Lecture 12: Bargaining with Incomplete Information or Reputation

Alexander Wolitzky

MIT

14.126, Spring 2024

Repeated Games vs. Bargaining

Repeated game models studied so far ask when/how players can cooperate to attain higher joint surplus than is possible with myopic play.

Reputation effects sometimes make sharper payoff predictions, e.g. about the division of the surplus.

Bargaining models focus on division of the surplus, as well as whether players' attempts to get a bigger share lead to some surplus being wasted.

Some of the same issues from repeated games and reputation arise (e.g. uniqueness vs. multiplicity, "insistent" play as source of equilibrium selection), but also some new ones.

Delay and Disagreement

A key new issue is understanding when agreements will be reached only after costly delays or conflicts, or not reached at all.

- E.g., strikes and other labor disputes, lawsuits, wars.

These Pareto-inefficient outcomes are "bargaining failures."

Why does bargaining fail?

- Cooperative game theory concepts like the core predict Pareto-efficient outcomes.
- In simple, complete-information bargaining models like Rubinstein-Stahl, bargainers reach agreement immediately, and the outcome is Pareto-efficient.
- Yet many if not most bargaining processes involve some delay.

Incomplete Information in Bargaining

A leading explanation for delay and other bargaining failures is **incomplete information**.

- If I don't know your reservation value, I might make an offer you find unacceptable.
- Even my offer is actually acceptable, you might reject it in the hope that I update and then make a more generous offer.

There is a large literature on incomplete information bargaining. Mostly theory, but also

- Empirical analysis of field data: e.g., can these models explain the duration of strikes or wars?
- Experiments: does play in the lab resemble the theory, which as we will see makes some strong predictions?

Some Empirical References

Classic field studies: Cramton-Tracy 92, 94, Kennan-Wilson 89.

- *Some findings:* In union contracts, strike incidence and duration seem to increase with uncertainty over private information (consistent with theory). More compromise settlements than war-of-attrition models predict (rarely see one side concede completely).

More recent field studies: Backus et al 20, Larsen 21,

- *Some findings:* With detailed back-and-forth offer data (e.g. eBay data), find significant delay and disagreement along lines of theory, but also more permanent disagreement or compromise agreements rather than concession. Also find more "split the difference" offers.

Some Empirical References (cntd.)

Classic lab studies: Roth-Murnighan 82, Forsythe-Kennan-Sopher 91, Seale-Daniel-Rapoport 01.

- *Some findings:* Introducing incomplete info has signficant effects on play. However, not strong evidence that play resembles sequential equilibrium.

Also much more recent work, e.g. Frechette and coauthors.

In general, bringing bargaining theory and empirics closer together is an active research area.

Who Has Private Information? Who Makes Offers?

Consider bargaining over the price of a single indivisible good.

Rubinstein: common knowledge that good is worth 0 to S, worth 1 to B, take turns making offers until agree on price.

- Result: unique SPE, immediate aggrement.

A natural generalization to incomplete info would keep the alternating-offers structure, but make both parties' valuations random and private information.

However, we would then have privately informed agents making offers, and hence a signaling game.

- These models have been studied, but there are lots of equilibria, so equilibrium refinements become critical and it's hard to get clean results.

To avoid this, much of the literature considers the case where one party (S) makes all the offers, and only the other party (B) has private information.

Two Interpretations

Bargaining interpetation: Single indivisible good worth 0 to S, worth $v \sim F$ to B, where v is B's private info. S repeatedly offers prices p.

Durable-good monopoly interpretation: S has large supply of goods, faces continuum of B's with unit demand and $v \sim F$ in the population. S repeatedly offers prices p.

Coase's conjecture: Preview

Coase 72 considered the problem of the US government selling federal land.

Since US gov't is the only seller of federal land, we might expect the monopoly outcome: in period 1, gov't sets monopoly price p_1^m, sells quantity $1 - F(p_1^m)$, and the game ends. However:

- If gov't sells $1 - F(p_1^m)$ at price p_1^m, next period tempted to set new monopoly price p_2^m $(< p_1^m)$ against truncated demand curve, and so on, cutting price each period.
- But buyers will anticipate future price cuts, so buyers with $v > p_1^m$ will **not** buy at p_1^m in period 1.
- Intuition: absent commitment power, the seller is not truly a monopoly but must compete against his own "future selves," so the outcome may be competitive rather than monopolistic.

Coase conjecture: when discounting between periods ("bargaining friction") vanishes, one-sided incomplete-info bargaining leads to the competitive outcome.

Plan

1. Linear example (Sobel-Takahashi 83)
2. General case and the Coase conjecture (Fudenberg-Levine-Tirole 85, Gul-Sonnenschein-Wilson 86)
3. Extensions
4. Reputational bargaining: combine reputation and Coase conjecture to give tractable model of bargaining with 2-sided incomplete info.

Model

- One indivisible good worth 0 to S, worth $v \sim F$ to B, v is B's private info.
- Assume $\operatorname{supp} F = [\underline{v}, \bar{v}]$ with $\underline{v} \geq 0$. (Without loss, as S never sells at price < 0.)
- Technical assumption: either $F(\underline{v}) > 0$ or F admits a strictly positive and continuous density at $\underline{v}$.
- In each period t, S names a price p_t, and B *accepts* or *rejects*.
- If type-v B accepts p_t in period t, payoffs are $\delta^{t-1} p_t$ for S and $\delta^{t-1}(v - p_t)$ for B, where $\delta \in (0, 1)$ is common discount factor.
- If B never accepts, payoffs are $(0, 0)$.

Solution concept: perfect Bayesian equilibrium (PBE).

- This is a multistage game with observed actions, so can define PBE as in FT.
- Only "belief" is S's belief about distribution of v. PBE requires that this is updated by Bayes' rule "when possible."

Linear Example (Sobel and Takahashi, 1983)

- Assume F is the uniform distribution on $[0, 1]$.
- A PBE is **linear** if there exist constants $\lambda > 1 > \gamma$ such that if S offers p_t then B accepts iff $v \geq \lambda p_t$; and if S's posterior is that $v \sim U[0, \hat{v}]$ then S offers $p_t = \gamma \hat{v}$.

Theorem (Linear Coase Conjecture)

The unique linear PBE is given by

$$\lambda = \frac{1}{\sqrt{1-\delta}} \text{ and } \gamma = \frac{\sqrt{1-\delta}\left(1-\sqrt{1-\delta}\right)}{\delta}.$$

In particular, as $\delta \to 1$, we have $\gamma \to 0$ and $\lambda \to \infty$. So, along the eqm path, $p_1 = \gamma \to 0$.

- Outcome is approximately competitive; also what would happen if B were known to have his "toughest" value $v = 0$.

Proof Sketch: Skimming Property

Key step is the **skimming property** (holds for any distribution F and any PBE): if $v < v'$ and type-v buyer accepts p_t with prob >0, then type-v' buyer accepts p_t with prob 1.

- Let V_v be continuation payoff of type-v buyer if rejects p_t. Since a type-v buyer can follow the strategy of a type-v' buyer, we have

$$V_v \geq V_{v'} - \left(v' - v\right).$$

- If type-v buyer accepts p_t with prob >0 then

$$v - p_t \geq \delta V_v.$$

- Hence,

$$\begin{aligned} v' - p_t &= v' - v + v - p_t \\ &\geq (1-\delta)\left(v' - v\right) + \delta\left(V_{v'} - V_v\right) + \delta V_v \\ &> \delta V_{v'}, \end{aligned}$$

so type-v' buyer accepts with prob 1.

Proof Sketch (cntd.)

- By skimming property, S's posterior is always a truncation, i.e. takes form $v \sim U[0, \hat{v}]$.
- Let $U(\hat{v})$ be S's continuation value when $v \sim U[0, \hat{v}]$. Then

$$U(\hat{v}) = \max_p p(\hat{v} - \lambda p) + \delta U(\lambda p).$$

- Taking FOC and using guess that $p^* = \gamma \hat{v}$ gives an equation in λ and γ.
- Since cutoff buyer type $v = \lambda p$ must be indifferent between buying at price p and waiting one period and then buying at price $\gamma \lambda p$, we have

$$\lambda p - p = \delta(\lambda p - \gamma \lambda p),$$

which is a second equation in λ and γ.
- Solving the two equations for λ and γ gives the result. (See FT Chapter 10 for details.)

Limitations of the Linear Example

- Specific prior (uniform) and equilibrium (linear).
  Even with uniform prior, there are other PBE.
- Simple proof, but hard to extract intuition from the algebra.
- With a general prior, can't hope to guess the equilibrium functional form like we did here.

Fortunately, Coase conjecture does not rely on uniform prior or linearity, and general argument will give more intuition.

General Case: Gap vs. No-Gap

There is a subtlety depending on whether $\underline{v} = 0$ or $\underline{v} > 0$.

$\underline{v} > 0$ is called the **gap case**.

- Economically, common knowledge of strict gains from trade.
- If posterior ever becomes concentrated around $\underline{v}$, S will price at $\underline{v}$ and sell for sure.
- Therefore, eqm can be found by backward induction and is unique, and all sales are made in finite time. (Fudenberg-Levine-Tirole 85)

$\underline{v} = 0$ is called the **no-gap case.**

- No common knowledge of strict gains from trade. (Can imagine that "really" have $\underline{v} < 0$, but ignore types below 0 since S never sells to them.)
- Since it's never optimal for S to price at exactly 0, sales trickle on forever. No backward induction.
- There can be multiple PBE, and the Coase conjecture holds only under a restriction to "stationary strategies." (Gul-Sonnenschein-Wilson 86)

General Case: Definitions

A sequence of equilibria (as $\delta \to 1$) **satisfies the Coase conjecture** if

1. Seller's expected payoff goes to $\underline{v}$, and
2. Probability buyer purchases in period 1 goes to 1.

An equilibrium is **stationary** if, whenever the current price p_t is lower than all past prices $p_{t'}$ $(t' < t)$, each buyer type's behavior is independent of the history of past prices.

- This is a weak form of Markov perfect. Full MPE would require that play at *every* history depends only on belief about B's type. This is too much to ask for, as sometimes S must mix with prob depending on yesterday's price to justify B's mixing yesterday.
- In the no-gap case stationarity rules out repeated game-like equilibria, where if S deviates then play switches to a different continuation equilibria. Such equilibria need not satisfy the Coase conjecture.

The Coase Conjecture

Theorem (Coase Conjecture)

1. *In the gap case: a PBE exists; for generic distributions F there is a unique PBE; and every sequence of PBE satisfies the Coase conjecture.*
2. *In the no-gap case: a stationary PBE exists; and every sequence of stationary PBE satisfies the Coase conjecture.*

We give a heuristic argument for the no-gap result, which gives the key intuition.

- See FT Ch. 10 for a more detailed sketch.
- See Gul-Sonnenschein-Wilson 86 for (complicated) full proof.

Heuristic Argument

- Re-normalize time so that the seller makes an offer every Δ units of time, and there is a constant real-time discount rate r, so $\delta = e^{-r\Delta}$. Taking $\delta \to 1$ is equivalent to taking $\Delta \to 0$.
- Fix any time ε. Since there are many periods before time ε when Δ is small, there can't be lots of sales in each of them: for any $\eta > 0$, there is a sufficiently small $\Delta > 0$ and a period $t < \varepsilon/\Delta$ (i.e., a time less than ε) such that, in equilibrium, fewer than η buyers buy at t.
- Let π^t be S's continuation payoff at t (given period t beliefs).
- Suppose S deviates in period t by "skipping" offer p^t: in period t, she offers p^{t+1} (the eqm price for period $t+1$) instead of p^t, and then offers p^{t+2} in period $t+1$, etc.
  - This is the key deviation: **cut prices faster.**
- By stationarity, all the buyer-types that previously accepted at t or $t+1$ now accept at t; all the buyer-types that previously accepted at $t+2$ now accept at $t+1$, etc.

Heuristic Argument (cntd.)

- S's "gain" from the deviation is that she reaps the profit π^t slightly faster: this is worth approximately $\pi^t r\Delta$.
  (Sell to **everyone** a little faster. First-order in Δ if π^t is bounded away from 0.)
- S's "loss" from the deviation is that the buyer-types who previously bought at price p^t at t now buy at price p^{t+1} at t. However, $p^t - p^{t+1} \leq \bar{v} r\Delta$ (otherwise, no one would buy at t), and by hypothesis at most η buyers buy at t, so total loss is at most $\eta\bar{v} r\Delta$. (Sell to a **few** buyers at a **slightly** lower price. "Second-order" in Δ.)
- Gain must be smaller than loss, so $\pi^t \leq \eta\bar{v}$.
- η was chosen arbitrarily. Taking $\eta \to 0$ implies that $\pi^t \to 0$, and hence $\pi^{\varepsilon/\Delta}$ (profit from time ε on) $\to 0$.
- $\pi^{\varepsilon/\Delta} \to 0$ implies that $p^{\varepsilon/\Delta} \to 0$: if instead $\pi^{\varepsilon/\Delta} \to 0$ due to long delay, B would get payoff close to 0, so S can get positive profit by cutting prices.
- Finally, $p^{\varepsilon/\Delta} \to 0$ implies $p^0 \to 0$, otherwise no one buys.

Remark 1: The Right to Remain Silent

In Coasian eqm, B gets best possible payoff

If B also gets to make offers (e.g., with alternating offers as in Rubinstein), need refinements to address signaling.

However, since B gets best possible payoff when he doesn't get to make offers, intuitively the same should be true when he can make offers, as long as he has the "right to remain silent" (i.e., make only unacceptable offers, without this triggering adverse belief updating by S).

Ausubel and Deneckere (1992) formalize this.

Note: The feature that it doesn't matter who makes the offers is a strength of the model, since in real-world negotiations we usually don't know the exact extensive-form.

Remark 2: Analogy with Reputation Models

The Coase conjecture is vaguely related to the Stackelberg payoff theorem of Fudenberg-Levine 89.

- Stackelberg payoff theorem: as $\delta \to 1$, a privately informed long-run player facing a series of short-run opponents does almost as well as she would if she were known to be her most favorable (Stackelberg) commitment type.
- Coase conjecture: as $\delta \to 1$, a privately informed bargainer facing an uninformed long-run opponent does almost as well as she would if she were known to be committed to the strategy of her most favorable ($v = \underline{v}$) payoff type.

There are also important differences. FL applies for any NE but doesn't work with a LR opponent. Coase conjecture requires stationary PBE.

The connection will be even stronger when we turn to reputational bargaining, where uncertainty is about whether a bargainer is committed to a certain strategy, rather than about her value.

Remark 3: Non-Stationary, Non-Coasian Equilibria

If non-stationary strategies are allowed, the Coasian eqm can be used as a threat point to restore commitment power for S.

- There are non-stationary equilibria of the form: "S is supposed to cut prices only very slowly, and if S tries to cut prices faster (e.g., by skipping a price), eqm reverts to Coase."

Ausubel-Deneckere 89 use this idea to prove a "folk theorem" for non-stationary eqm.

Theorem ("Folk Theorem")

Consider the no-gap case with a regularity condition on F. Let $V^m = \max_p p(1 - F(p))$. For any $\varepsilon > 0$, there exists $\delta(\varepsilon) < 1$ such that, for every $\delta > \delta(\varepsilon)$ and any $V \in [\varepsilon, V^m - \varepsilon]$, there is a PBE in the game with discount factor δ where S's payoff equals V.

- In the gap case, all equilibria are stationary for generic priors, as game solvable by backward induction.

Extensions of the Coase Conjecture

There are many, including:

- Ausubel and Deneckere (1987), Gul (1987): with two **competing** sellers, get folk theorem even in gap case, supported by threat of opponent cutting price.
- Conlisk, Gerstner, and Sobel (1984), Sobel (1991), Board (2008): with **inflow** of consumers, prices cycle with occasional "sales" to accumulated low-value buyers. Similar models are used in IO to model price dynamics.
- Admati and Perry (1987): with **strategically timed offers** where low-value buyer can commit to delay responding to offers, get separating eqm where delay screens buyer types, rather than pooling Coasean eqm.
- Hart (1989), Hörner and Samuelson (2011), Chen (2012), Fuchs and Skrzypacz (2013): with (different kinds of) **deadlines**, the probability of agreement can spike at the deadline. (Chen (2012) is a MIT JMP motivated by understanding the bargaining structure on Priceline.com.)

Extensions of the Coase Conjecture (cntd.)

Wang (1989), Board and Pycia (2014), Nava and Schiraldi (2019): failures or generalizations of Coase conjecture where buyers can accept **different contracts**.

E.g., if buyers can **exit** and get outside option $\bar{u} > 0$, then unique equilibrium is $p_1 = p^m$, buyer-types $v > p^m$ accept, buyer-types $v < p^m$ exit.

(There can't be an eqm where any buyer-types stick around, as if v^* is the lowest such type then S will never price below v^*, so type v^* should have exited.)

Wang, Board-Pycia present this as a failure of the Coase conjecture. Nava-Schiraldi present it as an instance of a generalized Coase conjecture where there are two goods and consumers have unit demand, interpreting the outside option as an inferior good.

Interdependent Values (Deneckere and Liang, 2006)

An especially rich extension is **interdependent values**: informed party's private info is payoff-relevant for both parties.

As in Akerlof's market for lemons, with interdependent values usually makes sense to call the informed party who receives offers the seller, call the uninformed party who makes offers the buyer.

- Seller has a good of *quality* $q \in [0, 1]$ (observed by S).
- A quality-q good is worth $c(q)$ to S and $v(q)$ to B.
- Assume $c(q) + \varepsilon < v(q)$ for all q (CK of strict gains from trade; "gap case").
- Timing: B makes a price offer p_t each period, S *accepts* or *rejects*, payoffs if agree at time t are $\delta^t (p_t - c(q))$ for S and $\delta^t (v(q) - p_t)$ for B.

Interdependent Values (cntd.)

Recall: in Akerlof, there is an efficient competitive equilibrium iff $E[v(q)] \geq c(1)$.

- [Uninformed B's willingness-to-pay] $>$ [Highest S cost]
- DL06 show that in this case the Coase conjecture holds, as in the private values case: $p_1 \longrightarrow c(1)$ as $\delta \longrightarrow 1$.
- However, this cannot happen if $E[v(q)] < c(1)$, because B would lose money if offered $c(1)$.
- DL show that in this case trade follows a complicated pattern with substantial delay.

More Papers on Interdependent Values

- Fuchs-Skrzypacz 10: Bilateral bargaining with **stochastic arrival** of a second buyer, at which point there's an auction between the two buyers. This makes the model one with interdependent values, because S's payoff in case an auction happens depends on B's value v. Delay occurs in eqm even as $\delta \to 1$, as in DL06.
- Hörner-Vieille 09: One seller, sequence of short-run buyers. Paper studies consequences of buyer offers being **public vs. private**. With public offers, rejecting offers is an opportunity to signal high quality, and no-trade can result. Efficiency is higher with private offers.
- Daley-Green 12,20: Exogenous, dynamic **arrival of information** in market for lemons. Eqm involves periods of delay while players "wait for news," as well as endogenous info at histories where low-quality S's sell with positive probability (so no-sale is good news).
  - Such models used to understand learning and signaling in markets with adverse selection, like labor and financial markets.

Reputational Bargaining

The main line of the literature on incomplete info bargaining considers incomplete info about valuations for or quality of a good.

A parallel line of literature (with roots in Nash 53; Schelling 56) models bargaining as a "struggle to establish commitments to favorable bargaining positions" (Schelling).

Two ways of modeling this:

1. Explicitly model "committing actions" in a complete info bargaining game.

- Crawford 82: 2-stage game where first each player demands a share of the pie, then each player privately learns her cost of backing down and decides whether to back down, and no-deal results if demands are incompatible and neither backs down.
- Related models: Muthoo 96, Ellingsen and Miettinen 08, 14, Dutta 12, 23.

Reputational Bargaining: Motivation (cntd.)

2. Stick with “standard” extensive form with only offers and accept/reject decisions, but add incomplete info about whether each player is a type who is committed to a certain demand. This is “reputational bargaining.”

- Key force: even a small prob of commitment types can make a big difference, because uncommitted players can “pretend” to be committed. (As in reputation effects in repeated games.)
- Early analysis in Kreps 90 and Myerson 91; seminal paper by Abreu and Gul 00.

Reputational Bargaining: Motivation (cntd.)

The popularity of reputational bargaining models comes from several nice features:

- Handle 2-sided incomplete info without signaling refinements: commitment-type behavior is pre-specified, so off-path offers must be attributed to the single rational type.
- As in the Coase conjecture, (mostly) doesn't matter which party makes offers.
- Equilibrium is typically unique and tractable, similar to incomplete-info war-of-attrition (as in Gang of Four, especially Kreps-Wilson 82).

The War of Attrition

Reputational bargaining is related to war of attrition. Start by covering WOA, which also has many other applications.

- 2 players are fighting over a prize.
- Each period, they simultaneously decide whether to *stop* or *continue*.
- When one player stops, the opponent gets the prize. If both stop, suppose no one gets it (not important).
- In the original version of the WOA (Maynard Smith 74), prize is worth v, fighting costs 1 per period.
- Unique symmetric eqm: stop each period with prob $p = 1/\left(1+v\right)$.
  - This stopping prob makes the opponent indifferent, as $pv - (1-p) = 0$.
- There are also asymmetric eqm, like "1 always stops, 2 always continues."

War of Attrition in Continuous Time

The WOA is convenient to analyze in continuous time.

- Each player chooses a cdf $\hat{F}_i$, where $\hat{F}_i(t)$=prob stops at or before t if opponent doesn't.
- Prize is worth v, fighting costs 1 per unit of time.
- Symmetric eqm: stop at constant hazard rate $\lambda = 1/v$.
  - This stopping rate makes the opponent indifferent, as $\lambda v - 1 = 0$.
- Can show that this eqm is the limit of the symmetric discrete-time eqm as $\Delta \to 0$.

Asymmetric Prizes

If winning prize is worth $v_i \neq v_j$ to player i, in the unique mixed eqm player i's stopping rate is $\lambda_i = 1/v_j$.

- Player with higher value stops faster, to keep the other player indifferent.
- Player who values the prize less wins it more often!

This is a paradoxical prediction. Somewhat like the paradoxes of backward induction we discussed in Lecture 5 (e.g. chain store paradox), it goes away if we add commitment types.

WOA with Commitment Types

Consider continuous time WOA with asymmetric prizes v_1, v_2, where each player i has prior prob z_i of being a commitment type that always continues. (Kreps-Wilson 82)

- Each (rational) player chooses a stopping time cdf $\hat{F}_i$.
- Player j's incentives are determined by player i's stopping time *unconditional on her type*, which is given by $F_i(t) = (1 - z_i)\hat{F}_i(t)$. (Since commitment type never stops.)

As in the case without commitment types, j is indifferent between stopping and continuing for another instant iff i's unconditional stopping rate $F_i'(t) / (1 - F_i(t))$ equals $1/v_j$.

However, if each player i concedes at unconditional rate $1/v_j$ and there are no discrete concessions, then generically one player's reputation hits 1 (i.e. $F_i(t)$ hits $1 - z_i$) strictly before the other's.

This is impossible, because if i learns that j is rational for sure at time T_j, i should immediately concede. Need discrete concession.

What Happens in Equilibrium?

Neccesary conditions:

"No gaps": There cannot be $t' > t$ s.t. F_i is constant on $[t, t']$ and but strictly increasing above t'. Otherwise, j would strictly prefer to concede at $t + \varepsilon$ rather than $t' + \varepsilon$. But then i would strictly prefer to concede at t' rather than $t' + \varepsilon$, a contradiction. So each F_i is strictly increasing up to $1 - z_i$.

"No jumps after $t = 0$": If F_i jumps up at $t > 0$, j would strictly prefer to concede at $t + \varepsilon$ rather than $t - \varepsilon$. This would contradict "no gaps."

"No mutual initial concession": If i concedes with prob >0 at $t = 0$, j strictly prefers to concede at $t = \varepsilon$ rather than $t = 0$.

Therefore, at most one player concedes with prob >0 at $t = 0$, and subsequently each player i concedes at rate $\lambda_i = 1/v_j$.

What Happens in Equilibrium? (cntd.)

"Simultaneous Hitting": $T_1 = T_2 = T^*$, where T_i is the time when $F_i(t)$ hits $1 - z_i$. Otherwise, opponent would concede for sure when first player hits.

Constant concession rates $\lambda_i = 1/v_j$ plus simultaneous hitting lets us determine who makes the initial concession, and with what probability.

Absent initial concession, i's reputation hits 1 at time T_i s.t.

$$z_i e^{\lambda_i T_i} = 1, \text{ or } T_i = -\frac{\ln z_i}{\lambda_i}.$$

If absent initial concession $T_i < T_j$, then j must make the initial concession, and must concede with unconditional prob $F_j(0)$ s.t.

$$\frac{z_j}{1 - F_j(0)} e^{\lambda_j T_i} = 1, \text{ where } T_i = -\frac{\ln z_i}{\lambda_i}.$$

This completes the characterization of equilibrium.

Small Commitment Probabilities

We've seen that player i immediately concedes with prob >0 iff

$$\frac{\ln z_i}{\ln z_j}\frac{\lambda_j}{\lambda_i} > 1.$$

The initial commitment probs enter with logs (unlike the v's), because reputation grows exponentially so the growth rates swamp the initial values.

If we fix the v's and take $z_i, z_j \to 0$ at the same rate (i.e., $z_i/z_j \in [1/K, K]$ for fixed $K \geq 1$), then if $\lambda_i < \lambda_j$, i immediately concedes with prob $\to 1$.

The player with the lower λ_i (i.e., lower v_i) immediately concedes with prob 1 in the limit, because a small initial commitment prob must be multiplied by a large factor to compensate for the fact that the other player's reputation will grow at a faster rate.

So, this somewhat complicated analysis gives the natural conclusion that the player who values the prize more wins it.

Extensions and Applications of WOA

Bliss-Nalebuff 84: N players waiting for 1 of them to volunteer for an unpleasant task. How does N affect delay?

Fudenberg-Tirole 86: WOA in a "declining industry," where eventually even a monopolist won't make a profit. Backward induction from the monopoly exit time pins down a unique solution.

Bulow-Klemperer 99: $N + K$ firms competing for N prizes.

Krishna-Morgan 97: WOA interpreted as an all-pay auction, with affiliated values.

...Now turn to reputational bargaining, starting with seminal paper of Abreu-Gul 00.

Abreu and Gul (2000)

AG study two possibly-committed players bargaining over a pie of size 1 in continuous time, with discount rate r.

Convenient to introduce the model in three steps:

1. Continuous-time "concession game," where each player is either rational or is committed to a particular, commonly known demand, and at each point in time each player either sticks with their demand or concedes (accepts opponent's demand). This is essentially the WOA again.
2. Concession game with multiple commitment types on each side and a prior stage where each player chooses which type to "mimic."
3. Discrete-time bargaining game (with offers and accept/reject decisions, not just "concession") with multiple commitment types. This is the ultimate game of interest, but AG show all equilibria converge to the equilibrium of the continuous-time concession game as $\Delta \to 0$, somewhat like Coase conjecture.

Concession Game with One Commitment Type on Each Side

- Two players must divide a dollar at some point in continuous time.
- Each player $i \in \{1, 2\}$ is a **commitment type** with indep prob z_i.
- Commitment type of player i always demands some fixed share of the pie $\alpha_i \in (0, 1)$, with $\alpha_1 + \alpha_2 > 1$.
- At each point in time, each player decides whether to stick with her commitment demand α_i or **concede** (accept opponent's demand).
- If the first concession is made by player i at time t, payoffs are $e^{-rt}(1 - \alpha_j)$ for i, $e^{-rt}\alpha_j$ for j.
- (If concede simultaneously, flip a coin between α_1 and α_2.)

Concession Rates

This is a WOA with commitment types, with the difference that instead of a prize v_i and a flow cost of fighting 1, we have a prize $\alpha_i + \alpha_j - 1$ and a flow cost of fighting $r_i (1 - \alpha_j)$.

- Prize is i's "capital gain" when j concedes, $\alpha_i - (1 - \alpha_j)$.
- Fighting cost is the lost interest on j's offer of $1 - \alpha_j$.

So, j's concession rate that makes i indifferent is given by

$$\begin{aligned} \lambda_j (\alpha_i + \alpha_j - 1) &= r_i (1 - \alpha_j), \text{ or} \\ \lambda_j &= \frac{r_i (1 - \alpha_j)}{\alpha_i + \alpha_j - 1}. \end{aligned}$$

Who Initially Concedes?

As in the WOA, absent initial concession, i's reputation hits 1 at time T_i such that

$$z_i e^{\lambda_i T_i} = 1, \text{ or } T_i = -\frac{\ln z_i}{\lambda_i}.$$

Therefore, i initially concedes with prob >0 iff $-\ln z_i / \lambda_i > -\ln z_j / \lambda_j$, or

$$r_i \left(1 - \alpha_j\right) \left(-\ln z_i\right) > r_j \left(1 - \alpha_i\right) \left(-\ln z_j\right).$$

Comparative Statics

What features make i more likely to "win the race," i.e. to satisfy

$$r_i\left(1-\alpha_j\right)\left(-\ln z_i\right) < r_j\left(1-\alpha_i\right)\left(-\ln z_j\right)?$$

1. i is more likely to win when she is more patient (small r_i) and her opponent is less patient (large r_j).
2. i is more likely to win when her prior commitment prob is higher (high z_i) and her opponent's is lower (low z_j).
3. i is more likely to win when her demand is less aggressive (low α_i) and her opponent's is more aggressive (high α_j).

Effect 3 will be key once we endogenize demands: reputation-formation pushes players to make modest demands, because if i makes an aggressive demand it reduces the concession rate λ_i that makes j indifferent, which disadvantages i in the reputational race.

Comparative Statics (cntd.)

Again, it's key that the prior commitment probs enter with logs.

- If we consider the complete-info limit where $z_i, z_j \longrightarrow 0$ at the same rate, any constant-factor difference between z_i and z_j is irrelevant for the outcome.
- If we fix the discount rates and demands and take $z_i, z_j \longrightarrow 0$ at the same rate, then whichever player loses the reputational race must concede with probability 1 in the limit.

Concession Game with Multiple Commitment Types

- Now suppose for each player i there's a finite set of commitment types $C_i \subset (0, 1)$ (with each type identified with its demand), with an independent prior over $C_1 \times C_2$.
- At the beginning of the game, first player 1 publicly announces a demand $\alpha_1 \in C_1$, then player 2 announces a demand $\alpha_2 \in C_2$.
- Then play continues into the same concession game as before.

Concession Game with Multiple Commitment Types

AG show there is still a unique Nash equilibrium outcome distribution.

- If i announces demand α_i, her initial reputation at the beginning of the concession game depends on the prior prob of commitment type α_i and her eqm announcement probs.
- Her initial reputation is decreasing in the eqm prob that she announces α_i, as this makes it less likely that she is truly committed conditional on announcing α_i.
- Since i benefits from having a high initial reputation, there is a kind of "substitutability": the more she's expected to announce α_i, the less she wants to announce α_i.
- This substitutability leads to a unique eqm distribution of announcements, and hence a unique eqm in the entire game.

The Complete-Info Limit

We typically view commitment probs as being small (or even just a perturbation/equilibrium selection device), so natural to consider the limit where $z_1, z_2 \to 0$.

Since z_1 and z_2 enter with logs in determining who wins the reputational race, the limit outcome doesn't depend on how z_1 and z_2 go to 0, so long as they go to 0 at the same rate ($z_1/z_2 \in [1/K, K]$ for fixed K).

Given analysis so far, follows that each player i can guarantee a limiting payoff of at least

$$\max\left\{\alpha_i \in C_i : \alpha_i \leq \frac{r_j}{r_i + r_j}\right\}.$$

- If i demands any $\alpha_i \leq \frac{r_j}{r_i+r_j}$ and j demands any $\alpha_j > 1 - \alpha_i$, then $\lambda_i > \lambda_j$, so j loses the reputational race. When $z_1, z_2 \to 0$ at the same rate, whichever player loses the reputational race must concede at $t = 0$ with prob$\to 1$.

The Complete-Info Limit (cntd.)

If each player i has a "rich" set of commitment types, they can each make a demand close to $\frac{r_j}{r_i+r_j}$.

- Therefore, if each player has a rich set of commitment types, payoffs in the complete-info limit are uniquely determined as $\left(\frac{r_2}{r_1+r_2}, \frac{r_1}{r_1+r_2}\right)$, independent of the prior distribution of commitment types on each side.
- Interestingly, these are exactly the same payoffs as in Rubinstein's complete-info, alternating-offers model. (Even though there aren't alternating offers in the current model.)
- Intuition: players' relative costs of delay determine payoffs in both models, in AG due to WOA structure, in Rubinstein due to alternating-offers.

Discrete-Time Bargaining

Finally, what does this analysis of concession games have to do with **bargaining**, where players can change their offers frequently (unlike in the concession game)?

Discrete-Time Bargaining: Intuition

Key insight (Kreps 1990, Myerson 1991): With (only) commitment types, eqm of a concession game with fixed demand and eqm of a bargaining game where players' can change demands are very similar, because

1. If a player changes her offer, this reveals that she's rational, and she is then in the position of an uninformed bargainer facing a possibly-committed opponent.
2. By logic similar to the Coase conjecture or the Stackelberg payoff theorem (details below), the revealed-rational player's continuation payoff at this point is very close to what she gets by accepting the opponent's demand.
3. Therefore, the consequences of changing one's offer in reputational bargaining are very similar to conceding. So the outcome of the game where players can change their offers frequently is very similar to the outcome of the game where the only choice is when to concede.

Discrete-Time Bargaining: Result

AG show that for any sequence of bargaining games (i.e., rules about who gets to make offers when) where each player gets to make at least one offer in every length-Δ interval of time, and for any sequence of PBE, as $\Delta \to 0$ the resulting distribution over outcomes converges to the unique Nash eqm outcome distribution of the concession game.

With rich sets of commitment types, in the iterated limit where first $\Delta \to 0$ and then $z_1, z_2 \to 0$ at the same rate, the Rubinstein (alternating-offers) outcome obtains even if one party makes offers 100 times as frequently as the other!

If instead the sets of commitment types are not rich, there is substantial delay in equilibrium, even as $\Delta \to 0$, as in the WOA.

- This gives a tractable bargaining model with two-sided incomplete info and substantial delay, which seems useful for applications. (However, the "no haggling" and sudden concession predictions are counterfactual.)

The Reputational Coase Conjecture

To complete our treatment of AG, we sketch the argument that if player 1 may be committed to demanding α and player 2 is known to be rational, in the $\Delta \to 0$ limit player 2's payoff cannot exceed $1 - \alpha$ in any PBE of any bargaining game where player 1 makes at least one offer in every length-Δ time interval.

- "Reputational Coase conjecture."

Reputational Coase Conjecture

First, there exists a finite time T such that, if P1 always demands α and never accepts, P2 accepts by T.

Proof is similar to FL89: if P2 doesn't accept, he must believe that P1 will deviate from "commitment behavior" soon; so if P1 doesn't deviate, P2 eventually becomes convinced that P1 will continue commitment behavior; and then accepts.

We now argue that the smallest such T goes to 0 as $\Delta \to 0$.

- Intuition: once reach $T - \varepsilon$, P2 can't gain much by holding out as P1 can get α by insisting for another ε units, so P2 should concede early.
- Unlike FL89, this step is not robust to addition of non-stationary commitment types: e.g., a type that always demands α^1, only accepts offers greater than α^1 before some time $\hat{T}$, but accepts any offer after time $\hat{T}$. These types are called "non-transparent" (Wolitzky 11).

The Reputational Coase Conjecture (cntd.)

- Suppose toward a contradiction that $T > \varepsilon$ for all Δ.
- Suppose we find ourselves at time $T - \varepsilon$, and P1 has insisted on α so far.
- Since P1 can get α by insisting for another ε units of time, P2 can't hope for more than $1 - e^{-r\varepsilon}\alpha$.
- Similarly, for any $\eta > 0$, once reach time $T - \eta\varepsilon$, P2 can't hope for more than $1 - e^{-r\eta\varepsilon}\alpha$.
- Hence, at time $T - \varepsilon$, P2 is willing to reject until $T - \eta\varepsilon$ only if there is a discrete prob (bounded away from 0) that P1 deviates from commitment behavior between $T - \varepsilon$ and $T - \eta\varepsilon$.
- Similarly, P2 is willing to reject from $T - \eta\varepsilon$ to $T - \eta^2\varepsilon$ only if there is a discrete prob that P1 deviates between $T - \eta\varepsilon$ and $T - \eta^2\varepsilon$ is high, and so on. (At step k, P2's waiting cost and the biggest prize he can hope for are both proportional to η^k, so the required concession prob is constant.)
- This is a contradiction, because the total prob that P1 deviates between $T - \varepsilon$ and T is at most 1.

Extensions: Abreu-Pearce 07

- Extend AG to a repeated game, where each period players choose an action and an offer (=share of a feasible continuation payoff). When an offer is accepted, it's binding on both players, so acceptance ends the game.
- (Equivalent interpretations: "repeated game with contracts," "bargaining with payoffs as you go")
- Commitment types can have complex, non-stationary strategies.
- Under some genericity assumptions, an argument similar to AG shows that in any PBE, each player can guarantee herself nearly her "Nash bargaining with threats" payoff by imitating a behavioral type that always plays the NBWT action.
- This shows that AG's equilibrium selection approach can be extended form bargaining to repeated games with contracts.
- Sharp prediction in stark contrast to folk theorem. However, again relies on no non-transparent types.

Extensions: Wolitzky 12

- As an alternative to AG's rather complex eqm reasoning, ask what payoff each player can guarantee herself assuming only that the opponent plays a best response.
- Paper first considers the case where only one player makes a commitment. Then shows that the same guarantees can be approximated in 2-sided reputation models when the commitment probabilities are small.
- Maximum guarantee is attained by a non-stationary commitment type that demands "compensation for delay": guarantee x_0 by announcing path of demands $x(t) = \min\{e^{rt}x_0, 1\}$.
- The guarantees vanishes as $z_1, z_2 \to 0$, but because z_1, z_2 enter the concession rates with logs the guarantees are much larger than z_1, z_2 (e.g., with $z = .001$, a player can guarantee 13% of the surplus).

Extensions: Fanning 16

- Reputational bargaining with an uncertain deadline continuously distributed over $[T - \varepsilon, T]$. As in AP, allows both simple and complex (but transparent) commitment types.
- As in AG, rational players concede for sure by some $T^* < T$, but new cost of delay given by deadline risk. As approach the deadline, there is a spike of agreements.
- Concession occurs with high prob at $t = 0$ (AG's initial concessions) and close to the deadline, with a lull in between. This "U-shaped agreement frequency" seems realistic.
- Non-stationary demands match those in Rubinstein alternating-offers bargaining with a deadline (which are non-stationary, as delay is more costly near the deadline).

Extensions: Abreu, Pearce and Stacchetti 15

- Consider a model with incomplete info about the discount factor of one player, in addition to commitment types.
- Since the privately informed player makes offers, without commitment types there are many PBE. However, adding a small prob of commitment types helps select a unique eqm.
- With stationary commitment types only, essentially select the complete information eqm when the informed player is known to be patient. (A Coase-type result.)
- With non-stationary (transparent) commitment types, there is a *different* unique eqm, now with substantial delay and a better payoff for the uninformed player.

Applications

- Compte and Jehiel 02 and Atakan and Ekmekci 13 introduce outside options, which in AE13 result from a search-and-matching model. AE13 show that endogenous outside options from being able to rematch lead to inefficiency even in the complete-info limit (in contrast to AG).
- Fanning 20 uses reputation bargaining to study **mediation**, where a mediator facilitates agreement by privately exchanges messages with the parties. The tractability of reputational bargaining allows a clean analysis of when/how mediation helps, without relying on signaling refinements.

Final Remarks and Open Questions

Reputational bargaining is a tractable framework for studying the implications of commitment and two-sided uncertainty in bargaining.

Open questions:

- Foundations for commitment types? (Abreu and Sethi 03: evolutionary approach. Weinstein and Yildiz 16: foundation via higher-order payoff uncertainty.)
- Reputation effects in related games, like multilateral bargaining or repeated games?
- There are some experiments on reputational bargaining (Embrey, Frechette, Lehrer 14), but remains to be seen if can be successfully applied empirically outside the lab.