Microeconomics 3

Spring 2021

Game Theory

Economics is all about interactions

- Competition between firms.
- Workers and managers.
- Buyers and sellers
- Experts
- Bargaining
- etc.

Game Theory

How to model decision making with multiple agents? We've previously focused on single agent settings

1. Micro 1: Single Agent - Maximize expected utility
2. Micro 2: Multiple Agents - Price takers, act as if they are single agents.
3. Macro: Model multiple agents as one agent

No strategic concerns. No one thinks about how their behavior influences others'

Game Theory

In this class, we will develop a framework to think about these.

Problems we have to solve:

1. How do we model a strategic setting?
2. What does rationality mean here? What does it mean to maximize utility?
3. How do we model what decision makers reason about other's reasoning?

Formalize the logic of "putting yourself in the other's shoes."

What do I think others will do? How will they change their behavior if they see me change mine?

Game Theory

Build a theory that accomodates many different types of games:

- Static (simultaneous move) games
- Dynamic
- (Im)Perfect Information
- (In)Complete information

Game

What is a game?

What do we need to describe Rock Paper Scissors?

Game

What is a game?

What do we need to describe Rock Paper Scissors?

- Who is playing?

Game

What is a game?

What do we need to describe Rock Paper Scissors?

- Who is playing?
- What they can do?

Game

What is a game?

What do we need to describe Rock Paper Scissors?

- Who is playing?
- What they can do?
- What happens after they've made their play?

Game

A normal form game is:

- N players, $i \in \{1, 2, \ldots N\}$.
- A set of strategies for each player S_i. $S = S_1 \times S_2 \times \ldots S_N$.
- Payoffs $u_i : S \to \mathbb{R}$

(We may occasionally consider infinitely many players.)

We've put no structure on S, which allows us to capture a lot (e.g. dynamics)

Example: Prisoner's Dilemma

- Two players simultaneously choose effort (E) or shirk (S).
- Output $= 6 \times \#$ of E, E costs 4.
- Output is shared equally

Example: Cournot Duopoly

- Two firms compete in an industry.
- Inverse demand: $p(q_1, q_2) = \max\{0, a - b(q_1 + q_2)\}$.
- Firms have constant marginal cost c, $a > c$, $b > 0$.

Example: 2nd Price Auction

- $n \geq 2$ bidders compete for one unit of an indivisible good.
- Each bidder i has valuation $v_i \geq 0$, unknown to others.
- Bids are submitted simultaneously.
- Highest bidder wins, pays 2nd highest bid.

Example: English Auction

- $n \geq 2$ bidders compete for one unit of an indivisible good.
- Each bidder i has valuation $v_i \in \mathbb{N}$, unknown to others.
- In round k, auctioneer announces price k, bidders each announce whether they would buy or not at that price.
- Repeat until one bidder is willing to buy.
- Award item to that bidder, he pays that price

Strategies

Strategies are contingent plans of action. Specify what to do in every situation that may arise!

- In static games (like examples), pretty clear.
- In dynamic games (e.g. Chess), a strategy specifies what move to make after every possible history of moves.
- So we could model these dynamic games as one shot (normal form) games.

Prisoner's Dilemma

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

Prisoner's Dilemma

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

No matter what other player does, S gives higher payoff than E.

Dominance

Definition

Let s_i,$s_i' \in S_i$.

1. s_i is strictly dominated by s_i' if for all $s_{-i} \in S_{-i}$
$$u_i(s_i, s_{-i}) < u_i(s_i', s_{-i}).$$
2. s_i is weakly dominated by s_i' if for all $s_{-i} \in S_{-i}$.
$$u_i(s_i, s_{-i}) \leq u_i(s_i', s_{-i})$$
and for some $s_{-i} \in S_{-i}$
$$u_i(s_i, s_{-i}) < u_i(s_i', s_{-i}).$$

Dominance

No matter how we model reasoning about others, rational players won't play strictly dominated strategies.

What about weak dominance?

- Less clear, there are beliefs that make a weakly dominanted strategy optimal.
- But a "cautious" player, who worried about others playing any strategy would not choose these.
- We'll formalize this later.

Second Price Auction

The SPA has no dominated strategies. Any bid b_i is a best response to $b_j > \max\{b_i, v_i\}$.

But, bidding your value is weakly dominant.

Cornout Duopoly

Problem: Most games don't have a dominant strategy

Cornout duopoly, if player i believes player j chooses q_j then best response is

$$\max q_i(a - b(q_i + q_j) - c)$$
$$\text{FOC } 0 = a - c - 2bq_i - bq_j$$
$$q_i = \frac{a - c - bq_j}{2b}$$

so any q_i between $\frac{a-c}{2b}$ and 0 are best responses to something, so not dominated.

All larger quantities are dominated

Cornout Duopoly

What if each player knows the other doesn't play dominated strategies?

Then $q_j \in [0, (a-c)/2b]$, any $q_i \notin [(a-c)/4b, (a-c)/2b]$ is dominated.

- Now assume player i knows that player j knows that player i knows the other doesn't play dominated strategies.
- We can eliminate more stuff.
- Common knowledge of rationality: Iterated deletion of strictly dominated strategies.

$q_i = (a-c)/3b$ is the only strategy that survives infinite iterations of this procedure.

$q_i + q_j = 2(a-c)/3b$ is greater than the monopoly quantity $(a-c)/2b$, but less than the competitive $(a-c)/b$. Competition decreases, but doesn't eliminate market power.

Iterated Weak Dominance

Iterated deletion of strictly dominated strategies is a clear implication of common knowledge. Moreover, clearly the order we eliminate strategies doesn't really matter.

Iterated weak dominance is less well founded. Consider

	L	R
T	$1,1$	$1,0$
B	$1,0$	$0,1$

We can eliminate B. Then R is dominated, so we eliminate R. But we justified eliminating B because R could be played!

Iterated weak dominance

Order matters!

	L	C	R
T	$1, 2$	$2, 3$	$0, 3$
M	$2, 2$	$2, 1$	$3, 2$
B	$2, 1$	$0, 0$	$1, 0$

Three possible orders of elimination give different predictions!

1. $TRBC$ gives ML
2. $BLCT$ gives MR
3. TCR gives pins down that player 2 plays L, player 1 plays M or B

Rationalizability

Is there anything else we can rule out?

- Suppose players are expected utility maximizers.
- Action they choose must be best response to something
- Eliminate never best responses.

Payoffs

Thinking about what is a never best response is trickier. How do we evaluate payoffs?

But from player i's perspective s_{-i} is unknown, not something they choose.

- How do we deal with this?
- Micro 1: put a probability distribution over it $\mu_{-i} \in \Delta(S_{-i})$ (beliefs).
- So u_i is the Bernoulli utility function, evaluate payoffs using expected utility.

Uncertainty is not entirely subjective. What sort of restrictions make sense to put on beliefs?

Beliefs

Now give each player belief $\mu_{-i} \in \Delta(S_{-i})$.

- Expected payoff of player i when choosing s_i is

$$E_\mu(u_i(s_i, s_{-i})) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \mu_{-i}(s_{-i})$$

- A strategy is a best-reply to belief μ_{-i} if

$$s_i \in \arg\max E_\mu(u_i(s_i, s_{-i})).$$

- A strategy is rationalizable if it survives iterated deletion of never best-replies.

Rationalizability vs. Dominance

Concepts seem related, strictly dominated strategies are clearly not rationalizable.

Is the converse true?

	L	R
T	$2, \cdot$	$-1, \cdot$
M	$0, \cdot$	$0, \cdot$
B	$-1, \cdot$	$2, \cdot$

Rationalizability vs. Dominance

	L	R
T	$2, \cdot$	$-1, \cdot$
M	$0, \cdot$	$0, \cdot$
B	$-1, \cdot$	$2, \cdot$

But, if we allow player 1 to play T w/ prob $1/2$, B o.w., then dominated.

Rationalizability vs. Dominance

Allowing for randomization lets us close the gap

Definition

A mixed strategy for player i is a probability distribution $\sigma_i \in \Delta(S_i)$

In general, when thinking about finite games, always consider mixed extension

- Players are allowed to "flip coins"
- Convexifies the strategy space.
- Can be justified by higher order uncertainty (purification)

Rationalizability vs. Iterated Dominance

Once we consider the mixed extension, these two concepts are the same!

Step 1: If a strategy is a best reply than it is not dominated

- Straightforward, if s_i is a best reply to μ_{-i} then

$$\sum_{s \in S_{-i}} \mu_{-i}(s) u_i(s_i, s) \geq \sum_{s \in S_{-i}} \mu_{-i}(s) u_i(s_i', s) \, \forall s_i' \in S_i$$

- Since expectations are linear, holds for all mixed strategies as well. By linearity, there must be some s_{-i} such that

$$u_i(s_i, s_{-i}) \geq u_i(\sigma_i, s_{-i})$$

Rationalizability vs. Iterated Dominance

Step 2: If a strategy s_i is not dominated, then it is a best reply

- Trickier: Suppose that s_i is not a best reply
- Define set:

$$X = conv\{x \in \mathbb{R}^{|S_i|-1} : \exists s_{-i} \in S_{-i} \, s.t. \, x_{s_i'} = u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})\}$$

- Corners of this set, gain from deviation, fixing s_{-i}.
- This is convex, compact, doesn't cross $\mathbb{R}_{-}^{|S_i|-1}$.
- SHT: $\exists \lambda$ s.t. $\forall x \in X$, $\lambda \cdot x > 0 > \lambda \cdot y$, $\mathbb{R}_{-}^{|S_i|-1}$
- All components of λ are positive, can renormalize to make it a probability. But

$$\sum_{s \in S_i \setminus s_i} \sum_{s_{-i} \in S_{-i}} \lambda_s \mu(s_{-i}) u_i(s, s_{-i}) > \sum_{s_{-i} \in S_{-i}} \mu(s_{-i}) u_i(s_i, s_{-i})$$

So these two concepts are the same

There are a few subtleties here

- When defining best replies, we took $\mu_{-i} \in \Delta(S_{-i})$
  - Correlation between strategies.
  - Natural consequence of expected utility theory w/ subjective uncertainty.
  - Mixed strategies were independent randomizations.
  - Needed this correlation to convexify X.
- If we required ind beliefs, only holds in $N = 2$ case.
- Otherwise, rationalizable strategies are subset of strategies that survive iterated deletion.

Rationalizability

Good news:

- We've characterized what common knowledge of rationality implies.
- It has bite in some games we care about.

Bad news:

- Anything goes in a lot of games (e.g. 2nd price auction)

	L	R
T	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

Equilibrium

	L	R
T	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

What if, in addition, we required beliefs to be correct?

Equilibrium

Definition

The best response correspondence $BR_i : \prod_{j \neq i} S_j \to S_i$ is defined by

$$BR_i(s_{-i}) = \arg\max_{s \in S_i} u_i(s, s_{-i}).$$

Definition

A strategy profile $s \in S$ is a Nash Equilibrium if for all i and for all $s_i' \in S_i$

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}).$$

(Equivalently, $s_i \in BR_i(s_{-i})$ for all i)

- Fixing others' strategies, no player has a **profitable deviation**.
- Straightforward to define for mixed extension.
- Where do correct beliefs come from?

Examples

	L	R
T	$\underline{2}, \underline{1}$	$0, 0$
B	$0, 0$	$\underline{1}, \underline{2}$

Cournot Oligopoly

Bertrand Oligopoly

$N \geq 2$ identical firms, marginal cost c. Compete by choosing prices.

Demand

$$D(p) = a - \min\{p_1, p_2, \dots p_n\}$$

Firm payoff

$$u_i(p_1, \dots p_n) = \begin{cases} (p_i - c)D(p_i) \text{ if } p_i < \min_{j \neq i} p_j \\ \frac{1}{|\{k : p_k = \min_j p_j\}|}(p_i - c)D(p_i) \text{ if } p_i = \min_{j \neq i} p_j \\ 0 \text{ o.w.} \end{cases}$$

Nash equilibrium: 2 or more firms price at cost, everyone else prices above.

Stag Hunt

	L	R
T	$\underline{1}, \underline{1}$	$0, 0$
B	$0, 0$	$\underline{0}, \underline{0}$

TL pareto dominates BR. Is it more compelling?

Stag Hunt

	L	R
T	$\underline{1}, \underline{1}$	$0, 0$
B	$0, 0$	$\underline{0}, \underline{0}$

TL pareto dominates BR. Is it more compelling?

	L	R
T	$\underline{1}, \underline{1}$	$-100, 0$
B	$0, -100$	$\underline{0}, \underline{0}$

BR risk dominates TL, is it more compelling?

Equilibrium

Concept gives us predictive power

- Predicts things that seem plausible-ish
  - Players aren't playing strictly dominated strategies.
  - Put self in other's shoes
- No one wants to unilaterally change their mind.
- Outcome of pre play coordination, repeated play w/ learning?

Uniqueness

- Not unique
- In some ways a feature
- Refinements: "focus" attention on more plausible set of equilibrium
  - Prediction vs. isolate economic forces

Do these exist?

Existence - Matching pennies

Unlike w/ rationalizability, not clear anything satisfies this criteria

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Mixed strategy equilibrium, both mix uniformly.

Mixed Strategies

Allowing for ind randomizations convexify strategy space, fixing existence problem.

Games can have pure and mixed strategies

	L	R
T	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

Mixed Strategies

Theorem

$\sigma \in \prod_{i=1}^{N} \Delta(S_i)$ is a Nash equilibrium of a finite game iff for all i

$$u_i(\sigma) = u_i(s_i, \sigma_{-i}), \; \forall s_i \in \operatorname{supp} \sigma_i$$

and

$$u_i(\sigma) \geq u_i(s_i, \sigma_{-i}) \forall s_i \in S_i \setminus \operatorname{supp} \sigma_i$$

(holds in infinite games as well w/ some stuff about measure 0 sets)

Gives us a way to solve for these.

Mixed Strategies

Do these make sense?

- Players only randomize to make the other player want to randomize.
- Some other justifications:
  - Steady state
  - Pure strategies in perturbed game
  - Beliefs

Existence

Theorem

Finite games have a mixed strategy Nash equilibrium

Proof:

- The best response correspondence is the arg max of a continuous function over a compact set (with no constraints).

$$BR_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Delta(S_i)} \sum_{s \in S} \left(\prod_{i=1}^{N} \sigma_j(s_j) \right) u_i(s)$$

Existence

Theorem

Finite games have a mixed strategy Nash equilibrium

Proof:

- The best response correspondence is the arg max of a continuous function over a compact set (with no constraints).

$$BR_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Delta(S_i)} \sum_{s \in S} \left(\prod_{i=1}^{N} \sigma_j(s_j) \right) u_i(s)$$

  - Maximum theorem: BR_i is UHC, non-empty, compact valued.

Existence

Theorem

Finite games have a mixed strategy Nash equilibrium

Proof:

- The best response correspondence is the arg max of a continuous function over a compact set (with no constraints).

$$BR_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Delta(S_i)} \sum_{s \in S} \left(\prod_{i=1}^{N} \sigma_j(s_j) \right) u_i(s)$$

  - Maximum theorem: BR_i is UHC, non-empty, compact valued.
  - Utility function is linear (quasiconcave) is σ_i. So BR_i is convex valued.

Existence

Theorem

Finite games have a mixed strategy Nash equilibrium

Proof:

- The best response correspondence is the $\arg\max$ of a continuous function over a compact set (with no constraints).

$$BR_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Delta(S_i)} \sum_{s \in S} \left(\prod_{i=1}^{N} \sigma_j(s_j) \right) u_i(s)$$

  - Maximum theorem: BR_i is UHC, non-empty, compact valued.
  - Utility function is linear (quasiconcave) is σ_i. So BR_i is convex valued.
- $\sigma \mapsto \prod BR_i(\sigma_{-i})$ is convex, compact valued, non-empty correspondence.
  - Kakutani's fixed point theorem: Convex, compact valued, non-empty correspondence that maps a compact space to itself has a fixed point.

Existence

Needed: Continuous, quasiconcave unity, “nice” strategy space.

- There are natural settings where our result doesn’t apply.
  - Bertrand oligopoly
  - Dynamic games w/ infinite periods.
- Argument is non-constructive.
- Other theorems can be used in discontinuous/non-quasiconcave settings (e.g. potential games)

Some odds and ends

Nash equilibrium are hard to compute

- Brouwer fixed point argument is non-constructive
- *BR* dynamics don't necessarily converge to a eqm.

A bit sketchy epistemically

- What additional assumptions on knowledge are we making.
- There are Nash equilibrium where everyone plays a weakly dominated strategy.

Correlated Equilibrium

It may make sense to allow players to correlate their actions.

Suppose players have access to a public randomization device.

	L	R
T	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

Now has a symmetric eqm where both players do better than mixed eqm. How?

- Players use public randomizing device to coordinate. No longer have to play BL or TR.
- Must coordinate on a Nash Equilibrium for each realization of randomizing device.

So achievable payoffs:

$$Conv\{(2,1),(1,2),(\frac{2}{3},\frac{2}{3})\}$$

Correlated Equilibrium

In some games we can do even better with a mediator (or arbitrary correlated strategies)

	L	R
T	$2, 7$	$6, 6$
B	$0, 0$	$7, 2$

Correlated Equilibrium

Mediator can make privately, correlated action recommendations, $a \in \Delta(S)$.

	L	R
T	$2,7$	$6,6$
B	$0,0$	$7,2$

- For a player to recommendation, they must not have a profitable deviation.
- Let a_i denote player i's recommended strategy.
- So if the mediator recommends player 1 plays T then

$$\underbrace{2Pr(L|a_1 = T) + 6Pr(R|a_1 = T)}_{\text{Payoff from T}} \geq \underbrace{7Pr(R|a_1 = T)}_{\text{Payoff from deviating}} .$$

- We can do a similar exercise for the rest of the strategies, any recommendations work that satisfy

$$2Pr(L|a_1 = T) \geq Pr(R|a_1 = T), Pr(R|a_1 = B) \geq 2Pr(L|a_1 = B)$$
$$Pr(T|a_2 = L) \geq 2Pr(B|a_2 = L), 2Pr(B|a_2 = R) \geq Pr(T|a_2 = R)$$

Correlated Equilibrium

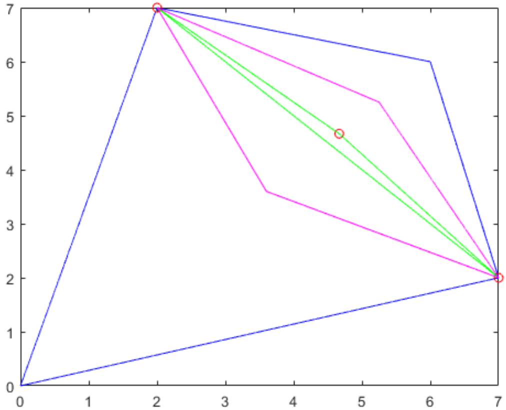

- Red circles: Nash equilibrium
- Area enclosed by green: Correlated equilibrium with public randomization
- Area enclosed by pink: Correlated eqm with mediator

Trembling Hand Perfection

Intuition: Robustness to small "mistakes".

	L	R
T	$1, 1$	$0, 0$
B	$0, 0$	$0, 0$

RB seems weird

Definition

An equilibrium σ of a finite normal form game is a normal form trembling hand perfect equilibrium if there exists a sequence $\{\sigma^k\}_{k=1}^{\infty}$ of completely mixed strategy profiles such that

- $\sigma^k \to \sigma$
- σ_i is a best reply to σ^k_{-i} for all k

Extensive Form Games

Consider the following game:

- An entrant considers entering a market with an incumbent firm.
- Entry costs 1, monopoly profits are 4.
- If the firm enters, incumbent can accommodate (split monopoly profits), or fight (both get 0).

	Fight if entry	Accommodate if entry
Enter	$-1, 0$	$1, 2$
Out	$0, 4$	$0, 4$

Two pure NE

Extensive Form Games

Timing seems like it should matter.

- Nash equilibrium can be supported with threats that don't seem credible.
- Normal form doesn't capture timing.

Extensive form

Extensive form allows us to visualize timing

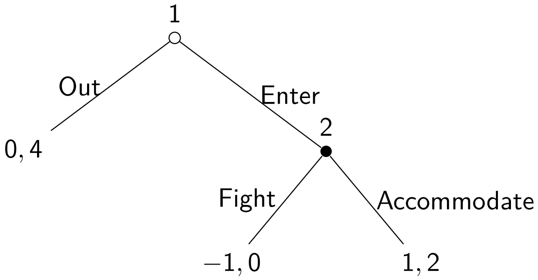

Extensive form

An extensive form game is:

- An arborescence: A tree with
  - A single initial node
  - For each node unique path connecting it to the root.
- Terminal nodes Z
- Decision nodes $X \setminus Z$
- $\iota : X \setminus Z \to \{1, 2, \ldots N\}$, player who owns each node.
- $A(x)$ set of actions available at node x.
- $u_i : Z \to \mathbb{R}$
- Info sets: Partition of nodes $H = \{h^1, \ldots h^k\}$.
  - Same player owns every node in info set.
  - Every info set has same actions

Extensive form

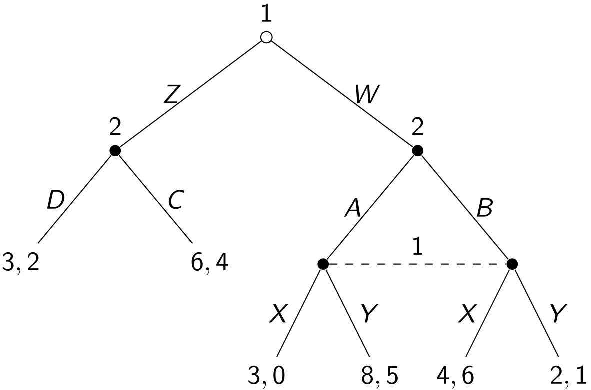

Info sets

Players can't tell which node in an info set they are at.

Throughout this course, assume perfect recall.

- Players don't forget their own past actions.
- Restriction on possible Hs

Extensive Form: Strategies

Strategies are tricky here!

- A strategy specifies what a player does at each info set.
- This includes what happens at nodes that aren't reached!

It's tempting to describe an equilibrium either by describing the outcome or on-path behavior. Don't do this!

As we'll see, behavior at unreached nodes is a crucial part of the eqm description.

Extensive Form

Clearly any static game can be written in extensive form.

Strategies in extensive form:

$$S_i = \times_{h_i \in H_i} A(h_i).$$

We can write extensive form as normal form. Do we lose important info by doing this?

Reduced Normal Form: Collapse equivalent strategies together

Example: Centipede game

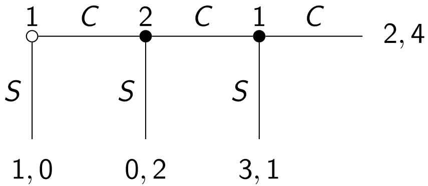

	C	S
CC	2, 4	0, 2
CS	3, 1	0, 2
SC	1, 0	1, 0
SS	1, 0	1, 0

	C	S
CC	2, 4	0, 2
CS	3, 1	0, 2
S	1, 0	1, 0

Reduced Normal Form

Can solve for eqm in the reduced normal form.

We've already seen nash eqm that seem plausible in normal form, but odd in extensive form

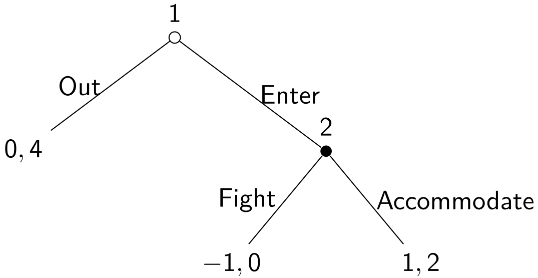

Behavior Strategies

Can define mixed extension as before.

- This is a bit messy, mixed strategy is over lots of actions.
- Conceptually easier to define things at each info set

Definition

A behavior strategy α_i is a distribution over actions at each of player i's info sets:

$$\alpha_i : H_i \to \cup_{h_i \in H_i} \Delta(A(h_i))$$

s.t. $\alpha_i(h_i) \in \Delta(A(h_i))$.

Players mix independently at each info set. Clearly behavior strategies correspond to some mixed strategy. The reverse is also true.

Theorem

(Kuhn) Every mixed strategy has a realization equivalent behavior strategy.

Backwards Induction

Definition

A game has perfect information if all information sets are singletons.

Can solve games of perfect information through backwards induction.

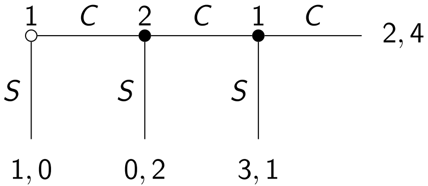

Backwards Induction

Theorem

Every game of perfect information has a pure strategy Nash Equilibrium.

Proof is mechanical. If no ties, backwards induction gives unique prediction.

Predation

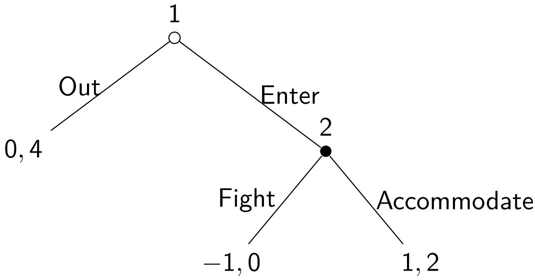

Predation

Backwards induction predicts (Enter, Accommodate), while there are 2 pure NE.

- Uncredible threats - If firm enters, will firm really fight?
- Connection to weakly dominated strategies

Stackelberg Game

Consider the Cornout game, but player 1 moves first.

- Solve through backwards induction:

$$q_2(q_1) = \frac{a-c}{2b} - \frac{q_1}{2}$$

- Player 1 maximizes payoffs knowing how q_2 responds

$$q_1 = (a-c)/(2b),\ q_2 = (a-c)/(4b)$$

- Player 1 makes $(a-c)^2/(8b)$, player 2 makes $(a-c)^2/(16b)$.

Subgame Perfection

Backwards induction needs perfect info.

But can apply same logic to imperfect info games.

Definition

A subgame of an extensive form game is a collection of nodes with

- There is a single node s.t. all other nodes are its successors.
- If a node is in the subgame, so are all its successors.
- If a node is in a subgame than so are all nodes in its info set

- Smaller trees contained in the extensive form.
- Root of a subgame must be a singleton info set.
- Every game has at least 1 subgame.
- In perfect info games, every decision node identifies a subgame

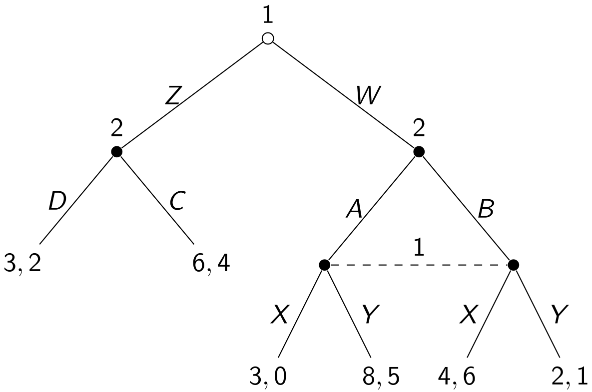
1
Z
W
2
2
D
C
A
B
3, 2
6, 4
1
X
Y
X
Y
3, 0
8, 5
4, 6
2, 1

Subgame Perfection

Definition

A behavior strategy profile α is a Subgame Perfect equilibrium if for every subgame the corresponding profile is a Nash Equilibrium.

- Every SPE is a NE.
- Stronger than NE, NE only requires equilibrium play in on path subgames.
- Can find using backwards induction in finite games w/ perfect info
- (W/ Perfect Info corresponds to iterated deletion of weakly dominated strategies, sort of)

Subgame Perfection

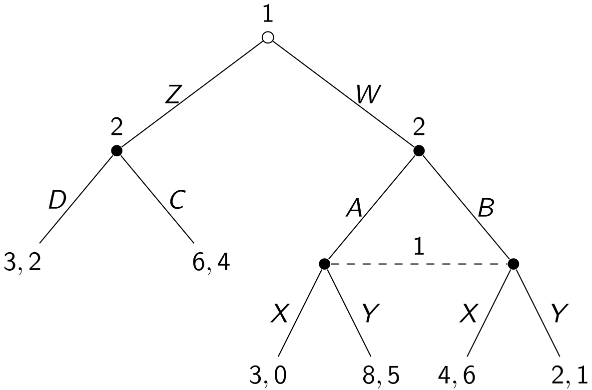

One Shot Deviation Principal

Once we don't have perfect info, backwards induction is harder to use.

Theorem

In a finite multi-stage game with observed actions, strategy profile s is a subgame perfect equilibrium if and only if there is no player i and no strategy s_i' that agrees with s_i except at a single t and h^t, and such that s_i' is a better response to s_{-i} than s_i conditional on history h^t being reached.

Example

Consider playing the following game twice

	E	S	P
E	$2, 2$	$-1, 3$	$-1, -1$
S	$3, -1$	$0, 0$	$-1, -1$
P	$-1, -1$	$-1, -1$	$-2, -2$

Is the following an SPE

$$s_i^1 = E$$

$$s_i^2(a_i, a_j) = \begin{cases} \text{S if } a_j = E \\ P \text{ o.w.} \end{cases}$$

Subgame Perfection

Now we have a dynamic solution concept. Some potential issues:

- Different extensive forms can have the same normal form.
- Different "equivalent" extensive forms can have sets of SPEs
- Some games have few or no subgames

We should maybe think more about games w/ imperfect info

- Should our solution concept treat all games with same reduced normal form the same?
  - Not clear, seems like we are losing something about role of dynamics
- But our concept should at least eliminate pathological cases, where dynamic games are clearly the same.

Repeated Games

	E	S
E	$2,2$	$-1,3$
S	$3,-1$	$0,0$

- Intuition: Maybe E because game is played multiple times.
- Predation: Maybe fight entrants to dissuade future entrants.
- What happens if we play games repeatedly.
- Dynamic incentives: Carrots + sticks

Repeated Games

Take a stage game G, and repeat it $K \in \mathbb{N}$ times.

- Strategies: What to play in stage game conditional on every history of stage game play.
- Can solve through backwards induction.
- If the stage game has a unique SPE, so does repeated game.

Example

Consider repeating the prisoners dilemma N times

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

Repeated Games

What if we repeat the game infinitely many times.

$$V_i(a) = (1-\delta)\sum_{t=0}^{\infty} \delta^t u_i(a_t)$$

(normalized discounted payoffs)

- Payoff is discounted sum of payoffs in stage game.
- Discount to capture impatience
- Info sets at time t identified by history of play.
- Repeated play of same action profile gives geometric series, recall

$$\sum_{t=0}^{\infty} \delta_t U = \frac{U}{1-\delta}$$

(we multiply payoffs by $1-\delta$ to normalize)

Repeated Games

Consider the repeated prisoners dilemma

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

Some problems:

- Strategies are hard to write down!
- No "last" subgame, can't work backwards.

We know an eqm exists. What can we say about the set of eqm?

Goal: Folk Theorem

Automata

Strategies say what to do after every possible history

- Have to specify infinite contingencies.
- Any strategy that we can tractably deal with has more structure.
- Easy way to visualize dynamic strategies: Flow charts!

Definition

An automaton is a collection of states Ω, an initial state $\omega_0 \in \Omega$, a mapping from states to actions $f : \Omega \to A$ and a transition rule $\tau : A \times \Omega \to \Omega$.

Automata

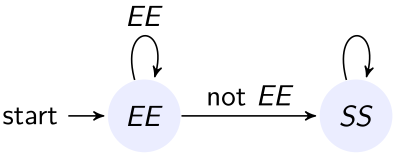

Immediately implies what happens after any contingency.

States are arbitrary (e.g. we define them). Can have multiple states w/ same recommended action

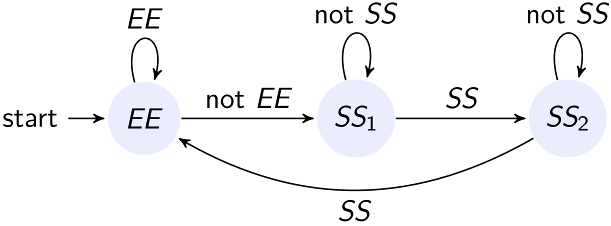

Automata

Easy to solve for payoffs assuming everyone follows automaton

$$V_i(\omega) = (1-\delta)u_i(a(\omega)) + \delta V_i(\tau(f(\omega),\omega))$$

Gives a system of linear equations, one for each state.

Example

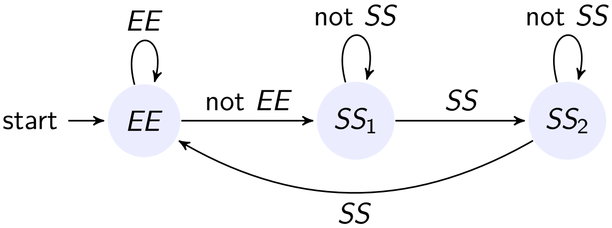

	E	S
E	$2, 2$	$-1, 3$
S	$3, -1$	$0, 0$

$$V_i(EE) = (1-\delta)2 + \delta V_i(EE)$$
$$V_i(SS_1) = (1-\delta)0 + \delta V_i(SS_2)$$
$$V_i(SS_2) = (1-\delta)0 + \delta V_i(EE)$$

SPE

Lemma

A strategy profile represented by an $(\Omega, \omega_0, f, \tau)$ *is a SPE iff for all states* ω *reachable from the initial state the strategy profile represented by* $(\Omega, \omega, f, \tau)$ *is a Nash equilibrium of the repeated game.*

The circles in the diagram tell us exactly what we need to check. Can we simplify more.

Definition

Player i has a profitable one-shot deviation from $(\Omega, \omega_0, f, \tau)$ if for some state ω and some action $a_i \in A$ such that

$$V_i(\omega) < (1 - \delta)u_i(a_i, f_{-i}(\omega)) + \delta V_i(\tau((a_i, f_{-i}(\omega)), \omega))$$

One Shot Deviation Principle

Definition

A strategy profile is subgame perfect iff no one shot deviations.

Example - Grim Trigger

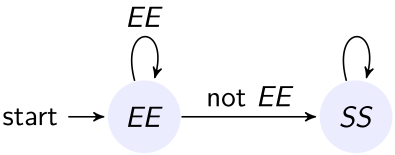

	E	S
E	$2, 2$	$-(1-\delta), 3(1-\delta)$
S	$3(1-\delta), -1(1-\delta)$	$0, 0$

	E	S
E	$2(1-\delta), 2(1-\delta)$	$-(1-\delta), 3(1-\delta)$
S	$3(1-\delta), -(1-\delta)$	$0, 0$

- Have to check whether autamaton describes eqm in each state.
- States are part of our construction of strategies, not part of the statement of the game!
- More complex strategies require more states, more deviations to check.

One Shot Deviation Principle - Finite Horizon

Suppose a finite stage game G is repeated T times.

WTS: No one shot deviations imply subgame perfection

- Consider a profile α that is not a SPE.
- Then there exists a subgame where some player i has a profitable deviation, $\hat{\alpha}_i$.
  - Let $\hat{\alpha}_i$ be the best such deviation.
- There must exist a time s and s length history h_s such that
  - History h_s occur w/ pos probability in the subgame.
  - For any history following h_s, $\hat{\alpha}_i = \alpha_i$ or i is indifferent
  - i strictly prefers $\hat{\alpha}_i(h_s)$ to $\alpha_i(h_s)$.
- Then following $\hat{\alpha}_i$ at h_s and α_i o.w. is a profitable deviation.

One Shot Deviation Principle - Infinite Horizon

Subgame perfection $\Rightarrow$ no one shot deviations is straightforward for all games.

Proof (for finite automaton) of other direction:

- Suppose not, $\exists$ a state ω and a player i s.t. i has a profitable deviation.
- Let $\bar{V}$ be the value from player i's best response.

$$\bar{V}_i(\omega) = \max_{a_i \in A_i}\{(1-\delta)u_i(a_i, f_{-i}(\omega)) + \delta \bar{V}_i(\tau((a_i, f_{-i}(\omega)), \omega))$$

Let $\bar{\omega}$ be the state that maximizes $\bar{V}_i(\omega) - V_i(\omega)$.

We've found where you gain the most from a deviation, now using this we construct a profitable one shot deviation

Proof Continued

- Let

$$V_i^{os}(\omega) = \underbrace{(1-\delta)u_i(a_i^\omega, f_{-i}(\omega))}_{\text{Deviate today}} + \delta \underbrace{V_i(\tau((a_i^\omega, f_{-i}(\omega)), \omega))}_{\text{Follow automaton tomorrow on}}$$

  where a_i^ω is the action chosen in the $\bar{V}$ problem.

- Note that

$$\bar{V}_i(\bar{\omega}) - V_i(\bar{\omega}) \geq \delta[\bar{V}_i(\omega) - V_i(\omega)] \, \forall \omega$$

  so

$$\begin{aligned}
\bar{V}_i(\bar{\omega}) - V_i(\bar{\omega}) &\geq \bar{V}_i(\tau((a_i^{\bar{\omega}}, f_{-i}(\bar{\omega})), \bar{\omega})) - V_i(\tau((a_i^{\bar{\omega}}, f_{-i}(\bar{\omega})), \bar{\omega}) \\
&> \delta[\bar{V}_i(\tau((a_i^{\bar{\omega}}, f_{-i}(\bar{\omega})), \bar{\omega})) - V_i(\tau((a_i^{\bar{\omega}}, f_{-i}(\bar{\omega})), \bar{\omega})] \\
&= \bar{V}_i(\bar{\omega}) - V_i^{os}(\bar{\omega})
\end{aligned}$$

One shot deviation principle

- Proof generalizes to autamoton with infinite states, compact A, etc.
- Similar results for finite games, games without recursive structure, incomplete info games (under right solution concept)
- Discounting important
  - Unlike in finite horizon.
  - Effectively allows us to ignore “far away” parts of deviations.
- There are nash equilibrium that do not have this property!

Folk theorem

For sufficiently patient players, (EE) was an eqm in Prisoner's Dilemma

- Patience + infinite horizon gave us rich enough set of rewards + punishments for cooperation.
- How general is this?

Definition

We say a payoff vector v is strictly Individually Rational if $v_i > \min_{\sigma_{-i} \in \prod_{j \neq i} \Delta(A_j)} \max_{\sigma_i \in \Delta(A_i)} u_i(\sigma_i, \sigma_{-i}) := \underline{v}_i$.

A payoff vector v is feasible if $v \in conv\{v \in \mathbb{R}^n : \exists a \in A \, s.t. \, u_i(a) = v_i \, \forall i\}$.

- Individually Rational: Players can always get at least this payoff
- Feasible: Exists a strategy profile that gives this payoff (as $\delta \to 1$)

Theorem (Folk Theorem)

Fix an infinitely repeated game with stage game G. If v is strictly individually rational and feasible then there exists an $\bar{\delta} < 1$ such that if $\delta > \bar{\delta}$ then there exists an SPE that gives payoffs v.

Folk Theorem

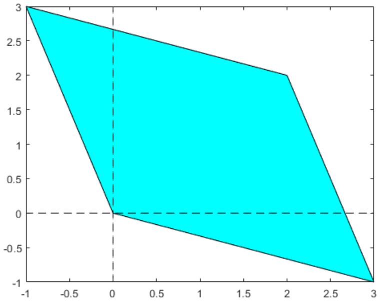

- Feasibility: Convex combination of stage game payoffs
  - Any feasible payoff can be achieved by some strategy profile as $\delta \to 1$.
- Individually rational: $\underline{v}_i$ is min max payoff
  - Best player i can do if everyone else is trying to minimize their payoffs.
  - Player i can always guarantee at least $\underline{v}_i$.

Folk theorem

Theorem

Folk Theorem Fix an infinitely repeated game with stage game G. If v is strictly individually rational and feasible then there exists an $\bar{\delta} < 1$ such that if $\delta > \bar{\delta}$ then there exists an SPE that gives payoffs v.

What's tricky about proving this?

- In prisoner's dilemma, logic straightforward.
  - Punish deviations w/ repeated play of static nash
  - More generally, punish by putting at approx $\min\max$ value.
  - Is that possible?
- Constructing on-path strategy that gives v is complicated.

What do these equilibria look like?

- Claim: there's a pure strategy SPE that gives both players a payoff of $\sqrt{2}$
- Potentially quite complicated
- Do we care?

Folk Theorem

Public randomization

- At the start of each stage game $x \sim U[0,1]$ drawn publicly.
- Makes stuff convex.
- Folk theorem in PD straightforward

Folk Theorem

How should we think about this?

- Good news: Repetition + Patience supports cooperation
- Inefficient SPEs help us understand how to provide incentives for efficient ones
- Bad news: Lots of eqm, many with complex behavior
  - Eqm without cycles in behavior.
  - Some equilibrium seem hard to compute.
  - How do people coordinate on an outcome?
- Prisoner's Dilemma had nice structure
  - Repeated play of static nash, lowest point in feasible, IR set.
  - Punishing deviations may be complicated.

Short Lived Players

Consider the following game (product choice)

	C	S
H	$2, 3$	$0, 2$
L	$3, 0$	$1, 1$

Player 2 is short-lived, new player 2 in each period.

- Equivalently: Continuum of player 2s

Natural modeling technique; captures intuitive features + simplifies analysis.

Short Lived Players

	C	S
H	$2,3$	$0,2$
L	$3,0$	$1,1$

Short lived players do not have dynamic incentives

- Play myopic optimum
- If think player 1 plays H, play C, if L than S
- In eqm those are only pure strategy profiles possible in stage game.

Short Lived Players

Can sustain cooperation

	C	S
H	$2, 3$	$0, 2$
L	$3, 0$	$1, 1$

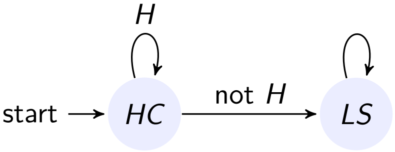

Short lived player:

- State HC

$$3 \geq 2$$

- State SL

$$1 \geq 0$$

Short Lived Players

Can sustain cooperation

	C	S
H	$2,3$	$0,2$
L	$3,0$	$1,1$

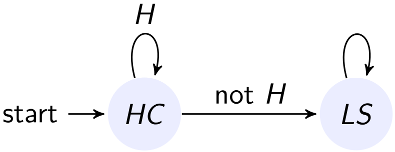

Long lived player:

- State HC

$$2 \geq 3(1-\delta) + \delta$$

- State SL

$$1 \geq 0$$

Short Lived Players

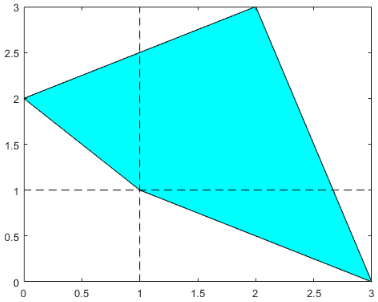

Player 1 can't get payoffs above 2.

- Could in game with long-lived player 2.
- Can still sustain payoffs above static nash.
- Firm establishing a reputation? Does that story make sense?

Another application: Repeated predation

Bargaining

Two players trying to split a dollar.

- In even periods, player 1 offers split $(x, 1-x)$, $x \in [0,1]$.
- Player 2 chooses to accept or reject.
  - Accept: game ends, player 1 gets x (discounted), player 2 gets $1-x$
  - Reject: Move to next period, player 2 proposes split $(x, 1-x)$...
- Player has discount rate δ_i.
- Both get 0 if no agreement.

Bargaining

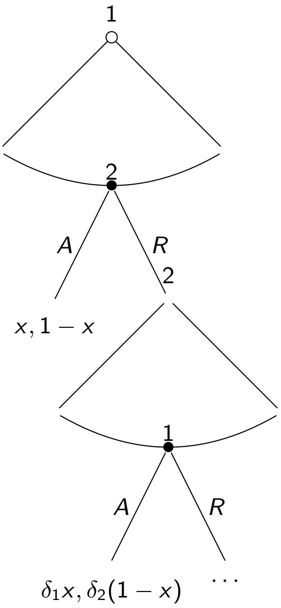

Bargaining

What are SPEs?

- Once again: Strategies are complicated.
- Stationary strategies: Player plays same strategy in each stage.

Now can work backwards, let V_i be stationary eqm payoff.

Bargaining

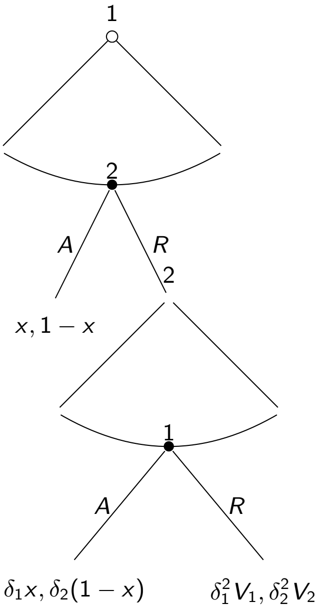

Bargaining

- Last subgame
  - Player 1 accepts iff $x \geq \delta_1 V_1$.
- 2nd to last
  - Player 2 offers $x = \delta_1 V_1$
- 3rd
  - Player 2 accepts iff $(1 - x) \geq \delta_2(1 - \delta_1 V_1)$
- Initial period
  - Player 1 offers $x = 1 - \delta_2(1 - \delta_1 V_1)$

If a stationary eqm of this form exists then:

$$V_1 = 1 - \delta_2(1 - \delta_1 V_1)$$
$$V_1 = \frac{1 - \delta_2}{1 - \delta_1 \delta_2}$$

Bargaining

Discount rate pins down bargaining power:

- First mover advantage, let $\delta_1 = \delta_2$ then

  $$V_1 = \frac{1}{1+\delta} > 1/2$$

  Vanishes as players get patient.
- As $\delta_2 \to 0$ or $\delta_1 \to 1$ player 1 gets whole pie.
- Time between periods
  - Isolate impact of impatience by taking time between period to 0.
  - Let Δ be the time between periods, $\delta_i = e^{-r_i \Delta}$.

  $$\lim_{\Delta \to 0} \frac{1 - e^{-r_2 \Delta}}{1 - e^{-(r_1 + r_2)\Delta}} = \frac{r_2}{r_1 + r_2}$$

  - Gets rid of 1st mover advantage.
  - Bargaining power depends on relative discount rates.

Bargaining - Uniqueness

Are there non-stationary *SPE*s? No

Proof:

- Let $M_i = \sup\{i\text{'s expected payoff in any SPE where i proposes initially}\}$.
- $m_i = \inf\{i\text{'s expected payoff in any SPE where i proposes initially}\}$.
- In any SPE
  1. i accepts any offer $> \delta_i M_i$
  2. i rejects any offer $< \delta_i m_i$
- Claim 1: $m_j \geq 1 - \delta_i M_i$.
- Claim 2 $M_j \leq 1 - \delta_i m_i$.
- Claim 1 + Claim 2:

$$M_j \leq 1 - \delta_i m_i \leq 1 - \delta_i(1 - \delta_j M_j)$$

Bargaining Uniqueness

Therefore

$$M_j \leq \frac{1-\delta_i}{1-\delta_i\delta_j}, \; M_i \leq \frac{1-\delta_j}{1-\delta_i\delta_j}$$

and from claim 1 again

$$m_i \geq 1 - \delta_j \frac{1-\delta_i}{1-\delta_i\delta_j} = \frac{1-\delta_j}{1-\delta_i\delta_j} = M_i$$

So all SPE give same payoff as the stationary eqm.

- Important features for this result:
  - Two players
  - No outside options
  - No breakdown
  - Alternating offers
- Solution is efficient, symmetric in limit when environment is symmetric
- Delay in bargaining?

Incomplete Information

Complete information seems demanding

- Players know exactly the game being played.
- Know that others know the game being played.
- ...

Most economically relevant games have incomplete info

- Cournot with private costs
- Auctions
- Signaling games

Incomplete info

How to model this? (Harsanyi)

- Give each player a type
- Type encodes beliefs about game, beliefs about others beliefs, beliefs about others beliefs about others beliefs...
- Types randomly drawn at start of game by nature.

Bayesian Game:

- Set of players $\mathcal{I} = \{1, 2, \ldots N\}$
- Types: $\theta_i \in \Theta_i$, $i \in \mathcal{I}$.
- Belief about others' types $p_i : \Theta_i \to \Delta(\prod_{j \neq i} \Theta_j)$.
- Strategies $s_i : \Theta_i \to A_i$
- Payoffs $u_i(a, \theta)$

Bayes Nash Equilibrium

Definition

The profile s^* is a Bayes-Nash equilibrium if for all i and $\theta_i \in \Theta_i$

$$E_{\theta_{-i}}(u_i(s^*(\theta), \theta)) \geq E_{\theta_{-i}}(u_i(a_i, s^*_{-i}(\theta_{-i}), \theta))$$

where the expectation is taken using probability $p_i(\theta_i)$

Standard to assume common prior

- Vector of types drawn from $p \in \Delta(\prod \Theta_i)$
- $p_i(\theta_i)$ is $p(\cdot|\theta_i)$

Incomplete Information

Types:

- Types + distribution over types capture all uncertainty.
- Define prob distribution over it, can now evaluate payoffs.
- Does this make sense?
  - Adding more structure to the game (types)
  - Maybe players don't know type structure
  - Not a problem (way beyond this course)

Common prior

- Standard assumption
- Does it make sense?

Examples - Cournot

- Demand is $a - q_1 - q_2$
- Firm 2 has marginal cost c
- Firm 1's costs are unknown to firm 2
  - With probability θ, c_h.
  - With probability $1 - \theta$ c_L.

Example - First Price auction

- N players
- Highest bidder wins, pays their bid
- Types are value for object, drawn iid from $U[0, 1]$

Example - Email Game

	A	B
A	$2, 2$	$1, -3$
B	$-3, 1$	$0, 0$

	A	B
A	$0, 0$	$1, -3$
B	$-3, 1$	$2, 2$

With probability $p < 1/2$, the right game describes payoffs.

- Only player 1 knows the game: $(A, A), A$ is unique BNE.
- Both know the game: 2 pure strategy eqm.

Note: A is the unique BR in both games if the probability the other player plays A is above $1/4$.

Example - Email Game

Now suppose there is close to common knowledge of the game. Specifically, before the game starts

- Player 1 learns the game.
- Player 1 sends a message to player 2 telling them the game if right game.
- Player 2 sends a message to player 1 telling them they got the message.
- Player 1 sends a message to player 2 telling them the previous message
- ...

Suppose messages go through with probability $1 - \epsilon$.

Example - Email Game

- Types: let Q_i be the number of messages player i has sent.
- Common prior:

$$Pr((Q_1, Q_2) = (q_1, q_2)) = \begin{cases} 1 - p & \text{if } q_1 = q_2 = 0; \\ p\epsilon(1-\epsilon)^{q_1+q_2-1} & \text{if } q_1 \geq 1 \text{ and } q_2 = q_1 - 1 \text{ or } q_2 = q_1; \\ 0 & \text{otherwise.} \end{cases}$$

- Posteriors:

$$Pr(Q_1 = Q_2 | Q_1) = \frac{\epsilon(1-\epsilon)}{\epsilon + (1-\epsilon)\epsilon} = 1 - Pr(Q_1 = Q_2 + 1) < 1/2$$

and

$$Pr(Q_1 = Q_2 | Q_2) = \frac{\epsilon}{\epsilon + (1-\epsilon)\epsilon} > 1/2$$

Player 1 and 2 always think it's more likely that their message didn't go through then that their message did and then the other player's message failed.

Example - Email Game

Claim: In equilibrium all types play A.

Proof:

- Base case: If $Q_1 = 0$, player 1 knows the game is the left game
- A is dominant strategy, so player 1 plays A.
- If $Q_2 = 0$, 2 believes $Q_1 = 0$ w. prob$> 1/2$, so 1 plays A w/ prob at least $1/2$, so A is BR
- Inductive hypothesis: Suppose $Q_2 = N - 1$ play A, w.t.s. $Q_1 = N$ plays A
- $Q_2 = N - 1$ with prob$> 1/2$, so A is best response for $Q_1 = N$.
- Inductive hypothesis: $Q_1 = N$ plays A, w.t.s $Q_2 = N$ plays A
- $Q_1 = N$ with prob$> 1/2$, so A is best response for $Q_2 = N$.

Example - Email Game

	A	B
A	$2, 2$	$1, -3$
B	$-3, 1$	$0, 0$

	A	B
A	$0, 0$	$1, -3$
B	$-3, 1$	$2, 2$

Types play A, no matter how close to common knowledge

- If game was known to be right game, 2 other eqm.
- Higher order uncertainty "selected" less risky eqm.
- Does this intuition generalize? (Global games)

Example - Signaling

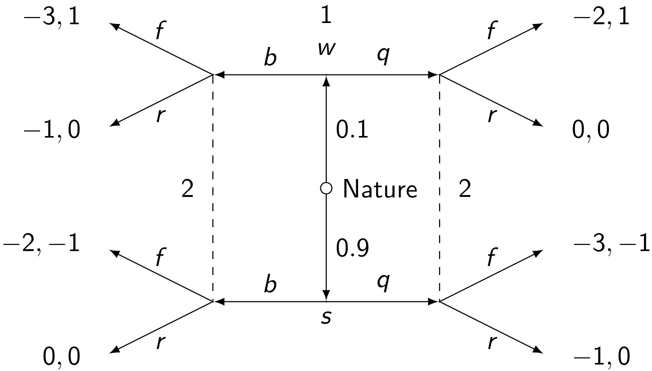

Dynamic games of incomplete information

Uh oh

Only subgame is entire game!

- Subgame perfection has no bite
- Dynamic features still seem important
- Can we come up with a "better" concept for games with imperfect info.

Want new solution concept.

Dynamic games

What do we want a solution concept to satisfy?

- Play should be "optimal" at each info set
- Subgame perfection
- Consistent behavior that respects dynamic properties of game (e.g. no incredible threats)

Imperfect Info

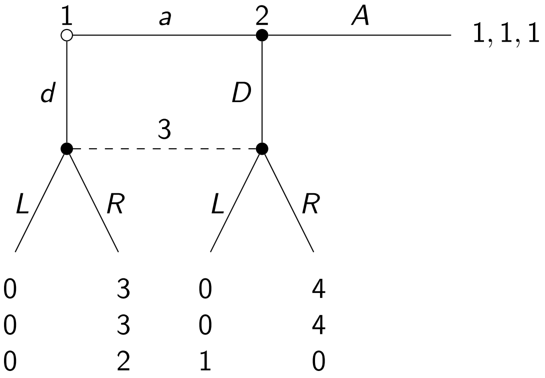

What are subgame perfect eqm?

- *dAR*
- *aA*, *L* with prob$\geq 3/4$

Sequential Rationality

First problem: Given an info set, how to evaluate continuation payoffs

Definition (Sequential Rationality)

A behavior strategy profile b is sequentially rational if at each information set there exists beliefs $\mu^h \in \Delta(h)$ such that if i moves at h

$$\sum_{x \in h} \mu^h(x) u_i(b|x) \geq \sum_{x \in h} \mu^h(x) u_i(b_i', b_{-i}|x)$$

where $u_i(b|x)$ is the utility player i receives have reached node x and future play follows behavior strategy b.

Imperfect Info

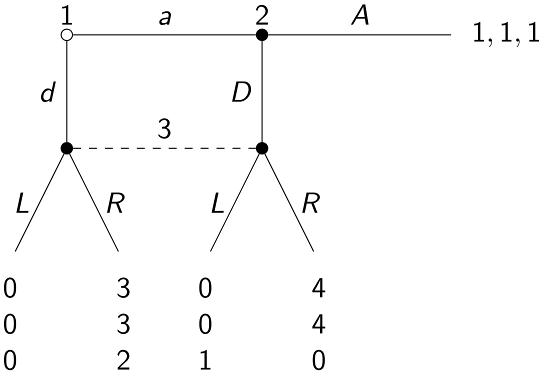

- *dAR* is not sequentially rational.
- *aAL* is
- So is *aDR*

Weak Perfect Bayesian Equilibrium

Stuff we want

- Seems like beliefs should at least behave like probabilities!
- One shot deviation principle
  - If a strategy profile is sequentially rational, then there are no one shot deviations.
  - But we don't yet have the converse
- Profiles should be subgame perfect.

Definition

A weak perfect bayesian equilbrium is a pair of strategies and beliefs (b, μ) such that b is sequentially rational given μ and μ is defined from b using Bayes' rule when possible, i.e.

$$\mu^h(x) = \frac{Pr^b(x)}{\sum_{y \in h} Pr^b(y)} \text{ whenever } x \in h, \sum_{y \in h} Pr^b(y) > 0.$$

PBE

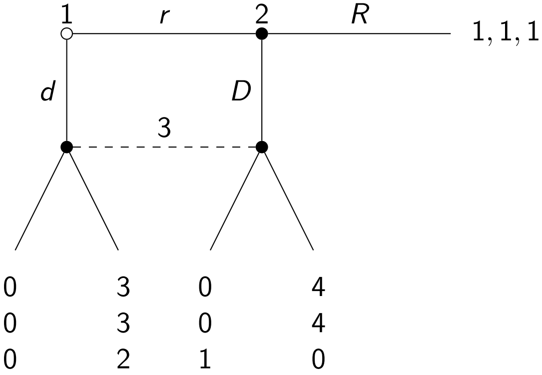

aA, *L* with prob$\geq 3/4$ are only weak PBE strategy profiles.

PBE

Does this do what we want?

Weak PBE is Nash eqm + sequential rationality

- Stronger than Nash: Restricts off-path behavior
- Players can "trick" themselves off path
- Does not imply subgame perfection.

PBE

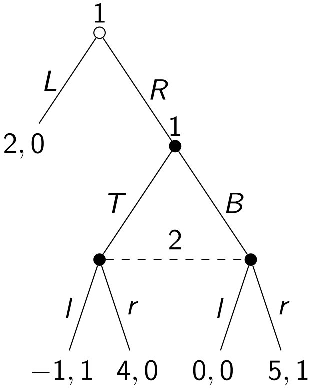

PBE

How to deal w/ probability 0 events?

- Require some sort of consistency
- Immediately after probability 0 event, beliefs can be arbitrary
- Subsequent beliefs formed using Bayes rule

Definition

A almost perfect bayesian equilibrium is a pair of beliefs (b, μ) such that

- b is sequentially rational given μ
- μ is defined using Bayes' rule when possible
- If information set h is reached with positive probability from information set h' under (b, μ) then beliefs at h are consistent with Bayes' rule applied to the beliefs at h'.

(almost because Perfect Bayesian Equilibrium generally assumes something additional, but no clear consensus on what that is)

PBE

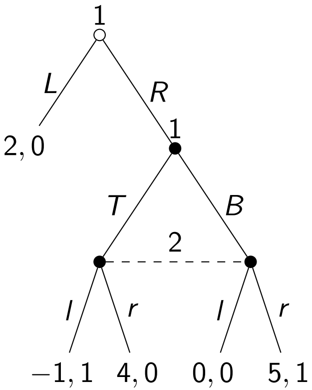

PBE

Off-path beliefs

- Still have some weird things (signaling what you don't know)

What should a player think if they see off path behavior?

- Failure of common knowledge?
  - Other player is playing a different eqm
  - Other player is not strategically sophisticated
- Other player made a mistake

We know how to think about mistakes!

Sequential Equilibrium

We can model mistakes with trembles!

Definition

A sequential equilibrium (μ, b) if there exists a sequence of completely mixed behavior strategy profiles $b_k \to b$ such that

- The associated beliefs (formed using bayes rule) $\mu_k \to \mu$.
- b is sequentially rational given μ

Note that we only require play to be optimal in the limit.

Sequential Equilibrium Example

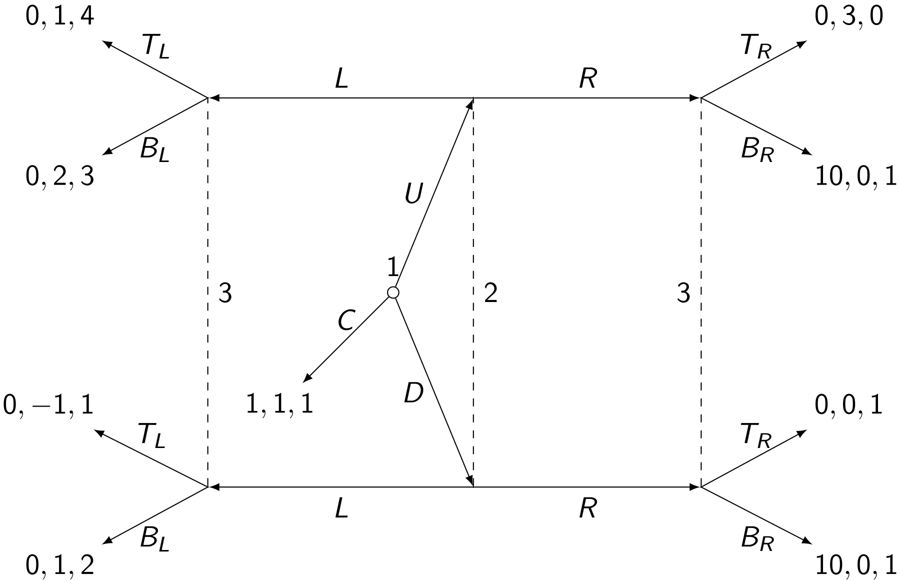

Sequential Equilibrium

- Exist in finite games
- Are aPBEs
  - Are subgame perfect
  - Satisfy the one shot deviation principle

In practice often use aPBE + get rid of implausible eqm

- Conceptually same thing in important class of incomplete info games (sender-receiver games)
- Not clear how to define sequences of beliefs, convergence, etc. outside of finite games
- (one less annoying thing to check)

Examples

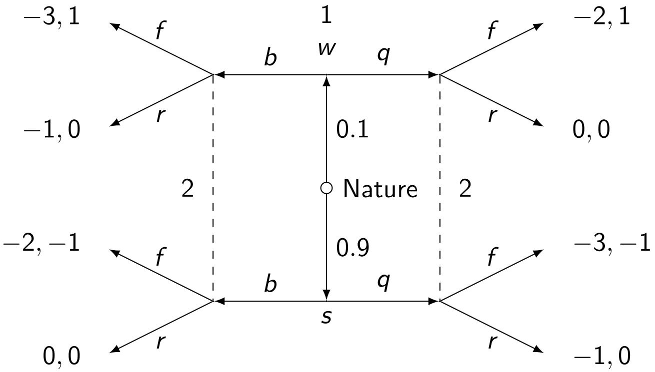

Spence Signaling

Worker is trying to signal ability to employer

- Worker type, $\theta \in \{1, 2\}$.
- $Pr(\theta = 2) = p > 0$
- Worker knows type, employer doesn't.
- 2 identical firms compete ala Bertrand for worker by choosing wage, receive payoff $\theta - w$.
- Before getting job, choose level of education $e \in \mathbb{R}_+$, cost $c(e) = e/\theta$.
- Timing:
  1. Worker sees private type
  2. Worker publicly chooses e
  3. Firms simultaneously make wage offers
  4. Worker chooses which offer to accept.

Signaling

Solution concept PBE

- Given beliefs $\mu(\theta|e)$, in last stage firm's offer $w = E(\theta|e)$, where the expectation is taken using μ.
- Worker problem becomes

$$\max_e E(\theta|e) - e/\theta$$

- Need to find behavior strategy $e(\theta)$ and eqm beliefs μ

Focus on two types of pure strategy eqm:

- Pooling eqm: $e(2) = e(1)$.
- Separating eqm $e(2) \neq e(1)$.

Pooling

Is it possible for no information to be transmitted?

- Fix strategy $e^* = e(1) = e(2)$.
- In equilibrium,

$$\mu(\theta|e^*) = 2p + (1-p) = 1 + p$$

- What about off path?
- Just want to show a pooling eqm exists, choose "worst" off path beliefs

$$\mu(1|e) = 1 \, \forall e \neq e^*.$$

- Finally, e^* is an eqm iff

$$(1+p) - e^*/\theta \geq \max_e 1 - e/\theta$$

- So any $e^* \in [0, p]$ is consistent with a pooling eqm.

Separating

Now suppose eqm $e(1) \neq e(2)$

- Now in eqm
$$\mu(\theta | e(\theta) = 1$$
- Similarly, $\mu(1|e) = 1$ for all other e.
- Firm doesn't want to deviate to off-path thing if
$$\theta - 1 \geq e(\theta)/\theta$$
- Now make sure a type doesn't want to pretend to be the other type
$$\theta - e(\theta)/\theta \geq \theta' - e(\theta')/\theta$$

Combining these:

$$\begin{aligned} e(1) &= 0 \\ 2 &\geq e(2) \\ 2 - e(2)/2 &\geq 1 \\ 1 &\geq 2 - e(2) \end{aligned}$$

So any strategies s.t. $e(1) = 0$ and $1 \leq e(2) \leq 2$ are consistent with an eqm

Pooling

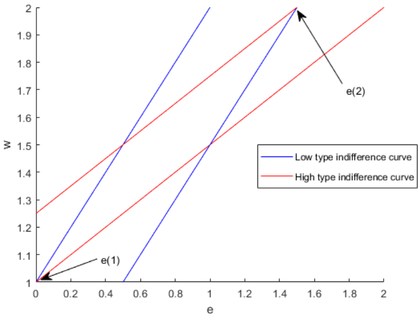

2
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1
w
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2
e
e(1)
e(2)
Low type indifference curve
High type indifference curve

Application: Reputation

Suppose a long-lived incumbent is playing predation against a sequence of T short-lived entrants.

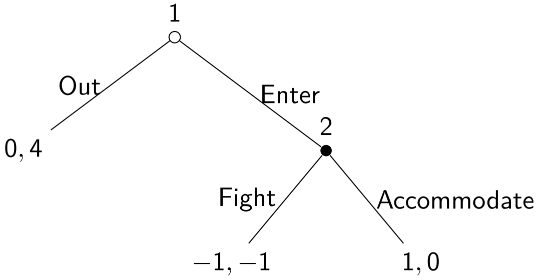

Entrant doesn't enter if probability of fight is greater then $1/2$.

No matter how large T is, unique SPE is (enter, accommodate) in every stage.

Reputation

What if, with small probability, incumbent is type that always plays fight. Let p_t be probability of commitment type at time t.

- In final period, if $p_T > 1/2$, entrant plays out.
- In previous period, let π_{T-1} be the probability strategic type fights.
- If $p_{T-1} \geq 1/2$, done, entrant never enters
- O.w. accommodate iff

$$0 \geq -1 + V_T.$$

Belief after seeing fight is

$$Pr(tough|fight) = \frac{p_{T-1}}{p_{T-1} + (1 - p_{T-1})\pi_{T-1}}.$$

If $\pi_{T-1} = 1$, short-lived guy never enters, if $\pi_{T-1} = 0$, incumbent deviates to fight after enter.

Reputation

So we need mixing. To mix $V_T = 1$, so

$$1/2 = \frac{p_{T-1}}{p_{T-1} + (1 - p_{T-1})\pi_{T-1}}.$$

so $\pi_{T-1} = p_{T-1}/(1 - p_{T-1})$.

Does the entrant enter? Probability of fight is

$$p_{T-1} + (1 - p_{T-1})\pi_{T-1} = 2p_{T-1}$$

So never enter if $p_{T-1} > 1/4$

Reputation

We can keep working backwards

In period $T - t$, eqm is such that

- Entrant plays Out, incumbent randomizes if $1/2^{t+1} < p < 1/2^t$.
- Entrant enters, incumbent randomizes if $p < 1/2^{t+1}$
- Entrant enters, incumbent fights if $1/2^t < p$.

So if the game is sufficiently long, eqm features:

- The entrant gets the market early in the game.
- Eventually entrant enters, incumbent either reveals its type, or continues to pool to deter future entry.
- Incumbent payoff is bounded away from static nash in all eqm

Reputation

Now let's think about the infinite horizon game.

- In any period where the short lived player enters, probability of fight must be less than $1/2$.
- So, given reputation p, whenever the entrant enters, if the incumbent fights

$$Pr(commitment|fight) = p \cdot Pr(fight|commitment)/Pr(fight) \geq 2p$$

- Every time the entrant enters, the incumbent reputation at least doubles if they fight and goes to 0 o.w.
- This must hold in any equilibrium.

In equilibrium, incumbent makes at least what they make from mimicking commitment type.

Use the previous observation to bound the payoff from playing fight forever.

- We can bound the number of periods where the short-lived player enters

$$
\begin{aligned}
2^k p_0 &= 1/2 \\
k &= -\log_2(p_0) - 1
\end{aligned}
$$

- So in infinite horizon game, eqm payoff bounded below by

$$
(1-\delta) \sum_{t=0}^{\lfloor -\log_2(p_0)-1 \rfloor} -1\delta^t + 4\delta^{\lfloor -\log_2(p_0)-1 \rfloor}.
$$

- In patient limit as $\delta \to 1$, firm payoff in all eqm goes to 4.

Similar logic works in many types of games.

In eqm, value of mimicking commitment type disciplines payoffs, because strategic type can always pretend to be them.

Social Choice Theory

How do societies consisting of many individuals make choices?

- Societies consist of many people with many different preferences.
- How can we aggregate those preferences?
- Are there rules that are "good" for combining them?
- How should we expect societies to make choices?

Social Choice Theory

Formally:

- Set of X of outcomes.
- $\mathcal{R}$, set of preferences over X.
- Vector of preferences: $R = (R_1, \ldots, R_n) \subseteq \mathcal{R}^n$.
  - P denotes strict part of preferences
- A social welfare function $f : \mathcal{R}^n \to \mathcal{R}$

Arrow's Theorem

Some natural properties for f to satisfy:

- Unanimity: For any R, $\{x, y\}$,

  $$xP_iy \, \forall i \in N \Rightarrow xPy$$

  where P is the strict part of $f(R)$.
- Independence of Irrelevant Alternatives: For any x, y and preferences R,R' , if $xR_iy \Leftrightarrow xR'_iy$ then $xf(R)y \Leftrightarrow xf(R')y$.
- Note: we are also implicitly assuming $f(R)$ is a preference relation and any R is possible.

Examples

Assume $|X| = M < \infty$.

Borda Count:

- Each agent assigns M points to most preferred alternative, $M - 1$ to second...
- $xf(R)y$ iff x has more points in total than y. Clearly a preference relation, satisfies unanimity, verify for yourselves it does not satisfy IIA.

Majority rule

- $xf(R)y$ iff the majority of preferences list x over y. Clearly satisfies unanimity + IIA
- But $f(R)$ is not a preference relation.

Arrow's Theorem

Theorem

If $|X| \geq 3$ and f satisfies Unanimity and Independence of Irrelevant alternative then f is dictatorial, i.e. there exists $d \in N$ such that for all R

$$xP_d y \Rightarrow xPy$$

for all x, y where P is the strict part of $f(R)$.

The only social choice functions that satisfies our "nice" axioms are those that just let one of the players make all choices.

Gibbard-Satterthwaite

We aren't thinking about incentives here, but our "nice" axioms seem to mimic things that incentives should imply.

- Restrict attention to the set of strict preferences, $\mathcal{P}$.
- Define a social choice function $\xi : \mathcal{P} \to X$.
- A social choice function is manipulable if there exists an agent i, a vector P and a P_i' s.t.

$$\xi(P_i', P_{-i}) P_i \xi(P)$$

  otherwise we say it is strategyproof
- Without loss, let's assume $\xi(\mathcal{P}) = X$

Strategyproofness requires reporting your true preferences to be a dominant strategy.

Gibbard-Satterthwaite

Theorem

Suppose $\xi : \mathcal{P} \to X$ *is strategyproof and* $|X| \geq 3$. *Then* ξ *is dictatorial.*

We can use Arrow's theorem to show this.

Proof

Lemma

If there exists a P and a P_i' s.t. $\xi(P) = x$, $\xi(P_i', P_{-i}) = y$ then if xP_iy and $xP_i'y$, ξ is manipulable.

Lemma

Suppose there exists a $B \subseteq X$ s.t. for all i and all $b \in B$, $a \in X \setminus B$, bP_ia then $\xi(P) \in B$.

Proof

With these, we can construct a social welfare function.

$xf(P)y$ iff x is chosen when we move x, y to the top of everyone's preferences without changing anything else, i.e. define P_i' as

1. $xP_i'y \Leftrightarrow xP_iy$.
2. If $a, b \notin \{x, y\}$ then $aP_i'b \Leftrightarrow aP_ib$.
3. If $a \notin \{x, y\}$ and $b \in \{x, y\}$ $bP_i'a$.

Then $xf(P)y \Leftrightarrow x = \xi(P')$.

We need to verify transitivity, IIA and unanimity (completeness and asymmetry are immediate).

Gibbard Satterthwaite

What does this mean:

- Revelation principle: Very hard to implement anything in dominant strategies.
- Nice games to aggregate preferences (e.g. things where the outcome of everyone playing a dominant strategy is Pareto efficient) don't exist in general.
- Not hopeless, we could relax some assumptions. For instance, restrict preferences.

Psychological Games

What if payoffs depend on factors beyond strategic choices

- Seems natural to allow for some sort of reciprocity
  - I want to treat people nice if I expect them to treat me nice. I want to treat people badly if I expect them to treat me badly.
- How can we formalize this?
- Payoffs now depend on both what others do, and I what I expect others to do.

Geanakoplos, Pearce, Stacchetti 1989

Psychological Games

Let S be the set of strategies in a normal form game. We can make payoffs depend on both the realized strategy profile and beliefs:

$$v_i : M_i \times S \to \mathbb{R}$$

where M_i is the set of beliefs over strategies of other players, M is the set of belief vectors (beliefs for each player) over play in the game.

Psychological Games

Definition

Given a normal form game, a Psychological Nash Equilbrium is a pair $(\mu^*, \sigma^*) \in M \times S$ where

1. μ_i puts probability 1 on σ^* for all i.
2. for all i and all $\sigma_i \in S_i$

$$v_i(\mu_i^*, (\sigma_i, \sigma_{-i}^*)) \leq v_i(\mu_i^*, \sigma^*)$$

- Changing strategies does not change beliefs. We've separated out expectations about others' play from actual play.
- Can be augmented to allow for mixed strategies. Fixed point arguments can be used to establish existence.

Example - Fairness and Reciprocity

We have a lot more flexibility now in terms of payoffs.

Rabin (1993) uses the following functional from to study many of the games we've already looked at:

$$v_i(b, \sigma) = u_i(\sigma) + \alpha_i(\mu)u_j(\sigma)$$

where $\alpha(b)$ captures how much to care about others' payoffs.

Battle of the Sexes

Remember Battle of the Sexes, our game where players were trying to coordinate on an activity

	B	S
B	$2, 1$	$0, 0$
S	$0, 0$	$1, 2$

Now let's think about this game where

$$\alpha_i(\mu_1, \mu_2) = \begin{cases} .8 \text{ if } \mu_1 = \mu_2 \\ -.8 \text{ o.w.} \end{cases}$$

Example

So now, if I think that B, B is being played, payoffs essentially become

	B	S
B	$2.8, 2.6$	$0, 0$
S	$0, 0$	$2.6, 2.8$

so the Nash equilbirium in this game effectively becomes more appealing, players really want to cooperate if they expect the other to be cooperating.

Example

On the other hand, if I think B, S is being played,

	B	S
B	$1.2, -.6$	$0, 0$
S	$0, 0$	$-.6, 1.2$

So B, S is an equilbrium of the psychological game. If each thinks the other is trying to selfishly force them to go to preferred activity, then both prefer to miscoordinate.

On the other hand S, B is not an equilibrium.