proposition as well as case (ii) if the axes of $\mathcal{P}$ are equal). We conjecture that $N^{0k} = N/K$. Then equilibrium prices are characterized by (30)–(33). Since $p^0 = p^*$, we have $\hat{p}^0(N) = p^*$. This fact, together with the assumption that p^* is the circumcenter of $\mathcal{P}$, implies that profits, given by (34), are equal for all $k \neq 0$, which verifies the conjecture that $N^{0k} = N/K$. Thus we can write

$$\Phi^{h_0}(N) = \frac{[z(N)]^2}{2\delta}\|p^k - p^0\|_2^2, \qquad k \neq 0. \tag{35}$$

The convergence statements in the proposition follow directly from equations (32) and (35), using (33).

If p^* is not the circumcenter of $\mathcal{P}$, then $\mathcal{P}$ is centrally symmetric with unequal axes. As we increase N from zero, the foregoing analysis applies, if we consider only the nodes that are endpoints of axes of maximal length. As hitherto shorter axes become maximal, the analysis applies again to a larger number of nodes (though the mass of arbitrageurs is lower on the hitherto shorter axes). Each axis shrinks uniformly to p^*. □

Proof of Proposition 8.2 We conjecture that $N^{1k} = N/K$, $k \neq 1$. Then equations (30)–(34) hold with node 0 replaced by node 1, and all the β^k's equal to β, i.e.

$$\hat{p}^1(N) = y(N)p^1 + [1 - y(N)]p^*, \tag{36}$$

$$\hat{p}^k(N) = z(N)p^k + [1 - z(N)]\hat{p}^1(N), \qquad k \neq 1, \tag{37}$$

$$\varphi^{1k} = \frac{[z(N)]^2}{2\delta}\|p^k - \hat{p}^1(N)\|_2^2, \qquad k \neq 1, \tag{38}$$

where

$$y(N) = \frac{\delta K}{\delta K + \beta N(K+1)}, \tag{39}$$

and $z(N)$ is given by (33). It follows that prices converge as stated in the proposition. We only need to verify that the profit φ^{1k} is the same for all $k \neq 1$, and is strictly decreasing in N.

Since $\beta^k = \beta$ for all k, p^* is given by (9). In particular, p^* is the circumcenter of the regular simplex $\mathcal{P}$. Suppressing the dependence of $\hat{p}^1$ and y on N, we have:

$$\begin{aligned} \|p^k - \hat{p}^1\|_2^2 &= \|p^k - p^1 + (1-y)(p^1 - p^*)\|_2^2 \\ &= \|p^k - p^1\|_2^2 + (1-y)^2\|p^1 - p^*\|_2^2 + 2(1-y)\langle p^k - p^1, p^1 - p^*\rangle_2, \end{aligned} \quad (40)$$

where $\langle \cdot, \cdot \rangle_2$ is the inner product associated with the $L^2(\Pi)$ norm $\| \cdot \|_2$. For vectors u, v and w in $\mathbb{R}^S$, we have

$$\|u - w\|_2^2 = \|u - v + v - w\|_2^2 = \|u - v\|_2^2 + \|v - w\|_2^2 + 2\langle u - v, v - w\rangle_2,$$

so that

$$2\langle u - v, v - w\rangle_2 = \|u - w\|_2^2 - \|u - v\|_2^2 - \|v - w\|_2^2. \quad (41)$$

Using this to evaluate the inner product in (40), we get

$$\|p^k - \hat{p}^1\|_2^2 = (1-y)^2\|p^1 - p^*\|_2^2 + y\|p^k - p^1\|_2^2, \quad (42)$$

where we have exploited the fact that p^* is the circumcenter of $\mathcal{P}$ and hence $\|p^k - p^*\|_2$ is invariant with respect to k. The (squared) ratio of the circumradius and the edge length of $\mathcal{P}$ is given by a standard formula (see Coxeter (1963), p. 292–295):

$$\frac{\|p^k - p^*\|_2^2}{\|p^k - p^\ell\|_2^2} = \frac{K}{2(K+1)}, \qquad k, \ell \in \mathcal{K}, k \neq \ell. \quad (43)$$

In particular, $\|p^1 - p^*\|_2^2 = \frac{K}{2(K+1)}\|p^k - p^1\|_2^2$. Substituting this into (42), we get

$$\|p^k - \hat{p}^1\|_2^2 = \left[\frac{K(1-y)^2}{2(K+1)} + y\right]\|p^k - p^1\|_2^2,$$

which does not depend on k, $k \neq 1$, because the edges of a regular simplex are congruent. Therefore, from (38), φ^{1k} is the same for all $k \neq 1$, so we can write

$$\Phi^{h_1}(N) = \frac{[z(N)]^2}{2\delta}\left[\frac{K[1-y(N)]^2}{2(K+1)} + y(N)\right]\|p^k - p^1\|_2^2, \qquad k \neq 1. \quad (44)$$

It is easy to check, using (33) and (39), that $\Phi^{h_1}(N)$ is strictly decreasing in N. □

Proof of Proposition 9.1 The following proof assumes that, if $\mathcal{P}$ is centrally symmetric, its axes are equal. The unequal axes case is straightforward to deal with and we omit the details.

Consider the complete architecture. Since $p^0 = p^*$ is the center of $\hat{\mathcal{P}}(N)$ for all N, it is also internal for all N. By Proposition 6.1, no trade occurs with exchange 0 and we can simply ignore it. Proposition 7.1 applies with K instead of $K+1$ nodes. From (13), equilibrium profits for the complete architecture are

$$\begin{aligned}
\Phi^c(N) &= [x(N)]^2\,\Phi^c(0) \\
&= \frac{1}{2\delta}[x(N)]^2\|p^k - p^\ell\|_2^2, \qquad k\ell \in \mathcal{A}^* \\
&\le \frac{1}{2\delta}[x(N)]^2\left[\|p^k - p^0\|_2 + \|p^\ell - p^0\|_2\right]^2, \qquad k\ell \in \mathcal{A}^* \qquad (45)\\
&= \frac{2}{\delta}[x(N)]^2\|p^k - p^0\|_2^2, \qquad k \neq 0, \qquad (46)
\end{aligned}$$

where (45) follows from the triangle inequality ((45) holds as an equality if and only if the polytope $\mathcal{P}$ is centrally symmetric), and (46) from the centrality of p^0. Note that $\mathcal{A}^*$ does not depend on N.

For the h_0-architecture, profits $\Phi^{h_0}(N)$ are given by (35). Comparing (46) with (35), $\Phi^{h_0}(N) > \Phi^c(N)$ if and only if $z(N) > 2x(N)$. The function $z(N)$ is given by (33). The function $x(N)$ and the number of active links $|\mathcal{A}^*|$ are respectively given by (17), (20) or (25), and (16), (19) or (24), depending upon the case under consideration, with $\bar{N} = 0$ and K replaced by $K-1$ (recall that we are applying Proposition 7.1 for K, not $K+1$, nodes). Therefore, $\Phi^{h_0}(N) > \Phi^c(N)$ if and only if, for the case of the simplex:

$$N > \frac{\delta}{2\beta}K(K-1) = \frac{\delta}{\beta}|\mathcal{A}^*|;$$

for the case of the odd polygon:

$$N > \frac{\delta K}{2\beta\cos(\pi/K)} = \frac{\delta|\mathcal{A}^*|}{2\beta\cos(\pi/K)};$$

and, for the centrally symmetric case:

$$N > \frac{\delta}{2\beta}K = \frac{\delta}{\beta}|\mathcal{A}^*|.$$

Now note that the condition for the odd polygon is most stringent for $K = 3$, in which case it is simply $N > \frac{\delta}{\beta}|\mathcal{A}^*|$, i.e. the same condition as for the other polytopes. $\square$

Proof of Proposition 9.2 For the complete and h_1-architectures, node 0 is inactive by Proposition 6.1. Therefore, we can ignore node 0 and apply our results for these architectures, replacing K by $K-1$.

Profits in the h_1-architecture are given by (44). Substituting for $y(N)$ and $z(N)$ from (39) and (33) respectively, and using (43) and the fact that $p^0 = p^*$ (and also replacing K by $K-1$),

$$\Phi^{h_1}(N) = \frac{\delta(K-1)^3[2\delta^2(K-1)+\beta NK(2\delta+\beta N)]}{4[\delta(K-1)+\beta N]^2[\delta(K-1)+\beta NK]^2}\|p^k-p^1\|_2^2 \tag{47}$$

$$= \frac{\delta K(K-1)^2[2\delta^2(K-1)+\beta NK(2\delta+\beta N)]}{2[\delta(K-1)+\beta N]^2[\delta(K-1)+\beta NK]^2}\|p^k-p^0\|_2^2, \tag{48}$$

for $k \notin \{0,1\}$. Comparing (48) with (35), $\Phi^{h_0}(N) > \Phi^{h_1}(N)$ if and only if

$$K[\delta(K-1)+\beta N]^2[\delta(K-1)+\beta NK]^2 > (K-1)^2(\delta K+\beta N)^2[2\delta^2(K-1)+\beta NK(2\delta+\beta N)],$$

or, equivalently,

$$K(2K-1)Q^4+2K(K+1)(K-1)Q^3 > (K-1)^2\Big[(K-2)Q^2+2K(2K-1)Q+K(K^2-1)\Big],$$

where $Q := \beta N/\delta$. The condition $N \geq \frac{\delta}{\beta}K$ in the statement of the proposition is equivalent to $Q \geq K$. In particular, $Q > K-1$, so it suffices to show that

$$K(2K-1)(K-1)^2Q^2+2K(K+1)(K-1)^2Q^2 > (K-1)^2\Big[(K-2)Q^2+2K(2K-1)Q+K(K^2-1)\Big],$$

or

$$2(2K^2+1)Q^2-2K(2K-1)Q-K(K^2-1) > 0. \tag{49}$$

It is easy to check that this expression is increasing in Q for $Q \geq K$, and is positive for $Q = K$. Hence (49) holds.

For the complete architecture, using (13) and (17), with $K-1$ instead of K, we get

$$\Phi^c(N) = \frac{\delta}{2}\left[\frac{K-1}{\delta(K-1)+2\beta N}\right]^2\|p^k-p^1\|_2^2,$$

for $k \notin \{0,1\}$. Comparing with (47), $\Phi^{h_1}(N) > \Phi^c(N)$ if and only if

$$(K-1)[2\delta^2(K-1)+\beta NK(2\delta+\beta N)][\delta(K-1)+2\beta N]^2 > 2[\delta(K-1)+\beta N]^2[\delta(K-1)+\beta NK]^2,$$

or, equivalently

$$2KQ^3-(K+3)(K-1)^2Q-2(K-1)^3 > 0.$$

The remainder of the proof is along the same lines as in the previous case. It suffices to show that

$$2K(K-1)Q^2-(K+3)(K-1)^2Q-2(K-1)^3 > 0,$$

or

$$2KQ^2-(K+3)(K-1)Q-2(K-1)^2 > 0.$$

Now we simply check that this expression is increasing in Q for $Q \geq K$, and is positive for $Q = K$. $\square$

Proof of Proposition 9.3 We omit the details for the case in which $\mathcal{P}$ is centrally symmetric with unequal axes. The following arguments apply in all the other cases. Using (10), (32) and (33):

$$d^{k,c}(N) = [x(N)]^2 \|p^k - p^*\|_2^2, \qquad k \in \mathcal{K}, \tag{50}$$

$$d^{k,h_0}(N) = \left[\frac{\delta K}{\delta K + \beta N}\right]^2 \|p^k - p^*\|_2^2, \qquad k \in \mathcal{K}. \tag{51}$$

These equations hold trivially for $k = 0$ since $\hat{p}^0 = p^0 = p^*$ (in the complete architecture, node 0 is inactive by Proposition 6.1, while in the h_0-architecture $\hat{p}^0 = p^0$ by Proposition 8.1). It is easy to check that $x(N) < \frac{\delta K}{\delta K+\beta N}$ in all the three cases studied in Proposition 7.1 (as in the proof of Proposition 9.1, we apply Proposition 7.1 for K nodes, not $K + 1$). The result follows. $\square$

Proof of Proposition 9.4 For the complete and h_1-architectures, node 0 is inactive by Proposition 6.1. So we apply our results for these architectures, replacing K by $K - 1$.

From (50) and (17),

$$d^{k,c}(N) = \left[\frac{\delta(K-1)}{\delta(K-1) + 2\beta N}\right]^2 \|p^k - p^*\|_2^2, \qquad k \in \mathcal{K}. \tag{52}$$

From (36) and (39),

$$\hat{p}^1(N) - p^* = \frac{\delta(K-1)}{\delta(K-1) + \beta NK}(p^1 - p^*), \tag{53}$$

so that

$$d^{1,h_1}(N) = \left[\frac{\delta(K-1)}{\delta(K-1) + \beta NK}\right]^2 \|p^1 - p^*\|_2^2. \tag{54}$$

For $k \notin \{0, 1\}$, we use (37), (33) and (53) to get:

$$\begin{aligned}\hat{p}^k(N) - p^* &= \frac{\delta(K-1)}{\delta(K-1) + \beta N}(p^k - p^*) + \frac{\beta N}{\delta(K-1) + \beta N}[\hat{p}^1(N) - p^*] \\ &= \frac{\delta(K-1)}{\delta(K-1) + \beta N}\left[(p^k - p^*) + \frac{\beta N}{\delta(K-1) + \beta NK}(p^1 - p^*)\right].\end{aligned}$$

Therefore,

$$\begin{aligned}d^{k,h_1}(N) = \left[\frac{\delta(K-1)}{\delta(K-1) + \beta N}\right]^2 \Bigg[&\|p^k - p^*\|_2^2 + \left(\frac{\beta N}{\delta(K-1) + \beta NK}\right)^2 \|p^1 - p^*\|_2^2 \\ &+ \left(\frac{2\beta N}{\delta(K-1) + \beta NK}\right) \langle p^k - p^*, p^1 - p^*\rangle_2\Bigg].\end{aligned} \tag{55}$$

Using (41) and (43), and the fact that all vertices of $\mathcal{P}$ are equidistant from p^*, we see that $\langle p^k - p^*, p^* - p^1 \rangle_2 = \frac{1}{K-1} \|p^k - p^*\|_2^2$, for $k \notin \{0, 1\}$. Substituting into (55), some algebraic manipulations give us:

$$\begin{aligned} d^{k,h_1}(N) = & \frac{\delta^2(K-1)}{[\delta(K-1)+\beta N]^2[\delta(K-1)+\beta NK]^2} \\ & \cdot [\delta^2(K-1)^3 + 2\delta\beta N(K-1)(K^2-K-1) + \beta^2 N^2(K^3-K^2-K-1)] \\ & \cdot \|p^k - p^*\|_2^2 \end{aligned} \tag{56}$$

for $k \notin \{0, 1\}$.

The rankings in the proposition can now be verified from (51), (52) and (54) for exchange 1, and from (51), (52) and (56) for exchanges $k \notin \{0, 1\}$. Monotonicity with respect to N can also be deduced from these equations. □

B Polytopes

In Section 6 we defined a polytope as the convex hull of a finite number of points in Euclidean space. In that discussion we had sidestepped some formal definitions, which we briefly summarize below. The reader may consult Coxeter (1963) and Grünbaum (2003) for further details.

Let $A \subset \mathbb{R}^d$, and $\{x_i\}$ a finite number of points in A. An *affine combination* of $\{x_i\}$ is a linear combination $\sum \nu_i x_i$ in which the weights $\{\nu_i\}$ add up to one. The points $\{x_i\}$ are *affinely independent* if none of these points can be expressed as an affine combination of the other points. An *affine subspace* is a translate of a linear subspace, i.e. of the form $x + V$, where x is a point in $\mathbb{R}^d$ and V is a linear subspace of $\mathbb{R}^d$. The *affine hull* of A is the smallest affine subspace containing A; it is the set of all affine combinations of points in A. The dimension of A is the dimension of the affine hull of A, which is defined to be the dimension of the corresponding linear subspace.

Let P be a d-polytope. A *face* of P is the intersection of P with a supporting hyperplane. Each face is itself a polytope. The 0-faces are called *vertices*, the 1-faces are called *edges*, and the $(d-1)$-faces are called *facets*.[17] Thus a 1-polytope is a line segment, a 2-polytope is a polygon whose facets (which are also its edges) are segments, a 3-polytope is a three-dimensional solid, whose facets are polygons and whose edges are segments, and so forth. A *simplex* is a polytope whose vertices are affinely independent.

If the midpoints of the edges incident at a vertex v of P lie on a hyperplane, then these midpoints are the vertices of a $(d-1)$-polytope called the *vertex figure* of P at v. The notion of a *regular polytope* can be defined inductively as follows. A

[17] As mentioned in footnote 3, we employ the terms "vertex" and "edge" as is standard in the theory of polytopes. We do not use these terms in the graph-theoretic sense.

regular polygon is a polygon that is equilateral and equiangular. A regular polytope is a polytope with regular facets and vertex figures. This definition implies that the facets are in fact congruent and so are the vertex figures.

References

Agnes, P. (2000). The "End of geography" in financial services? Local embeddedness and territorialization in the interest rate swaps industry. *Economic Geography*, pages 347–366.

Allais, M. (1967). Les conditions de l'efficacité dans l'economie. Fourth International Seminar, Centro Studi e Ricerche su Problemi Economico-Sociali, Milan.

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

Berman, A. and Plemmons, R. (1979). *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York.

Border, K. (1989). *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press.

Coxeter, H. (1963). *Regular Polytopes*. Macmillan, 2nd edition.

Culbertson, J. (1957). The term-structure of interest rates. *Quarterly Journal of Economics*, 71:485–517.

Diestel, R. (2005). *Graph Theory, 3rd edition*. Springer.

Drobny, S. (2000). *Inside the House of Money: Top Hedge Fund Traders on Profiting in the Global Markets*. Wiley.

Fique, J., Page, F., and Resende, J. (2013). *Strategic Network Formation Games and Endogenous Network Dynamics*. Book in preparation.

Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks*. Princeton University Press.

Grünbaum, B. (2003). *Convex Polytopes*. Springer-Verlag, 2nd edition.

Horn, R. and Johnson, C. (1991). *Topics in Matrix Analysis*. Cambridge University Press.

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

Modigliani, F. and Sutch, R. (1966). Innovations in interest rate policy. *American Economic Review*, 56(2):178–197.

Rahi, R. and Zigrand, J.-P. (2009). Strategic financial innovation in segmented markets. *Review of Financial Studies*, 22(8):2941–2971.

Rahi, R. and Zigrand, J.-P. (2013). Financial innovation in arbitrage networks. Mimeo, London School of Economics.

Zigrand, J.-P. (2004). A general equilibrium analysis of strategic arbitrage. *Journal of Mathematical Economics*, 40(8):923–952.

Zigrand, J.-P. (2006). Endogenous market integration, manipulation and limits to arbitrage. *Journal of Mathematical Economics*, 42(3):301–314.