Lecture 18: Bargaining and Intermediation in Networks

Alexander Wolitzky

MIT

6.207/14.15: Networks, Spring 2022

Bargaining and Market Power in Networks

As we mentioned Lecture 1, an important social science question is how a player's position in a network conveys social or economic "power".

One can think of different kinds of "power", but one important and readily formalized one is economic "bargaining power": what share of the value created by economic transactions that one is involved in does one appropriate oneself, and what share goes to one's partners?

Bargaining and Market Power in Networks (cntd.)

To analyze this, build on matching market model from last class.

- Bipartite network of buyers I and sellers J.
- Each buyer wants to consume at most 1 good; each seller has 1 good for sale.

Last class, we showed that a competitive equilibrium always exists and is efficient (maximizes total value) in this market.

But we didn't say much about how the value is divided among the players—that is, what network positions are most "powerful."

- To isolate role of network in determing division of value, today we assume that $v_{ij} \in \{0, 1\}$ for each buyer i and seller j. One interpretation: the goods are identical but there are restrictions on who can trade with whom.
- Given "unit-demand/unit-supply," each player will get a payoff between 0 and 1.
- This payoff is a measure of the bargaining power of that player's network position.

A Simple Example

- Suppose network consists of one seller (S) connected to two buyers (B_1 and B_2).
- S has one good, which is worth 0 to her and 1 to each of the buyers.
- What outcome do we expect will result?

Market-Clearing Prices

What are market-clearing prices in this example?

- At any price $p < 1$, both buyers demand the good. Demand$>$Supply.
- At any price $p > 1$, neither buyer demands the good. Demand$<$Supply.
- Unique market-clearing price is $p = 1$.
  - Both buyers are indifferent between purchasing and not.
  - To clear market, one buyer purchases at price 1, other buyer does not purchase.
  - Both buyers get payoff 0; seller gets payoff 1.

In this simple network, the seller has "all the bargaining power" and gets the maximum possible payoff of 1; the buyers have "no bargaining power" and get the lowest possible payoff of 0.

Stable Outcomes

We can reach the same conclusion that the seller has "all the bargaining power" by forgetting about prices and instead directly looking at "stable outcomes" of exchange.

Intuitively, an outcome is "stable" if every subset of agents receives at least much value as it can create by trading on its own.

- "Unstable" outcomes cannot persist because some group of agents can split off and do better on their own.
- We'll see that this notion of stability is more basic than competitive equilibrium: it doesn't involve prices, yet in every competitive equilibrium the players' values must be stable in this sense.

Stable Outcomes (cntd.)

Formally, for each subset of agents $M \subseteq N$, let $v(M)$ be the maximum value they can create by trading on their own.

- With unit-demand/unit-supply and $v_{ij} \in \{0, 1\}$ for all i, j, this is simply the total number of trades that can be executed among buyers and sellers in M—that is, the number of links in a maximal matching in the subnetwork M (where ij are linked iff $v_{ij} = 1$).

An outcome is **stable** if the vector of final payoffs $(u_i)_{i \in N}$ are such that, for *every* subset $M \subseteq N$,

$$\sum_{i \in M} u_i \geq v(M).$$

- The set of all stable payoff vectors $(u_i)_{i \in N}$ is called the **core** of the game.

Stable Outcomes in the Example

With one seller and two buyers, the only stable payoff vector is $(u_S = 1, u_{B_1} = u_{B_2} = 0)$.

- This matches the competitive equilibrium, so in this example competitive equilibrium and the core coincide.

We prove this by establishing some general facts about stable outcomes in matching markets.

Note: This can also be proved by showing that the payoffs of any competitve equilibrium are always stable, which follows by a similar argument to the proof last class that competitive equilibria are always efficient. But the proof we'll give here is more useful for the rest of today's lecture.

Stable Outcomes in the Example (cntd.)

Fact 1: If player i is unmatched in **some** maximal matching, she gets payoff 0 in **every** stable outcome.

Proof:

- If i is unmatched in some maximal matching then $v(N) = v(N \setminus \{i\})$.
- At a stable outcome, the players in $N \setminus \{i\}$ must receive at least the value they can create on their own, which equals $v(N \setminus \{i\})$.
- Since $v(N \setminus \{i\}) = v(N)$, this leaves nothing for player i.

A player who is unmatched in some maximal matching is called **under-demanded**.

- In the example, the buyers are under-demanded, and hence each get payoff 0.

Stable Outcomes in the Example (cntd.)

Fact 2: If player j is linked to some under-demanded player i, he gets payoff 1 in every stable outcome.

Proof:

- Since i and j are linked, $v(\{i,j\}) = 1$.
- Hence, at a stable outcome, $u_i + u_j$ must be at least 1.
- Since i is under-demanded, $u_i = 0$.
- Hence, $u_j = 1$.

A player who is linked to an under-demanded player is called **over-demanded**.

- In the example, the seller is over-demanded, hence gets payoff 1.

More General Networks

With 1 seller linked to 2 buyers, every node is either under-demanded (the buyers) or over-demanded (the seller).

So, in this example, the facts that under-demanded players get 0 and over-demanded players get 1 are enough to pin down everyone's payoff.

The same logic applies to any network in which every node is either under-demanded or over-demanded.

However, in many networks some nodes are neither under-demanded nor over-demanded: e.g. 1 seller linked to 1 buyer.

We need to analyze this case further to determine the payoffs/bargaining power of nodes that are neither under-demanded nor over-demanded.

Dulmage–Mendelsohn Decomposition

Nodes that are not neither under-demanded nor over-demanded are called **perfectly matched**.

This terminology is explained by the following important graph theory: the **Dulmage–Mendelsohn decomposition**.

Theorem

Fix a bipartite network G with under-demanded, over-demanded, and perfectly matched sets U, O, and P.
In every maximal matching in G,

1. *Every node in O is matched to a node in U.*
2. *Every node in P is matched to another node in P.*

Intuition: don't "waste" over-demanded nodes by matching them to perfectly matched nodes.

- Proof is based on alternating and augmenting paths, similarly to Hall's theorem. We skip the details.

Perfectly Matched Players

What payoffs do we expect perfectly matched players to receive?

Consider 1 seller linked to 1 buyer: this situation is called **bilateral bargaining**.

Trade at any price $p \in [0, 1]$ is a stable outcome.

- Each player's payoff is at least 0, which is what she could create on her own.
- The players jointly receive payoff 1, which equals what they can create on their own.

Similarly, any price $p \in [0, 1]$ is a market-clearing price.

Thus, general theories like competitive equilibrium or stability can predict that under-demanded players get payoff 0 and over-demanded players get payoff 1, but they do not predict anything about perfectly matched players' payoffs (other than that the efficient value is created and split among them in some way).

Bilateral Bargaining: Remark

The inability to predict what price will result from bilateral bargaining is a famous hole in standard economic theory, dating back to the 19th century where Francis Ysidro Edgeworth (one of the founders of formal economics) called it the "indeterminacy of contract."

- According to standard economic theory, in bilateral bargaining the price is indeterminate.

Bilateral Bargaining

While classical economic theory cannot predict the price in bilateral bargaining, game theory models can, at least if we're willing to make assumptions about the details of the "bargaining game" that the buyer and seller play.

We consider two versions of the bargaining game:

- **Ultimatum bargaining,** where one party makes a take-it-or-leave it offer to the other.
- **Alternating-offer bargaining**, where the parties take turns making offers until they agree on a price.

We will argue that alternating-offers bargaining is usually more realistic (and also doesn't bake a huge asymmetry in who gets to make the offer), so that's the version we'll use in building our general theory of bargaining on networks.

Ultimatum Bargaining

The ultimatum bargaining game is as follows:

- First, the seller names a price $p \in [0, 1]$.
- Then, the buyer says either *Accept* or *Reject*.
- If he says *Accept*, the parties trade at price p: seller's payoff is p, buyer's payoff is $1 - p$.
- If he says *Reject*, the game ends and both parties get payoff 0.

(Of course, there's a symmetric version of the game where it's the buyer who gets to make the take-it-or-leave-it offer.)

Intuitively, what will happen in this game?

- Arguing by **backward induction** (a key tool for solving dynamic games like this one), the buyer will accept any price $p < 1$.
- Anticipating this, the seller will name a price "close" to 1.
- Seller will get payoff ≈ 1, Buyer will get payoff ≈ 0.

Aside: Subgame-Perfect Equilibrium

The main game-theoretic solution concept for dynamic games with complete information is called **subgame-perfect equilibrium (SPE)**.

It says that the strategy profile is a Nash equilibrium (mutual best responses) at the beginning of the game, and in addition remains a Nash equilibrium conditional on reaching any point in the game (or "subgame").

E.g. In ultimatum bargaining, this says that the seller must choose an optimal price given the buyer's strategy (a function from p to $\{Accept, Reject\}$), and in addition the buyer's strategy must be optimal for **any** price the seller might name (not just the one she actually names in equilibrium; this is required by subgame perfection because each price leads to a different subgame).

We skip the formal definition of SPE and just use it in an intuitive way.

Ultimatum Bargaining

Theorem

The ultimatum bargaining game has a unique SPE. In it, Seller's strategy is to offer $p = 1$, and Buyer's strategy is to accept any price.

Intuitive proof: by backward induction (which finds SPE since it finds best responses in each subgame), buyer accepts any price $p < 1$, so seller names $p \approx 1$.

Ultimatum Bargaining (cntd.)

More formal proof: (shows buyer also accepts when $p = 1$)

- By "Nash in the subgame starting with B's decision", B must accept any $p < 1$. B can accept or reject if p exactly equals 1.
- Clearly, it is a SPE for S to offer $p = 1$ and B to accept every p, including $p = 1$.
  (Neither player has a profitable deviation in any subgame.)
- This is the only SPE, because if B's strategy is to accept any $p < 1$ and reject $p = 1$ (or even reject it with positive probability), then S's optimal action would be to offer "the largest p strictly less than 1," which is impossible.

Ultimatum Bargaining: Aside

The ultimatum game is one of the most widely studied games in experimental economics, where human subjects are put together in labs to play games.

A typical result in these experiments is that splits are often close to 50:50, with proposals more extreme than 70:30 typically rejected.

There are many possible explanations for this pattern.

- Perhaps the simplest: stakes in lab experiments are often small, and vengeance can be a powerful motivator. If you're offered only \$10 out of a \$100 pot, may be worth \$10 to you to "punish" the other player for making such a low offer.

Reminder that real-world bargaining involves a richer set of considerations than those in our game-theoretic model.

- Nonetheless, simple models are a useful ingredient for analyzing more complicated situations, like bargaining in networks.

Alternating-Offer Bargaining

The assumption in ultimatum bargaining that the game ends if B rejects is usually unrealistic: in reality, wouldn't B come back with a counter-offer, or wouldn't S try again with another offer?

Also, it's very "unfair" that we only let one player make offers. To study the influence of network position on bargaining power, we should adopt a more equal bargaining procedure.

Both of these problems are addressed by **alternating-offer bargaining.**

Alternating-Offer Bargaining

- In even periods ($t = 0, 2, 4, \ldots$), S offers a price p. Then B *Accepts* or *Rejects*. If B *Accepts*, they trade at price p and the game ends. If B *Rejects*, move to the next period.
- In odd periods ($t = 1, 3, 5, \ldots$), the roles are reversed: first B offers a price p, then S *Accepts* or *Rejects*. If S *Accepts*, they trade at price p and the game ends. If S *Rejects*, move to the next period.
- If the game ends with trade at price p in period t, payoffs are $\delta_S^t p$ for S, and $\delta_B^t (1 - p)$ for B, where $\delta_S, \delta_B \in (0, 1)$ are the players' **discount factors** (i.e. how patient they are).

Note: if $\delta_S = \delta_B = 0$, we'd be back to ultimatum bargaining.

Alternating-Offer Bargaining

With ultimatum bargaining, it was intuitively fairly clear that the outcome should be $p = 1$ following by *Accept*.

With alternating-offers bargaining, it's not at all obvious what happens in SPE, or even if the SPE is unique.

A famous theorem due to Ariel Rubinstein (1982) says that there is indeed a unique SPE, and the resulting price depends on the players' discount factors and who gets to make the first offer in a natural way.

Alternating-Offer Bargaining: Theorem

Theorem

The alternating-offers bargaining game has a unique SPE: S always offers price $p_S = \frac{1-\delta_B}{1-\delta_S\delta_B}$ *and accepts all prices greater or equal to* $p_B = \delta_S p_S$*; similarly, B always offers price* p_B *and accepts all prices less than or equal to* p_S*.*

Alternating-Offer Bargaining: Remarks

Recall: $p_S = \frac{1-\delta_B}{1-\delta_S\delta_B}$, $p_B = \delta_S p_S$

- In the unique SPE, trade occurs in the first period at price p_S.
- If we fix δ_B and take $\delta_S \to 1$ (i.e. S is very patient), then $p_S, p_B \to 1$.
  If we fix δ_S and take $\delta_B \to 1$ (i.e. B is very patient), then $p_S, p_B \to 0$.
  Thus, being more patient than your opponent is an advantage in bargaining.
- If $\delta_S = \delta_B$, then $p_S = \frac{1}{1+\delta}$ and $p_B = \frac{\delta}{1+\delta}$.
  When the players are equally patient, S gets a slight advantage from making the first offer.
- If we then take $\delta \to 1$, we have $p_S, p_B \to \frac{1}{2}$.
  When the players are equally patient and both very patient, they split the gains from trade equally.
  This reflects the fact that their bargaining positions are almost symmetric in this case.

Proof Sketch: Existence

Let's first argue that the proposed strategy profile is **one** SPE.

Intuitively, when B makes an offer, he should offer the lowest price that S prefers to waiting one period and then offering p_S (since this is S's best alternative to accepting). That is,

$$p_B = \delta_S p_S.$$

Similarly, when S makes an offer, she should offers the highest price that B prefers to waiting one period and then offering p_B. That is,

$$1 - p_S = \delta_B (1 - p_B).$$

Solving this system of equations gives

$$p_S = \frac{1 - \delta_B}{1 - \delta_B \delta_S} \text{ and } p_B = \frac{\delta_S (1 - \delta_B)}{1 - \delta_B \delta_S}.$$

as claimed. So, the proposed strategy profile is an SPE.

Proof Sketch: Uniqueness

This SPE is the only one where players use "stationary" strategies that involve making the same price offer in every period and using the same threshold for accepting the opponent's offer in every period.

- By a similar logic as in the ultimatum game, each player's offer must make the opponent indifferent between accepting and rejecting.
- We just showed that p_S and p_B are the only pair of (stationary) offers that satisfy this property.

In fact, this is the unique SPE even without the restriction to stationary strategies (so the theorem is true as stated).

- That proof is slightly harder. We omit it.

Bargaining on General Networks

Finally, we can put everything together to formulate an alternating-offers bargaining model where the following theorem holds:

Theorem

In every bipartite network of buyers and sellers, for every discount factor δ, there is an efficient SPE in which under- and over-demanded players get payoffs 0 and 1 (respectively), while perfectly matched players get payoffs $\frac{1}{1+\delta}$ or $\frac{\delta}{1+\delta}$, depending on whether they are on the side that offers first or second.

- "Efficient" means maximum total value is created: a maximal matching of agents trade immediately in period 0 (i.e. as many agents trade as possible, and no value is lost to delay).

In the 5-node example, $u_1 = 1$, $u_3 = u_4 = 0$, and u_2 and u_5 equal $\frac{1}{1+\delta}$ and $\frac{\delta}{1+\delta}$ (in some order).

Alternating-Offers Bargaining in Networks

The bargaining model we consider is the following generalization of alternating-offers bargaining:

- In even periods, each S offers a price. Then each B says which of the offered prices he's willing to accept (if any).
  - Now consider the subnetwork where two nodes are linked iff they are linked in the original network (i.e. are able to trade) and the S's price is acceptable to the B.
  - Select an arbitrary maximal matching of this subnetwork. Agents in this subnetwork trade at the accepted price and exit the network.
- In odd periods, the roles are reversed (i.e. B's offer prices).
- Game continues until everyone has traded.
  - If an S trades at price p in period t, gets payoff $\delta^t p$.
  - If a B trades at price p in period t, gets payoff $\delta^t (1 - p)$.

Equilibrium

Theorem

In every bipartite network of buyers and sellers, for every discount factor δ, there is an efficient SPE in which under- and over-demanded agents get payoffs 0 and 1 (respectively), while perfectly matched players get payoffs $\frac{1}{1+\delta}$ or $\frac{\delta}{1+\delta}$, depending on whether they are on the side that offers first or second.

Behavior in the SPE:

- Over-demanded S's offer $p = 1$, only accept $p = 1$.
- Under-demanded B's offer $p = 1$, accept any price.
- Over-demanded B's offer $p = 0$, only accept $p = 0$.
- Under-demanded S's offer $p = 0$, accept any price.
- Perfectly matched S's offer $p = \frac{1}{1+\delta}$ accept any $p \geq \frac{\delta}{1+\delta}$.
- Perfectly matched B's offer $p = \frac{\delta}{1+\delta}$, accept any $p \leq \frac{1}{1+\delta}$.

Intuition for the Theorem

Consider an under-demanded buyer, $i \in U$.

This buyer is supposed to accept price 1 in period 0. Does he have a profitable deviation?

- Since $i \in U$, by definition, all his partners are over-demanded S's.
- They all offer $p = 1$. So the buyer doesn't receive a better price offer in period 0.
- Suppose the buyer rejects in period 0.
- Since $i \in U$, there exists a maximal matching that doesn't include this buyer.
- It matches each over-demanded S to an under-demanded B.
- Since in period 0 all over-demanded S's offer $p = 1$ and all (other) under-demanded B's accept any price, all over-demanded S's will trade in period 0 even if the buyer we're focusing on rejects.
- Hence, the buyer will be isolated and receive payoff 0 starting in period 1. So there is no profitable deviation.

Intuition for the Theorem (cntd.)

The argument for other under- or over-demanded agents is similar.

For a perfectly matched agent:

- By definition, she is not linked to any under-demanded agent.
- Hence, her only hope is to trade with either an over-demanded agent or another perfectly matched agent.
- Over-demanded agents are more demanding (S's require $p = 1$; B's require $p = 0$), so better to trade only with other perfectly matched agents.
- If a perfectly matched player does not trade today, all other perfectly matched agents but one will trade, so the player who doesn't trade will still be perfectly matched tomorrow.
- Hence, among perfectly matched agents, the same strategies as in alternating-offers bargaining form a SPE, by the same argument.

Experimental Evidence

This model has been tested in the lab with human subjects.

The findings conform remarkably closely to the theory (at least for the small networks that were studied).

- Usually a maximal matching of agents trade with very little delay (>90% of trades within first 2 rounds of offers).
- Over-demanded agents do well, under-demanded agents do badly, perfectly matched agents get about $\frac{1}{2}$.

One major difference from theory: payoffs for over-/under-demanded agents are not as extreme as predicted.

- Over-demand agents get more like .8 or .9, under-demanded agents more like .05 or .1.

However, given that experimental studies of ultimatum bargaining rarely show divisions more extreme than 70-30, this is already quite extreme. (Other models of bargaining on networks predict less extreme splits and thus fit the experimental data even better.)

Intermediation on Networks

A topic closely related to bargaining on networks in **intermediation on networks**: trade on networks where the seller and buyers cannot trade directly, so the good must be re-sold through intermediaries.

- Bargaining on networks: multiple buyers and sellers, a link means that two agents can directly trade a good, no possibility of re-selling.
- Intermediation on networks: (usually) one seller and one buyer, not directly linked, good must be re-sold through intermediaries.

Intermediation is important in many markets.

- Financial markets
- Agricultural markets
- Markets for illegal goods (drugs, contraband, stolen goods, etc.)

Intermediation: Modeling Framework

Directed network G, initial node $s \in N$ with only out-links, final node $b \in N$ with only in-links, all other nodes have both in-links and out-links, no cycles.

- In each period $t = 0, 1, 2 \ldots$, some node is the current **owner** of the good.
  - s is the initial owner.
- Each period, current owner chooses a downstream neighbor, sells good to her at price determined by some bargaining process.
  - Can imagine different versions of the model depending on the bargaining process.
- Once good reaches b, b consumes and receives utility 1.

Example: A Line Network

Analyzing intermediation on arbitrary networks can be complicated, but some key ideas can be seen from considering the line network:

$$s = i_0 \longrightarrow i_1 \longrightarrow i_2 \longrightarrow \dots \longrightarrow i_{n-1} \longrightarrow i_n = b$$

Let v_i be player i's equilibrium payoff in the subgame where player i is the current owner of the good.

Taking all the v_i's as given, if i and $i+1$ do not trade, they each get payoff 0; if they trade at price p, i gets p and $i+1$ gets $v_{i+1} - p$.

- Thus, bargaining between i and $i+1$ is like bargaining over a surplus of size v_{i+1}.

We assume the bargaining process is such that this surplus is split equally: $p = \frac{v_{i+1}}{2}$.

Therefore, the v_i's must satisfy the recursive equation $v_i = \frac{v_{i+1}}{2}$.

Line Network (cntd.)

Since the good is worth 1 to the buyer, we have

$$\begin{aligned} v_n &= 1, \\ v_{n-1} &= \frac{1}{2}, \ldots \\ v_i &= \frac{1}{2^{n-i}}. \end{aligned}$$

Payoff for each player i is

$$u_i = \frac{1}{2^{n-i+1}}.$$

The price of the good doubles at each step along the intermediation chain!

- Downstream intermediaries get higher payoffs than upstream intermediaries, because moving the good one step closer to the buyer is more valuable when fewer intermediaries remain to take a cut of the profits.

Empirical Evidence

A study by Olken and Barron (2009) finds empirical support for this theory in an unlikely place: trucking in Aceh, Indonesia.

In Aceh, the main trucking route (the only big road) has a number of checkpoints and weigh stations between the two biggest cities.

The toll collectors at the checkpoints are usually corrupt and accept bribes.

Olken and Barron sent research assistants to ride along with dozens of truckers to record how much was paid in bribes.

Empirical Evidence (cntd.)

Letting the truck pass is like passing on the good in the line example: the "good" (truck) gets one node closer to the node where it creates economic value.

Theory predicts that bribes will be higher at checkpoints closer to the destination city.

- This is exactly what Olken and Barron find.

But could this be explained by other differences between the checkpoints, e.g. checkpoints closer to cities may be staffed by higher-ranking soldiers?

No, because the same checkpoint gets different bribes depending on which way the truck is going!

- For example, a checkpoint close to city 1 gets small bribes from trucks leaving city 1 and gets large bribes from trucks entering city 1.
- This is strong evidence in favor of the theory.

Summary

- Bargaining theory studies how agents share the economic value created by their transactions.
- Stability says no subset of agents receives less total value than it could create on its own. This is compelling but often is not a very sharp prediction.
- Game-theoretic models of bargaining make sharper predictions but depend on the details of how bargaining is modeled.
- The Dulmage–Mendelsohn Decomposition tells us which nodes in a bargaining network are in strong, weak, or balanced positions in the network. In a simple bargaining model, this is the only network information that's relevant for determining an agent's payoff.
- In intermediated markets, downstream intermediaries are in a stronger position than upstream intermediaries.

MIT OpenCourseWare
https://ocw.mit.edu

14.15 / 6.207 Networks
Spring 2022

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.