Government 2005: Formal Political Theory I

Lectures 5 & 6

Instructor: Tommaso Nannicini
Teaching Fellow: Jeremy Bowles

Harvard University

September 28 & October 5, 2017

Overview

- Dynamic games with complete and perfect information
  - Market entry
  - Sequential voting
  - Centipede game
  - Stackelberg competition
- Dynamic games with complete but imperfect information
  - Market entry (modified)
  - Bank runs
- Formal definitions and useful general results
  - Extensive-from representation
  - Subgame-perfect Nash equilibrium
  - Single-deviation principle
- Sequential bargaining
  - Finite (with patient or impatient players)
  - Infinite (with impatient players)

Finite games in extensive form (informal definition)

A finite (dynamic) game in extensive-form consists of:

- A finite set of players
- A finite set of histories/nodes (with some of them terminal)
- Player functions (i.e., who gets to decide at each non-terminal node)
- The set of available strategies of every (deciding) player in each non-terminal node
- Payoff functions (i.e., who gets what at terminal nodes)
- Information sets that individuals have in every node

Sequential rationality (informal definition)

In dynamic games, the idea of credible vs non-credible threats (or promises) is crucial

The solution concept of Nash equilibrium is not equipped to deal with this (see next examples)

We need to introduce some concept of sequential rationality to remove irrational (off-equilibrium) choices

Two solution concepts that will take care of that:

- **Backward induction**, which we can apply in games with complete and perfect information
- **Subgame perfection**, which we can apply in contexts with complete but imperfect information

A simple example of market entry

- To illustrate how backward induction works, consider the following simple game with perfect information
- At time zero, a potential "entrant" (E) decides whether to enter the market or not ("in" vs "out")
- At time one, after observing the entry decision, the "incumbent" firm (I) decides whether to accommodate the entry ("soft") or to fight it back with an aggressive market strategy ("hard")
- The history of the game, the decision nodes, the information sets, and the payoffs associated with the terminal nodes are captured by the following game tree

A simple example of market entry (contd.)

- Extensive-form representation of the game

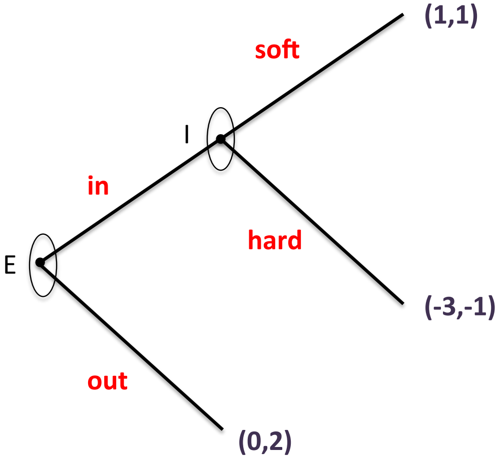

A simple example of market entry (contd.)

- Strategic-form representation has got four strategy profiles, as combinations of $\{$in, out$\}$ and $\{$soft if in, hard if in$\}$
  - Two Nash equilibria: (out, hard if in) and (in, soft if in)
  - The first is based on the non-credible threat of retaliation
- Backward-induction solves this problem:
  - Starting from last decision node, we see that playing "hard" is irrational for the incumbent if we got there
  - The entrant anticipates it and plays "in" at the initial decision node
  - (in, soft if in) is the only Nash equilibrium consistent with this sequential rationality

Sequential voting

- Suppose players 1, 2, and 3 are legislators deciding whether or not to vote themselves a pay-raise
- Each player prefers that the pay-raise passes, but would like to vote against it to avoid resentment from his constituents
- So, for example, player 1's payoffs are:

$$\pi_1 = \begin{cases} 1 & \text{if } v_1 = N,\ v_2 = Y, \text{and } v_3 = Y \\ 1 - c & \text{if } v_1 = Y, \text{ and } v_2 = Y \text{ or } v_3 = Y \text{ or both} \\ 0 & \text{if } v_1 = N, \text{ and } v_2 = N \text{ or } v_3 = N \text{ or both} \\ -c & \text{if } v_1 = Y,\ v_2 = N \text{ and } v_3 = N \end{cases}$$

  where $0 < c < 1$ represents the "constituent resentment" cost
- Analogous expressions hold for players 2 and 3
- Suppose the players vote in a fixed order, 1 then 2 then 3 (and they cannot change their votes)

Sequential voting (contd.)

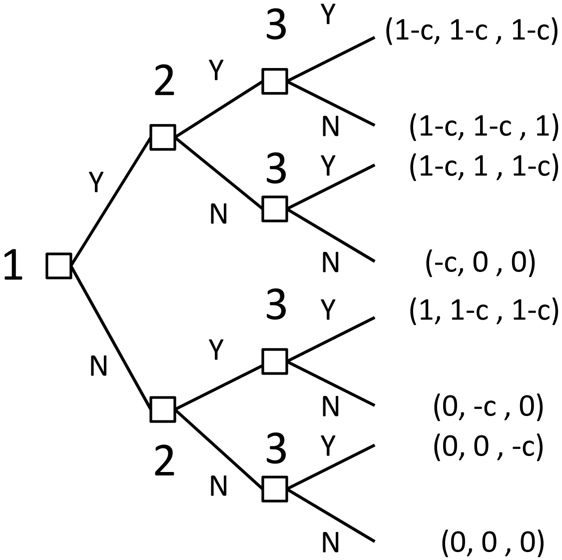

Sequential voting (contd.)

- Using backward induction, we can solve the game as follows
- If players 1 and 2 have both voted Y, then player 3 will vote N and receive his highest payoff: The bill passes and she gets to vote against it
- If 1 and 2 have both voted N, then 3 will vote N and the bill fails
- If 1 and 2 have split their votes, then 3 votes Y and the bill passes
- Note that, if players 1 votes N, the votes of players 2 and 3 are **strategic complements**

Sequential voting (contd.)

- Thus, from player 2's point of view, the game really looks like this when it is her turn to move:

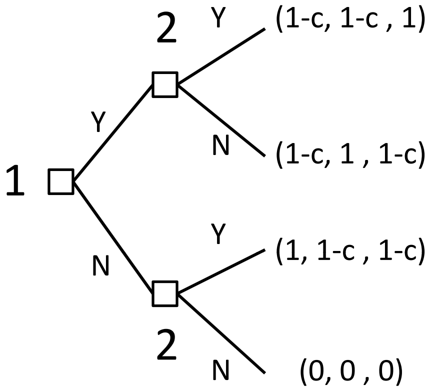

Sequential voting (contd.)

- Continuing the backward induction:
  - If player 1 votes Y then player 2 will vote N
  - If 1 votes N, then 2 will vote Y

- So, from player 1's point of view, the game is as follows when it is her turn to move:

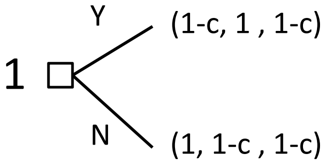

Sequential voting (contd.)

- Player 1 then votes N
- By voting N she *forces* players 2 and 3 to do the "dirty work" of passing the pay-raise
- She can anticipate that since it is in their own best interests to do this, and they will
- The (unique) backward-induction equilibrium is (N, Y, Y)
- There aren't **non-credible threats** off the equilibrium path
- What are the Nash equilibria if the legislators vote simultaneously? (We solved this in class)

Centipede game

- Consider the following game
- There is a pile of money on a table, which grows over time
- Players A and B take turns deciding whether to take the money that is already there, or to let the pile grow larger and hope to take it at a later date
- The game lasts 4 turns and therefore the game tree is as follows:

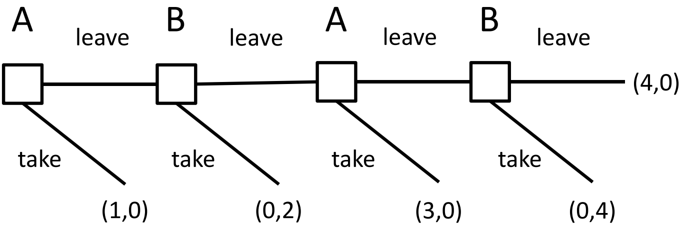

Centipede game (contd.)

- Solving by backward induction, we see that player A must take whatever money there is on the table at the beginning of the game
- If the game reaches the final turn, then B will definitely play "take"
- Anticipating this, if the game reaches turn 3, then A will definitely play "take"
- Anticipating this, if the game reaches turn 2, then B will definitely play "take"
- Anticipating this, A will play "take" at turn 1

Centipede game (contd.)

- This behavior raises an important and difficult issue about backward induction and sequential rationality
- As implemented, backward induction requires that each player assumes that all other players will be rational in all future moves, even if they have played irrationally in the past
  - Suppose the centipede game has reached turn 2
  - Then, A has played *irrationally* at turn 1
  - If B assumes that A will play rationally at turn 3 (and therefore will play "take"), then B should play "take" at turn 2
  - But, should B make such an assumption?
  - After all, A has demonstrated *strange* behavior once already; maybe she will do it again
  - And if she does, then B will get to play t at turn 4 and receive a payoff of 4 rather than 2
- The same puzzle arises for A if the game reaches turn 3

Stackelberg duopoly

- Consider the Cournot duopoly model, but suppose firm 1 chooses its quantity first, and firm 2 observes firm 1's choice before choosing its quantity
- The game tree is:

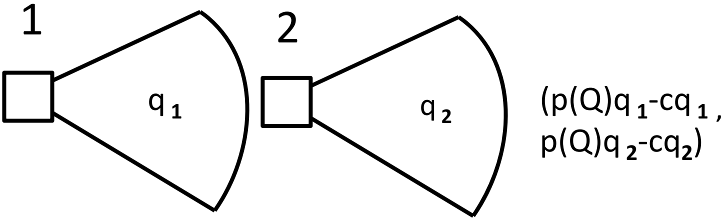

Stackelberg duopoly (contd.)

- To solve the game, consider firm 2's optimization problem given any choice of q_1. Firm 2's profits are

$$\pi_2(q_1, q_2) = (a - c - q_1)q_2 - q_2^2$$

- Differentiating and setting $\pi_2' = 0$ yields the best-response curve for firm 2:

$$\tilde{q}_2(q_1) = \frac{a - c - q_1}{2}$$

  assuming $q_1 < a - c$
- If $q_1 > a - c$, there is a corner solution with $\tilde{q}_2(q_1) = 0$
- Also, note that SOC $-2 < 0$ and thus $\tilde{q}_2(q_1)$ is an argmax

Stackelberg duopoly (contd.)

- Firm 1 can solve this problem also, and therefore *anticipates* what firm 2 will do
- Applying backward induction, firm 1 will choose q_1 to maximize:

$$\begin{aligned}
\tilde{\pi}_1(q_1) &= \pi_1(q_1, \tilde{q}_2(q_1)) \\
&= [a - c - \tilde{q}_2(q_1)]q_1 - q_1^2 \\
&= [a - c - (a - c - q_1)/2]q_1 - q_1^2 \\
&= \tfrac{1}{2}[(a-c)q_1 - q_1^2]
\end{aligned}$$

- Differentiating this and setting $\tilde{\pi}_1' = 0$ yields firm 1's optimal choice:

$$q_1^* = \frac{a - c}{2}$$

- And, firm 2's best-response to this choice is

$$q_2^* = \tilde{q}_2(q_1^*) = \frac{a-c}{4}$$

Stackelberg duopoly (contd.)

- $(q_1^*, q_2^*) = (\frac{a-c}{2}, \frac{a-c}{4})$ is the **backward-induction equilibrium** of the dynamic game
- The equilibrium quantity/price and each firm's profits are:
  $Q^S = \frac{3}{4}(a-c) > \frac{2}{3}(a-c) = Q^C$
  $p(Q^S) = \frac{1}{4}(a+3c) < \frac{1}{3}(a+2c) = p(Q^C)$
  $\pi_1^S = \frac{1}{8}(a-c)^2 > \frac{1}{9}(a-c)^2 = \pi_1^C$
  $\pi_2^S = \frac{1}{16}(a-c)^2 < \frac{1}{9}(a-c)^2 = \pi_2^C$
- The solution differs from the Nash equilibrium of the simultaneous-move game analyzed earlier in several ways
- Total output is higher, so price and total profits are lower
- Firm 1's profits are higher, but firm 2's profits are lower
- There is a **first-mover advantage** for firm 1

Games of complete but imperfect information

- In the games above, players never move *simultaneously* at any point of their sequential interaction
- To allow for simultaneous moves, we must allow for **imperfect information**
- This means that, at some nodes in the game, players do not know what other players are doing
- In other words, *information sets* are no longer singletons
- We need to expand the equilibrium concept from pure backward induction to **subgame perfection** (which can be seen as a generalized concept of backward induction)

Subgame perfection (informal definition)

- A **subgame-perfect Nash equilibrium** is a strategy profile in which players' strategies constitute a Nash equilibrium in every subgame of the game

- A subgame is a part of the game that is itself a game
- A subgame has all of the necessary components of a game, and could be analyzed as a game by itself
- Subgames start at each information set containing a single decision node, contain all subsequent nodes and only them, and cannot cut across information sets
- Instead of backward induction through every node of the game, we will apply backward induction through every subgame of the game

A simple example of market entry (modified)

- Let's modify the market entry game to allow for simultaneous actions after entry:

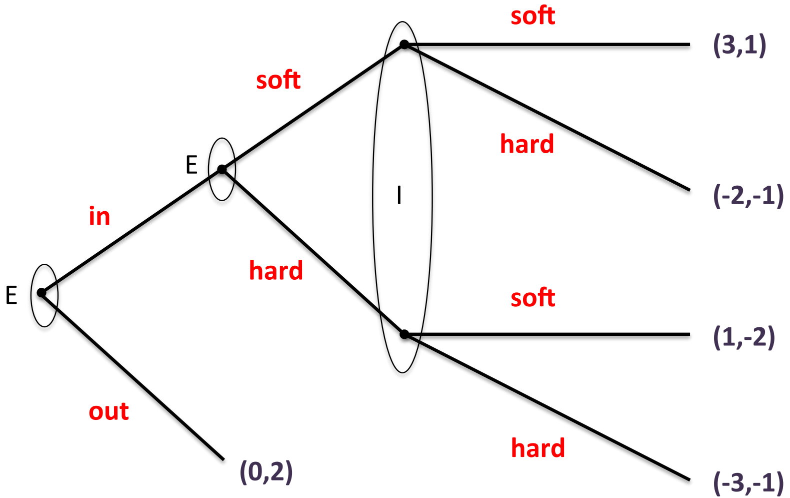

A simple example of market entry (modified and contd.)

- This game contains two subgames: The entire game itself and the simultaneous-move game taking place after entry

		Incumbent	
		Soft	Hard
Entrant	Soft	(3,1)	(-2,-1)
	Hard	(1,-2)	(-3,-1)

- Starting from the last subgame, we realize that it has a unique Nash equilibrium: (soft, soft)
- Anticipating this at the initial decision node, the entrant will definitely play "in"
- The (unique) subgame-perfect Nash equilibrium is (in, soft if in) from E and (soft if in) from I

A simple example of market entry (modified and contd.)

- If we represent the game in strategic form

		Incumbent	
		Soft if In	Hard if In
Entrant	Out, Soft if In	(0,2)	(0,2)
	Out, Hard if In	(0,2)	(0,2)
	In, Soft if In	(3,1)	(-2,-1)
	In, Hard if In	(1,-2)	(-3,-1)

- We find three NE: (out, soft if in) & (hard if in); (out, hard if in) & (hard if in), (in, soft if in) & (soft if in)
- But the first two are based on the non-credible threat of "hard if in" from incumbent

- What if multiple equilibria in a subgame?
- What if ties in single-node individual decision problems?

Bank runs

- Two players have each put \$3 in a bank, and the bank has invested their funds in a three-period project
- The project pays \$10 (out of \$6 of initial investment) if reaches the third period
- Each player must decide whether to withdraw (W) her money in the first and second period or not (D)
- If neither players withdraw her funds in period 1 the project matures, otherwise only \$4 can be recovered
- If neither players withdraw her funds in period 2 the project cashes fully
- However, if any player withdraws at any period she can get her full amount (the deposit in period 1 and the capital gain in period 2) while the other just gets the residual

Bank runs (contd.)

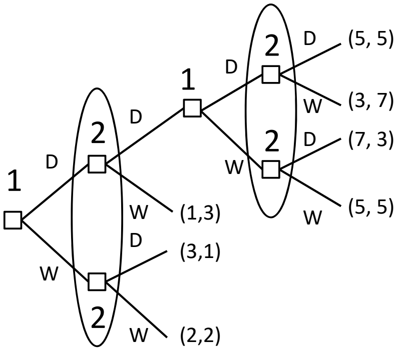

- Not all information sets are singletons $\Rightarrow$ imperfect information
- Two subgames: the game itself and the simultaneous-move game at period 2

Bank runs (contd.)

- We find subgame perfect equilibria by solving backward (one subgame at the time)
- The second period, which begins at the second node for player 1, is a subgame of the whole game
- This subgame has a unique Nash equilibrium: (W,W)

	W	D
W	(5,5)	(7,3)
D	(3,7)	(5,5)

- So, in the first period both players expect that if the game reaches the second period the final payoffs will be (5,5)

Bank runs (contd.)

- The first period can then be reformulated as:

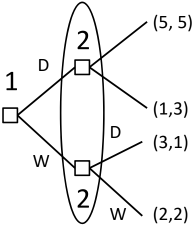

- In this reformulated game there are two Nash equilibria: (W,W) and (D,D)
- This is easily seen by looking at the strategic form:

	W	D
W	(2,2)	(3,1)
D	(1,3)	(5,5)

Bank runs (contd.)

- So, there are two subgame-perfect Nash equilibria to the whole game: (W,W,W,W), and (D,D,W,W)
- We have a bank run equilibrium and normal-functioning-of-banks equilibrium
- It all really depends on players' expectations about what other players will do

Extensive-form representation (reloaded)

Definition. Define $\Gamma^E = \langle N, H, I, p(.), u(.)\rangle$ as the extensive-form representation of the game, which must contain these elements:

- set of players, N;
- set of histories/nodes, H, containing also the terminal histories/nodes H^T and the initial history/node $H^0 = \emptyset$;
- player function, mapping each decision node h to the player who takes action there, $p(h) : H \setminus H^T \to N$;
- information sets, $I \subseteq H \setminus H^T$, such that $p(h) = p(h')$ if h and h' are in the same information set;
- payoff functions, $u_i(h) : H^T \to \mathbb{R}$ for each player $i \in N$.

Subgame-perfect Nash equilibrium (reloaded)

Definition. Define H_i as the subset of histories for which $p(h) = i$ and $A(h)$ as the set of actions available at h. Then, for each player i we define a strategy profile as: $s_i(h) : H_i \rightarrow A(h)$. [Of course, $s_i(h) = s_i(h')$ if h and h' are in the same information set.]

Definition. A subgame of Γ^E is a subset of the game such that:

1. it begins with information set containing single decision node, and contains all subsequent decision nodes (and only them);
2. if h is in the subgame, every $h' \in I(h)$ is also in the subgame, where $I(h)$ is the information set that contains h.

Definition. Given Γ^E, $s(.)$ is a **subgame-perfect Nash equilibrium** if in every subgame of Γ^E the restricted strategy profile $s(.)$ to the subgame is a Nash equilibrium of the subgame.

Useful general results

Theorem. Every finite Γ^E in which every information set is a singleton has a subgame-perfect Nash equilibrium. The equilibrium is unique if no player is indifferent between any two histories.

Theorem. In every finite Γ^E in which every information set is a singleton, the set of subgame-perfect Nash equilibria coincides with the subset of Nash equilbria that can be derived through backward induction.

Theorem (single-deviation principle). A strategy profile $s(.)$ is subgame-perfect if and only if no player has an incentive to deviate at any single information set.

Discounting 101

- In dynamic games, discounting of future values is often a key ingredient
- Define δ as the one-period discount factor, that is, if $\delta = 0.95$ one dollar tomorrow is worth 95 cents today
- The discount factor is linked to the discount rate: $\delta = 1/(1+r)$
- After t periods: $\delta^t = 1/(1+r)^t$
- If you are discounting an infinite flow of constant one-period amounts, it's important to remember these limits:
  - $\sum_{t=0}^{\infty} \delta^t = 1/(1-\delta)$
  - By substituting $\delta = 1/(1+r)$: $\sum_{t=0}^{\infty} \delta^t = (1+r)/r$

Ultimatum game

- Two players: $i \in \{1, 2\}$
- 1 makes offer on how to split sum equal to \$1 (x for herself)
- 2 either accepts (A) or rejects (R) the offer
- In case of rejection, they both get zero
- Histories: (x, Z) with $0 \leq x \leq 1$ and $Z \in \{A, R\}$
- Player function: $p(\emptyset) = 1$, $p(x) = 2$ for any x
- Payoff functions:
  - $U_1 = x$ if $Z = A$, zero otherwise
  - $U_2 = 1 - x$ if $Z = A$, zero otherwise

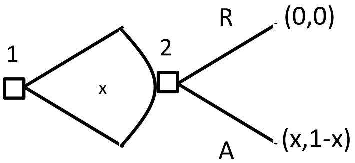

Ultimatum game (contd.)

- There are two subgames, which start at each decision node
- In the last subgame:
  - either 2 accepts every $1 - x \geq 0$
  - or 2 accepts $1 - x > 0$ and rejects $1 - x = 0$
- Going backward to 1's decision:
  - 1 offers $x = 1$ and 2 accepts every $1 - x \geq 0 \Rightarrow$ this is SPNE
  - if 2 accepts $1 - x > 0$ and rejects $1 - x = 0 \Rightarrow$ no offer by 1 is optimal $\Rightarrow$ the above SPNE is unique

- What happens if \$1 is made up of indivisible units/cents?
- What happens if 2 cannot reject the offer? (Dictator game)
- What happens if 1 gets x even if 2 rejects? (Impunity game)
- Does it change if 2 can make costly investment to increase \$1 before the ultimatum game starts? (Holdup problem)

Ultimatum game (contd.)

- Before, we assumed only one offer ($k = 1$)
- What if 2 gets the chance of making an offer too ($k = 2$)
- After rejecting x, 2 can offer y to 1 and take $1 - y$ for herself

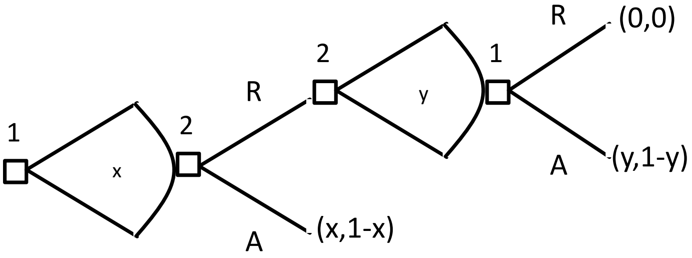

Ultimatum game (contd.)

- The subgame starting with 2 making an offer is an ultimatum game itself
- Its solution is that 2 offers $y = 0$ and 1 accepts every y
- Reasoning backward, 2 is going to reject any $x > 0$
- In all SPNE, 2 gets $1 and 1 nothing
- In this situation (or for any even k), 1 is powerless because 2 is the last guy with the chance to make an offer (vice versa with $K = 1$ or any odd k)

Finite sequential bargaining

- Before, we assumed patient players (no discounting)
- Let's assume player i has discount factor δ_i (impatient players)

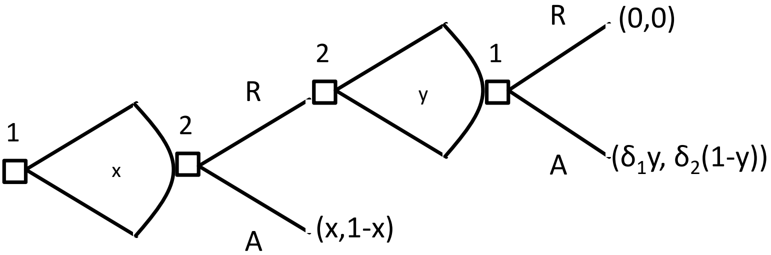

Finite sequential bargaining (contd.)

- Again, in the last ultimatum game, 2 offers $y = 0$ and 1 accepts every y
- Reasoning backward, 2 accepts any x s.t. $(1 - x) \geq \delta_2(1 - y) = \delta_2$
- Reasoning backward, 1 offers $x = 1 - \delta_2$ (offering more for herself is irrational because she ends up with 0, offering less is irrational too)
- The final outcome is $(1 - \delta_2, \delta_2)$ and no rejection takes place in equilibrium
- More precisely, the SPNE is:
  - 1 offers $x = 1 - \delta_2$
  - 2 accepts any $1 - x \geq \delta_2$
  - 2 offers $y = 0$
  - 1 accepts any $y \geq 0$
- The larger δ_2, the more powerful is player 2

Finite sequential bargaining (contd.)

- Now, assume $k = 3$: that is, 1 is the first and last guy who alternates making an offer
- The last offer would lead to: $(\delta_1^2, 0)$
- Therefore, one stage before, 2's offer would lead to: $(\delta_1^2, \delta_2(1 - \delta_1))$
- Therefore, at the first stage, 1's offer is: $(1 - \delta_2(1 - \delta_1), \delta_2(1 - \delta_1))$
- This offer (and anyone providing a larger amount to 2, for what it's worth) is accepted by 2
- Again, patience pays off

Recursivity of this problem (with k odd)

Proposing player	Player 1 discounting	Offered to player 1	Player 2 discounting	Offered to player 2
1	δ_1^k	1	δ_2^k	0
2	δ_1^{k-1}	δ_1	δ_2^{k-1}	$(1-\delta_1)$
1	δ_1^{k-2}	$1-\delta_2(1-\delta_1)$	δ_2^{k-2}	$\delta_2(1-\delta_1)$
2	δ_1^{k-3}	$\delta_1[1-\delta_2(1-\delta_1)]$	δ_2^{k-3}	$1-\delta_1[1-\delta_2(1-\delta_1)]$
1	δ_1^{k-4}	$1-\delta_2\{1-\delta_1[1-\delta_2(1-\delta_1)]\}$	δ_2^{k-4}	$\delta_2\{1-\delta_1[1-\delta_2(1-\delta_1)]\}$

If offers by 1 could go on indefinitely:

$$1-\delta_2+\delta_1\delta_2-\delta_1\delta_2^2+\delta_1^2\delta_2^2-\ldots$$

$$\sum_{m=0}^{M}(\delta_1\delta_2)^m-\delta_2\sum_{m=0}^{M}(\delta_1\delta_2)^m$$

Taking the limit with $M\to\infty$, we get:

$$\frac{1}{1-\delta_1\delta_2}-\frac{\delta_2}{1-\delta_1\delta_2}=\frac{1-\delta_2}{1-\delta_1\delta_2}$$

Infinite sequential bargaining

- Let's assume the 2-player bargaining goes on up to infinity (this makes a lot of sense once you keep in mind that Chuck Norris has counted to infinity, twice)

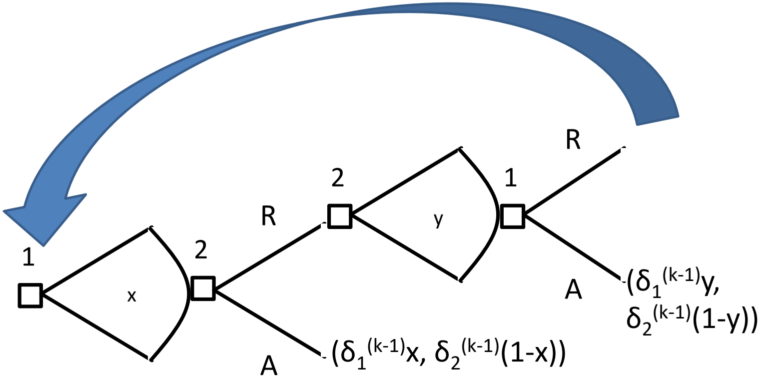

Infinite sequential bargaining (contd.)

- We can no longer use backward induction (even generalized)
- But we can exploit the stationary structure of the problem by assuming a **stationary** solution (i.e., one where choices are the same in each time-period)
- Let's prove that the following is the (unique) SPNE

SPNE of infinite sequential bargaining.

- 1 proposes $(x, 1-x)$ with
  - $x = (1-\delta_2)/(1-\delta_1\delta_2)$
  - $1-x = \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$
- 2 accepts at least $1-x$
- 2 proposes $(y, 1-y)$ with
  - $y = \delta_1(1-\delta_2)/(1-\delta_1\delta_2)$
  - $1-y = (1-\delta_1)/(1-\delta_1\delta_2)$
- 1 accepts at least y

Infinite sequential bargaining (contd.)

How did we come out with those numbers above?

- First way: form the recursive argument discussed above
- Second way: define v_i as the continuation value of the game when i is the proposer (they don't depend on time in a stationary equilibrium)
- 1 must offer: $x = 1 - \delta_2 v_2$, $x = \delta_2 v_2$
- As the offer is accepted: $v_1 = 1 - \delta_2 v_2$
- 2 must offer: $y = \delta_1 v_1$, $1 - y = 1 - \delta_1 v_1$
- As the offer is accepted: $v_2 = 1 - \delta_1 v_1$
- As a result:
  - $v_1 = (1 - \delta_2)/(1 - \delta_1 \delta_2)$
  - $v_2 = (1 - \delta_1)/(1 - \delta_1 \delta_2)$

Infinite sequential bargaining (contd.)

Let's prove that this is indeed a SPNE

- We use the single-deviation principle: No player can make profitable deviation in one single period
- Consider 1 (offering x): She cannot get more than x, and if her offer is rejected she gets $\delta_1^2 x < x$
- Consider 2 (pondering the offer of x): Accepting it, she gets $1 - x = \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$. But refusing it, she also gets $\delta_2(1-y) = \delta_2(1-\delta_1)/(1-\delta_1\delta_2)$
- Similar arguments apply to player 2 offering y, and player 1 pondering that offer

Infinite sequential bargaining (contd.)

Let's prove that this SPNE is unique

- Define $\underline{v_1}, \overline{v_1}$, respectively, as the lowest and highest continuation value that 1 can get in any SPNE (starting at a period when she makes an offer)
- Assume 2 makes an offer: She can get at least $(1 - \delta_1 \overline{v_1})$ and at most $(1 - \delta_1 \underline{v_1})$
- Assume 1 makes an offer: She must offer at least $\delta_2(1 - \delta_1 \overline{v_1})$ to player 2
- So it must hold: $\overline{v_1} \leq 1 - \delta_2(1 - \delta_1 \overline{v_1})$
- Therefore: $\overline{v_1} \leq (1 - \delta_2)/(1 - \delta_1 \delta_2)$
- By the same token, 2 will certainly accept more than $\delta_2(1 - \delta_1 \underline{v_1})$
- So it must hold: $\underline{v_1} \geq 1 - \delta_2(1 - \delta_1 \underline{v_1})$
- Therefore: $\underline{v_1} \geq (1 - \delta_2)/(1 - \delta_1 \delta_2)$
- It follows that: $\overline{v_1} \leq (1 - \delta_2)/(1 - \delta_1 \delta_2) \leq \underline{v_1}$
- But as $\underline{v_1} \leq \overline{v_1}$ by definition, we must have:

$$\overline{v_1} = \underline{v_1} = (1 - \delta_2)/(1 - \delta_1 \delta_2) \quad \text{(Q.E.D.)}$$

Infinite sequential bargaining (contd.)

Discussion of the equilibrium characteristics

- This equilibrium turns out to be efficient, i.e., no time is wasted in bargaining
- It pays to be patient:
  - Payoff of 1 increasing in δ_1 and decreasing in δ_2
  - As $\delta_1 \to 1$, also $x \to 1$
- First-mover advantage:
  - If $\delta_1 = \delta_2 = \delta$: $x = 1/(1+\delta) > \delta/(1+\delta) = 1 - x$