Econ 200C

Joel Sobel

April 30, 2010

INTRODUCTION

Joel Sobel, Room 311, Available after class and by appointment.

IMPORTANT INFORMATION

- I intend to post these slides (for as long as I create them). I urge you to use this as a reason to take fewer notes (instead, think and listen in class).
- No class on Wednesday of this week.
- Class on Friday (8:30-10:30) April 2 instead.

September 8, 2009

- I apologize for not completing 205.
- I learned that there are things worse than teaching math camp.
- Thank you for your support.
- I have accumulated a lot of new scars and some permanent hardware in my right leg.
- Otherwise, I am essentially recovered.

OUTLINE

- Subject: Non cooperative game theory and "information economics"
- Unlike 200B: Cannot summarize class with one model and two basic results. Instead, techniques, definitions, and a few basic models.
- Paternalism: Work as many problems as you can, carefully and seriously.
- Requirements: Homework, Midterm, Final

TOPICS

First half: game theory foundations: equilibrium concepts, extensive-form games, repeated games, incomplete information.
Second half: signaling, screening, agency, mechanism design

Warning

I have never taught class with slides before.

Advantages:

- It will look like I am prepared.
- You'll have access to the slides.
- I won't need to stand for two hours.
- Nageeb will know what to do in case of accident.

Disadvantages:

- I might not sustain the energy to prepare notes.
- Typos.
- Tendency to go too fast.

Game: Definition

- Set of players, $i = 1, \ldots, I$. (Need not be finite.)
- Strategy set $S = S_1 \times \cdots \times S_I$.
- Payoff functions $u_i : S \rightarrow \mathbb{R}$.

Comments on Strategy Set

Notation

- $s_i \in S_i$ is a strategy for Player i.
- $s = (s_1, \ldots, s_I)$ is a strategy profile (sometimes: outcome).
- $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_I)$

Comments:

- Concept of strategy is extremely subtle and general.
- We'll often need to extent to "mixed strategies" or "correlated strategies" replacing S by $\Delta(S_1) \times \cdots \times \Delta(S_I)$ in the first case and by $\Delta(S)$ is the second case. (Details later.)

Comments on Utility Functions

- Note: Player i's utility depends on s_{-i}. This is the essence of game theory.
- Foundations of utility functions from 200A.
- Especially: Extension to mixed or correlated strategies by linearity (von Neumann-Morgenstern)

Game Theory for Social Science

- Start with strategic problem.
- Translate problem into game.
- Solve game.
- Translate solution back to motivating problem.

First step is hard (finding interesting problem). Second step requires making choices. It is something of an art. Third step is both technical and philosophical. What is a solution? Fourth step is easy (although there is a tendency to exaggerate). The first part of the class is primarily about the third step. The second part gives examples of the first step.

GAMES AND ECONOMIC BEHAVIOR

	c_1
UP	$1, 1$
DOWN	$1, 0$

Note: Matrix representation of two-player games. One row for each strategy of Player 1 (ROW). One column for each strategy of Player 2 (COLUMN). In each cell of payoff matrix the first number is ROW's payoff, the second is COLUMN's payoff.
This game is a decision problem. It is trivial in the sense that R gets the same payoff no matter what he does. But C cares about R's choice of action.

Matching Pennies

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

Players simultaneously decide whether to reveal the Head (H) or Tail (T) of a penny. If the pennies match, Row wins; if not, Column wins.

Pure conflict ("zero sum"). Intuition (symmetry) suggests that each player should expect a zero payoff, but there no zeros in the payoff matrix.

Coordination

	Call	Wait
Call	$0,0$	$2,2$
Wait	$2,2$	$0,0$

Players receive the same payoff (no conflict of interest).
Friends talk on phone. The line goes dead. They want to resume the conversation. If both try to call, they get busy signals. If neither call, then they cannot talk.
Two sensible predictions.

Dominance

	Left	Right
Up	$10, 1$	$2, 0$
Down	$5, 2$	$1, 100$

No matter what Player 2 does, Player 1 players his top strategy. If Column understands this, she'll play left even though that gives her no chance to get her favorite payoff.

Prisoner's Dilemma

	Left	Right
Up	$4, 4$	$0, 5$
Down	$5, 0$	$1, 1$

Here both players have dominant strategies. If both play them, result is inefficient.

Battle of the Sexes

	Left	Right
Up	$4, 1$	$0, 0$
Down	$0, 0$	$1, 4$

Players want to coordinate, but have different views of the best place to coordinate.

Matching Pennies with a Cheater

	HH	HT	TH	TT
H	$1, -1$	$1, -1$	$-1, 1$	$-1, 1$
T	$-1, 1$	$1, -1$	$-1, 1$	$1, -1$

Like matching pennies, but now Column can observe Row's move before deciding what to do.
Note difference in strategies.

DEFINITIONS

- s_i is (strictly) dominated (by s_i^*) if for all s_{-i} $u_i(s_i^*, s_{-i}) > u_i(s)$.
- s_i is weakly dominated (by s_i^*) if for all s_{-i} $u_i(s_i^*, s_{-i}) \geq u_i(s)$ with strict inequality for at least one s_{-i}.
- s_i^* is a (weakly) dominant strategy if it (weakly) dominates all other $s_i \in S_i$.
- Security level: $\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$.

SIGNIFICANCE

"Rational" players avoid (strictly) dominated strategies.
Security Level provides a lower bound to a player's payoffs.
How do these concepts apply to examples?

Mixed Strategies

Given any game, one can extend the game by replacing S_i by $\Delta(S_i)$ and by extending u_i to the new, larger domain.
An element $\sigma_i \in \Delta(S_i)$ is called a mixed strategy.
$\sum_{s_i \in S_i} \sigma_i(s_i) = 1$, $\sigma_i(s_i) \geq 0$ for all $s_i \in S_i$.
$\sigma_i(s_i)$ is interpreted as the probability that Player i selects (pure) strategy s_i under σ_i.
$U_i(\sigma) \equiv \sum_{s \in S} \left(\Pi_{i=1}^{I} \sigma_i(s_i) \right) u_i(s)$
Expected utility justifies this.

Why Mix?

1. Raises security level (to 0 in matching pennies).
2. Increases the power of dominance:

	Left	Right
Up	$10, 0$	$0, 10$
Middle	$0, 10$	$10, 0$
Down	$a, 1$	$a, 1$

Down is dominated by a pure strategy when $a < 0$. It is dominated by a mixed strategy when $a < 5$. It is not dominated (by a pure strategy) when $a \geq 0$. It is dominated (by a mixed strategy) when $a < 5$. It is not dominated (by a pure or mixed strategy) when $a \geq 5$.)

Unless otherwise noted from now on, "strategy" means mixed strategy.

Interpretation of Mixed Strategy

- Conscious Randomization.
- Opponent uncertain ("purification").

Correlated Strategies

$\Pi_{i-1}^{I} \Delta(S_i)$ is smaller than $\Delta(S)$
(The first looks at only "product measures" or independent distributions.)
For example:

	Left	Right
Up	.5	0
Down	0	.5

(numbers are probabilities, not payoffs)

Implications of Correlation

- None for security level (somewhat subtle, uses linearity of payoffs)
- Makes it harder for a strategy to be "dominated" (when $I > 2$)
- Useful idea in repeated games, mechanism design, equilibrium theory.

Iterated Dominance

Let $S_i^0 \equiv S_i$. Form Game G^k: same players, same utilities, strategy set S^k. Let $S_i^k = \{s_i \in S_i : s_i \text{ is undominated in } G^{k-1}\}$.
Properties:

1. S^k contained in S^{k-1}.
2. S^k nonempty.
3. There is k^* such that $S^k = S^{k^*}$ for all $k > k^*$.

Call limit set $S^* = \cap_k S^k$.

Rationale

"Rational" agents avoid dominated strategies. "Rational" agents expect their "rational" opponents to do so. Hence the process iterates.

Suggests a solution method for some games (like the dominance and prisoner's dilemma examples).

Technicalities

- Order of deletion does not matter.
- Order would matter for iterative weak dominance.
- Reason for difference: a strictly dominated strategy stays strictly dominated. A weakly dominated strategy may no longer be weakly dominated after other strategies are deleted.

Examples

- Linear Cournot is Dominance Solvable. (Since strategy set is continuous, need infinite number of rounds).
- Guessing Games.

A Different Approach

Rational agent best responds to beliefs about what other players are doing.
Beliefs of player i: probability distribution over $\Pi_{j \neq i} \Delta(S_j)$.
Rational players assume that opponents only play best response to beliefs.
Generates iterative process.

Rationalizability

Let $R_i^0 \equiv R_i$. Form Game G^k: same players, same utilities, strategy set R^k. Let
$R_i^k = \{s_i \in R_i : s_i \text{ is a best response to some beliefs in } G^{k-i}\}$.
Properties:

1. R^k contained in R^{k-1}.
2. R^k nonempty.
3. There is k^* such that $R^k = R^{k^*}$ for all $k > k^*$.

Call limit set $R^* = \cap_k R^k$.
Strategies in R^* are rationalizable.

Rationalizable and IDDS

1. $R^* \subset S^*$.
2. $R^* = S^*$ for two player games.
3. $R^* \neq S^*$ in general.
4. $S^* = T^*$ in general, where T^* is the set obtained by iterated deletion of strategies are not best responses to correlated strategies. (Like rationalizability, but it is easier to survive because you can best respond to more things.)

Reasons

1. Fewer rationalizable strategies because strictly dominated strategy cannot be a best response. (This part is easy.)
2. Second part requires the separating hyperplane theorem. It says that any strategy that is not dominated must be a best response (to some mixed strategy).
3. The separating hyperplane argument generalizes to more than 2 player games except that the proper generalization is that any strategy that is not dominated must be a best response to some correlated strategy. (When there are two players, there is one opponent so correlation does not matter.)

Equilibrium

A strategy profile s^* is a (Nash) equilibrium if for each i, s_i^* is a best response to s_{-i}^*.
The definition has two parts: players best respond to beliefs; beliefs are accurate.

Properties of Equilibrium

1. Nash equilibrium does not exist (in pure strategies).
2. Nash equilibrium exists (in mixed strategies); S finite.
3. Nash equilibrium is rationalizable and survives IDDS.
4. Hence dominance solvable implies unique Nash.
5. Games have multiple Nash equilibria.
6. Nash equilibria need not be efficient.
7. Nash equilibrium payoffs are at least equal to security level.

Existence Proof

Let $\phi : \Pi_i \Delta(S_i) \rightarrow \Pi_i \Delta(S_i)$ by

$$\phi_i(\sigma) = \{\sigma_i^* : \sigma_i^* \text{ is a best response to } \sigma_{-i}\}.$$

Check that ϕ satisfies assumptions of Kakutani's Fixed Point Theorem and that a fixed point is a Nash equilibrium.

Lessons of Existence Proof

- Depends on Expected Utility (convex valued).
- Can be extended to existence of pure strategy equilibria in games with continuous stractegy spaces if u_i strictly quasi-concave.

Two-Player Zero-Sum Games

$I = 2$, $u_1(s) + u_2(s) = 0$ for all $s \in S$.
(Matching pennies.) n general,

$$\min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma) \geq \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma).$$

- What is the importance of mixed strategies?
- Interpretation of LHS?
- Why is inequality true?
- Inequality is an equation for zero-sum games.
- Combined with existence of Nash equilibrium, this is the "Fundamental Theorem of Two-Player Zero-Sum Games."

Refinements

- Nash removes strictly dominated strategies (and iterates).
- What about weakly dominated strategies?

	Left	Right
Up	4, 4	0, 0
Down	0, 0	0, 0

- Can we "trick" Nash into removing weakly dominated strategies?

Perturbed Game

Given game G, a number $\varepsilon \in (0, 1)$ and a totally mixed strategy $\tilde{\sigma}$ ($\tilde{\sigma}_i(s_i) > 0$ for all $s_i \in S_i$), form a perturbed game that has the same player set, the same strategy set, and payoffs:

$$\tilde{u}_i(s) = \varepsilon u_i(s_i, \tilde{\sigma}_{-i}) + (1 - \varepsilon) u_i(s).$$

Weakly dominated strategy in G becomes strictly dominated in perturbed game.

Trembling Hand Perfection

A strategy profile σ^* is trembling-hand perfect if there exists a totally mixed strategy $\tilde{\sigma}$ and a sequence of mixed strategies $\{\sigma(n)\}$, such that $\{\sigma(n)\}$ is a Nash equilibrium of the $\tilde{\sigma}$, $\varepsilon(n)$ perturbation; $\lim_{n\to\infty} \varepsilon(n) = 0$; and $\lim_{n\to\infty}\{\sigma(n)\} = \sigma^*$.

- Perfect Equilibria Exist (finite games, in mixed strategies)
- Players do not use weakly dominated strategies in perfect equilibrium.
- There are equivalent definitions.
- Perfect equilibria are discontinuous in data of problem.
- Perfect equilibria may not iteratively delete weakly dominated strategies.
- There is a stronger definition ("totally perfect") that replaces "there exists" by "for all" in above definition.

Observations

- Definition and existence problems with perfect equilibrium in continuous games.
- Existence problems with totally perfect equilibria.
- Examples.

Extensive Forms

Games trees are a way to depict vividly the dynamic structure of games.
Looking at the trees for many games is insightful.
But
Describing trees formally is tedious.
There is reason to believe that all information needed to think about a game can be captured in "normal" or strategic form.

The tedious part

Extensive Game

1. Game is: finite set of nodes: $\mathcal{X}$; actions $\mathcal{A}$; players $\{1, \dots . I\}$.
2. $p : \mathcal{X} \to \{\mathcal{X} \cup \emptyset\}$. $p(x)$ is the immediate precedessor of x; nonempty for all but one node. x_0, the initial node is the exception. Successors of x, $s(x) \equiv p^{-1}(x)$. Terminal nodes, $T \subset \mathcal{X}$ have no successors.
3. $\alpha : \mathcal{X} \setminus \{x_0\} \to \mathcal{A}$ is the action that leads to x.
4. Choices at x: $c(x) \equiv \{a \in \mathcal{A} : a = \alpha(x') \text{ for some } x' \in s(x)\}$. It takes different actions to get to different nodes: $x', x'' \in s(x)$, $x' \neq x''$, implies $\alpha(x') \neq \alpha(x'')$.
5. A partition $\mathcal{H}$ of $\mathcal{X}$ and a function $H : \mathcal{X} \to \mathcal{H}$. Associated with each node x is an information set $H(x) \in \mathcal{H}$. $c(x) = c(x')$ if $H(x) = H(x')$ (otherwise one could distinguish elements in an information set). $C(H) \equiv \{a \in \mathcal{A} : a \in c(x) \text{ for some } x \in H\}$ are the choices at H.
6. $\iota : \mathcal{H} \to \{0, 1, \dots, I\}$ identifies a player to each information set. $\mathcal{H}_i = \{H \in \mathcal{H} : i = \iota(H)\}$.

Features

1. Strategy: complete contingent plan (what to do at every information set).
2. Mixed strategy: probability distribution over pure strategies.
3. Outcome: probability distribution over terminal nodes.
4. Behavior strategy: independent probability distribution at each information set.
5. Can turn extensive game into strategic game. (And conversely)
6. Nature is a trick that can be used to incorporate incomplete information.
7. Perfect information game: all information sets are singletons.
8. Perfect Recall: formalization of "no player forgets."

Two Classical Results

1. Perfect information games have pure strategy equilibria.
2. In games with perfect recall, any outcome induced by mixed strategies can be induced by behavior strategies.

Subgames

Given an extensive form game, a subgame is a subset of the nodes for which the restrictions of predecessor function is still well defined; there is a unique initial node; and in which if x is in the subgame so are all nodes in $H(x)$.
(Start with an arbitrary singleton information set and include all successors. If you can do this and you hit all nodes in all successor information sets, the you have a subgame.)
Subform: Same as subgame except that you can start with a non-trivial information set. (If so, then add initial move of nature.)

Solution Concepts

1. Backward Induction.
2. Subgame perfection.
3. Bayesian Perfect Equilibrium.

Subgame Perfection

A strategy profile σ of an extensive-form game G is a subgame perfect Nash equilibrium if it induces a NE in every subgame of G. The key concept is "induces": A strategy for G is a complete contingent plan. It specifies an action at each information set. Consequently, it generates (induces) a strategy for any subgame of G. This is a tangible reason for the complicated definition of a strategy. It guarantees that once you specify a strategy for a big game you automatically specify a strategy for all subgames.
Properties:

1. SGPNE exist in finite extensive games.
2. A SGPNE induces a SGPNE in every subgame.
3. If G has no proper subgames, then any NE of G is a SGPNE.
4. The backward-induction equilibrium of a perfect information game is subgame perfect.
5. A NE σ induces a NE on any subgame that is reached with positive probability by σ.

Continued

The last three observations are trivial. The first follows either from a version of "backward induction" or the observation that a trembling-hand equilibrium of the normal form derived from an extensive form is a SGPNE (this follows from the last observation).

Sequential Rationality

There are not enough subgames for SGPNE to destroy all "incredible" threats. Generalize notion of subgame to subform: A subset of nodes (and associated actions, information sets, and payoffs) is a subform if it contains precisely the successors of a given information set (not necessarily a singleton) and either all or none of the nodes in any information set of the original game. A subform becomes a game if one assigns beliefs (a probability distribution) on the nodes in the initial information set. (That is, first nature moves and determines a probability distribution over the initial information set, and the game follows the structure of the original game.)
Formally, a system of beliefs μ for an extensive game G is a specification of a probability $\mu(x) \geq 0$ for each decision node in G such that $\sum_{x \in H} \mu(x) = 1$ for every information set H.

Belief-based solution concepts

A belief based equilibrium (not a technical term) for G is a pair (σ, μ) (strategy profile, system of beliefs) such that (a) μ satisfies a consistency condition; (b) σ induces an equilibrium on every subform of G (where subforms become subgames using μ).

(b) is called sequential rationality. (a) comes in many forms. A minimal restriction is that μ is consistent with σ whenever possible. ($\mu(x)$ is computed by Bayes's Rule: the prob of reaching x from σ divided by the probability of reaching the information set containing x – this makes sense when the denominator is positive). This leads to Weak Perfect Bayesian Equilibrium.

A stronger consistency restriction is to require μ to be derived from a sequence of completely mixed (putting positive probability on all pure strategies) strategies converging to σ. This leads to sequential equilibrium.

Connections

1. Trembling hand perfect equilibria must be sequential; sequential must be weak perfect bayesian, subgame perfect, and Nash.
2. Weak perfect Bayesian and Subgame Perfect must be Nash.
3. Weak perfect Bayesian may fail to be Subgame Perfect.
4. Subgame Perfect may fail to be weak perfect Bayesian.

Reasons

1. The totally mixed "approximating" strategies generate consistent beliefs.
2. Sequential rationality includes Nash (because the entire game is a subgame/subform).
3. Example: Player I moves "in" or "out". If "out" both players get 10. If "in" they play matching pennies. Represent matching pennies as: first I moves, then II moves (not knowing I's move). A WPBE is, for I: "out" and "H" if "in" fand "H" for Player II along with beliefs that Player I always moves T. This is NE, II best responds to beliefs at his information set. Player *I* best responds to strategies at both information sets. But the players don't play an equilibrium in matching pennies.

More

The reason that you can have a WPBE that is not NE is that II's beliefs are not consistent with I's strategies off the equilibrium path (II thinks that player I always moves T, this is permissible according to the definition of consistency used in WPBE because player II's information set isn't reached when Player I moves out). But it is not consistent with what Player I does in the subgame.

4. Player I goes out or left or right. If he goes out, the game ends; payoffs are $(1, 10)$. Player II does not know whether Player II goes left or right. In either case, Player II can go left or right. If Player II goes left then payoffs are $(0, 0)$ if Player I goes left or right. Payoffs are $(5, 5)$ if Player II goes right (and Player I doesn't go out). There is a subgame perfect equilibrium in which Player 1 goes out and Player II goes left, but in any WPBE Player II must go right.

Repeated Games

Given a game G and a discount factor $\delta \in (0,1)$ and a positive integer (possibly infinite), one can form a new game, the n-time repetition of G with discount factor δ. Loosely, players play G, observe the outcome, play G again, and so on. Payoffs are he discounted sum of the payoffs from each "stage."

Formally, the repeated game has the same player set. The strategy sets are complicated. Let $\mathcal{H}_0 = \emptyset$ and $\mathcal{H}_t = S^t$ for $t > 0$. The strategy set for player i consists of a sequence of functions $f_i^t : \mathcal{H}_{t-1} \to S_i$. The interpretation is that $f_i^t(h_{t-1})$ is the action that i takes in the tth repetition following history h_{t-1}. The payoff function for the repeated game given strategy profile f is:

$$U_i(f) = (1 - \delta^n)^{-1}(1 - \delta) \sum_{t=1}^{n} \delta^{t-1} u_i(f^t(h_{t-1}))$$

where $f^t = (f_1^t, \ldots, f_I^t)$ and h_t is the sequence of histories induced by f $(h_t = h_{t-1} \times f^t(h_{t-1})))$

Some literature on "time average" criterion: roughly $\delta = 1$ case (which causes technical problems in infinite horizon games).

Comments

1. Strategy Space is HUGE.
2. $(1-\delta^n)^{-1}(1-\delta)\sum_{t=1}^{n}\delta^{t-1} = 1.$
3. Hence $(1-\delta^n)^{-1}(1-\delta)$ is a normalization factor (designed to make feasible payoff set independent of n).

Definitions

1. Feasible Set $F(G)$: convex hull of $\{u(s) : s \in S\}$.
2. Minmax: $\underline{v}_i^i = \min_{\sigma_{-i} \in \Pi_{j \neq i} \Delta(S_j)} \max_{s_i \in S_i} u_i(s_i, \sigma_{-i})$. (Let $\underline{v}_j^i$ be the payoff that j obtains for strategies that lead to the payoff $\underline{v}_i^i$ for i.)
3. Individually Rational Set of Payoffs: $IR(G) = \{v : v_i \geq \underline{v}_i^i \text{ for all } i\}$.

The set of strictly individually rational payoffs is:

$$SIR(G) = \{v : v_i > \underline{v}_i^i \text{ for all } i\}.$$

The feasible set consists of all possible payoffs that one could obtain from correlated strategies in the given game.
$\underline{v}_i^i$ is a lower bound to the payoff that player i can obtain in a Nash equilibrium. (The definition allows i to best respond to the strategy of the other players, but for the other players to pick the strategy that makes i worse off.)

Questions to test understanding

1. How would $\underline{v}_i^i$ change if i could select $\sigma_i \in \Delta(S_i)$ (instead of $s_i \in S_i$)?
2. How would $\underline{v}_i^i$ change if $-i$ could select $\sigma_{-i} \in \Delta(S_{-i})$ (instead of $\sigma_{-i} \in \Pi_{j \neq i} \Delta(S_j)$?
3. What is the relationship between $\underline{v}_i^i$ and security level?
4. What is the relationship between $\underline{v}_i^i$ and $\underline{v}_i^j$?

Maxmin versus Minmax

Consider

$$\min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma).$$

This is the "min max".

In this quantity Player i can pick σ_i to BR to σ_{-i} (the inside maximization), but then the other agents can pick their strategies to hurt Player i (the outside minimization). This is the quantity that is relevant for repeated games.

This quantity is greater than or equal to

$$\max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma),$$

which is the security level. This is the "max min" that is relevant for zero-sum games. In the second quantity, the other players can pick a strategy that hurts Player i as much as possible (this is the inside minimization), but Player i can select the strategy that is the best possible under these circumstance. Note that when you solve the second problem and get σ^*, there is no guarantee that σ_i is a best response to σ_{-i}

Inequality

.

$$\min_{\sigma_{-i}} \max_{\sigma_i} u_i(\sigma) \geq \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma)$$

The reason is that for any σ'_{-i} and σ_i,

$$u_i(\sigma_i, \sigma'_{-i}) \geq \min_{\sigma_{-i}} u_i(\sigma)$$

and so

$$\max_{\sigma_i} u_i(\sigma_i, \sigma'_{-i}) \geq \max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma)$$

and since the left-hand side is greater than or equal to the right hand side for all σ'_{-i} it is true when you minimize with respect to σ'_{-i}, which gives the desired result.

Largest Possible Equilibrium Payoff Set

The set of equilibrium payoffs must be contained in the set of feasible and individually rational payoffs. This is true for G and for any repeated version of G (with proper normalization).

Folk Theorem

Let G be given. Let $V = F(G) \cap SIR(G)$. If G satisfies a regularity condition (sufficient: $F(G)$ is I dimensional), then for each $v \in V$ and $\varepsilon > 0$ there is a $\delta_0 > 0$ such that the infinitely repeated version of G with discount factor $\delta > \delta_0$ has a subgame perfect equilibrium payoff v^* such that $|v - v^*| < \varepsilon$.
Informally: in an infinitely repeated game between patient players any individually rational and feasible payoff can be a subgame perfect equilibrium payoff.

Assumptions

- Patience
- Solution concept
- Regularity

Easy Folk Theorem, I: Nash Equilibrium

Normal phase: play to get target. Punishment phase: max min deviator.
Problem: Punishment may not be "credible" (not an equilibrium in a subgame)

Easy Folk Theorem, II: Nash Threats

Let $V' = \{v \in F(G)$: for each i there exists a NE of G with payoffs $\tilde{v}^i$, such that $v_i > \tilde{v}^i_i\}$.
V' is smaller than V (it may be empty).

Proof

Normal phase: play to get target.
Punishment phase: play worst NE for deviator.

Optimal Penal Codes

Let $\tilde{v}^i$ be the lowest subgame perfect equilibrium payoff to player i in the infinitely repeated discounted version of G. You can check to see whether something is a subgame perfect equilibrium payoff by checking whether it is attractive to deviate if punishment is to worse equilibrium.

This requires the existence of worst equilibrium for each player. It does not depend on the discount factor being close to one, although in practice it is hard to compute worst payoff.

Variations

- Finite Games when G has multiple equilibria.
- Incomplete information.
- Incomplete Monitoring
- Infinite Horizon Games that Aren't Repeated Games
- Limited Interaction

Key Insight

Folk Theorems work because one can have more than one continuation payoff. This enables the future payoff to depend on today's actions. Hence punishment is possible.

Sketch of Proof

Need to think about three types of payoff:

1. The target payoff v^*.
2. The minmax payoffs $\underline{v}^i$ (one payoff vector associated with each player).
3. Conditional reward payoffs, $v'(i)$ defined to be

$$v'(i) = (\tilde{v}_1 + \nu, \ldots, \tilde{v}_{i-1} + \nu, \tilde{v}_i, \tilde{v}_{i+1} + \nu, \ldots, \tilde{v}_I + \nu)$$

   where $\tilde{v}$ such that $\tilde{v}_i \in (\underline{v}_i, v_i)$, $\nu > 0$, $v'(i)$ feasible for all i.

- Phase I: Equilibrium Path: play to get target payoffs v^*.
- Phase II: Punishment Path: If there is a deviation by more than one player, ignore it. If there is a deviation by Player i, minmax i for N periods.
- If anyone deviates during punishment, punish the deviator.
- Phase III: If not, play to get payoffs $v'(i)$.

Technicality

If there is a "public randomization device," then players can correlate strategies that v^* and $v'(i)$ and $\underline{v}^i$. Otherwise, they can cycle through pure strategies in a way that approximates v^*. Doing so makes the proof messier and a bit trickier, but is not interesting. I'll ignore the problem.

Phase I

The strategies generate the desired payoff (essentially by definition). The problem is to confirm that there are no incentives to deviate.

In Phase I playing according to equilibrium yields v_i^*. A single deviation yields (potentially) a one shot payoff of M, the largest payoff in G, followed by N periods of $\underline{v}_i^i$, followed by v_i forever. Hence we need:

$$v_i^* \geq (1-\delta)M + \delta(1-\delta^N)\underline{v}_i^i + \delta^{N+1}v_i'.$$

As long as $v_i^* \geq v_i$, this inequality will hold for N large and δ close to one.

Phase II

Must check that neither i nor $j \neq i$ want to deviate. For i the comparison is (when k periods of punishment remain):

$$(1-\delta^{k+1})\underline{v}_i^i + \delta^{k+1} v_i'(i) \geq (1-\delta)\underline{v}_i^i + \delta(1-\delta^N)\underline{v}_i^i + \delta^{N+1} v_i'(i)$$

For j the comparison is:

$$(1-\delta^{k+1})\underline{v}_j^i + \delta^{k+1} v_i'(j) \geq (1-\delta)M + \delta(1-\delta^N)\underline{v}_j^j + \delta^{N+1} v_j'(j)$$

The first inequality guarantees that a deviator does not gain from deviating from her punishment. The intuition is that the punishment phase is finite. A deviator would prefer to end the punishment quickly rather than take an immediate reward and re-start the punishment.

The second inequality guarantees that a punisher is willing to participate in the punishment. This is the critical point at which subgame perfection comes into play. The potential problem is that in order to punish i, another player must settle for a payoff that is lower than his individually rational payoff. This could only happen if latter he is rewarded for punishing. That is why we need a phase three payoff that is relatively high for the punishers.

Phase II Algebra

The verification is straightforward. In the case of a deviation by i, there is a potential one-period gain, but at the expense of (a) longer punishment and (b) postponing the final period.
In the case of a deviation by j, there is a potential one period gain and a potential gain because punishing may lead to a payoff worse than j's minmax. These gains are dominated for patient players because the final phase gives a uniformly higher payoff to j if she doesn't deviate.

Phase III

The comparison from Phase I establishes that deviations are not attractive here.

Implicit Assumption

I assumed that if there exists a profitable deviation, then there exists a profitable deviation involving just one change. In discounted games it is not hard to prove this "one-shot deviation principle," but it does require an argument (and the property does not hold for undiscounted games).

Regularity Condition

Where did we need a regularity condition?
In order to reward punishers you need to have the possibility of phase three payoffs that reward a punisher (but not a rewarder).

Why the epsilon in the statement of the proof?