Economics of Networks
Networked Markets

Evan Sadler

Massachusetts Institute of Technology

Agenda

Perfect matchings

Bargaining

Competitive equilibrium in a two-sided market

Supply networks and aggregate volatility

Suggested Reading:

- EK chapters 10 and 11; Jackson chapter 10
- Manea (2011), "Bargaining in Stationary Networks"
- Acemoglu et al. (2012), "The Network Origins of Aggregate Fluctuations"

Buyer-Seller Networks

Often assume trade is unrestricted: any buyer can costlessly interact with any seller

Not true in practice:

- Product heterogeneity
- Geographic proximity
- Search costs
- Reputation

Develop theory of buyer-seller networks

- Connections to bargaining, auctions, market-clearing prices

Questions:

- Can every buyer (seller) find a seller (buyer)?
- Do market clearing prices exist?
- Is the outcome of trade efficient?

Perfect Matchings

A simple model:

- Disjoint sets of buyers and sellers B and S, $|B| = |S| = n$
- Bipartite graph G (all edges connect a buyer to a seller)
- Write $N(A)$ for set of neighbors of agents in A
- A *matching* is a subset of edges such that no two share an endpoint

Say i and j are matched if the matching contains an edge between them

A matching is *perfect* if every buyer is matched to a seller and vice versa

- Contains $\frac{n}{2}$ edges

Buyer-Seller Networks

Often assume trade is unrestricted: any buyer can costlessly interact with any seller

Not true in practice:

- Product heterogeneity
- Geographic proximity
- Search costs
- Reputation

Develop theory of buyer-seller networks

- Connections to bargaining, auctions, market-clearing prices

Perfect Matchings

A simple model:

- Disjoint sets of buyers and sellers B and S, $|B| = |S| = n$
- Bipartite graph G (all edges connect a buyer to a seller)
- Write $N(A)$ for set of neighbors of agents in A
- A *matching* is a subset of edges such that no two share an endpoint

Say i and j are matched if the matching contains an edge between them

Perfect Matchings

Theorem

The bipartite graph G has a matching of size $|S|$ if and only if for every $A \subseteq S$ we have $|N(A)| \geq |A|$

Clearly necessary (why?), sufficiency is harder

Call a set $A \subseteq S$ with $|A| > |N(A)|$ a constricted set

Elegant alternating paths algorithm to find maximum matching and constricted sets (EK, 10.6)

Rubinstein Bargaining

A seemingly unrelated problem...

Consider one buyer and one seller

- Seller has an item the buyer wants
- Seller values at 0, buyer at 1
- At time 1, seller proposes a price, buyer accepts or rejects
- If accept, game ends, realize payoffs
- If reject, proceed to time 2, buyer makes offer

Bargaining with alternating offers

Players are impatient, discount future at rate δ

The One-Shot Deviation Principle

Game has infinite time horizon, cannot use backward induction

- Payoff is a discounted infinite sum

Useful fact: one-shot deviation principle

Theorem (Blackwell, 1965)

In an infinite horizon game with bounded per-period payoffs, a strategy profile $\mathbf{s}$ is a SPE if and only if for each player i there is no profitable deviation s_i' that agrees with s_i everywhere except at a single time t.

HUGE simplification: only need to check deviations in a single period

- Proof is beyond our scope

Rubinstein Bargaining

Consider a profile of the following form:

- There is a pair of prices (p_s, p_b)
- The seller always proposes p_s and accepts any $p \geq p_b$
- The buyer always offers p_b and accepts any $p \leq p_s$

Suppose the seller proposes in the current period

Acceptance earns the buyer $1 - p_s$, rejection earns $\delta(1 - p_b)$

- Incentive compatible if $p_s \leq 1 - \delta + \delta p_b$

Similarly, when buyer proposes, acceptance is incentive compatible if $p_b \geq \delta p_s$

Rubinstein Bargaining

In equilibrium, seller proposes highest acceptable price

- $p_s = 1 - \delta + \delta p_b$

Similarly, buyer offers lowest acceptable price

- $p_b = \delta p_s$

Solving yields

$$p_s^* = \frac{1}{1+\delta}, \qquad p_b^* = \frac{\delta}{1+\delta}$$

Theorem (Rubinstein 1982)

The alternating offers bargaining game has a unique SPE with offers (p_s^*, p_b^*) that are immediately accepted.

Bargaining in a Bipartite Network

Let's extend this framework to a bipartite graph G connecting sellers S to buyers B

At time 1, sellers simultaneously announce prices

- A buyer can accept a single offer from a linked seller
- All buyers who accept offers are cleared from the market along with their sellers
- In case of ties, social planner chooses trades to maximize total number of transactions

Others proceed to time 2, when buyers make offers

- Alternating offers framework as before
- Previous model equivalent to a single buyer linked to a single seller

Example: Two Sellers, One Buyer

Suppose there are two sellers linked to a single buyer

Buyer will choose seller who offers lowest price

- If sellers offer same $p > 0$, buyer randomizes
- Profitable deviation: offer $p - \epsilon$ to ensure a sale

In unique SPE, both sellers offer $p = 0$

- Logic is reminiscient of Bertrand competition
- The "short" side of the market has all the bargaining power

Bargaining in Networks

What if there are two buyers and one seller?

- Same logic applies, sells at price 1

What if there are two buyers, each linked to same two sellers?

Work backwards, what happens if one pair trades and exits the market?

- Bargaining power is the same as in the one buyer one seller example

Bargaining in Networks

Less clear what happens in more complicated graphs

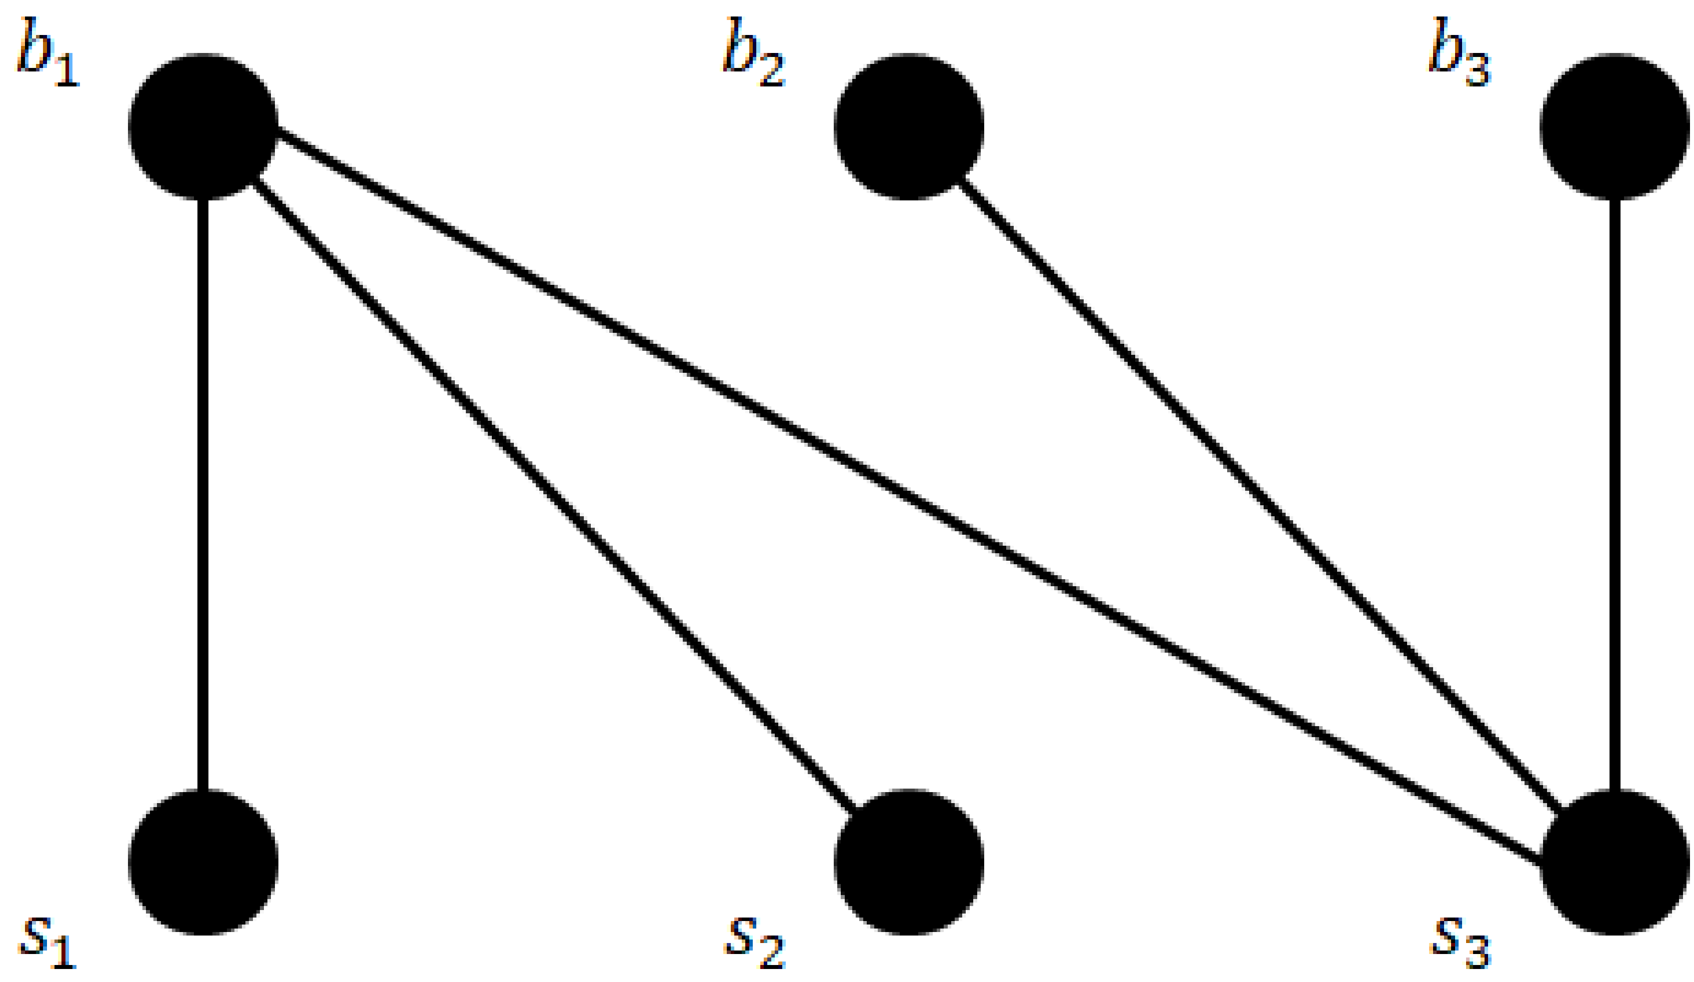

Redundant Links

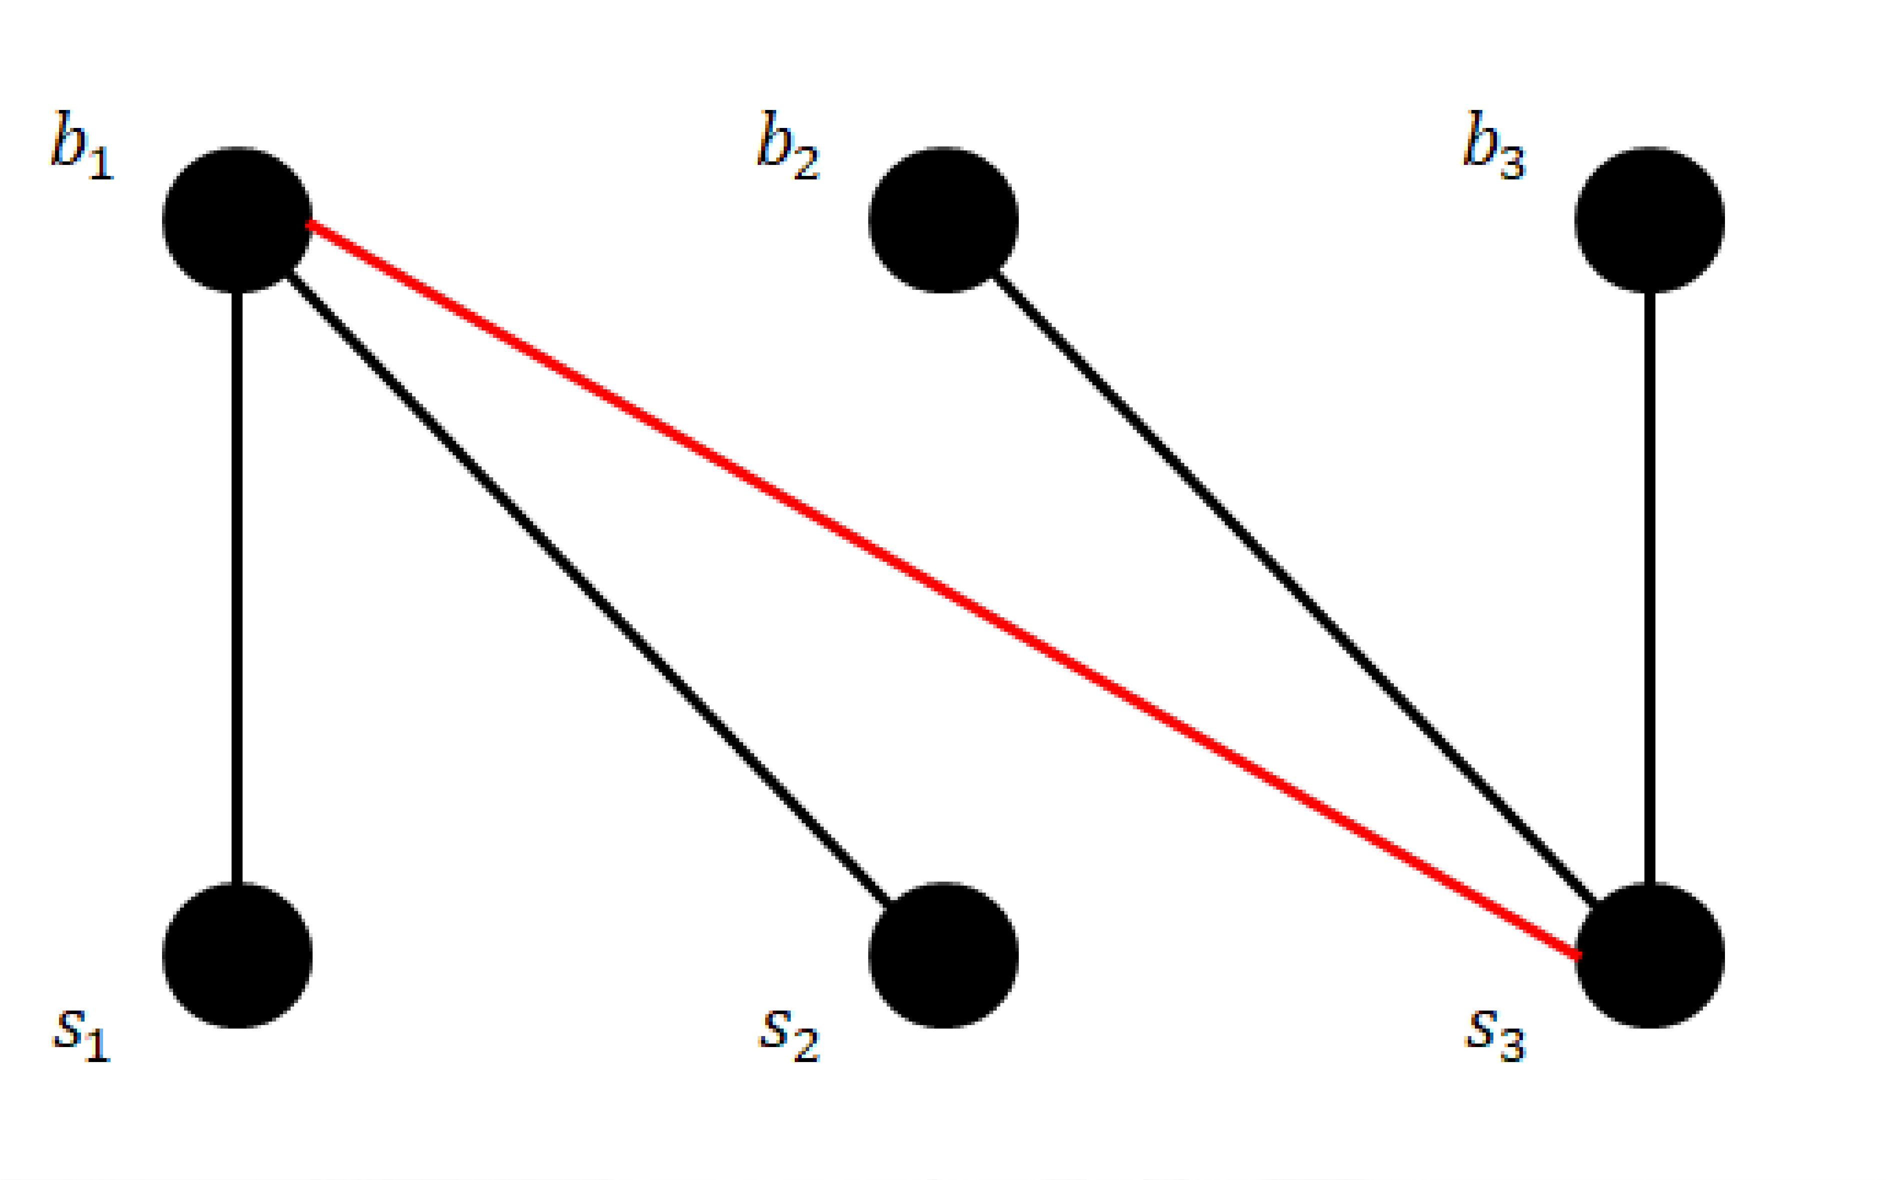

Bargaining in Networks

Existence of perfect matching ensures near-equal bargaining power

Once we eliminate redundant links, reduction to three cases

- Price 0, 1, or close to $\frac{1}{2}$

Decomposition algorithm, three sets G_S, G_B, G_E initially empty

- First, identify sets of two or more sellers linked to a single buyer, remove and add to G_S
- Next, identify remaining sets of two or more buyers linked to a single seller, remove and add to G_B
- Repeat: for each $k \geq 2$, look for sets of $k+1$ or more sellers (buyers) linked to k buyers (sellers); remove and add to the corresponding sets
- Add remaining players to G_E

Decomposition Example

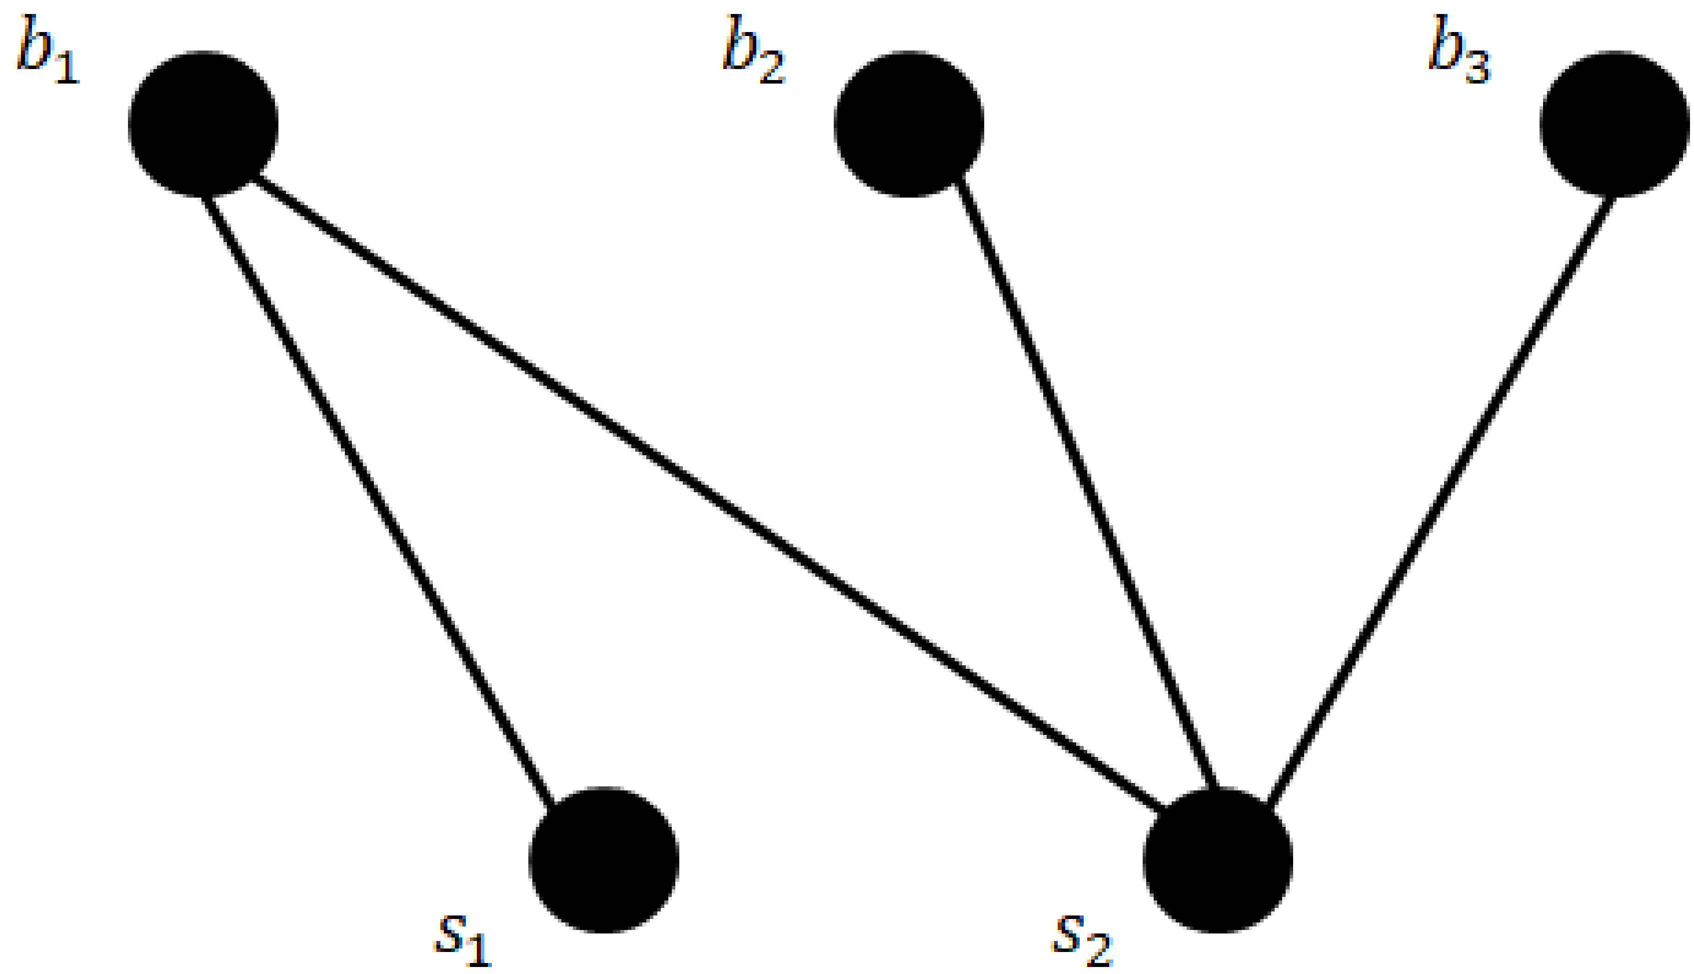

Decomposition Example

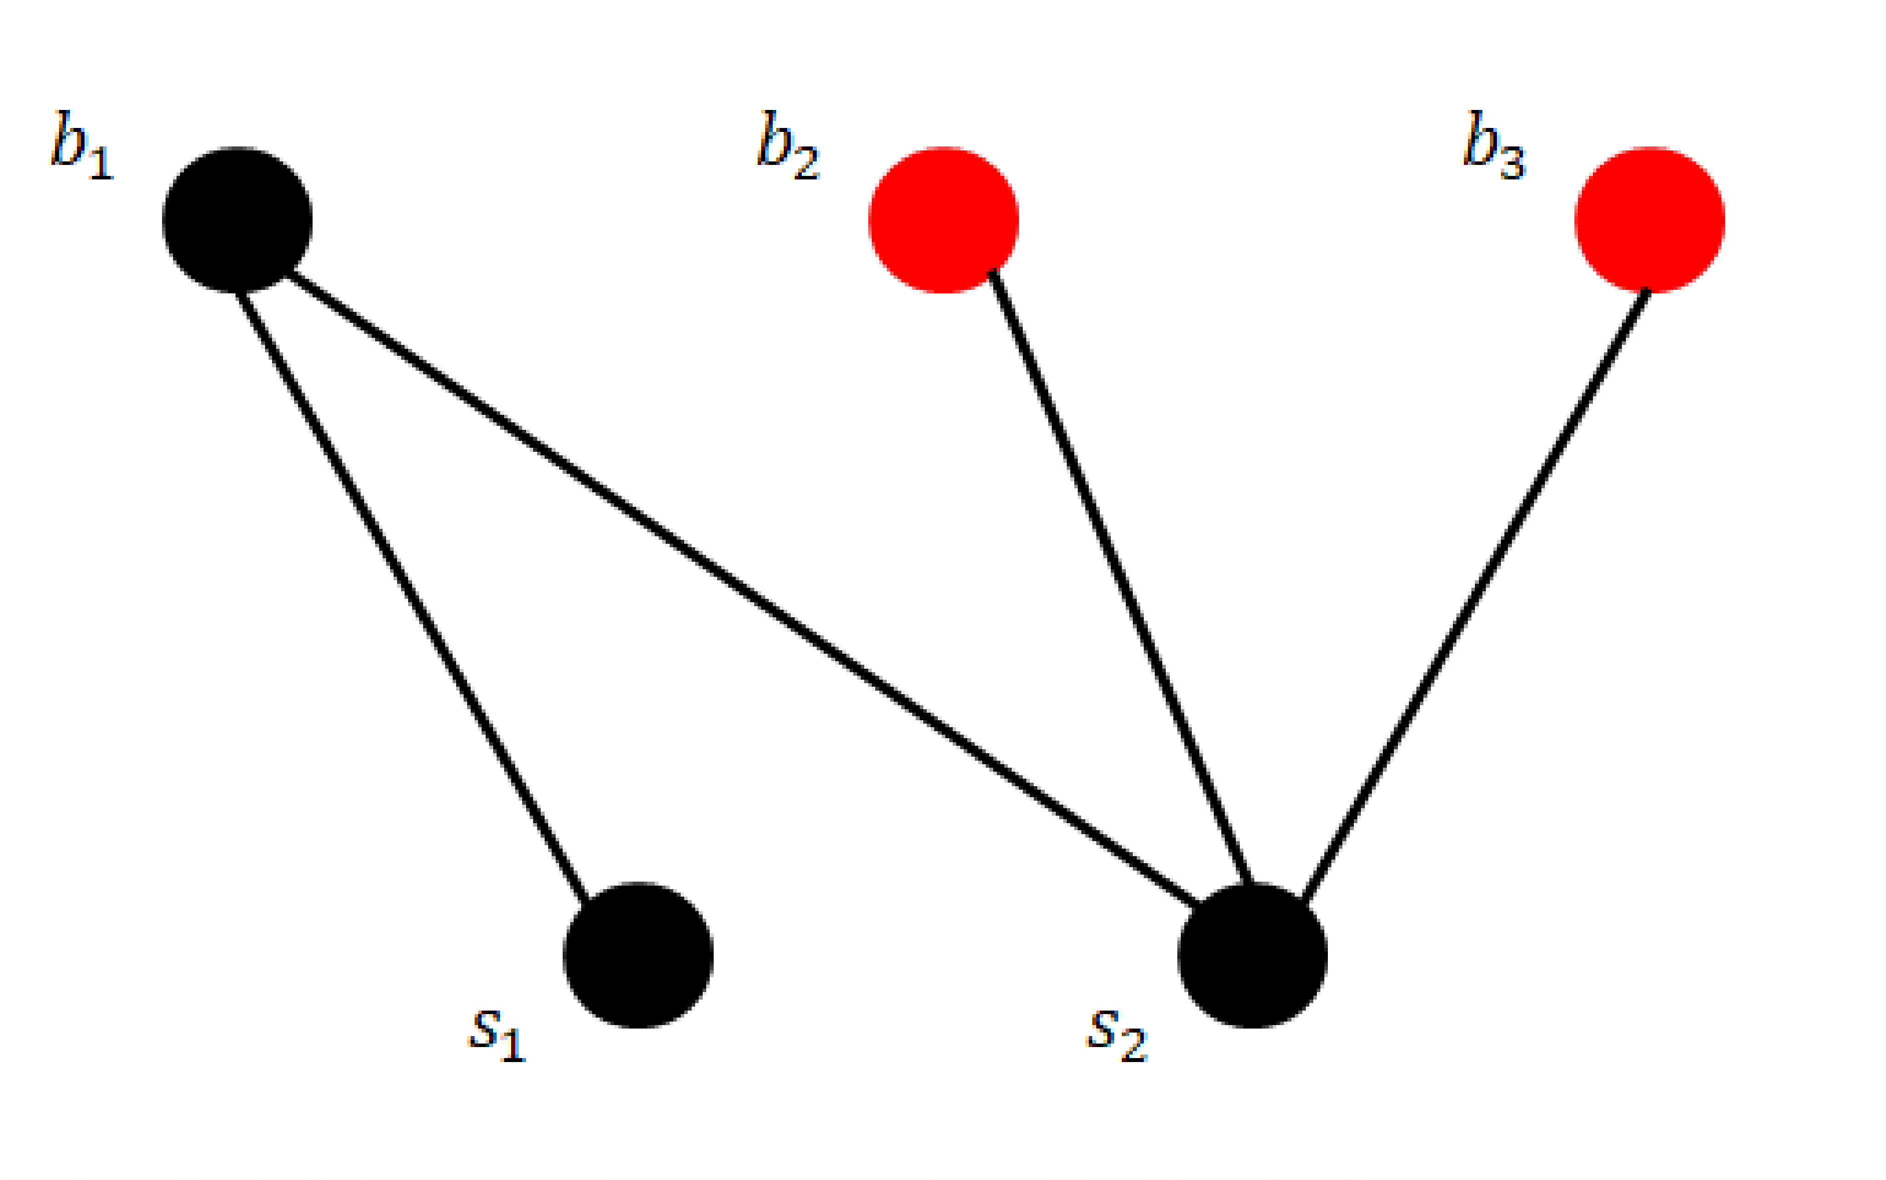

Decomposition Example

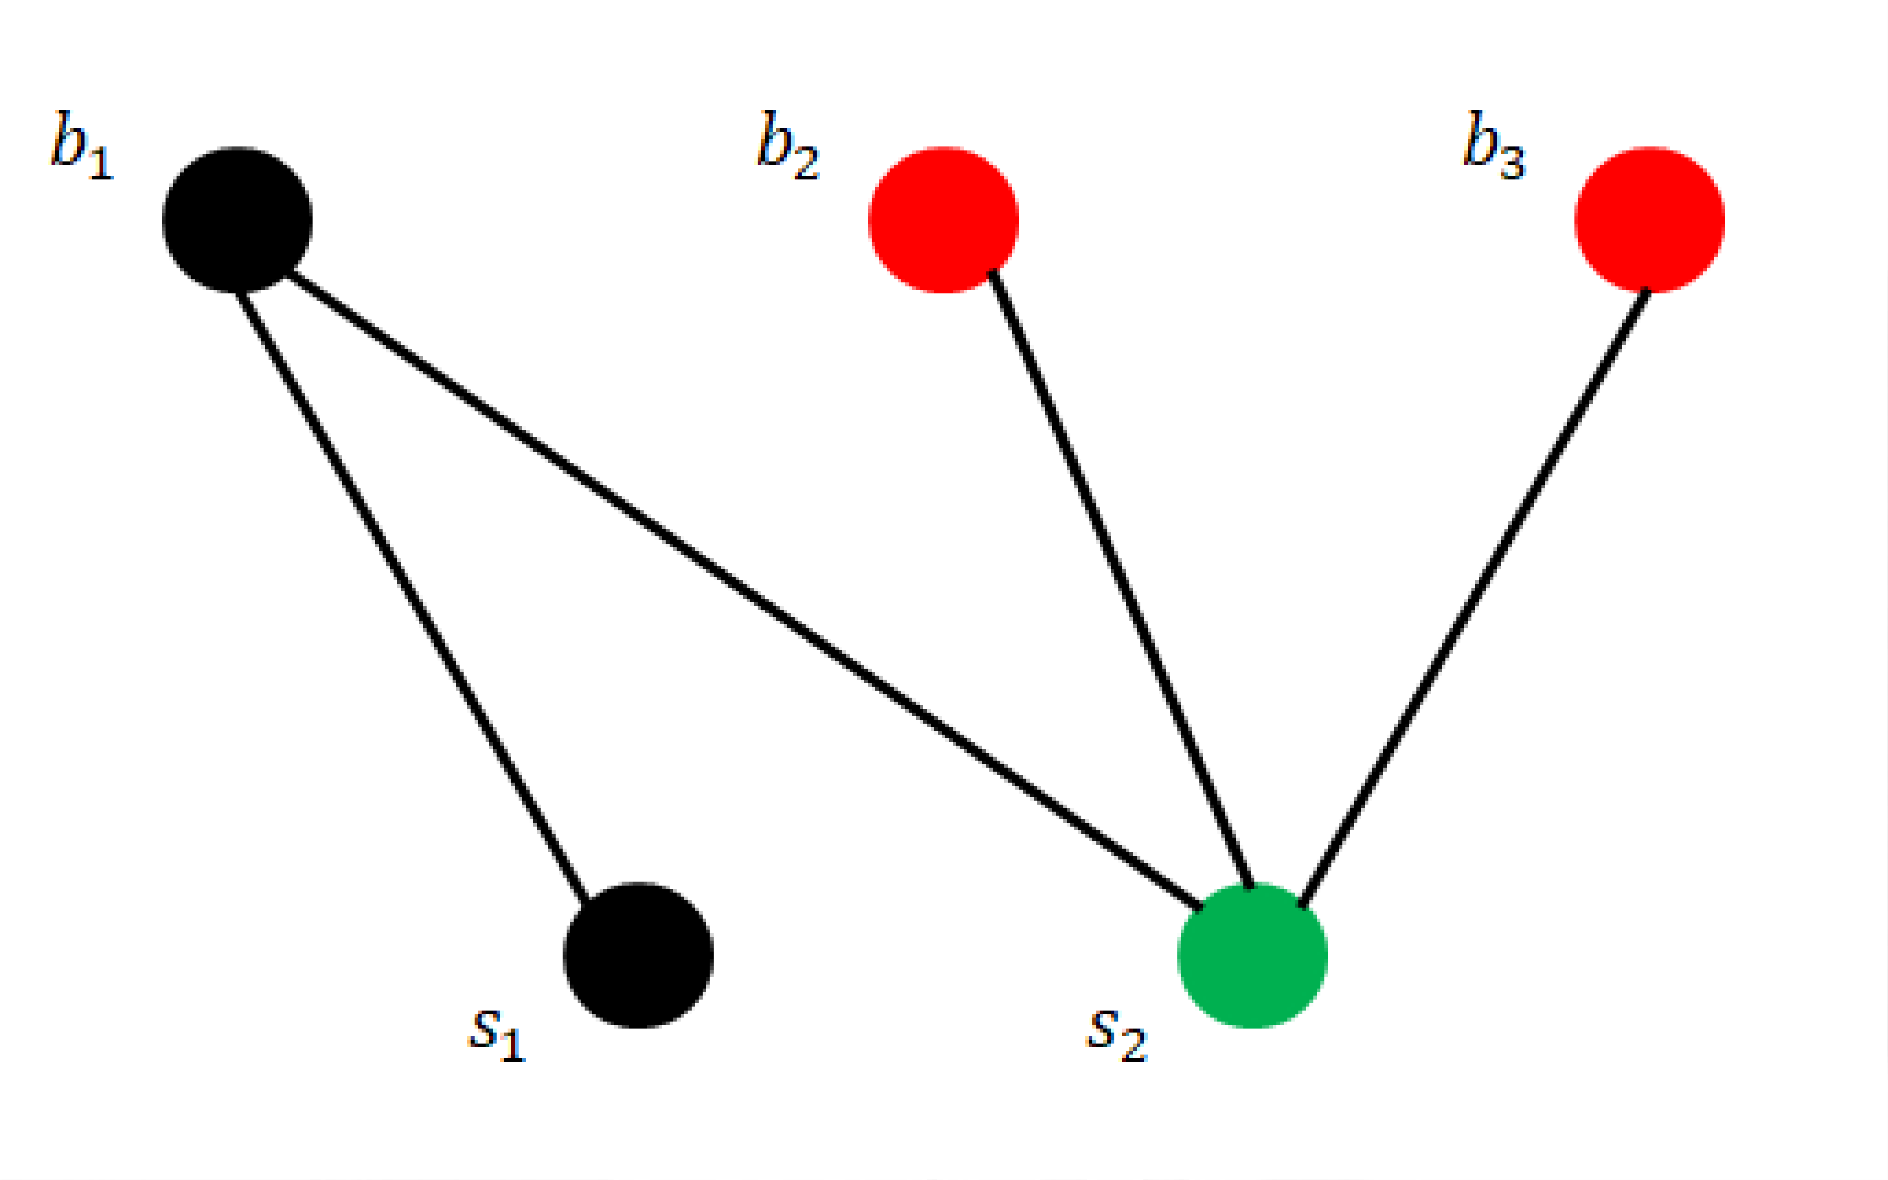

Decomposition Example

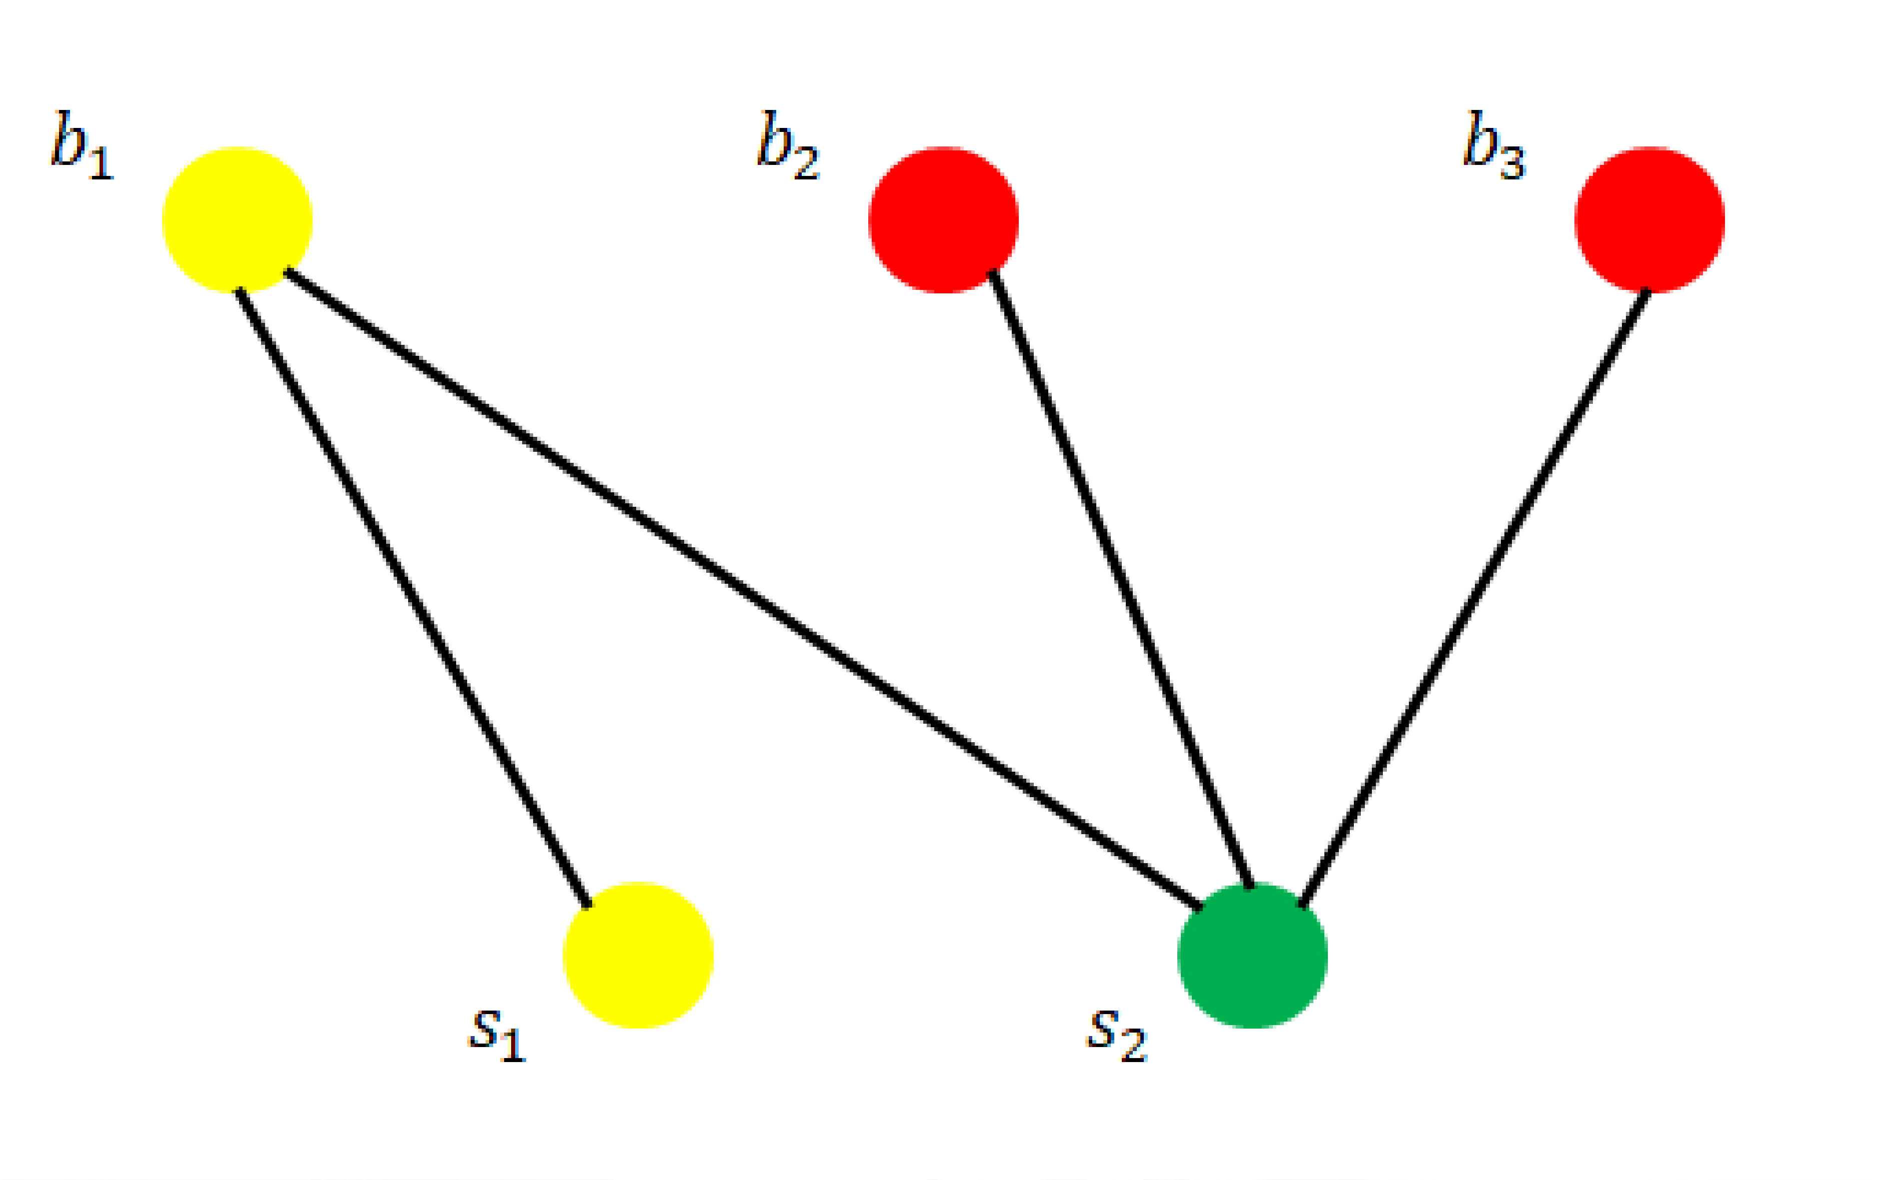

Bargaining in Networks

This simple algorithm pins down bargaining payoffs

Theorem

There exists a SPE in which:

- *Sellers in* G_S *get* 0, *buyers in* G_S *get* 1
- *Sellers in* G_B *get* 1, *sellers in* G_B *get* 0
- *Sellers in* G_E *get* $\frac{1}{1+\delta}$, *buyers in* G_E *get* $\frac{\delta}{1+\delta}$

Prediction matches well with experimental findings

Valuations and Prices

Suppose now that buyers have heterogeneous values for different sellers' products

- Each seller has an item, values it at zero, wants to maximize profits
- Posted price

Buyer i values seller j at v_{ij}, wants at most one object

- Buy from seller j, pay $p_j \geq 0$
- Buyer utility $v_{ij} - p_j$, seller utility p_j

The transaction generates surplus v_{ij}

Valuations and Prices

For a buyer i, set of *preferred sellers* given prevailing prices $\mathbf{p}$

$$D_i(\mathbf{p}) = \{j \, : \, v_{ij} - p_j = \max_k [v_{ik} - p_k]\}$$

Preferred seller graph contains edge ij if and only if $j \in D_i(\mathbf{p})$

A perfect matching in the preferred seller graph means we can match every buyer to a preferred seller, and no item is allocated to more than one buyer

- Note, whether such a matching exists will depend on the prices

Valuations and Prices

Who sells to whom?

Definition

A price vector $\mathbf{p}$ is *competitive* if there is an assignment $\mu : B \rightarrow S \cup \{\emptyset\}$ such that $\mu(i) \in D_i\mathbf{p}$, and if $\mu(i) = \mu(i')$ for some $i = i'$, then $\mu(i) = \emptyset$ (i.e. buyer i is unmatched). The pair $(\mathbf{p}, \mu)$ is a *competitive equilibrium* if $\mathbf{p}$ is competitive, and additionally if seller j is unmatched in μ, then $p_j = 0$.

Competitive equilibrium prices are market-clearing prices

- Equate supply and demand
- Corresponds to perfect matching in preferred seller graph

Existence and Efficiency

Theorem (Shapley and Shubik, 1972)

A competitive equilibrium always exists. Moreover, a competitive equilibrium maximizes the total valuation for buyers across all matchings (i.e. it maximizes total surplus).

Proof beyond our scope

More general versions of this result are known as the First Fundamental Theorem of Welfare

Bargaining in Stationary Networks

What if there are multiple opportunities to trade over time?

Simplest stationary model:

- Set of players $N = \{1, 2, ..., n\}$
- Undirected graph G
- Common discount rate δ
- No buyer-seller distinction, any pair can generate a unit surplus

In each period, a directed link ij is chosen uniformly at random

- Player i proposes a division to player j
- Player j accepts or rejects

If accept, players exit the game and are replaced by new, identical players

Bargaining in Stationary Networks

Theorem

There exists a unique payoff vector $\mathbf{v}$ such that in every subgame perfect equilibrium, the expected payoff to player i in any subgame is v_i. Whenever i is selected to make an offer to j, we have

- *If $\delta(v_i + v_j) < 1$, then i offers δv_j to j, and j accepts*
- *If $\delta(v_i + v_j) > 1$, then i makes an offer that j rejects*

Proof is beyond our scope; for generic δ, always have $\delta(v_i + v_j) \neq 1$

Intuition for strategies: $\delta(v_i + v_j)$ is the joint outside option

- Players make a deal if doing so is better than the outside option for both

Bargaining in Stationary Networks

Can place bounds on payoffs in *limit equilibria*

- As $\delta \to 1$, equilibrium payoff vectors converge to a vector $\mathbf{v}^*$

Let M denote an independent set of players (no two linked)

- Let $L(M)$ denote set of players linked to those in M

Theorem

For any independent set M, we have

$$\min_{i \in M} v_i^* \leq \frac{|L(M)|}{|M| + |L(M)|}, \qquad \max_{j \in L(M)} v_j^* \geq \frac{|M|}{|M| + |L(M)|}$$

Manea (2011) provides an algorithm to compute the payoffs

Supply Networks

During the financial crises, policy makers feared that firm failures could propagate through the economy

- The president of Ford lobbied for GM and Chrysler to be bailed out
- Feared that common suppliers would go bankrupt, disrupting Ford's operations

Such cascade effects are not a feature of standard theory

- In a perfectly competitive market with many firms, the effects of a shock to one are spread evenly across the others
- A failure has a small effect on aggregate output

Structure of supply networks can help tell us when cascade effects are possible and how severe they might be

Supply Networks: A Model

Variant of a multisector input-output model

- Representative household endowed with one unit of labor
- Household has Cobb-Douglas preferences over n goods:

$$u(c_1, c_2, \ldots, c_n) = A \prod_{i=1}^{n} (c_i)^{1/n}$$

- Each good i produced by a competitive sector, can be consumed or used as input to other sectors
- Output of sector i is

$$x_i = z_i^{\alpha} l_i^{\alpha} \prod_{j=1}^{n} x_{ij}^{(1-\alpha)w_{ij}}$$

- l_i is the labor input, x_{ij} is the amount of commodity j used to produce commodity i, w_{ij} is the input share of commodity j, z_i is a sector productivity shock (independent across sectors)

Supply Networks: A Model

Output:

$$x_i = z_i^\alpha l_i^\alpha \prod_{j=1}^n x_{ij}^{(1-\alpha)w_{ij}}$$

Assumption: $\sum_{j=1}^n w_{ij} = 1$

- Constant returns to scale

Input-output matrix W with entries w_{ij} captures inter-sector relationships

- Can think of W as a weighted network linking sectors

Define weighted out-degree $d_i = \sum_{j=1}^n w_{ji}$, and let F_i be the distribution of $\epsilon_i = \log z_i$

Economy characterized by a set of sectors N, distribution of sector shocks $\{F_i\}_{i\in N}$, network W

Equilibrium Output

Acemoglu et al (2012) show that the output in equilibrium (i.e. when the representative consumer maximizes utility and firms maximize profits) is given by

$$y \equiv \log(GDP) = \sum_{i=1}^{n} v_i \epsilon_i$$

where $\mathbf{v}$ is the *influence vector*

$$\mathbf{v} = \frac{\alpha}{n} \left[I - (1-\alpha) W' \right]^{-1} \mathbf{1}$$

Influence vector is closely related to Bonacich centrality

Shocks to more central sectors have a larger impact on aggregate output

Aggregate Volatility

Let σ_i^2 denote the variance of ϵ_i

We can compute the standard deviation of aggregate output as

$$\sqrt{var(y)} = \sqrt{\sum_{i=1}^{n} \sigma_i^2 v_i^2}$$

If we have a lower bound on sector output variances $\underline{\sigma}$, then this implies

$$\sqrt{var(y)} = \Theta(\|v\|_2)$$

Volatility scales with the Euclidean norm of the influence vector

Example

Suppose all sectors supply each other equally

- $w_{ij} = \frac{1}{n}$ for all i, j

The influence vector then has $v_i = \frac{c}{n}$ for some c and all i

Also assume $\sigma_i = \sigma$ for all i

Aggregate volatility is then

$$\sqrt{var(y)} = \sigma \sqrt{\sum_{i=1}^{n} v_i^2} = \frac{\sigma c}{\sqrt{n}}$$

Goes to zero as number of sectors becomes large

Example

Suppose we have a dominant sector 1 that is the only supplier to all others

- $w_{1j} = 1$ for all j

This implies $v_1 = c$ for some c, independent of n

This implies a lower bound on aggregate volatility

$$\sqrt{var(y)} \geq \sigma_1 c$$

Volatility does not shrink with n

Asymptotics

Can interpret economy with large n as more disaggregated

- Increased specialization
- Might expect less volatility

For economy with n sectors, define the coefficient of variation

$$CV(d^{(n)}) = \frac{STD(d^{(n)})}{\bar{d}}$$

Theorem

Consider a sequence of economies with increasing n. Aggregate volatility satisfies

$$\sqrt{var(y)} \geq c\frac{1 + CV(d^{(n)})}{\sqrt{n}}$$

for some c.