Economics 703: Advanced Microeconomics
Peter Cramton

Matthew Chesnes

Updated: December 12, 2006

1 Lecture 1: August 31, 2006

1.1 Introduction

- **Definition:** Game theory is the study of mathematical models of conflict and cooperation between rational (expected utility maximizers) and intelligent (everyone understand the game and knows other players are rational) decision makers.

- Player: $i \in N = \{1, \dots, n\}$. Strategy: $s_i \in S_i$. Strategy vector (profile):

  $$s = (s_1, \dots, s_n) \in S = S_1 x \dots x S_n.$$

  Payoff function: $u_i(s) : S \mapsto \Re$.

- Normal form (strategic form) game:

  $$\Gamma = \{S_1, \dots, S_n; u_1, \dots, u_n\}.$$

- Extensive form of a game includes the following information: who plays when, actions available at each node, information available to each player when they have the move, and resulting payoffs.

- **Definition:** Strategy - a complete plan of action or what to do in EVERY contingency.

- **Definition:** Information set - a collection of decision nodes for a player such that the player has the move at all nodes in the set and the player does not know which node has actually been reached.

- **Definition:** Perfect Information - each information set is a singleton. Finite games of perfect information can be solved by backward induction or elimination of weakly dominated strategies.

2 Lecture 2: September 5, 2006

- **Definition:** Imperfect information: at some point in the tree, some player is not sure of the complete history of the game so far.
- Can solve prisoners dilemma (PD) problem by IESDS. Yields a "Dominant Strategy Equilibria."
- **Definition:** Strictly dominated strategy: x strictly dominates y if the player gets a higher payoff from x compared to y, regardless of what the other players do.
- **Definition:** Weakly dominated strategy: x weakly dominates y if the player gets at least as a great a payoff from playing x compared to y, regardless of what the other players do.
- Most games do not have a dominate strategy equilibrium, so we use our workhorse, the Nash Equilibrium.
- **Definition:** For an n-person game in normal form, a strategy profile, $s^* \in S$ is a Nash Equilibrium (NE) in pure strategies if for all i,

  $$u_i(s^*) \geq u_i(s_i; s^*_{-i}) \ \forall\ s_i \in S_i.$$

  So we have a set of mutual best responses (BR).
- In matching pennies, the strategies of the players are diametrically opposed. It's a game a pure conflict. There is no equilibrium in pure strategies, so there must be a NE in mixed strategies.
- **Definition:** Zero-sum game: the sum of the payoffs is zero, regardless of outcome. This is a game of pure conflict which is rare in economics. Subset of constant-sum games.
- **Definition:** Mixed strategy: a randomization over pure strategies.
- **Theorem 1:** The existence of Nash Equilibria in finite games. Consider a normal form game,

  $$\Gamma = (S^1, \ldots, S^n; u^1, \ldots, u^n),$$

  with pure strategy profile, $s = \{s^1, \ldots, s^n\} \in S = S^1 x \cdots x S^n$, and mixed strategy profile:

  $$\sigma = \{\sigma^1, \ldots, \sigma^n\} \in \Delta(S^1) x \cdots x \Delta(S^n),$$

  where $\sigma^i : S^i \mapsto [0,1]$ and $\sigma^i(s^i) = \Pr(\text{player } i \text{ plays pure strategy } s^i)$. Define player i's expected payoff as:

  $$v^i(\sigma) = \sum_{s \in S} u^i(s) \prod_{j=1}^{n} \sigma^j(s^j).$$

  Note this assumes that players randomize INDEPENDENTLY! They each take other player's strategies as given. Define a Nash Equilibrium in mixed strategies (of which pure strategy NE are a subset) as an n-tuple, $\sigma = (\sigma^1, \ldots, \sigma^n)$ such that for every player i,

$$v^i(\sigma) \geq v^i(\hat{\sigma}^i; \sigma^{-i}) \; \forall \; \hat{\sigma}^i \in \Delta(S^i).$$

So, again, we have mutual best responses. Note the size (or contents) of the strategy sets need not be the same across players. Then, we have the theorem which says that every finite game has a Nash equilibrium in mixed strategies.

- **Theorem 2:** Consider an n-person game, $\Gamma = \{D^1, \ldots, D^n; v^1, \ldots, v^n\}$ where D^i are the pure strategies available to player i and v^i is player i's payoff function. If each D^i is a compact convex subset of a Euclidean space, and each v^i is continuous and quasiconcave in d^i, then Γ has a NE in pure strategies.

- More next time on definitions and proofs of these 2 theorems.

3 Lecture 3: September 7, 2006

3.1 Existence of Nash Equilibria

- Before proving the theorems from last time, we need a couple definitions and another theorem.

- **Definition:** A function, v^i, is Quasiconcave if:

$$v^i(\alpha d^i + (1-\alpha)\hat{d}^i, d^{-i}) \geq min\{v^i(d^i, d^{-i}), v^i(\hat{d}^i, d^{-i})\}.$$

  See G-3.1.

- **Definition:** A correspondence, $\phi : T \mapsto V$, is upper hemicontinuous at every point $x \in T$, if given two sequences, $x_r \to x$ and $y_r \to y$ such that $y_r \in \phi(x_r)$ for all r, we have,

$$y \in \phi(x).$$

  Note BRs need not be unique so that's why we're using correspondences instead of functions.

- The payoff function that we are maximizing must actually have a maximum, so we assume the payoff function is closed and bounded (or compact) and continuous.

- **Theorem: Kakutani**. If T is a nonempty compact convex subset of a Euclidean space, and ϕ is an upper hemicontinuous nonempty convex valued correpondence from T to T, then ϕ has a fixed point (FP), ie there is an x in the domain of ϕ such that $x \in \phi(x)$. See G-3.2 and G-3.3 for a proof by picture of Kakutani.

- Note that usually we have an ODD number of equilibria. Even numbers of equilibria are not generic.

- **Proof of Theorem 2: existence of pure strategy NE.** For each player i, define a best response correpondence, $\phi^i : D \mapsto D^i$ as:

$$\phi^i(d) = \{d^i \in D^i | v^i(d^i, d^{-1}) \geq v^i(d^{i'}, d^{-i}) \; \forall \; d^{i'} \in D^i\}.$$

  Note that ϕ^i is nonempty since D^i is compact and v^i is continuous. ϕ^i is convex-valued since v^i is quasiconcave in d^i. And ϕ^i is upper hemicontinuous since v^i is continuous. Let ϕ be the collection of best response correspondences for all players. Since each D^i is compact, D is also a compact subset of Euclidean space. Since each ϕ^i is an upper hemicontinuous nonempty convex-valued correpondence, so is ϕ. Thus by kakutani, ϕ has a fixed point. But a fixed point of ϕ, ie a set of mutual best responses, is just a NE of Γ. Q.E.D.

- **Proof of Theorem 1: existence of a mixed strategy NE in finite games.** Let $D^i = \Delta(S^i)$, ie the complete set of all mixed strategies. Each D^i is a compact convex subset of Euclidean space and each v^i is continuous and quasiconcave in d^i (actually linear). Thus, by theorem 2, ∃ NE in which each player chooses a "pure" strategy from $\Delta(S^i)$. Q.E.D.

- Note that you can get existence without quasiconcavity (see Dasgupta and Maskin, RES 1986). We cannot use Kakutani in this case. See the Glicksberg theorem.

3.2 Correlated Equilibria

- The idea of playing a mixed strategy NE in a one shot game (like in BoS) is not intuitively pleasing because given that your opponent is playing the MSNE, you are indifferent between your two strategies. However the MSNE relies on the fact that you both mix. This is more reasonable for repeated games. Short of actually coordinating explicitly on an equilibria, one might consider a situation where the two players can talk to a mediator who can tell the two players how to play the game.

- In BoS, the two NE yield payoffs $(2, 1)$ and $(1, 2)$. Suppose a mediator flips a coin and if it comes up heads, he tells the players to play (U, L), and if it comes up tails they should play (D, R), yielding each of the two payoffs. On average, each player obtains a payoff of $\frac{3}{2}$.

- Compare this with the unique MSNE of the game, where player one randomizes on $(2/3, 1/3)$ and two randomizes on $(1/3, 2/3)$. This yields an average payoff of $\frac{2}{3}$ for each player.

- So the correlated equilibia allows for a much higher average (or expected) payoff to the players. The reason is that we have eliminated the possibility of falling into the "bad outcomes" of (U, R) and (D, L) which yield payoffs of zero to both players.

4 Lecture 4: September 12, 2006

4.1 A Note on Quasiconcavity

- See G-4.1. This function is not concave, but it is Q-concave. Once it starts going up (even if it levels off for a while), it has to keep going up. This means there is a top of the mountain so to speak (or many if the function plateaus).
- In our proof of theorem 2, we had that the BR correspondence, ϕ^i, was convex valued since the payoff function, v^i, was Q-concave in d^i. See G-4.2 (from new slides) for a picture of this.

4.2 Perfect Recall

- A game has perfect recall if each player knows whatever he knew previously, including his previous actions. Ie, players do not forget anything.
- In Bridge, 2 players play against 2 other players. To model the game as a two player game would mean that the teammates were acting together (ie, they could see each other's cards). Since they don't, the game only has perfect recall if you model it as a 4 player game with the payoffs being identical to each set of teammates.
- All games of perfect information have perfect recall (but not vice versa). Some games of imperfect information have perfect recall.
- **Definition:** A behavioral strategy specifies a probability distribution over feasible actions at each information set. This is more useful in extensive games.

4.3 Correlated Equilibria Example

- See last lecture for definitions. Consider the game in G-4.3. There are two NE in pure strategies of (D, L) and (U, R) yielding payoffs of $(7, 2)$ and $(2, 7)$ respectively.
- There is also a MSNE of $(2/3U, 2/3L)$ yielding a payoff of $(14/3, 14/3)$.
- The players could do better though. Suppose 3 balls labeled A, B, and C are placed in an urn and a mediator draws them out at random. If ball A is drawn, player 1 is told to play D. If ball C is drawn, player 2 is told to play R. If players are not told anything, they play U and L respectively. Thus we end up in each of the three outcomes $(2, 7)$, $(6, 6)$, and $(7, 2)$ one third of the time each. This yields an average payoff of $(5, 5)$ to each player. So two questions:
  - Why is this an equilibrium? Consider player 1 (and symmetrically for player 2). If player 1 is told to play D, he knows the A ball has been drawn which means that 2 is told nothing and will play L. Thus the outcome is $(7, 2)$ and player 1 cannot profitably deviate. If balls B or C are drawn, player 1 is told nothing

and hence plays U. Player 1 knows that A was NOT chosen and there is an equal chance that B or C were chosen. Thus, his expected payoff from playing U is $0.5*6+0.5*2=4$. By deviating to D, his expected payoff is $0.5*7+0.5*0=3.5$. So playing U is a best response. Thus both players are playing a BR so we have a correlated equilibrium.

  - Is this the absolute best the players can do in a correlated equilibrium? No, there are other probabilities (say placing more weight on the payoff $(6,6)$) which will yield an expected payoff to each player of 5.25. This can be found by solving a simple constrainted maximization problem. Maximize the expected payoff to each player by allocating weights to the different balls, subject to the constraint that each player cannot profitably deviate.

- **Definition:** For a normal form game, a correlated strategy is a probability distribution, $p(s)$, over the set of pure-strategy n-tuples, S.

- **Definition:** The correlated strategy, $p(s)$, is a correlated equilibrium of the mediated game if for each player, i, and for all s_i^* such that $p(s_i^*)>0$,

$$\sum_{s_{-i}^*\in S_{-i}} u_i(s_i^*;s_{-i}^*)p(s_{-i}^*|s_i^*) \geq \sum_{s_{-i}^*\in S_{-i}} u_i(s_i;s_{-i}^*)p(s_{-i}^*|s_i)\ \forall\ s_i\in S_i.$$

- **Theorem:** Every point in the convex hull of the NE payoffs is a correlated-equilibrium payoff. So plot all the NE (pure and mixed), connect the dots, and everything inbetween is achievable. This is done with a mutually observable randomizing device.

- **Theorem:** The correlated equilibrium payoffs are a convex polyhedron defined by *linear* inequalities, unlike the $(n-1)^{th}$ degree equations that determine NE.
Proof: The linear inequalities (those that define correlated equilibria – see above), determine a convex polyhedron in the space of correlated strategies. These in turn determine a convex polyhedron of correlated equilibrium payoffs. I'm not convinced.

5 Lecture 5: September 14, 2006

5.1 Dynamic Games

- Missed Lecture.
- **Problem:** Recall a NE involves each player acting optimally given the other players' strategies. Ie, players all play a BR. The problem is that the optimality condition is at the beginning of the game. Hence some NE of dynamic games involve non-credible threats.
- **Definition:** Consider a game, Γ, of perfect information consisting of a tree, T, linking information sets, $i \in I$ (all single nodes), and payoffs at each terminal node of T. A subtree, T_i, is a tree beginning at the information set i, and the subgame Γ_i is the subtree T_i and the payoffs at each terminal node of T_i.
- **Definition:** A NE of Γ is subgame perfect (SPNE) if is specifies NE strategies in every subgame of Γ. Ie, players act optimally at every point in the game.
- Note that subgames may NOT cut information sets.

5.2 Repeated Games

- Consider a stage game, $G = \{A_1, \ldots, A_n; u_1, \ldots, u_n\}$ and an outcome of G:

$$a = (a_1, \ldots, a_n) \in A = A_1 x \ldots x A_n.$$

- Denote the game where G is played T times, $G(T)$, a repeated game.
- Denote a^t the outcome of the t^{th} repetition of G.
- The history prior to the t^{th} repetition is denoted:

$$h^{t-1} = (a^1, \ldots, a^{t-1}) \in A^{t-1}.$$

- A strategy for player i of $G(T)$ is:

$$\sigma_i = (\sigma_i^1, \ldots, \sigma_i^t, \ldots, \sigma_i^T),$$

  where $\sigma_i^t : A^{t-1} \mapsto A_i$, ie the strategy maps the history up to $t-1$ into an action.
- Stage game payoffs are denoted:

$$\{u_i(a^1), \ldots, u_i(a^T)\}.$$

- Average payoffs for $G(T)$:

$$U_i(h^T) = \frac{1}{T}\sum_{t=1}^{T} u_i(a^t).$$

- Discounted aggregate payoff for $G(T)$:

$$U_i(h^T) = \sum_{t=1}^{T} \delta^{t-1} u_i(a^t),\ \delta \in [0,1].$$

- So players condition their behavior in period t on the history up until that period. This may allow for reputation effects or cooperation that is not possible in the one period stage game.

- **Definition:** A strategy profile is a subgame perfect Nash equilibrium (SPNE) of $G(T)$ if:

  - (1) σ is a NE of $G(T)$, and
  - (2) $\forall\ t < T$ and $\forall h^t \in A^t$, $\sigma[h^t]$ is a NE of $G(T-t)$ where $\sigma[h^t]$ is the strategy profile for the game $G(T-t)$ specified by σ following the history h^t.

- In stage games that have a unique NE, the game $G(T)$, $T < \infty$, has the same NE (ie, play the stage game NE in every period). Cooperation is not sustainable in finitely repeated games (ie, because there is a last period).

6 Lecture 6: September 19, 2006

6.1 More on Finitely Repeated Games

- Consider the game in G-6.1. The game has two pure strategy NE of $(x1, x2)$ and $(z1, z2)$. When we have multiple NE of the stage game, the finitely repeated game can have more equilibria than just playing the stage game Nash in every period.

- Consider the following. Both play y in all periods and both play z in the last period. If any deviation occurs by either player, both play x until the end of the game.

- Equilibrium path payoffs for say $G(2)$ are $4+3=7$. A deviation from y in period 1 yields $5+1=6$. Thus playing y in period 1 and z in period 2 is a NE of $G(2)$.

- Formally, the strategies are:

$$\sigma_i^1 = y_i, \text{ and } \sigma_i^2(h^1) = z_i \text{ if } h^1 = (y_1, y_2) \text{ and } \sigma_i^2(h^1) = x_i \text{ otherwise.}$$

- Due to Blackwell, we only need to consider the single BEST unilateral deviation to see if the game has a NE. We don't need to check every possible history which is nice.

6.2 Infinitely Repeated Games

- See Fudenberg and Maskin (Econometrica 1986).

- Consider the stage game, $G = \{A_1, \dots, A_n; u_1, \dots, u_n\}$ and the infinitely repeated game, $G(\infty)$.

- A strategy in $G(\infty)$ is:

$$\sigma_i = (\sigma_i^1, \dots, \sigma_i^t, \dots), \text{ where } \sigma_i^t : A^{t-1} \mapsto A_i.$$

- The action set, A_i, is the set of all mixed strategies and strategies based on public randomization devices (strong assumption).

- Define the minimax payoff as:

$$v_j^* = min_{a_{-j}}\ max_{a_j}\ u_j(a_j, a_{-j}).$$

  So v_j^* is the absolute worst that player j can do if everyone else gangs up on him: he can always guarantee himself v_j^*.

- Define the minimax strategy as:

$$M_{-j}^j \in arg\ min_{a_{-j}}\ [max_{a_j}\ u_j(a_j, a_{-j})].$$

  So if the others play M_{-j}^j, they can hold player j to no more (less?) than v_j^*.

- The minimax strategies will define boundries on what we can achieve in the (infinitely) repeated game setting.

7 Lecture 7: September 21, 2006

7.1 More on Infinitely Repeated Games

- A note on minimax strategies and payoffs: a group of players, $-j$, that play, M^j_{-j}, play the strategies that punish player j the hardest. However, given these harsh strategies, j can still acheive v^*_j.
- **Definition:** Payoffs $v_j \geq v^*_j$ are said to be individually rational (IR) for player j. In any equilibrium of $G(T)$, j's expected payoff must be at least v^*_j. So no player can ever do any worse than his minimax payoff.
- In what follows, we will always normalize $v^*_i = 0$ for all i.
- Define the set of attainable (feasible) payoffs in the entire game as:

$$U = \{v_1, \ldots, v_n)| \; \exists \; a \in A \; s.t. \; u_i(a) = v_i \; \forall \; i\}.$$

- Define V^* to be the set of individually rational and feasible payoffs achievable via correlated strategies. V^* is a convex hull of U.
- See G-7.1 for the V^* region of the PD game.
- Note the minimax payoff will be the NE payoff only if the stage game 1) has a unique dominant strategy NE or 2) the game is zero-sum.
- Since, (if $a^t = a \; \forall \; t$):

$$U_i(h^\infty) = \sum_{t=0}^{\infty} \delta^t u_i(a^{t+1}) = \frac{1}{1-\delta} u_i(a),$$

then the average payoff each period is:

$$u_i(a) = (1-\delta) U_i(h^\infty).$$

7.2 Folk Theorem

- **Theorem:** For any $\{v_1, \ldots, v_n\} \in V^*$, if players discount the future sufficiently little (ie, are patient enough), (ie $\exists \; \delta^* \in (0,1)$ s.t. $\forall \; \delta \in (\delta^*, 1)$), then there exists a NE of $G(\infty)$ where for all players, i, the average payoff to player i is v_i.
- **Proof:** Consider a vector of actions that sustains v_i for all players (equilibrium path payoffs). If anyone deviates in any period, play the minimax strategy forever.
- **Problem:** Not Subgame Perfect ! Playing minimax forever is not optimal! Minimax strategies are not mutual best responses (except for a dominant strategy unique NE). The punishers would want to eventually deviate. Try again ...

7.3 Rubinstein Perfect Folk Theorem

- **Theorem:** For any $\{v_1, \ldots, v_n\} \in V^*$, $\exists$ a SPNE of $G(\infty)$ with NO discounting where for each player i, player i's expected payoff is v_i. So players do NOT punish forever, they eventually revert. The punishment must fit the crime which is a nice property.
- **Proof Sketch:** Consider a vector of actions that sustains v_i for all players (equilibrium path payoffs). If anyone deviates in any period, play the minimax strategy ONLY long enough to wipe out the gains for the deviator. Then revert. The punishers themselves are threatened with punishment if they deviate from the punishment.
- **Problem:** No discounting! Not realistic. Try again ...

7.4 Friedman Folk Theorem with Nash Threats

- **Theorem:** If $\{v_1, \ldots, v_n\} \in V^*$ pareto dominates the payoffs of a NE of G, then if the players discount the future sufficiently little, $\exists$ a SPNE of $G(\infty)$ where for each player i, player i's expected payoff is v_i.
- **Proof:** Same as the Folk Theorem proof except punishment is now a NE so no reason to deviate. This is nice for games like the PD. We can sustain cooperation with minimal levels of patience.
- **Problem:** No real problems here except we revert to NE forever. Last one ...

7.5 Fudenberg-Maskin Folk Theorem

- **Theorem:** Consider the case where $n = 2$. For any $(v_1, v_2) \in V^*$, $\exists$ $\delta^* \in (0, 1)$ such that there exists a SPNE of $G(\infty)$ in which player i's average payoff is v_i when players discount the future according to $\delta \geq \delta^*$.
- **Proof:** Let M_i be player i's minimax strategy against player j. Ie M_i is the strategy that player i would play to stick it to player j in the worst way possible. So,

  $$M_i \in arg\ min_{a_i}\ [max_{a_j}\ u_j(a_i, a_j)].$$

  Consider strategies for each player where they play the strategies that yield v_i on the equilibrium path, but if a deviation has occured, then revert to minimax strategies for N periods. If a deviation occurs by a player during the punishment, start the N periods over again, forget about the original deviant, and start punishing the most recent deviator.
  Define the best possible payoff for player i as:

  $$\overline{v}_i = max_{a_1, a_2}\ u_i(a_1, a_2).$$

  So this is player i's lucky payoff. It suffices to show that $\exists$ N and $\delta^* \in (0, 1)$ s.t. for each player i and $\delta \geq \delta^*$:

– (1) $\frac{v_i}{1-\delta} > \overline{v}_i + \frac{\delta}{1-\delta} p_i$.

– (2) $\frac{p_i}{1-\delta} = \frac{1-\delta^N}{1-\delta} u_i(M_1, M_2) + \frac{\delta^N}{1-\delta} v_i > 0$.

So if p_i is the "continuation value" of the deviator after deviating for one period (ie, it's the punishment phase payoff plus the payoff once the players revert to the equilibrium path), then condition (1) says that playing the equilibrium path is optimal (ie deviating is not optimal). What about condition (2)? Note that:

$$\sum_{i=0}^{N} \delta^i = \frac{1-\delta^{N+1}}{1-\delta}.$$

We want the payoff in the punishment phase to be better than deviating. Since everyone plays minimax strategies in the punishment phase, a unilateral best response payoff will be v_i^* which we have normalized to zero. So the LHS of (2) is the PDV of the continuation value, ie getting p_i forever. What does this equal? The first term on the RHS is the payoff during the punishment phase. M_1 and M_2 are the minimax strategies. Note that $u_i(M_1, M_2)$ is NOT the minimax payoff for player i! BOTH players are playing minimax strategies (ie they are "sticking it to each other"), so $u_i(M_1, M_2)$ is worse than the minimax payoff, so:

$$u_i(M_1, M_2) < 0.$$

The second term is what the player gets after we revert to the equilibrium path (gravy train) after N periods of punishment. This whole thing has to be greater than zero which is the normalized minimax payoff. If it is, not deviating from the punishment strategies is optimal.
Note that conditions (1) and (2) can be rewritten:

– (1) $v_i > (1-\delta)\overline{v}_i + \delta p_i$.

– (2) $p_i = (1-\delta^N) u_i(M_1, M_2) + \delta^N v_i > 0$.

- **Problem:** It seems pretty good. What are the key ideas? First of all, we only punish for a finite amount of time. Ie, we don't add punishment periods to punishment periods which causes problems if players are a bit impatient (ie they don't care about a punishment that doesn't happen for a long time). The punishment strategies of M_1 and M_2 are fairly draconian and sustain cooperation for smaller δ's.

- So we didn't actually prove this theorem above, but it's clear that there will exist a δ and an N that satisfies (1) and (2). Since $u_i(M_1, M_2)$ is negative and v_i is positive, then for large enough δ, $p_i > 0$. Since $\overline{v}_i$ is very large and p_i is smaller than v_i, then also for big δ, the RHS of (1) will be larger than the LHS. Sort of Q.E.D.

8 Lecture 8: September 26, 2006

8.1 Fudenberg-Maskin n-Player Folk Theorem

- Note before we proved the F-M theorem for 2 players. Here we generalize to n players.
- **Theorem:** Assume the dimensionality of V^* is n. Ie, the feasible IR region must have a dimension equal to the number of players in the game. Then for any $v \in V^*$, $\exists\, \delta^*$ such that for $\delta \in (\delta^*, 1)$, $\exists$ a SPNE of $G(\infty)$ in which player i's average payoff is v_i when players have discount factor δ.
- Why do we need this dimensionality condition? It's because we need to have the flexibility to create strategies that punish some players (say from deviating) and reward others (say from successfully carrying out a punishment of a deviator). If the dimension of V^* is not big enough, we won't be able to punish and reward to sustain as many cooperative equilibria.
- See notes for a 3 player example where $dim(V^*) = 1$.

8.2 Applications of Dynamic Games

Cartel Maintenance - Porter (JET 1983)

- Consider a static oligopoly model with N firms, indexed by i that play a Cournot game on a homogeneous good.
- Denote $q = (q_i, q_{-i})$ and $Q = \sum_{i=1}^{N} q_i$.
- Inverse demand is equal to $\theta p(Q) = \theta(a - bQ)$ where $\theta \sim F(\theta)$ on $[0, \infty)$ and $E[\theta] = \mu$. Firms do not perfectly observe draw of θ each period.
- Linear costs: $C(q_i) = c_0 + c_1 q_i$ where $0 < c_1 < \mu a$.
- Expected profits are thus:

$$
\begin{aligned}
E[profits] = \pi_i(q) &= E[P(Q)]q_i - c_0 - c_1 q_i \\
&= \mu[a + bQ]q_i - c_0 - c_1 q_i \\
&= \mu[a + b(q_i + Q_{-i})]q_i - c_0 - c_1 q_i
\end{aligned}
$$

- So how do we write this is formal game theoretic notation? Consider the game:

$$\Gamma = \{S_1, \ldots, S_n; \pi_1, \ldots, \pi_n\},$$

where $S_i \in [0, \infty)$ for all players, i. $q^* = (q_1^*, \ldots, q_n^*)$ is a Cournot NE if:

$$\pi_i(q_i^*, q_{-i}^*) \geq \pi_i(q_i, q_{-i}^*)\ \forall\ q_i \in S_i.$$

- Each firm maximizes their expected profit by choosing q_i:

$$FOC(q_i) = 0 = \mu(a - 2bq_i - bQ_{-i}) - c_1.$$

Imposing symmetry,

$$c_1 = \mu(a - 2bq - b(N-1)q).$$

$$q_i = \frac{\mu a - c_1}{\mu b(N+1)} = \frac{A}{B(N+1)} \equiv s \;\forall\; i.$$

- So s is the Cournot NE quantity choice.

- Could the firms do better? Yes, they could cooperate, produce the monopoly output level and split the monopoly level of profits. However, choosing the monopoly quantity is NOT a NE so we need to devise a trigger strategy to sustain this outcome. Note splitting the monopoly outcome implies:

$$q_i = \frac{A}{2BN} \equiv r.$$

Clearly $r < s$.

- So, how close to the cooperateive outcome can we get in $G(\infty)$ with appropriate trigger strategies. Firms do NOT observe output choices exactly due to this θ randomness. However, even if it is a bad draw of θ that is driving down the price, firms must STILL PUNISH each other (enter a price war) for some period of time to sustain the cartel.

- Each firm in $G(\infty)$ maximizes:

$$E\{\sum_{t=0}^{\infty} \beta^t \pi_i(q^t)\},\; \beta \in (0,1).$$

Note this game has NO proper subgames because no firm ever reaches a singleton information set due to the uncertainty about demand.

- So what's our trigger strategy? A strategy is a triplet, (q, p, T) where q is the cooperative output choice, p is the trigger price, and T is the duration of the punishment phase. All firms play q in all periods if no deviation has occured (ie price has been above p), and if in any period, the price drops below p, revert to Cournot Nash for T periods. Then return to cooperation.

- So each firm has the following value function:

$$V_i(q) = \pi_i(q) + Pr(p_t \geq p)\beta V_i(q) + Pr(p_t < p)\left[\sum_{t=1}^{T-1} \beta^t \pi_i(s) + \beta^T V_i(q)\right].$$

- Solve for $V_i(q)$:

$$V_i(q) = \underbrace{\frac{\pi_i(s)}{1-\beta}}_{Cournot\ Profits} + \underbrace{\frac{\pi_i(q)-\pi_i(s)}{1-\beta+(\beta-\beta^T)Pr(p_t<p)}}_{Trigger\ Strategy\ Premium},$$

where:

$$Pr(p_t < p) = Pr(\theta_t p(Q) < p) = Pr(\theta_t < \frac{p}{p(Q)}) = F(\frac{p}{p(Q)}).$$

- So this is nice. But what about (q, p, T)? Which to choose? Note that a trigger-price strategy (q, p, T) is a NE if for all players, i:

$$V_i(q_i, q_{-i}) \geq V_i(\tilde{q}_i, q_{-i}) \ \forall \ \tilde{q}_i \geq 0.$$

- Consider the following series of proposition which may (or may not) be proved. Consider the optimal cooperative quantity, $q^*(p,T) = \{q_1^*(p,T), \ldots, q_n^*(p,T)\}$, which results in value function for player i: $V_i^*(p,T) \equiv V_i(q^*(p,T))$. Then:
  - **Proposition 1**. The Cournot Nash quantity $q_i^* = s$ for all players is a NE for all p, T.
  - **Proposition 2**. Given p and T, $q_i^* = q_j^*$ for all players i, j.
  - **Proposition 3**. For all p and T, $q_i^* \in (s/N, s]$ where $s/N < r < s$.
  - **Proposition 4**. If $F(\theta)$ is convex (say uniform), then $V_i(q)$ is concave in q_i so the FOC is sufficient for finding the optimal q^*.

More next time.

9 Lecture 9: September 28, 2006

9.1 More on Porter's Repeated Cournot with Demand Shocks

- Note in this model we have demand uncertainty (through the θ), as well as unobservability of the actual quantities choosen by the players.
- Last time we talked about 4 propositions that any equilibrium of the game would satisfy. In general, the optimal trigger price and punishment, (p, T), have to satisfy (for an interior solution):

$$\frac{\partial V_i^*(q^*(p,T))}{\partial p} = \frac{\partial V_i^*(q^*(p,T))}{\partial T} = 0.$$

- So if we can take those partials (see proposition 4 above), then we can solve for the optimal q^*, p^* and T^*. See Cramton notes for details but I think the point is that it all depends on the distribution of θ. If the variance is high, ie lots of demand shocks, we are more constrained in the set of cooperative quantities that we can sustain. Ie, for some distributions, there may NOT exist an optimal trigger price and punishment length that will sustain the joint monopoly outcome. It simply will require too many price wars or too long of a punishment and firms will find it optimal to deviate and we end up back at Cournot-Nash.
- Note as the uncertainty associated with θ goes to zero, the optimal quantity that we can sustain approaches the joint monopoly outcome.

9.2 Sequential Bargaining - Rubinstein (Econometrica 1982)

- Consider a game where two players are bargaining over a pie. Players make alternating offers with player 1 going first. The game continues forever unless a player agrees to the other's offer. Players discount the future at rate δ_1 and δ_2 respectively.
- Offers are always made in terms of player 1's payoff, so if player 2 offers x, the payoffs would be $(x, 1-x)$.
- The discount factors account for bargaining costs, the time value of money, or even the potential of "breakdown". This is the idea that the pie might vanish with some probability in the future as one of the players decides to quit or possibly attains a better offer.
- See G-9.1 for the game in extensive form.
- Note that EVERY division of the pie is a Nash equilibrium. You can always write down strategies that support any division of the pie, including the very unfair divisions like $(0, 1)$ and $(1, 0)$. Note if the length of the game is finite, it can be solved (uniquely) by backward induction.

Uniqueness

- As long as players discount future payoffs ($\delta < 1$), then the game has a UNIQUE SPNE in which trade occurs immediately. This is the efficient bargaining outcome.
- So what's the solution to the infinite period game? Suppose there exists an equilibrium, $(s, 1-s)$. Note that period 3 and period 1 are identical in terms of who has the move. Thus in period 3, when player 1 gets to accept or reject player 2's offer, player 1 will accept if:
$$\delta_1 s_2 \geq \delta_1^2 s.$$
Player 2 will offer just enough for acceptance so:
$$s_2 = \delta_1 s.$$
In period 2, player 2 will accept if:
$$1 - s_1 \geq \delta_2(1 - \underbrace{\delta_1 s}_{s_2}).$$
Player 1 will offer just enough for acceptance so:
$$s_1 = 1 - \delta_2 + \delta_1\delta_2 s. \quad (*)$$
- So what have we done? We supposed that there existed an equilibrium $(s, 1-s)$ and we have found another one: $(s_1, 1-s_1)$! Note that the $(s, 1-s)$ equilibrium supports the $(s_1, 1-s_1)$ division. We could also rearrange (*) to get:
$$s = \frac{s_1 - 1 + \delta_2}{\delta_1\delta_2}.$$
So if $(s_1, 1-s_1)$ is an equilibrium, then it support $(s, 1-s)$ as well! And we could continue this. However, if we plot (*), as in G-9.2, we see that since $\delta_i < 1$ for both players, the function has slope less than one and only intersects the 45 degree line once. Hence there is a unique fixed point. Solving (*) for that fixed point, we have:
$$s^* = 1 - \delta_2 + \delta_1\delta_2 s^*$$
$$s^* = \frac{1-\delta_2}{1-\delta_1\delta_2}.$$
- So s^* supports itself! This is our unique SPNE.

Existence

- Consider the following strategies:
$$\text{Player 1 always offers: } \; x = \frac{1-\delta_2}{1-\delta_1\delta_2}.$$

$$\text{Player 2 always offers: } y = 1 - \frac{1-\delta_1}{1-\delta_1\delta_2}.$$

Player 1 always accepts y or more.

Player 2 always accepts x or less.

Note that when 2 offers y, the payoffs are $(y, 1-y)$.

- Note also that as $\delta_i \to 0$, ie player i gets more impatient, their own share of the pie shrinks. Intuitive.
- Next time we'll show that these strategies guarantee existence.

10 Lecture 10: October 3, 2006

10.1 Proof of Existence in the Rubinstein Bargaining Model

- A nice way to show existence is to consider the necessary condition for an equilibrium. Both players must be indifferent between accepting and rejecting any current offer. If they STRICTLY prefer to accept an offer that you make, you're a moron and should lower your offer until they are just indifferent. However, in practice, it makes sense to be a bit generous to guarantee acceptance.
- Thus, player 2 should be indifferent between $1 - x$ today and $1 - y$ tomorrow.
- And player 1 should be indifferent between y today and x tomorrow. So we have:

  $$1 - x = \delta_2(1 - y),$$

  $$y = \delta_1 x.$$

  Solving these two equations and two unknowns:

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

  as stated last time. This is our SPNE.

10.2 Mechanism Design

- We now consider games of imperfect or incomplete information, Bayesian games, and the Revelation Principal.

Imperfect vs Incomplete Information

- **Definition:** A game of imperfect information involves one or more players being uncertain about the full history of the game. Eg, firms not knowing the quantity choices of other firms in the Cartel maintenance game.
- **Definition:** A game of incomplete information involves one or more players having different private information about their preferences or abilities. Eg, firms not knowing the cost levels of other firms in the Cartel maintenance game.
- Note perfect information implies complete information. Why?
- The key to analyzing games of incomplete information is to transform them into games of imperfect information by letting nature move first! Ie, we randomly select each player's payoff function.

- Example 1. See G-10.1. Here we have a game where nature moves first and selects the cost levels for two firms (high or low). Each of the four possibilities occurs with equal probability (1/4). Each player then chooses quantity simultaneously as shown in the extensive form game with information sets for each player. Profits are then realized.

- The CRITICAL ASSUMPTION of all of this is that each player in the game agrees on the probability distribution over draws from nature. Ie, beliefs must be consistent between the players.

- We can also represent this game in normal form as in G-10.2. The strategies of each player looks like (w, z) where w is the quantity choice (high or low) when costs are realized to be low and z is the quantity choice when costs are realized to be high. Thus each player has four strategies. Refering to the payoffs in G-10.1, we see, eg,

$$x = E[u_1|(q_L^1, q_H^1), (q_L^2, q_H^2)] = \frac{1}{4}[a + b + c + d].$$

  A similar calculation can be used to fill in the entire table.

- So we have taken the incomplete information in the game (not knowing the other's cost level) and transformed it into imperfect information about nature's initial draw at the beginning of the game.

- Example 2. Consider a first priced seal bid (FPSB) auction with two bidders, each with valuation,

$$v_i \sim U[0, 1],$$

  independently distributed. Nature first chooses the v_i for each player and then players choose a bid, $b_i(v_i)$. The player with the highest bid wins the item and pays his bid. We'll come back to this in a second.

Bayesian Games (Harsanyi, Management Science 1967)

- **Definition:** A Bayesian game is a normal form game of actions and payoffs but we add to it the types of each player and probability distributions over types:

$$\Gamma = \{A_1, \ldots, A_n; T_1, \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\}.$$

  A_i is the strategy set, T_i is the type space for player i so:

$$t = (t_1, \ldots, t_n) \in T = T_1 x \ldots x T_n.$$

  Beliefs are represented by a p_i such that

$$p_i(t_{-i}|t_i)$$

is player i's beliefs about the types of the other players, conditional on his own type, t_i. All the p_i's come from the SAME joint distribution. The payoff for each player,

$$u_i(a, t),$$

may depend on the actions and types of ALL the players.

- So what about these beliefs? We say that the beliefs, $(p_1, \ldots, p_n)$, are consistent if they can be derived from Bayes' rule from a common joint distribution $p(t)$ on T. So there exists $p(t)$ such that:

$$p_i(t_{-i}|t_i) = \frac{p(t_1, \ldots, t_n)}{\sum_{t_{-i} \in T_{-i}} p(t_{-i}, t_i)} \; \forall \; i, t_i.$$

  Recall the conditional is just the joint over the marginal, eg $f(y|x) = f(x, y)/f(x)$. For a two player game:

$$p_1(t_2|t_1) = \frac{p(t_1, t_2)}{\sum_{t_2 \in T_2} p(t_1, t_2)},$$

  which is the joint distribution of (t_1, t_2) divided by the marginal of t_1.

- So beliefs are consistent if nature moves first and types are determined according to $p(t)$ and each player i is only informed of t_i.

- Back to example 2. Types are valuations, $t_i \in [0, 1]$. Actions are bids, $a_i \in [0, \infty)$. Given the uniform (and independent) valuations,

$$p_1(t_2|t_1) = p_2(t_1|t_2) = 1.$$

  Payoffs:

$$u_i(a, t) = \begin{cases} t_i - a_i & if \; a_i > a_j \\ \frac{1}{2}(t_i - a_i) & if \; a_i = a_j \\ 0 & if \; a_i < a_j \end{cases}$$

- Solution to example 2. Player 1 maximizes:

$$\begin{aligned} \pi_1(a_1, a_2, v_1, v_2) &= (v_1 - a_1) * Pr(a_1 > a_2) \\ &= (v_1 - a_1) * a_1 \\ &= a_1 v_1 - a_1^2 \\ FOC &\Rightarrow v_1 = 2a_1 \\ a_1(v_1) &= \frac{1}{2} v_1 \end{aligned}$$

  Players bid half their valuation.

- **Definition:** A strategy for player i is a complete plan of action for each of player i's possible types. That is:

$$\sigma_i : T_i \mapsto A_i.$$

- **Defintion:** A strategy profile, $\sigma = (\sigma_1, \ldots, \sigma_n)$, is a Bayesian equilibrium (BE) of Γ if:

$$\sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(\sigma(t), t) \geq \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i([\sigma_{-i}(t_{-i}), a_i], t) \ \forall \ i, a_i \in A_i.$$

  Note that the p_i is the same on each side of the equation. So we are looking for no deviations in actions, beliefs stay the same no matter what.

Existence of a Bayesian Equilibrium

- Since we can turn the game into a normal form game (by removing the incomplete information and adding imperfect information), we can apply the Nash Existence theorem from before. As long as we have finite players and finite strategies, there will always be a BE.
- With consistent beliefs (Harsanyi's assumption), a BE of Γ is simply a NE of the game with imperfect information in which nature moves first.
- Any game of incomplete information with consistent beliefs can be tranformed into a standard normal form game.

11 Lecture 11: October 5, 2006

11.1 Revelation Principle

- Myerson (Econometrica 1979) and others.
- An equilibrium of a Bayesian game, Γ, can be represented by a simple equilibrium of a modified Bayesian game, Γ', as follows. Consider:

$$\Gamma = \{A_1, \ldots, A_n; T_1 \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\},$$

and the modifed game:

$$\Gamma' = \{A'_1, \ldots, A'_n; T_1 \ldots, T_n; p_1, \ldots, p_n; u'_1, \ldots, u'_n\},$$

such that:

$$A'_i = T_i \text{ and } u'_i(a', t) = u_i(\sigma(a'), t)\ \forall\ i.$$

- So each player reports (potentially dishonestly) her private type in the play of Γ'. By reporting type t_i, you get the payoff that t_i gets by playing the equilibrium strategy $\sigma_i(t_i)$ in Γ.
- Consider an example. In the FPSB auction with two bidders, the optimal strategy is to bid $b_i(v_i) = \frac{1}{2}v_i$. Instead of playing that game, suppose that the two players revealed their valuations to the auctioneer. In this game, the player that reported the highest value wins and the auctioneer tells them to pay $\frac{1}{2}$ of their reported value. The outcome of the two games is equivalent, however in the latter, you have players reporting their valuations, and as it turns out, they report their TRUE valuations in equilibrium.
- **Definition:** The Revelation Principal states that for any Bayesian equilibrium of Γ, reporting your true type is a Bayesian equilibrium of Γ' defined above.
- Note in Γ, we have:

$$\text{Types } (T) \underbrace{\Longrightarrow}_{Strategy\ \sigma} \text{Actions } (A) \underbrace{\Longrightarrow}_{Payoff\ u} \text{Outcomes } (\Re^n)$$

- and in Γ' we have:

$$\text{Types } (T) \underbrace{\Longrightarrow}_{Identity\ I} \text{Actions } (A) \underbrace{\Longrightarrow}_{Payoff\ u\ o\ \sigma} \text{Outcomes } (\Re^n)$$

- Γ' is called the "Direct Revelation Game." The identity mapping means $I(t) = t$, ie report your true type.
- A direct mechanism, Γ', in which truthful reporting is a BE is called "incentive compatible."

- So the revelation principal states that WLOG, the analysis of BE can be restricted to incentive compatible direct mechanisms (ICDM). The key to an ICDM is that what players reveal about their private types is NOT used against them. So a proportional tax is not incentive compatible because the more income a person reports, the more they are taxed.

11.2 Provision of a Public Good

- Groves (Econometrica 1973).
- Suppose three households are considering contributing money to pave their road. The cost of paving is c. Each has independent private value, $v_i \sim F_i$. No subsidies are available so the 3 households must cover the complete cost of the road if it is to be built.
- Suppose the three households simultaneously announce their willingness to pay, b_i. If $b = b_1 + b_2 + b_3 \geq c$, the road is built and each pays:

$$\frac{b_i}{b} * c.$$

This may seem fair but each has an incentive to underbid and let the others cover the cost of the road. The more honest are your neighbors, the more you should lie. In the end the road may not be built even though $v = v_1 + v_2 + v_3 \geq c$.

- Consider an alternative mechanism: the "Groves Mechanism." Each household makes a bid $b_i(v_i)$ and if $b > c$, the road is built. However, each contributes the amount that the other players' bids fall short of c. That is:

$$u_i(b_1, b_2, b_3, v_i) = \begin{cases} 0 & if\ b < c \\ v_i - (c - b_j - b_k) & if\ b \geq c\ and\ b_j + b_k < c \\ v_i & if\ b \geq c\ and\ b_j + b_k \geq c \end{cases}$$

- What's nice about this is that my bid does NOT influence how much I pay (at least on the margin). My bid only affects the decision of whether or not the road will get built. The problem with this is that the road may not be funded. We could have $b > c$ but the revenue raised:

$$(c - b_2 - b_3) + (c - b_1 - b_3) + (c - b_1 - b_2) = 3c - 2b \leq c.$$

Thus we would need a subsidy from elsewhere to cover the difference.

- There is also an incentive to collude among the bidders. If everyone overstates their value, the shortfall of everyone is small so each pays very little which leave a lot to be covered by some outside subsidy.
- So in this game, there is actually a unique DOMINANT strategy equilibrium which is bidding truthfully, $b_i(v_i) = v_i$. The reasoning is this (from David Givens):

- Case 1: $b_{-i} > c$. In this case, $u_i = v_i$, regardless of i's bid (i.e. - you pay nothing). So $b_i(v_i) = v_i$ is a BR.
- Case 2: $b_{-i} < c$. In this case, there are two possibilities.
  * The first possibility is that $v_i + b_{-i} > c$. In this case, either i bids $b_i \geq c - b_{-i}$ (the shortfall) and the road is built ($u_i = v_i - (c - b_{-i})$) or i bids less than the shortfall and the road doesn't get built ($u_i = 0$). Since, by assumption, $v_i - c + b_{-i} > 0$, bidding $b_i \geq c - b_{-i}$ dominates bidding $b_i < c - b_{-i}$. And since i's payoff only depends on his valuation and not his bid, bidding $b_i(v_i) = v_i$ is a BR.
  * The second possibility is that $v_i + b_{-i} < c$. In this case, building the road is inefficient. Suppose i bids his valuation, $b_i = v_i$. In that case, the road isn't built and $u_i = 0$. Suppose i bids below his valuation, $b_i < v_i$. In that case the road isn't built and $u_i = 0$. Now, we're at the interesting scenario. Suppose i bids above his valuation, $b_i > v_i$. One of two outcomes will result. Either $b_i < c - b_{-i}$ and the road doesn't get built, and $u_i = 0$. Or, $b_i > c - b_{-i}$, and the road does get built. But notice in the latter case, i's payoff is $u_i = v_i - (c - b_{-i}) < 0$, by assumption (first sentence of this paragraph). So, overbidding is weakly dominated by truthful bidding.

  Therefore, truthful bidding is a dominant strategy equilibrium.

11.3 Second Price Sealed Bid Auction

- n bidders, highest bidder wins, pays second highest bid.
- Under usual argument, bidding truthfully is a (weakly) dominate strategy.
- Similar to the Groves mechanism, my bid does not affect what I pay, only if I win or not.
- "My profits are 100% of the incremental value that I create by participating in the auction."

12 Lecture 12: October 10, 2006

12.1 Bilateral Trading Mechanisms

War of Attrition

- Two animals fighing over a prize (prey). Valuations are private, $v_i \sim iid\ F(\cdot)$ with density $f(\cdot)$ on $[0,1]$.
- Each unit of time means a cost of c for each.
- What is the optimal concession time, $t_i(v_i)$ for each animal?
- Suppose $t_i'(v_i) > 0$. See G-12.1. Assume $t_i(0) = 0$.
- Define the inverse function $v_i = x_i(t_i)$. So $x(t)$ is the valuation of an animal that concedes at time t.
- Since ties happen with measure zero, the payoff for each animal is:

$$u_i(v_1, v_2, t_1, t_2) = \begin{cases} v_i - ct_j & if\ t_j \le t_i \\ -ct_i & if\ t_j > t_i \end{cases}$$

- Calculus note. Recall:

$$\frac{d}{dx}\int_{u(x)}^{v(x)} f(x,t)dt = v'(x)f(x,v(x)) - u'(x)f(x,u(x)) + \int_{u(x)}^{v(x)} \frac{d}{dx}f(x,t)dt.$$

- Thus, each animal, i, maximizes their expected utility given animal j's strategy, $t_j(v_j)$. So i solves:

$$Max_{t_i}\ \left\{ \underbrace{\int_0^{x_j(t_i)} [v_i - ct_j(v_j)]f(v_j)dv_j}_{i\ wins} \underbrace{-ct_i[1 - F(x_j(t_i))]}_{i\ loses} \right\}.$$

Note that $x_j(t_i)$ is the valuation of player j evaluated at i's concession time, t_i.

- Derive:

$$\begin{aligned} 0 &= x_j'(t_i)[v_i - c\underbrace{t_j(x_j(t_i))}_{t_i}]f(x_j(t_i)) - c[1 - F(x_j(t_i))] - ct_i[-f(x_j(t_i))x_j'(t_i)] \\ &= x_j'(t_i)v_i f(x_j(t_i)) - ct_i f(x_j(t_i))x_j'(t_i) - c[1 - F(x_j(t_i))] + ct_i f(x_j(t_i))x_j'(t_i) \\ &= x_j'(t_i)v_i f(x_j(t_i)) - c[1 - F(x_j(t_i))] \end{aligned}$$

- Impose symmetry (only after you take FOC!):

$$x'(t)vf(x(t)) - c[1 - F(x(t))] = 0$$

$$x'(t) = \frac{c[1 - F(x(t))]}{vf(x(t))}.$$

- Since $v = x(t)$ and since x is the inverse of t, $x'(t) = 1/t'(v)$, we have:

$$t'(v) = \frac{vf(v)}{c[1 - F(v)]}.$$

- This last equation is a solvable differential equation since we can impose the initial condition, $t(0) = 0$. Thus,

$$t(v) = \int_0^v \frac{zf(z)}{c[1 - F(z)]} dz.$$

  Notice this is strictly increasing as we would assume.

Simultaneous Offers

- Chatterjee and Samuelson (Operations Research 1983).
- Suppose a seller and buyer are considering trading an object which is worth s to the seller and b to the buyer. These values are private to each though they both know the distributions from which they are drawn:

$$s \sim F \ on \ [\underline{s}, \overline{s}]$$

$$b \sim G \ on \ [\underline{b}, \overline{b}]$$

- Strategies for each are function, $p(s)$ and $q(b)$. Trade occurs if $q(b) \geq p(s)$ and the trade price is:

$$P = \frac{1}{2}(p(s) + q(b)).$$

  Ex-post efficiency requires trade iff $s < b$.

- Payoffs:

$$u_{seller}(p, q, s, b) = \begin{cases} P - s & if \ p \leq q \\ 0 & if \ p > q \end{cases}$$

$$u_{buyer}(p, q, s, b) = \begin{cases} b - P & if \ p \leq q \\ 0 & if \ p > q \end{cases}$$

  Note we are evaluating the change in assets as payoffs.

- Example. Suppose F and G are both $U[0, 1]$. Each player will choose their offer (p for the seller and q for the buyer) to maximize their expected payoff conditional on their private valuation and the bidding strategy of their opponent. Since we don't know this, we'll integrate over what we don't know.

- Define inverse functions $b = y(q)$ and $s = x(p)$. Eg, a seller who offers p has valuation $x(p)$.

- Seller's Problem (note $y(p)$ is the valuation of the buyer at offer p):

$$\begin{aligned} Max_p \left\{ E[u(p,q,s,b)|s,q(b)] \right\} &= Max_p \left\{ \int_{y(p)}^{1} \frac{1}{2}(p+q(b)) - s \underbrace{f(b)}_{1} db \right\} \\ FOC: 0 &= -y'(p)[(p+\underbrace{q(y(p))}_{p})/2 - s] + \int_{y(p)}^{1} \frac{1}{2} db \\ &= -y'(p)[p-s] + \frac{1}{2}[1-y(p)] \end{aligned}$$

- Buyer's Problem (note $x(q)$ is the valuation of the seller at offer q):

$$\begin{aligned} Max_q \left\{ E[u(p,q,s,b)|b,p(s)] \right\} &= Max_q \left\{ \int_{0}^{x(q)} b - \frac{1}{2}(p(s)+q) \underbrace{f(s)}_{1} ds \right\} \\ FOC: 0 &= x'(q)[b - \frac{1}{2}(\underbrace{p(x(q))}_{q} + q)] - \frac{1}{2}x(q) \\ &= x'(q)[b-q] - \frac{1}{2}x(q) \end{aligned}$$

- Consider FOCs:

$$\begin{aligned} -y'(p)[p-s] + \frac{1}{2}[1-y(p)] &= 0 \\ x'(q)[b-q] - \frac{1}{2}x(q) &= 0 \end{aligned}$$

Now substitute $s = x(p)$ and $b = y(q)$ and multiply through by 2:

$$\begin{aligned} -2y'(p)[p-x(p)] + 1 - y(p) &= 0 \qquad (a) \\ 2x'(q)[y(q)-q] - x(q) &= 0 \qquad (b) \end{aligned}$$

Solving (b) for $y(q)$ and evaluating at p:

$$\begin{aligned} y(q) &= \frac{x(q)}{2x'(q)} + q \\ y(p) &= \frac{x(p)}{2x'(p)} + p \qquad (c) \end{aligned}$$

Differentiate (c):

$$\begin{aligned} y'(p) &= 1 + \frac{1}{2}\frac{x'(p)x'(p) - x(p)x''(p)}{(x'(p))^2} \\ &= 1 + \frac{(x'(p))^2 - x(p)x''(p)}{2(x'(p))^2} \\ &= \frac{2x'(p)^2 + x'(p)^2 - x(p)x''(p)}{2(x'(p))^2} \\ &= \frac{3x'(p)^2 + -x(p)x''(p)}{2(x'(p))^2} \\ &= \frac{3}{2} - \frac{x(p)x''(p)}{2(x'(p))^2} \qquad (d) \end{aligned}$$

Now plug (c) and (d) into (a):

$$\begin{aligned} 0 &= -2[\frac{3}{2} - \frac{x(p)x''(p)}{2(x'(p))^2}][p - x(p)] + 1 - y(p) \\ 0 &= [x(p) - p][3 - \frac{x(p)x''(p)}{(x'(p))^2}] + 1 - [\frac{x(p)}{2x'(p)} + p] \\ 0 &= [x(p) - p][3 - \frac{x(p)x''(p)}{(x'(p))^2}] + 1 - p - \frac{x(p)}{2x'(p)} \end{aligned}$$

- So what do we do next? We have this differential equation (2nd order) and we lack an initial condition. So we'll guess a linear solution (with, as in all bargaining situations, buyers understating their value and sellers overstating their value). Then we'll verify that our solution is a good one.
- More next time.

13 Lecture 13: October 12, 2006

13.1 Bilateral Trading Mechanisms

More on Simultaneous Offers

- Recall we solved for the single second order differential equation last time:

$$[x(p) - p][3 - \frac{x(p)x''(p)}{(x'(p))^2}] + 1 - p - \frac{x(p)}{2x'(p)} = 0.$$

- Suppose we consider a linear solution: $x(p) = \alpha p + \beta$. Substituting,

$$[\alpha p + \beta - p][3 - 0] + 1 - p - \frac{\alpha p + \beta}{2\alpha} = 0.$$

  This is just one equation and two unknowns so there are a lot of α's and β's that solve (need terminal conditions to solve – see below). One combination yields:

$$x(p) = \frac{3}{2}p - \frac{3}{8}.$$

- Plugging this into equation (c) from last lecture:

$$y(q) = \frac{3}{2}q - \frac{1}{8}.$$

- Inverting these two equations we get the offer functions for the buyer and seller as a function of their value, b and s:

$$q(b) = \frac{2}{3}b + \frac{1}{12}$$

$$p(s) = \frac{2}{3}s + \frac{1}{4}$$

- See G-13.1. This shows that buyers understate their true value and sellers overstate. Consider the seller's offer curve, $p(s)$. When the seller values the object very little, he overstates a lot. This is because there is a fairly good probability of trade and he wants to maximize his expected gains (which only are positive if there is trade). As the seller's value increases, he inflates less and less up to the point where the buyer will never offer anything higher. The probability of trade falls as the seller's valuation increases so he has to be more and more honest. We have assumed the buyer never offers more than 0.75 (a terminal conditon used to solve the above diffeq). Beyond 0.75, no trade will occur so the seller can offer anything.

- The opposite reasoning applies to the buyer which yields the $q(b)$ curve. Again, we have assumed (as a terminal condition) that the seller never offers less than 0.25.

"The buyer's incentive to understate vanishes at the point where the seller will not bid below."

- Due to having to make these assumptions, bargaining games have multiple of (infinitely many) solutions. We could focus on just the symmetric solution and would find that it is unique.
- So when does trade occur?

$$\begin{aligned} q(b) &\geq p(s) \\ \frac{2}{3}b + \frac{1}{12} &\geq \frac{2}{3}s + \frac{1}{4} \\ \frac{2}{3}(b-s) &\geq \frac{1}{4} - \frac{1}{12} \\ b - s &\geq \frac{3}{2} * \frac{1}{6} = \frac{1}{4} \end{aligned}$$

- So we wanted trade to take place whenever $b > s$ (ex-post efficient), but we get this inefficiency that the gains from trade must be at least a quarter for trade to take place. The outcome is inefficient but it's the best we can do.

Public Choice Problem

- Suppose there are two members of society, say drivers and a bridge building company. There is a project, say building a bridge, and let $d = 1$ if the bridge is built, $d = 0$ else.
- Members have privately known values of the project, $v_i \in (-\infty, \infty)$, so this could be the value of driving to the drivers and the profits from the bridge project to the builder.
- Ex-post efficiency requires $d(v_1, v_2) = 1$ if $v_1 + v_2 \geq 0$. Otherwise, don't build.
- We seek a mechanism to implement an efficient choice rule.
- By the revelation principle, we can restrict our attention to incentive compatible direct mechanisms (ICDMs). Not to be confused with ICBMs. So our mechanism takes reported valuations and determines a decision variable, $d(v_1, v_2)$ and transfers, $t(v_1, v_2)$. Suppose $t_i(v_1, v_2)$ is the transfer that player i receives.
- We would like both a socially efficient mechanism, and if possible, one that is dominant strategy incentive compatible, ie:

$$v_i^* \in \ argmax_{v_i} \ \{v_i^* d(v_i, v_j) + t_i(v_i, v_j)\} \ \forall \ v_j,$$

where v_i^* is the TRUE valuation of player i. So truth-telling is a dominant strategy (optimal for any v_j).

- So one thing we could try is a Groves Mechanism. Build if the reported types, $v_1+v_2 \geq 0$, ie the net benefit of the project to the two members of society is positive. Let the transfer be:

$$t_i(v) = d(v_i, v_j)v_j + h_i(v_j),$$

where h is some function that does not depend on i's reported value.

- So consider i's problem:

$$\begin{aligned} & Max_{v_i}\ \{v_i^* d(v_i, v_j) + t_i(v_i, v_j)\} \\ \equiv\ & Max_{v_i}\ \{v_i^* d(v_i, v_j) + d(v_i, v_j)v_j + h_i(v_j)\} \\ \equiv\ & Max_{v_i}\ \{(v_i^* + v_j)d(v_i, v_j)\} \end{aligned}$$

And this implies that $v_i = v_i^*$ is optimal since $d = 1$ only if $v_i^* + v_j \geq 0$. So truth-telling is a dominant strategy.

- HOWEVER, Groves still suffers from the usual budget balance problem. We cannot design a tranfer mechanism to ensure the bridge is built whenever socially efficient and it gets funded.

- So consider another mechanism which is NOT a dominant strategy equilibrium, but truth-telling is still a Bayesian equilibrium.

- Now require:

$$v_i^* \in argmax_{v_i}\ E_{v_j}[v_i^* d(v_i, v_j) + t_i(v_i, v_j)|v_i].$$

So given we don't know v_j, we integrate out over what we don't know and want truth-telling to be a best response.

- Budget balance will be satisfied: $t_1(v) + t_2(v) = 0\ \forall\ v$.

- So we want a mechanism to satisfy ex-post efficiency, incentive compatibility, and budget balance.

- Consider a (guessed) transfer function $t_i(v) = g_i(v_i) - g_j(v_j)$, where:

$$g_i(v_i) = \int_{-\infty}^{\infty} v_j d(v_i, v_j) f_j(v_j) dv_j.$$

Clearly this implies $t_1(v) + t_2(v) = 0$.

- First check that we satisfy incentive compatibility. Player i solves:

$$Max_{v_i} \int_{-\infty}^{\infty} [v_i^* d(v_i, v_j) + t_i(v_i, v_j)] f_j(v_j) dv_j,$$

and substituting for the transfer:

$$Max_{v_i} \int_{-\infty}^{\infty} [v_i^* d(v_i, v_j) + g_i(v_i) - g_j(v_j)] f_j(v_j) dv_j$$

$$\equiv Max_{v_i} \int_{-\infty}^{\infty} [v_i^* d(v_i, v_j) + \int_{-\infty}^{\infty} v_j d(v_i, v_j) f_j(v_j) dv_j - \int_{-\infty}^{\infty} v_i d(v_i, v_j) f_i(v_i) dv_i] f_j(v_j) dv_j$$

$$\equiv Max_{v_i} \int_{-\infty}^{\infty} [v_i^* + v_j] d(v_i, v_j) f_j(v_j) dv_j$$

because (I guess) the second (double) integral evaluates to zero. Again, by inspection of this last expression, player i would want to truthfully report $v_i = v_i^*$ in order to maximize for the same reason as above.

- The problem with all this is that individual rationality may not be satisfied:

$$U_i(v_i) = \int_{-\infty}^{\infty} [v_i d(v_i, v_j) + t_i(v_i, v_j)] f_j(v_j) dv_j \geq 0 \; \forall \; v_i \in V_i,$$

may NOT hold for some densities f_j. So we either get IC or IR, but it's hard to satisfy both at the same time.

14 Lecture 14: October 24, 2006

14.1 Bilateral Trading Mechanisms - General Results of Mechanism Design

- We now consider mechanisms that characterize all possible equilibria in a particular class of games.
- We would like to know what sort of mechanisms that we can come up with that will satisfy certain properties, namely Incentive Compatibility (IC), Individual Rationality (IR), and Ex-Post Efficiency.
- The key insight is that we can use the Revelation Principal to help classify the set of possible equilibria. Recall this says that ANY Bayesian Equilibrium of a Bayesian game can be represented by an equilibrium of the direct revelation game (where players report their types) that is both IC and IR.

Model Setup

- Myerson and Satterthwaite, JET 1983
- Consider a bilateral exchange between a buyer and seller with independent private values:

$$\text{Seller: } s \sim F \text{ with positive PDF } f \text{ on } [\underline{s}, \overline{s}].$$

$$\text{Buyer: } b \sim G \text{ with positive PDF } g \text{ on } [\underline{b}, \overline{b}].$$

- Note that F and G are common knowledge but actual values are not. In our direct revelation game, traders report their values and an outcome is selected (by the mechanism designer). So given reports (s, b), we will have two outcome functions:
  - (1) $p(s, b)$: the probability of trade given s and b.
  - (2) $x(s, b)$: the payment from the buyer to the seller given s and b.

So we seek a mechanism, $< p, x >$, that is both IC and IR. Then we'll check for efficiency at the end.

Ex-Post Utilities

- We assume that the seller and buyer have quasi-linear utility functions:

$$\text{Seller: } u(s, b) = x(s, b) - s * p(s, b)$$

$$\text{Buyer: } v(s, b) = b * p(s, b) - x(s, b)$$

So note the seller gets the payment less his value weighted by the probability of trade. Note the probability does NOT multiply the payment. This allows for a more general class of mechanisms where the cash transfer could happen independent of the transfer of the good.

- So everything is extremely general so far expect that we are not modelling risk aversion (these utilities imply risk neutrality) and there are no income effects.

Expected Payments and Probabilities

- The seller does not know the buyer's type so he must integrate out over this information:

$$\text{Seller's Expected Payment: } X(s) = \int_{\underline{b}}^{\overline{b}} x(s,b)g(b)db.$$

$$\text{Seller's Expected Prob of Trade: } P(s) = \int_{\underline{b}}^{\overline{b}} p(s,b)g(b)db.$$

And similarly for the buyer:

$$\text{Buyer's Expected Payment: } Y(b) = \int_{\underline{s}}^{\overline{s}} x(s,b)f(s)ds.$$

$$\text{Buyer's Expected Prob of Trade: } Q(b) = \int_{\underline{s}}^{\overline{s}} p(s,b)f(s)ds.$$

Interim (Ex-ante) Utilities

- Given our expected payments and probabilities, we have:

$$\text{Seller's Ex-ante Utility: } U(s) = X(s) - sP(s)$$

$$\text{Buyer's Ex-ante Utility: } V(b) = bQ(b) - Y(b)$$

Incentive Compatibility and Individual Rationality

- We want our mechanism $< p, x >$ to be IC and IR.
- IC requires that for all s, b, s' and b':

$$U(s) \geq X(s') - sP(s')$$

$$V(b) \geq bQ(b') - Y(b')$$

So telling the truth is always at least as good as lying.

- IR requires that for all $s \in [\underline{s}, \overline{s}]$ and for all $b \in [\underline{b}, \overline{b}]$,

$$U(s) \geq 0$$

$$V(b) \geq 0$$

Lemma 1 (Mirrlees, Myerson)

- The mechanism, $< p, x >$, is IC iff $P(\cdot)$ is decreasing, $Q(\cdot)$ is increasing, and:

$$U(s) = U(\overline{s}) + \int_s^{\overline{s}} P(t)dt$$

$$V(b) = V(\underline{b}) + \int_{\underline{b}}^{b} Q(t)dt$$

Call these two conditions (IC'). So the ex-ante utility is equal to the integral of the expected probability of trade up to a constant.

- PROOF.
  - (Only If) Lets assume IC holds and show P is decreasing, Q is increasing and IC' holds. We'll do it for the seller but it's all symmetric for the buyer. Given $U(s) = X(s) - sP(s)$ and $U(s') = X(s') - s'P(s')$ by definition. IC requires (substituting for $X(s)$ and $X(s')$):

$$U(s) \geq X(s') - sP(s') = U(s') + s'P(s') - sP(s') = U(s') + (s' - s)P(s')$$

$$U(s') \geq X(s) - s'P(s) = U(s) + sP(s) - s'P(s) = U(s) + (s - s')P(s)$$

Rewriting these two inequalities:

$$U(s) - U(s') \geq (s' - s)P(s')$$

$$U(s') - U(s) \geq (s - s')P(s)$$

So putting them together, we have:

$$(s' - s)P(s) \geq U(s) - U(s') \geq (s' - s)P(s')$$

Divide by $s' - s$:

$$P(s) \geq \frac{U(s) - U(s')}{s' - s} \geq P(s')$$

Let $s' \rightarrow s$,

$$\frac{dU}{ds} = -P(s) \quad (1)$$

Now see G-14.1. Here we have $U(s)$ plotted and a tangent line at s. Assume $s' > s$. Since IC holds, we must have the point z on the graph:

$$z = U(s) - P(s)(s' - s) \leq U(s').$$

So $U(s)$ lies above these tangents (supporting hyperplanes). This implies $U(s)$ is decreasing and convex. Thus $P(\cdot)$ is decreasing. $Q(\cdot)$ will be increasing by the same argument for the buyer. How do we get to IC'? Integrate (1):

$$\int_{\overline{s}}^{s} \frac{dU}{dt} dt = -\int_{\overline{s}}^{s} P(t)dt$$

$$U(s) - U(\overline{s}) = \int_{s}^{\overline{s}} P(t)dt$$

$$U(s) = U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt$$

which is (the seller's part of) IC'.

- (If) Lets assume P decreasing, Q increasing, and IC' holds and show that IC holds. Recall IC says:

$$U(s) \geq X(s') - sP(s')$$

Given $U(s) = X(s) - sP(s)$,

$$X(s) - sP(s) \geq X(s') - sP(s')$$

$$s[P(s) - P(s')] + X(s') - X(s) \leq 0 \quad (2)$$

So lets show that (2) holds. Recall, using IC', we have:

$$U(s) = U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt,$$

and substituting for $U(s)$:

$$X(s) - sP(s) = U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt$$

$$X(s) = sP(s) + U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt$$

So substituting this into (2), we can instead show the following:

$$0 \geq s[P(s) - P(s')] + s'P(s') - sP(s) + \int_{s'}^{\overline{s}} P(t)dt - \int_{s}^{\overline{s}} P(t)dt$$

Or,

$$0 \geq (s'-s)P(s') + \int_{s'}^{s} P(t)dt = -P(s')\int_{s'}^{s} dt + \int_{s'}^{s} P(t)dt = \int_{s'}^{s} [P(t) - P(s')]dt \quad (3)$$

So note that (3) holds because $P(\cdot)$ is decreasing. The proof is similar for the buyer. QED.

Lemma 2 (Mirrlees, Myerson)

- An incentive compatible mechanism, $< p, x >$, is individually rational iff:

$$U(\overline{s}) \geq 0, \text{ and } V(\underline{b}) \geq 0.$$

Call this condition IR'. Ie, the worst types (highest value seller and lowest value buyer) must be willing to participate in the mechanism.

- PROOF. Clearly IR' is necessary for the mechanism to be IR because if IR is satisfied for the worst guys, it is satisfied for all. By Lemma 1, since $U(\cdot)$ is decreasing (see graph), IR' is sufficient as well.

- It will turn out that the worst types will be the only types that we don't have to "bribe" in order for them to tell the truth in a IC/IR mechanism. Everyone else will have to receive a bit more to make them honest. This is what will lead to the inefficiency (ie, trade won't take place at times when it is ex-post efficient).

Corollary

- An IC/IR mechanism, $< p, x >$, satisfies:

$$U(\overline{s}) + V(\underline{b}) = \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left[b - \frac{1 - G(b)}{g(b)} - s - \frac{F(s)}{f(s)} \right] p(s, b) f(s) g(b) ds db \geq 0.$$

- So this is really an accounting exercise. The aggregate payoff of those worst off equals this expression which is non-negative. The non-negativity comes from the IR constraints above. The whole thing is approximately the expected gains from trade except we have these ratios of CDF and PDF terms inside. These terms represent the bribes to the buyer and seller to induce truth-telling.

- Note there is NO $x(s, b)$ in this condition! It only depends on the probability of trade, $p(s, b)$.

- PROOF. Recall IC' for the seller:

$$U(s) = U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt.$$

Substitute in for $U(s)$:

$$X(s) = sP(s) + U(\overline{s}) + \int_{s}^{\overline{s}} P(t)dt.$$

Substitute in definition of X and P:

$$\int_{\underline{b}}^{\overline{b}} x(s,b)g(b)db = s\int_{\underline{b}}^{\overline{b}} p(s,b)g(b)db + U(\overline{s}) + \int_{s}^{\overline{s}} [\int_{\underline{b}}^{\overline{b}} p(t,b)g(b)db]dt.$$

Now take the expectation with respect to s:

$$\int_{\underline{s}}^{\overline{s}}\int_{\underline{b}}^{\overline{b}} x(s,b)f(s)g(b)dbds = \int_{\underline{s}}^{\overline{s}}\int_{\underline{b}}^{\overline{b}} sp(s,b)f(s)g(b)dbds + U(\overline{s}) + \int_{\underline{s}}^{\overline{s}}\int_{s}^{\overline{s}}\int_{\underline{b}}^{\overline{b}} p(t,b)f(s)g(b)dbdtds.$$

Or,

$$\int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb = U(\overline{s}) + \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} sp(s,b)f(s)g(b)dsdb + \int_{\underline{b}}^{\overline{b}} [\int_{\underline{s}}^{\overline{s}}\int_{s}^{\overline{s}} p(t,b)f(s)dtds]g(b)db.$$

Last term in brackets (see G-14.2):

$$\begin{aligned}
\int_{\underline{s}}^{\overline{s}}\int_{s}^{\overline{s}} p(t,b)f(s)dtds &= \int_{\underline{s}}^{\overline{s}}\int_{\underline{s}}^{t} p(t,b)f(s)dsdt \\
&= \int_{\underline{s}}^{\overline{s}} p(t,b)\int_{\underline{s}}^{t} f(s)dsdt \\
&= \int_{\underline{s}}^{\overline{s}} p(t,b)F(t)dt \\
&= \int_{\underline{s}}^{\overline{s}} p(s,b)F(s)ds
\end{aligned}$$

Plug this into the equation above:

$$\int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb = U(\overline{s}) + \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} sp(s,b)f(s)g(b)dsdb + \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} p(s,b)F(s)g(b)dsdb.$$

- So as an exercise, lets show that the same condition holds for the buyer. IC' for the buyer:

$$V(b) = V(\underline{b}) + \int_{\underline{b}}^{b} Q(t)dt$$

Substitute in for $V(b)$:

$$bQ(b) - Y(b) = V(\underline{b}) + \int_{\underline{b}}^{b} Q(t)dt$$

$$Y(b) = bQ(b) - V(\underline{b}) - \int_{\underline{b}}^{b} Q(t)dt$$

Substitute in definition of Y and Q:

$$\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)ds = b\int_{\underline{s}}^{\overline{s}} p(s,b)f(s)ds - V(\underline{b}) - \int_{\underline{b}}^{b}[\int_{\underline{s}}^{\overline{s}} p(s,t)f(s)ds]dt.$$

Now take the expectation with respect to b:

$$\int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb = -V(\underline{b}) + \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} bp(s,b)f(s)g(b)dsdb - \int_{\underline{b}}^{\overline{b}}\int_{\underline{b}}^{b}\int_{\underline{s}}^{\overline{s}} p(s,t)g(b)f(s)dsdtdb.$$

The outside double-integral of the last term can be simplified:

$$\begin{aligned}
\int_{\underline{b}}^{\overline{b}}\int_{\underline{b}}^{b} p(s,t)g(b)dtdb &= \int_{\underline{b}}^{\overline{b}}\int_{t}^{\overline{b}} p(s,t)g(b)dbdt \\
&= \int_{\underline{b}}^{\overline{b}} p(s,t)\int_{t}^{\overline{b}} g(b)dbdt \\
&= \int_{\underline{b}}^{\overline{b}} p(s,t)(1-G(t))dt \\
&= \int_{\underline{b}}^{\overline{b}} p(s,b)(1-G(b))db
\end{aligned}$$

Plug this into the equation above:

$$\int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb = -V(\underline{b}) + \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} bp(s,b)f(s)g(b)dsdb - \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} p(s,b)f(s)[1-G(b)]dsdb.$$

- So we have these two giant expressions for the expected transfer from the buyer to the seller, one from each player's perspective. They must be equal! So:

$$\begin{aligned}
E[X(s)] &= E[Y(b)] \\
\int_{\underline{b}}^{\overline{b}} x(s,b)g(b)db &= \int_{\underline{s}}^{\overline{s}} x(s,b)f(s)ds \\
\int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb &= \int_{\underline{b}}^{\overline{b}}\int_{\underline{s}}^{\overline{s}} x(s,b)f(s)g(b)dsdb
\end{aligned}$$

Equating the RHS of each of these expressions (from above):

$$U(\overline{s}) + \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} sp(s,b)f(s)g(b)dsdb + \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} p(s,b)F(s)g(b)dsdb$$

$$= -V(\underline{b}) + \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} bp(s,b)f(s)g(b)dsdb - \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} p(s,b)f(s)[1-G(b)]dsdb$$

More simply:

$$U(\overline{s}) + \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} sp(s,b)f(s)g(b) + p(s,b)F(s)g(b)dsdb$$

$$= -V(\underline{b}) + \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} bp(s,b)f(s)g(b) - p(s,b)f(s)[1-G(b)]dsdb$$

And finally,

$$\begin{aligned}
& U(\overline{s}) + V(\underline{b}) \\
= & \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} bp(s,b)f(s)g(b) - p(s,b)f(s)[1-G(b)] - [sp(s,b)f(s)g(b) + p(s,b)F(s)g(b)]dsdb \\
= & \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} [bp(s,b) - p(s,b)\frac{1-G(b)}{g(b)} - sp(s,b) - p(s,b)\frac{F(s)}{f(s)}]f(s)g(b)dsdb \\
= & \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left[b - \frac{1-G(b)}{g(b)} - s - \frac{F(s)}{f(s)} \right] p(s,b)f(s)g(b)dsdb \geq 0
\end{aligned}$$

QED. This is where we started a day and a half ago.

- Some intuitition. The expression is approximately the expected gains from trade though the PDF and CDF terms reflect bribes to the players to guarantee truth-telling.
- Thus for ANY mechanism you give me, $p(s,b)$, I can tell you if the mechanism is IC and IR. I know that ONE exists for sure. Let $p(s,b) = 0$ always (no trade). The corollary equation is trivially satisfied.
- What about a nice, ex-post efficient mechanism like $p^*(s,b) = 1$ if $s \leq b$ and $p^*(s,b) = 0$ else. In turns out (and we'll prove it soon) that this VIOLATES the corollary equation. This mechanism does NOT satisfy ex-post efficiency.
- **Theorem:** If it is not common knowledge that gains exist (ie the supports of the traders' valuations have a non-empty intersection), then NO IC/IR trading mechanism can be ex-post efficient. If the supports do not overlap, then one might exist.

- See G-14.3 for the no overlap case. Clearly $p(s, b) = 1$ for all s and b with:

$$x(s, b) = \frac{\overline{s} + \underline{b}}{2},$$

will satisfy the corollary equation and be ex-post efficient.

- See G-14.4 for the overlap case. The corollary will be violated and no ex-post efficient IC/IR mechanism will exist (because bribing will be required to get IR and IC).

- We'll prove the theorem next time.

15 Lecture 15: October 31, 2006

15.1 More on Efficiency of IC/IR Mechanism

- Recall the theorem from last time, which was based on the corollary equation:

$$U(\overline{s}) + V(\underline{b}) = \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left[b - \frac{1-G(b)}{g(b)} - s - \frac{F(s)}{f(s)} \right] p(s,b) f(s) g(b) ds db \geq 0$$

  which said that if it is NOT common knowledge that there are gains from trade, then no IC/IR trading mechanism exists that is ex-post efficient. The idea is that if the support of the distributions overlap, it may be the case that $s > b$ and gains do not exist.

- So in summary, if we have risk neutrality (quasi-linear utility) and independent private values, the corollary being satisfied means there does exist a Bayesian Game with the mechanism $< p, x >$ as a Bayesian Equilibrium. We just need the net gains (gross gains less bribes) to be non-negative.

- We can show that the ex-post efficiency requirement, ie always trade if $s \leq b$, cannot be satisfied with an IC/IR mechanism. See Cramton notes for a tedious IBP technique to prove this. The proof shows the corollary equation is negative if $\underline{b} < \overline{s}$, ie the supports overlap.

- So what to do next? Since ex-post efficiency is unattainable, we need a weaker efficiency criterion with which to measure a mechanism's performance.

- Recall the definition of Pareto optimality: an allocation such that you can't make at least one party strictly better off without making some other parties worse off.

- In games of incomplete information, we have the following definition of efficiency: **Efficiency:** "A decision rule is efficient iff no other feasible decision rule can be found that may make some individuals better off without ever making any other individual worse off." Note we have added the concept of feasibility and a timing element because there is uncertainty about the future.

- The key issue in all of this will be: On what information should the expectation be conditioned? Three possibilities:

  - (1) Ex-ante information: a planner's (designer's) information at the beginning of the game (no knowledge of player's types).
  - (2) Interim information: a player's private information at the beginning of the game (they know their type and the distributions of their competitors).
  - (3) Ex-post information: all the private information (say what a judge and jury could act on).

- So one problem is that suppose we find an ex-ante efficient decision rule. What if, after realizing their private information, a player proposes a new decision rule that is guaranteed to be accepted. Do we use that one instead of the ex-ante efficient rule?

- So some theory. Consider a Bayesian game, $\Gamma = \{A_1, \ldots, A_n; T_1, \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\}$ with finite A_i and T_i. Beliefs are consistent and let D be the set of probability distributions over actions.

- A decision rule (or direct mechanism), $\delta : T \mapsto D$, maps reports (true valuations via the revelation principal) into a randomization over feasible actions. Then,

$$u_i(d, t) : D \times T \mapsto \Re,$$

maps decisions and types into payoffs.

- A decision rule, $\delta \in \Delta \equiv \{\delta : T \mapsto D\}$ is incentive compatible if,

$$\sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(\delta(t), t) \geq \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(\delta(t_{-i}, t_i'), t) \; \forall \; i \text{ and } t_i' \in T_i.$$

So let $\Delta^* = \{\delta : T \times D \mapsto \Re\}$ be the set of all IC decision rules. By the revelation principal, we can restrict our attention to $\delta \in \Delta^*$.

- So given a decision rule, $\delta(\cdot)$, the expected utility at each of the three information stages is:

  - (1) Ex-ante utility: $U_i(\delta) = \sum_{t \in T} p(t) u_i(\delta(t), t)$.
  - (2) Interim utility: $U_i(\delta|t_i) = \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(\delta(t), t)$.
  - (3) Ex-post utility: $U_i(\delta|t) = u_i(\delta(t), t)$.

  So, in economics, we would like to base our judgements on ex-ante efficiency considerations. However, in the real world, says for judges in a court of law, decisions are based on ex-post considerations and mess up the ex-ante efficiency requirements. An additional alternative would be to ignore our Bayesian approach where we assume we know the distributions $F(\cdot)$ and $G(\cdot)$, and find decision rules that are optimal for ANY distributions. This essentially collapses to a worse-case scenario analysis.

- **Definition:** Domination.

  - (1) A decision rule, γ, ex-ante dominates δ iff $U_i(\gamma) \geq U_i(\delta) \; \forall \; i$ with at least one strict inequality.
  - (2) A decision rule, γ, interim dominates δ iff $U_i(\gamma|t_i) \geq U_i(\delta|t_i) \; \forall \; i$ and $t_i \in T_i$ with at least one strict inequality.
  - (3) A decision rule, γ, ex-post dominates δ iff $U_i(\gamma|t) \geq U_i(\delta|t) \; \forall \; i$ and $t \in T$ with at least one strict inequality.

- **Definition:** Efficiency.

- (1) A decision rule, δ, is ex-ante (incentive) efficient iff $\nexists\ \gamma \in \Delta^*$ that ex-ante dominates δ.
- (2) A decision rule, δ, is interim (incentive) efficient iff $\nexists\ \gamma \in \Delta^*$ that interim dominates δ.
- (3) A decision rule, δ, is ex-post (classically) efficient iff $\nexists\ \gamma \in \Delta$ that ex-post dominates δ.

Note we don't require the γ to be IC in the last definition. The ex-post efficiency requirement is a nice first best definition, but as we have shown, is not always attainable.

Ex-ante Efficiency in the Bilateral Trading Game

- Returning now to Myerson and Satterthwaite. Suppose we want to find the ex-ante efficient mechanism that maximizes the expected gains from trade, ie,

$$Max\ \{\int_{\underline{s}}^{\overline{s}} U(s)f(s)ds + \int_{\underline{b}}^{\overline{b}} V(b)g(b)db\}.$$

Here we assume equal weight on the buyer and seller but clearly you could make other assumptions.

- So in the bilateral trading game, we are looking to solve:

$$Max_{p(s,b)} \left\{ \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} (b-s)p(s,b)f(s)g(b)dsdb \right\},$$

Subject to:

$$\int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} [b - \frac{1-G(b)}{g(b)} - s - \frac{F(s)}{f(s)}]p(s,b)f(s)g(b)dsdb \geq 0.$$

- Consider the lagrangian:

$$\begin{aligned}
\mathcal{L} &= \int_{\underline{b}}^{\bar{b}}\int_{\underline{s}}^{\bar{s}}\left[(b-s)+\lambda[b-\frac{1-G(b)}{g(b)}-s-\frac{F(s)}{f(s)}]\right]p(s,b)f(s)g(b)dsdb \\
&= \int_{\underline{b}}^{\bar{b}}\int_{\underline{s}}^{\bar{s}}\left[(1+\lambda)(b-s)+\lambda[-\frac{1-G(b)}{g(b)}-\frac{F(s)}{f(s)}]\right]p(s,b)f(s)g(b)dsdb \\
&= \int_{\underline{b}}^{\bar{b}}\int_{\underline{s}}^{\bar{s}}\left[(b-s)+\frac{\lambda}{1+\lambda}[-\frac{1-G(b)}{g(b)}-\frac{F(s)}{f(s)}]\right]p(s,b)f(s)g(b)dsdb \\
& \text{Let } \alpha=\frac{\lambda}{1+\lambda} \\
&= \int_{\underline{b}}^{\bar{b}}\int_{\underline{s}}^{\bar{s}}\left[b-s+\alpha[-\frac{1-G(b)}{g(b)}-\frac{F(s)}{f(s)}]\right]p(s,b)f(s)g(b)dsdb \\
&= \int_{\underline{b}}^{\bar{b}}\int_{\underline{s}}^{\bar{s}}\left[\Bigg(\underbrace{b-\alpha\frac{1-G(b)}{g(b)}}_{d(\alpha,b)}\Bigg)-\Bigg(\underbrace{s+\alpha\frac{F(s)}{f(s)}}_{c(\alpha,s)}\Bigg)\right]p(s,b)f(s)g(b)dsdb
\end{aligned}$$

Thus, by inspection, to maximize this expression, we want to set $p(s,b)=1$ whenever $d(\alpha,b)\geq c(\alpha,s)$, and $p(s,b)=0$ otherwise.

- So our ex-ante efficient decision rule (for the probability of trade) is:

$$p^{\alpha}(s,b)=\begin{cases}1 & if\ d(\alpha,b)\geq c(\alpha,s)\\ 0 & if\ d(\alpha,b)<c(\alpha,s)\end{cases}$$

- Note that our DR still depends on α. So we need to choose the α that just satisfies the corollary equation, ie:

$$U(\bar{s})=V(\underline{b})=0.$$

So we just want the constraint to bind.

- **Remark 1** If $\alpha=0$, then p^{α} is ex-post efficient (all the weight is being placed on the objective function).

- **Remark 2** If $\alpha=1$, then p^{α} maximizes the corollary equation.

- **Remark 3** The ex-ante efficient trading rule has the property that, given the reports, trade either occurs with probability one or not at all! This is important. Note that in the dynamic setting, say the repeated bargaining game, this is equivalent to saying that players either agree in the first round or not at all. There won't be an ex-ante efficient mechanism that involves players bargaining for a few periods and then settling on a division of the pie.

- So back to our $U[0,1]$ example. Note that $F(s) = s$, $G(b) = b$, and $f(s) = g(b) = 1$. Thus,

$$c(\alpha, s) = s + \alpha s = s(1+\alpha).$$

$$d(\alpha, b) = b - \alpha(1-b) = b(1+\alpha) - \alpha.$$

Trade occurs ($p(s,b) = 1$) if:

$$d \geq c \Longleftrightarrow b - s \geq \frac{\alpha}{1+\alpha} \equiv \phi.$$

To find what the RHS equates to, go back to the corollary equation. Note that $b - s \geq \phi$ yields an area like G-15.1. I use this diagram to find my limits of integration.

$$\begin{aligned}
0 &= U(\overline{s}) + V(\underline{b}) \\
&= \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} [b - (1-b) - s - s] p(b,s) f(s) g(b) ds db \\
&= \int_{\underline{b}}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} [2(b-s) - 1] p(b,s) ds db \\
&= \int_{0}^{1} \int_{0}^{1} [2(b-s) - 1] p(b,s) ds db \\
&= \underbrace{\int_{\phi}^{1} \int_{0}^{b-\phi} [2(b-s) - 1] * 1 \; ds db}_{trade\ occurs} \\
&= \int_{\phi}^{1} \int_{0}^{b-\phi} [2b - 2s - 1] ds db \\
&= \int_{\phi}^{1} [2bs - s^2 - s]_0^{b-\phi} db \\
&= \int_{\phi}^{1} 2b(b-\phi) - (b-\phi)^2 - (b-\phi) db \\
&= \int_{\phi}^{1} 2b^2 - 2b\phi - b^2 - \phi^2 + 2b\phi - b + \phi db \\
&= \int_{\phi}^{1} b^2 - \phi^2 - b + \phi db \\
&= [\frac{b^3}{3} - b\phi^2 - \frac{b^2}{2} + \phi b]_{\phi}^{1} \\
&= \frac{1}{3} - \phi^2 - \frac{1}{2} + \phi - \frac{\phi^3}{3} + \phi^3 + \frac{\phi^2}{2} - \phi^2 \\
&= \frac{2\phi^3}{3} - \frac{3\phi^2}{2} + \phi - \frac{1}{6}
\end{aligned}$$

And magically (matlab magic), this cubic solves to $\phi = \frac{1}{4}$, as we have shown previously.

- So if traders cannot commit to walking away from gains from trade, the mechanism will fail. Pretty strong assumption. Traders want to act sequentially rational but also realize that if they do, they lose some efficiency.
- There is a tension between ex-ante efficiency and sequential rationality.
- So what have we done so far:
  - (1) We have a complete characterization of the set of all BE of all bargaining games in which players' strategies map their private valuations into a probability of trade and a payment from buyer to seller.
  - (2) We have shown that ex-post efficiency is unattainable if it is uncertain that gains from trade exist.
  - (3) We have found the set of ex-ante efficient mechanisms.
  - (4) And we have shown that ex-ante efficiency is incompatible with sequential rationality.

Durability of Trading Mechanisms

- We don't like this tradeoff between efficiency and rationality, so are there further restrictions on the feasible set of decision rules that would (and could) be made ?
- Consider the notion of Durability. Is it ever the case that a player, knowing her type, could suggest an alternative decision rule, γ, that others would surely prefer? It turns out the answer is Yes! We say that a decision rule is durable if such an alternative does not exist.
- Consider an example. Each of two players, 1 and 2, is one of two types, a or b. The players' utilities as a function of their types and the three possible decisions (X, Y, Z) are as follows:

	1a	1b	2a	2b
d = X	2	0	2	2
d = Y	1	4	1	1
d = Z	0	9	0	-8

- The ex-ante efficient decision rule (satisfies IC) that maximizes the sum of the players' payoffs is:

	2a	2b
1a	X	Y
1b	Z	Y

PROOF: Suppose $u(p,q)$ is the expected utility of a player of type p that says he's of type q. So:

$$
\begin{aligned}
u(1a,1a) &= \frac{1}{2}*2+\frac{1}{2}*1 = 1.5 &\geq&\ 0.5 = u(1a,1b) = \frac{1}{2}*0+\frac{1}{2}*1 \\
u(1b,1b) &= \frac{1}{2}*4+\frac{1}{2}*9 = 6.5 &\geq&\ 2.0 = u(1b,1a) = \frac{1}{2}*0+\frac{1}{2}*4 \\
u(2a,2a) &= \frac{1}{2}*2+\frac{1}{2}*0 = 1.0 &\geq&\ 1.0 = u(2a,2b) = \frac{1}{2}*1+\frac{1}{2}*1 \\
u(2b,2b) &= \frac{1}{2}*1+\frac{1}{2}*1 = 1.0 &\geq&\ -3.0 = u(2b,2a) = \frac{1}{2}*2+\frac{1}{2}*(-8)
\end{aligned}
$$

So this DR is incentive compatible. No outsider could suggest an alternative DR that would make some type better off without making another type worse off.

- The problem is, what about an INSIDER! Suppose player 1 is of type 1a. Player 1 suggests to 2 that decison X be always chosen. Note that player 2 (types 2a or 2b) strictly prefers X to either Y or Z, and player 1a also prefers X to either Y or Z. So player 2 would surely accept! Thus the decision rule above (the ex-ante (incentive) efficient) DR) is NOT DURABLE!
- A DR is durable iff the players would NEVER unanimously approve a change to any other DR.
- Next time some extensions beyond the bargaining environment to auctions and dissolving partnerships.

16 Lecture 16: November 7, 2006

16.1 Mechanism Design Example: Dissolving a Partnership

- Cramton, Gibbons, and Klemperer (CGK), 1987.
- Consider n traders, indexed by i, each with a share, $r_i \geq 0$ of an asset. Note $\sum_i r_i = 1$.
- As in Myerson/Saterthwaite (MS), a player i's valuation for the entire asset is v_i, which is private to that player.
- The utility from owning share r_i is $r_i v_i$ where $v_i \sim iid\ F(\cdot)$ on $[\underline{v}, \overline{v}]$.
- A partnership, (r, F), is fully described by the vector of ownership right, r, and the traders' beliefs, F, about valuations.
- MS Case. Two traders (seller and buyer) with shares $r = \{1, 0\}$. We have the result that there does NOT exist a BE, σ, of the trading game such that:
  - (1) σ is (interim) individually rational and
  - (2) σ is ex-post efficient.

  (This assumes the distributions of valuations overlap.)
- CGK Case. If the ownership shares are not too unequal, ie one player doesn't own too much, then it is possible to satisfy both (1) and (2), ie satisfy IC, IR, ex-post efficiency, and balanced budget. If the shares are very different, the mechanism will require a lot of bribing to get IC so we won't get efficiency.
- The idea is that with fairly symmetric shares, the incentive to misrepresent is weakened, the bribes required are smaller, and thus, we're golden.
- So a partnership, (r, F), can be dissolved efficiently if there exists a Bayesian Equilibrium, σ, of a Bayesian trading game such that σ is interim individually rational and ex-post efficient.
- **Theorem:** The partnership, (r, F), can be dissolved efficiently if and only if:

$$\sum_{i=1}^{n} \left[\int_{v_i^*}^{\overline{v}} [1 - F(u)] u dG(u) - \int_{\underline{v}}^{v_i^*} F(u) u dG(u) \right] \geq 0, \qquad (1)$$

  where v_i^* solves $F(v_i)^{n-1} = r_i$ and $G(u) = F(u)^{n-1}$.
- So equation (1) is the same as the corollary equation in MS. The difference is that the worse off type is now the guy who is completely unsure about if he wants to be a net seller or a net buyer. Suppose r_i is my share and I'm the guy with $r_i = F(v_i)^{n-1} = Pr\{$ my value is the highest $\}$. Thus r_i is the probability, under ex-post efficiency, that

I should be awarded the whole good. So with probability r_i, I should (efficiently) purchase $1 - r_i$, ie the rest, of the good. Thus,

$$E[purchases] = (1 - r_i)r_i.$$

However, there is a probability $1 - r_i$ that I am NOT the highest valuation guy and in that case, I should (efficiently) sell my entire share, r_i. Thus,

$$E[sales] = r_i(1 - r_i).$$

So $E[purchases] = E[sales]$ for player i, I have NO idea if I'm supposed to overstate or understand my share as we did in the MS setting. Thus, I have no incentive to misrepresent and thus I am the worst off player. I don't require any bribe to bid / report truthfully.

- The proof of this theorem is the same as in MS, and thus, is left as an exercise to someone other than me.

- Consider an example. Suppose $n = 3$ and $F(v_i) = v_i$, uniform on $[0, 1]$. Note v_i^* solves $v_i^2 = r_i$, or $v_i^* = \sqrt{r_i}$. Equation (1) reduces to:

$$
\begin{aligned}
\sum_{i=1}^{3}\left[\int_{\sqrt{r_i}}^{1}[1-u]ud[u^2]-\int_{0}^{\sqrt{r_i}}u*ud[u^2]\right] &\geq 0\\
\sum_{i=1}^{3}\left[\int_{\sqrt{r_i}}^{1}[u-u^2]*2udu-\int_{0}^{\sqrt{r_i}}u^2*2udu\right] &\geq 0\\
\sum_{i=1}^{3}\left[\int_{\sqrt{r_i}}^{1}2u^2-2u^3du-\int_{0}^{\sqrt{r_i}}2u^3du\right] &\geq 0\\
\sum_{i=1}^{3}\left[\frac{2}{3}u^3-\frac{2}{4}u^4\right]_{\sqrt{r_i}}^{1}-\sum_{i=1}^{3}\left[\frac{2}{4}u^4\right]_{0}^{\sqrt{r_i}} &\geq 0\\
\sum_{i=1}^{3}[\frac{2}{3}-\frac{1}{2}-\frac{2}{3}r_i^{3/2}+\frac{1}{2}r_i^2-\frac{1}{2}r_i^2] &\geq 0\\
\sum_{i=1}^{3}[\frac{1}{6}-\frac{2}{3}r_i^{3/2}] &\geq 0\\
\frac{1}{2}-\sum_{i=1}^{3}\frac{2}{3}r_i^{3/2} &\geq 0\\
-\sum_{i=1}^{3}r_i^{3/2} &\geq -\frac{1}{2}*\frac{3}{2}\\
\sum_{i=1}^{3}r_i^{3/2} &\leq \frac{3}{4}
\end{aligned}
$$

See G-16.1 for the range of allocations that satisfy equation (1).

- **Proposition:** For any distribution, F, the one-owner partnership, $r=\{1,0,0,\ldots,0\}$, cannot be dissolved efficiently. Note the one-owner partnership can be interpreted as an auction. We cannot get ex-post efficiency because the seller's valuation is private information. The seller (as we will show) finds it in her best interest to set the reserve price above her value. Thus there will be circumstances where the object is not traded if the highest buyer has a value below the reservation price (though above the seller's value).

- An optimal auction maximizes the seller's expected revenue over the set of feasible (ex-post inefficient) mechanisms.

- **Theorem:** If a partnership, (r,F), can be dissolved efficiently (ie, equation (1) is satisfied), then the unique symmetric equilibrium of the following bidding game is interim individually rational and achieves ex-post efficiency. The game: given an arbitrary minimum bid, R,

  - Players choose $b_i \in [R,\infty)$

- The good is allocated to the highest bidder
- Each bidder (even the losers) pay:

$$p_i(b_1, \ldots, b_n) = b_i - \frac{1}{n-1} \sum_{j \neq i} b_j.$$

  So this is like an all pay auction where the revenues received are recycled.
- Each player also receives a side-payment, independent of the bidding (used to satisfy IR):

$$c_i(r_1, \ldots, r_n) = \int_{\underline{v}}^{v_i^*} u dG(u) - \frac{1}{n} \sum_{j=1}^{n} \int_{\underline{v}}^{v_j^*} u dG(u).$$

I don't have any good intuition for the form of the side payment.

- So this all works for an auction, but can we do even better? Yes!

16.2 Mechanism Design Example: Optimal Auctions

- Myerson, 1981.

- Suppose there are n buyers indexed by i. Buyer's willingness to pay / valuation is:

$$t_i \in [\underline{a}_i, \overline{a}_i],$$

  where t_i is independently drawn from $f_i(\cdot)$. Note we do not assume symmetry of types / valuations.

- The seller's type, however, denoted t_0, IS common knowledge to all.

- **Definition:** A Bayesian Auction consists of bids spaces $\{B_1, \ldots, B_n\}$ and outcome functions,

$$\tilde{p}_i : B \mapsto [0,1] \text{ and } \tilde{x}_i : B \mapsto \Re.$$

  where $\tilde{p}_i$ is the probability that player i gets the object and $\tilde{x}_i$ is the payment made from player i to the seller. Note for each vector of bids, $b \in B$,

$$\sum_{i=1}^{n} \tilde{p}_i \leq 1.$$

  So we are allowing for the case where no buyer get the object and the seller retains.

- Utility functions are quasilinear in money:

$$\tilde{u}_i(b,t) = t_i \tilde{p}_i(b) - \tilde{x}_i(b),$$

  where t is a vector of types. The fungibility of money becomes crucial.

- A strategy for bidder i is a bid function mapping his type into a bid. A strategy profile, $b = \{b_1, \ldots, b_n\}$ is a Bayesian Equilibrium (BE) if for each $t_i \in T_i$, the prescribed bid $b_i(t_i)$ is a best response to the $n-1$ other strategies, b_{-i}. Nothing new here.

- So what are we trying to solve (as mechanism designers) ? We want to choose $\{B_i, \tilde{p}_i, \tilde{x}_i\}$ to maximize the expected revenue:

$$ER = E_b \Bigg\{ \underbrace{\Bigg[1 - \sum_{i=1}^{n} \tilde{p}_i(b)\Bigg] t_0}_{\textit{seller keeps good}} + \sum_{i=1}^{n} \tilde{x}_i(b) \Bigg\},$$

subject to:

$$E_{b_{-i}} \Bigg\{ t_i \tilde{p}_i(b_i, b_{-i}) - \tilde{x}_i(b_i, b_{-i}) | b_i = b_i(t_i) \Bigg\} \geq 0,$$

where this constraint means the bidders are playing optimally and truthfully (IR and IC).

- So via, the revelation principal, we can consider the following problem of the seller:

$$Max_{p_i(t), x_i(t)} \Bigg\{ ER = \int_T \Bigg(\Bigg[1 - \sum_{i=1}^{n} p_i(t)\Bigg] t_0 + \sum_{i=1}^{n} x_i(t) \Bigg) f(t) dt \Bigg\},$$

subject to:

  - (1) Feasibility: $\forall\, i, \forall\, t$, $p_i(t) \geq 0$ and $\sum_i p_i(t) \leq 1$.
  - (2) IC: $V_i(t) \equiv v_i(t_i, t_i) \geq v_i(\tau_i, t_i) \;\forall\, i, \forall\, \tau_i, t_i \in T_i$.
  - (3) IR: $V_i(t_i) \geq 0 \;\forall\, i$.

(Note $v_i(x, y)$ is the value to player i of type y who reports x.)

- **Definition:** Conditional probability that player i gets the object when i's type is t_i:

$$P_i(t_i) \equiv \int_{T_{-i}} p_i(t) f_{-i}(t_{-i}) dt_{-i}.$$

- **Lemma 1:** A mechanism, $\{p_i(\cdot), x_i(\cdot)\}$, satisfies (IC) and (IR) iff $\forall\, i$,

  - (1) Probability of trade: $P_i(t_i)$ is weakly increasing
  - (2) Interim value: $V_i(t_i) = V_i(\underline{a}_i) + \int_{\underline{a}_i}^{t_i} P_i(\tau_i) d\tau_i \;\forall\, t_i \in T_i$
  - (3) Worst off IR: $V_i(\underline{a}_i) \geq 0$.

- **Lemma 2:** If a mechanism, $\{p_i(\cdot), x_i(\cdot)\}$, satisfies (IC) and (IR) then the expected revenue becomes:

$$ER' = t_0 - \sum_{i=1}^{n} V_i(\underline{a}_i) + \int_T \Bigg[\sum_{i=1}^{n} \Bigg\{ t_i - \frac{1 - F_i(t_i)}{f_i(t_i)} - t_0 \Bigg\} p_i(t) \Bigg] f(t) dt.$$

Note that $t_i - \dfrac{1-F_i(t_i)}{f_i(t_i)}$ is called a bidder's "Virtual Valuation".

- The proof of lemma 2 is just like in MS. Omitted. Note for a fixed $\{p_i(\cdot)\}$, choose the following payment in the optimal auction:

$$x_i(t) = t_i p_i(t) - \int_{\underline{a}_i}^{t_i} p_i(t_{-i}, \tau_i) d\tau_i,$$

  and set $V_i(\underline{a}_i) = 0$ for all players.

- So to solve the problem of optimal auction design, consider choosing $p_i(\cdot)$ to maximize ER' subject only to the feasibility constraint. Define these virtual valuation things to be:

$$c_i(t_i) = t_i - \frac{1-F_i(t_i)}{f_i(t_i)}.$$

  Let j be the maximium c_i over all bidders. So bidder j has the highest virtual value. Proceed via point-wise maximization. Thus for each fixed type vector, t, if $c_j(t_j) - t_0 > 0$, then set $p_j(t) = 1$ and $p_i(t) = 0$ for all $i \neq j$. If $c_j(t_j) - t_0 \leq 0$, set $p_i(t) = 0 \; \forall \; i$. This defines the optimal auction if condition (1) of Lemma 1 above holds. Ie we need to make sure the probability of trade is weaklying increasing in a bidder's type / valuation.

- For regular cases (uniform and normal distributions for example) it holds. No problems. Otherwise we just correct for it in some way. See Cramton notes for a graph.

- What about (ex-post) efficiency of the optimal auction? To maximize expected revenue when the buyers know their types but the seller does not, the seller may need to design an auction that sometimes FAILS to award the object to the player with the highest willingness to pay. Ie, the optimal auction won't be ex-post efficient.

- Example 1. Suppose $t_0 = 0$ and $F_i(t_i) = t_i$ for t_i on $[0,1]$, uniform. Virtual valuation is:

$$c_i(t_i) = t_i - \frac{1-F_i(t_i)}{f_i(t_i)} = t_i - \frac{1-t_i}{1} = 2t_i - 1.$$

  So,

$$c_i(t_i) - t_0 > 0 \Longleftrightarrow 2t_i - 1 - t_0 > 0 \Longleftrightarrow t_i > \frac{1}{2}.$$

  So the seller sets the reservation price at $R = \frac{1}{2}$, thus detering half the bidders. Clearly this (could be) ex-post inefficient. Ie, the reserve price maximizes the seller's ex-ante expected revenue, it also may lead to no-trade when gains exist. "One interesting feature of this example is that the reservation price, which is essential when $n = 1$ to get any profits, does not go to zero as $n \to \infty$. Indeed, the reservation price of 1/2 does not depend on n! So what happens is that as you increase the number of bidders, the reservation value just becomes meaningless as competition drives the bids up.

- Example 2. See Cramton notes for an example with 2 bidders with asymmetric uniform type supports. This example illustrates the second inefficiency: (from Givens) "which is related to the seller's desire to price discriminate when faced with asymmetric bidders. The inefficiency that results here is that sometimes the good is traded (good!) but NOT to the highest valuation bidder (bad!)." The seller withholds the object from low types of bidder 2 in order to extract more from the high types.

- In a second price sealed bid auction, the expected payment is just the second highest type/value.

- **Theorem:** Revenue Equivalence Theorem. One form of this theorem states that if:

  - (1) Lowest types get zero: $V(\underline{a}_i) = 0 \; \forall \; i$, and
  - (2) Ex-post efficiency: for each t, $p_j(t) = 1$ if $t_j > max_{j \neq i} t_i$,

  then the seller's expected revenue is just the expected value of the second-highest type. Given alternative auctions that all assign the good to the same guy, these must yield the same expected revenue.

- **Remark:** In the absence of reservation prices, symmetric equilibria of English, Dutch, FPSB and SPSB auctions all satisfy the revenue equivalence theorem above.

17 Lecture 17: November 9, 2006

17.1 Mechanism Design: Optimal Auctions - Final Remarks

- The key result of all this auction design material is that the revenue of the seller is going to depend on whom the good is allocated to, not on the particular form of the auction. So since FPSB, SPSB, English and Dutch all allocate the item to the bidder with the highest value, they all yield the same expected revenue to the seller.
- So why study auctions? Because we have been relying on the assumption of private values, which is usually violated in practice. We'll come back to this after we discuss aspects of dynamic games of imperfect information.

17.2 Dynamic Games of Imperfect Information: Refinements to Subgame Perfection

- Once we relax the perfect information assumption, the equilibrium concept of sub-game perfect NE is no longer sufficient.
- Consider the extensive form game and its strategic form in G-17.1. The extension to SPNE eliminates the non-credible threat, (L,r).
- Same definitions of subgame and subgame perfection as before.
- Given our Nash existence theorem, we have a similar result (proved the same way) that every finite extensive form game has a SPNE.
- Now consider the game in G-17.2. The game only has one subgame so NE = SPNE. The strategic form shows two NE at (L,l) and (R,r). However player 2 has a dominate strategy: to play l! So the second NE is implausible. The problem with SPNE (and NE) is that there is NO requirement to play optimally "off the equilibrium path".
- So SPNE does not prevent the use of strictly dominated strategies. This is unsatisfactory.
- Note we could alter the game with an inessential transformation that would eliminate this implausible equilibrium as shown in the notes. So this is also disconcerting.
- **Definition:** Sequential Rationality. To get around these problems, we introduce the concept of sequential rationality (SR). It requires:
  - (1) Every decision must be part of an optimal strategy for the remainder of the game.
  - (2) At every juncture, the players' subsequent strategy must be optimal wrt some assessment of the probabilities of all uncertain events. Ie, we need to find beliefs that could rationalize the strategies. In G-17.2, none exist. Player 2 will always play l if she gets the move. So, we require optimality off the equilibrium path (for some beliefs). NE did not require this.

- Thus, in G-17.2, the only NE that is SR is (L,l).
- Denote $A(h)$, with $h \in H_i$ to be the set of actions available to player i at information set h. Note the same set of actions must be feasible at every node in an information set.

Sequential Equilibrium (SE)

- **Definition:** Behavioral strategy. Denote a behavioral strategy:

$$\pi_i \equiv \{\pi_h^i(a)\}_{h \in H_i},$$

where the π_i's are non-negative and sum to one. So π_i is a probability distribution that describes the player's behavior at information set h. Denote a strategy profile:

$$\pi = (\pi^1, \dots, \pi^n).$$

- **Definition:** Beliefs are denoted $\mu_h(x)$, which is the probablity that player i assesses that node $x \in h$ has been reached. Again they are non-negative and sum to one over all nodes in the information set. The entire set of beliefs throughout the play of the game is denoted:

$$\mu \equiv \{\mu_h(x)\}_{h \in H}.$$

- **Definition:** Call the beliefs-strategies pair (μ, π) an assessment.
- An assessment, (μ, π), is sequentially rational if given the beliefs μ, no player i prefers at ANY information set $h \in H_i$ to change her strategy π_h^i given the others' strategies π^{-i}. So each player is playing a best response given her beliefs at each information set.
- So the nice thing about SE is that it:
  - (1) Eliminates dominated strategies from consideration off the equilibrium path.
  - (2) It also elevates beliefs to the importance of strategies.
- Now we consider the merits of competing sequentially rational equilibria.
- Consider now the game in G-17.3. Note that player 2 no longer has a dominate strategy at her information set. The game has 2 NE (L,l) and (R,r). But what about Sequential Equilibria. We require sequential rationality. Suppose player 2 has beliefs $(p, 1-p)$ that player 1 played L and M respectively. So clearly (L,l) with $p = 1$ is a SE via Bayes rule and the strategy (which is known) of her opponent. But what about (R,r). Now, player 2's information set is off the equilibrium path so player 1's strategy does

not put a restriction on player 2's beliefs. But if she is to play r, we require:

$$\begin{aligned} E_2[r] &\geq E_2[l] \\ -2p - 1(1-p) &\geq -p - 2(1-p) \\ 1 &\geq 2p \\ p &\leq \frac{1}{2} \end{aligned}$$

So another sequential equilibrium exists with (R, r), $(p, 1-p)$, $p \leq 0.5$.

- So what do we think about this second equilibrium? Look at player 1's strategies and payoffs. She always does better playing R versus playing M. So, via FORWARD INDUCTION, player 2 shouldn't put ANY weight on being at the node following M. Player 2's beliefs should be $(1, 0)$ which means the only equilibrium is (L,l).

- So sequential equilibrium also has its problems. It allows for implausible beliefs. Thus we may want to impose some restrictions on players' beliefs.

- **Definition:** Formally, a sequential equilibrium is an assessment (μ, π) that is both sequentially rational and consistent. Without consistency (next), we have a perfect bayesian equilibrium (PBE) or aka a weak sequential equilibrium (WSE).

- **Definition:** A strategy profile, π, is totally mixed if it assigns strictly positive probability to each action $a \in A(h)$ for each $h \in H$. So EVERYTHING happens with postive probability. There is no "off the equilibrium path" to worry about.

- **Definition:** An assessment (μ, π) is consistent if there exists a sequence of totally mixed strategies, π_n, and corresponding beliefs, μ_n derived from Bayes' rule, such that:

$$lim_{n\to\infty}(\mu_n, \pi_n) = (\mu, \pi).$$

  Note this is almost ALWAYS satisfied so we usually work with PBE instead of SE.

- **Theorem:** For every finite extensive-form game, there exists at least one SE. Note that SE are a subset of SPNE.

- How about weakly dominated strategies. How does the sequential equilibrium concept do? Not well. Consider the game in G-17.4. Two NE at (L,l) and (R,r). Note that l weakly dominates r for player 2. Suppose player 2 has beliefs $[1/n, 1-1/n]$ at the two nodes in her information set. Note at the left node, she plays l and at the right node, she is indifferent. So if there is positive weight on L (ie some small chance that player 1 has played L), it is optimal for player 2 to play l. However, let $n \to \infty$. In the limit, player 2's beliefs become $[0, 1]$. Now all of a sudden, playing r is a best response. So with beliefs $[0, 1]$, there would exists a SE at (R,r), even though for even the slightest preturbation of those beliefs, (R,r) would fail to be sequentially rational.

- So SE has (essentially) no cutting power on weakly dominated stratgies. We'll use trembling hand perfection to strengthen this.

- See notes for something about Structural Consistency which he did not focus on. Fill in these note if this become important. Essentially this caused more problems than solutions.

- More next time.

18 Lecture 18: November 14, 2006

18.1 Dynamic Games of Imperfect Information: Trembling Hand Perfection

- So we have moved from Nash equilibria in static games of perfect information, to Subgame Perfect equilibria (Selten) in dynamic games of perfect information where we required optimality in all subgames, to Bayesian equilibria (Kreps / Wilson) in static games of imperfect information where we required optimality at each information set, and finally we've added beliefs at those information sets in dynamic games of imperfect information.
- We have shown that those beliefs, particular among sequential equilibria (SE) are often implausible because they may involve one players believing that another may play a strictly dominated strategy.
- So we would like to do better.
- Consider G-18.1. We would like to eliminate unstable equilibria like the one at the top of the hill and only keep the stable equilibria, those that are invariant to preturbations.
- Preturbations will be one of two types: slight changes in strategies of players and slight changes in the players' payoffs.
- **Definition:** The behavioral strategy, π, is a trembling-hand perfect (THP) equilibrium if $\exists$ a sequence of totally mixed strategies, π_n, such that:
  - (1) $\pi_n \rightarrow \pi$, and
  - (2) for each player i and each element of the sequence, π_n, π^i is a best response to π_n^{-i}.

  So item (2) is the innovation in THP. A strategy is only THP is it is optimal for EVERY strategy along the way to the limit, not just at the limit as in a SE.
- If π is THP, it is robust to trembles.
- **Theorem:** For every finite extensive game, there exists at least one THP equilibrium.
- Note that THP $\Longrightarrow$ SE. So if we find a π that is THP, it is also a SE.
- Thus we can define THP equilibria as equilibria that are 1) sequentially rational, 2) consistent, and 3) (the innovation) optimal along the way to the limit.
- What the advantage of THP? **THP eliminates weakly dominated strategies (WDS).** The disadvantage is that it is harder to verify than the requirements of a SE.

- Recall the game in G-18.2 which we studied last lecture. Note that the game has two SE:
$$SE = \{(L, l); (p, 1-p)|p = 1 \text{ and } (R, r); (p, 1-p)|p = 0\},$$
Consistency is satisfied because *at the limit*, (R,r) is optimal. However, r is weakly dominated by l for player 2. Clearly the second SE above is not trembling-hand perfect because all along the way to the limit, l is the unique best response of player 2.

- **Remark:** Extensive form THP equilibria $\not\Rightarrow$ Normal form THP equilibria. See G-18.3. There are two THP equilibria, (Lr,r) and (Rr,r). However in the normal (strategic) form, we see that r and Rr are WDS. So that only leaves (Rr,r) as the THP equilibrium in the normal form.

- **Remark:** Normal form THP equilibia $\not\Rightarrow$ Extensive form THP equilibria. See G-18.4. (Lr) and (Rr) are both SE, but only (Lr) is THP because if the player is playing (Rr), she might "hit the wrong button" and tremble to (Rl). In the normal form, (Ll), (Lr), and (Rr) are all THP.

Strategic Stability

- So we would like to define a notion of strategic stability. When are equilibria stable and given some finite game, must there always exist some stable equilibria? First some definitions.

- **Definition:** Define the agent normal form of a game tree as the normal form of the game between agents, obtained by letting each information set be manned by a different agent where we give each agent the payoff of the player that he represents. Ie, if two agents represent player 2 (at player 2's two information sets), then these two agents both recieve whatever payoff would go to player 2 in the actual (non-agent) game.

- See G-18.5 for an extensive form, normal form, and agent normal form example.

- Behavioral strategies and the notion of a sequential equilibium are the same in agent form games except now the players are agents.

- **Definition:** An ϵ-perfect equilibrium of a normal form game is a totally mixed strategy vector, such that any pure strategy which is NOT a best reply has weight less than ϵ. So if you're mixing between all strategies, you can't place too much weight on strategies that aren't best responses.

- **Definition:** An ϵ-perfect-proper equilibrium of a normal form game is a totally mixed strategy vector, such that whenever some pure strategy, s_1 is a WORSE reply than some other pure strategy, s_2, then the weight on s_1 is smaller than ϵ times the weight on s_2. See G-18.6 for an example of how you assign these weights. Essentially this is just a THP equilibrium, but we have given priority to certain trembles (ie, the ones that you are more likely to tremble to!).

- **Definition:** Finally, we have that a perfect-proper equilibrium of a normal form game is a limit (as $\epsilon \to 0$) of the ϵ-perfect-proper equilibrium. And a perfect-proper equilibrium of a tree is a perfect proper-equilibrium of the tree's agent normal form.

- **Remark:** Recall that the problem with SE is that they are NOT invariant to inessential transformations of a game. Sequential rationality is never a problem because it holds for SE and:

  $$\text{Sequential Equilibria} \supseteq \text{Trembling-Hand Perfect Equilibria} \supseteq \epsilon\text{ -perfect-proper equilibria},$$

  so all of these satisfy SR as well.

- Note that the problem with SE is not that players play strictly dominated strategies, it's that players may believe that their opponent may play a strictly dominated strategy. Thus the problem is in the beliefs, not in the strategies.

Invariance

- Requirement 1: So one requirement that we would like to put on our set of equilibria is invariance. Strategically stable equilibia should be invariant to inessential tranformations of the game tree. Does this mean we sacrifice existence by imposing this?

- **Proposition 0:** A proper equilibrium of a normal form is sequential in any tree with that normal form. So if we are just interested in the SE of a game tree, we can just look at the proper equilibria of the normal form! Surprising result. By converting to the normal form, you lose the dynamic element of the game tree, but still we get the same equilibria. Intuition: successive elimination of weakly dominated strategies (as in proper-equilibria) is equivalent to backward induction to invoke sequential rationality of an extensive form game. Proof omitted.

- Requirement 2: We would also like strategically stable equilibria to depend only on the reduced normal form, determined by eliminating all pure strategies that are convex combinations of other pure strategies. See an example of this in Cramton's notes.

Admissibility and Iterated Dominance

- We would also like to find admissible strategies, ie those that are NOT weakly dominated.

- Both THP and proper equilibria satisfy admissibility (dominated strategies are never played with positive probability). However, as we have noted, sequential equilibria may not eliminate WDS so it may not satisfy admissibility.

- Iterated dominance means that we want strategies that survive iterated elimination no matter the order in which we eliminate them. See G-18.7 for an example where IEWDS will eliminate ALL equilibria depending on the order of elimination. So requiring iterated dominance may cost us existence.

Desirable Properties of a Stable Equilibrium

- Lets try to find characteristics of equilibria that do not involve balancing balls on mountain tops.
- We could consider a set-valued definition of stability. Desirable properties include:
  - (1) Existence: Every game has at least one solution.
  - (2) Connectedness: Equilibria are connected sets in the simplex of mixed strategies. See G-18.8. What would a non-connected set look like?
  - (3) Backwards Induction: usual definition.
  - (4) Invariance: Solution to the game is invariat to inessential transformations.
  - (5) Admissibility: Only involves non-dominated strategies.
  - (6) Iterated Dominance: A solution of a game, G, contains a solution of any game, G', obtained from G by deletion of a dominated strategy.
  - (7) Forward Induction: A solution must remain so after deletion of a strategy which is an inferior response for all strategies contained in the solution.
- **Proposition 1:** The set of NE of any finite game has finitely many connected components. At least one is such that for any equivalent game (ie, with the same reduced normal form) and for any preturbation of the normal form, there is a NE close to this component.
- So how should we define stability? Two definitions (included in the notes) that fall short are hyper-stability and full-stability. These both are fairly good though they both violate admissibility.
- **Definition:** So a better definition of stability that satisfies all but connectedness and backward induction (??) is defining equilibria as stable if they are robust to ANY totally mixed strategy. Note with THP equilibria, the equilibrium is only robust to SOME totally mixed strategy. Thus stability, as defined here is stronger than THP.

Forward versus Backward Induction

- Consider the game in G-18.9. The NE from the strategic form are (QM,R), (TB,L), and (TM,L). The unique SPNE is (TB,L) with a payoff of (2,0). Now consider the forward-induction argument. The only way that player 1 will initially play Q is if he is committing to playing M in his next move. Since player 2 observes this choice of Q, he knows that by playing L, he gets zero while if he plays R, he gets 1. So the forward induction equilibrium is (QM,R) with a payoff of (3,1). Forward and backward induction equilibria are often at odds with each other. Correct reasoning ?

19 Lecture 19: November 16, 2006

19.1 Signaling Games and the Intuitive Criterion

- Cho and Kreps (1987).
- Consider the standard signaling game with a sender (S) of type, $t \in T$, who observes his type (drawn by nature) and sends a message, $m \in M$, to a reciever (R) who observes the message, but not S's type, and then takes an action, $a \in A$.
- Assume T, M, and A are all finite so we get nice existence.
- Denote payoffs to S and R respectively as: $U^s(t, m, a)$ and $U^r(t, m, a)$.
- Everything but S's type is common knowledge.
- So what we're going to do now is identify unreasonable sequential equilibria (SE). SE required that strategies and beliefs were sequentially rational and consistent. It did not require that the beliefs were reasonable or plausible. The intuitive criterion will place some more structure on what beliefs are plausible.
- Consider a signaling game with two types, t1 and t2, sender messages L and R, and receiver actions u and d at each of his two information sets. Consider a potential equilibrium of pooling on R. This implies that the beliefs of the receiver following the R message are just the initial distribution from nature (common knowledge). Following an off the equilibrium path message (L), we impose beliefs on the receiver such that the sender will not want to deviate. What IF it was the case that a t1 sender is never better off by deviating to L but a t2 sender would actually like to deviate if the receiver would also play a different strategy (following L) than is specified in the pooling SE. The idea is that since the t1 sender would NEVER deviate, any strange (ie, off the equilibrium path) signal that the receiver sees, is more likely coming from a t2 sender. Hence the beliefs of the receiver should be such that the strange signal came from a t2 sender for sure. Thus, we would say that this "pooling on R" SE does NOT satisfy the intuitive criterion.
- The receiver is attempting to rationalize the deviation and knowing that one type of sender would never find it optimal to deviate, *intuitively* it makes sense for the receiver to think that the signal is coming from the other sender.

Beer and Quiche

- See G-19.1. Incumbent is either surly (prob 0.9) or a wimp (prob 0.1) and sends a signal by choosing his breakfast, beer or quiche. The entrant then reacts by choosing to duel or not to duel. The breakfast choice is just a signal to the entrant of the incumbent's type. The entrant might think only a surly incumbent would have beer for breakfast so I should not duel against him, but a wimpy incumbent might prefer quiche and so he'll choose to duel (eg, enter the market).

- It is easy to verify that separating on (Q,B) and separating on (B,Q) are not SE. In both cases, one of the incumbents will want to deviate.

- How about pooling equilibria? Consider pooling on beer. If (B,B) is played, then $p = 0.1$, and the entrant plays "not duel" following a beer signal. The surly sender gets his maximum payoff of 3 after drinking his beer and not dueling so he has no reason to deviate. The wimpy sender gets 2 from having a beer and not dueling, so he will stick to his strategy as long as:

$$E_{entrant}[Duel|Quiche] \geq E_{entrant}[Not\ Duel|Quiche]$$

$$2q \geq 1 \Rightarrow q \geq \frac{1}{2}.$$

So we have a SE at $\{(B, B), (Not\ Duel, Duel), (p, 1-p), (q, 1-q)|p = 0.1, q \geq 0.5\}$.

- How about pooling on quiche? (Q,Q) implies $q = 0.1$. Therefore following a quiche signal, the entrant plays not duel. The wimp gets his highest payoff so has no reason to deviate. The surly incumbent gets 2 from playing the SE, but could deviate and get either 1 or 3. So the surly sender requires:

$$E_{entrant}[Duel|Beer] \geq E_{entrant}[Not\ Duel|Beer]$$

$$2p \geq 1 \Rightarrow p \geq \frac{1}{2}.$$

So we have a SE at $\{(Q, Q), (Duel, Not\ Duel), (p, 1-p), (q, 1-q)|p \geq 0.5, q = 0.1\}$.

- So which one violates the Intuitive criterion?
  - Consider pooling on beer. The surly sender is getting his highest payoff so would never want to deviate. The wimpy sender is getting 2 in equilibrium but would get 3 if he deviated and the receiver didn't duel. So the wimpy sender is more likely to try to deviate. This is consistent with the receiver's beliefs following a quiche signal, $q \geq 0.5$. The receiver has to be fairly certain that the deviation came from a wimpy sender, and indeed it is more likely that it did! So the intuitive criterion is satisfied.
  - Consider pooling on quiche. The wimpy sender is getting his highest payoff so would never want to deviate. The surly sender is getting 2 in equilibrium but would get 3 if he deviated and the receiver didn't duel. So the surly sender is more likely to try to deviate. This is *inconsistent* with the receiver's beliefs following a beer signal, $p \geq 0.5$. The receiver has to be fairly certain that the deviation came from a wimpy sender, but in fact deviations are more likely to come from a surly sender! So the intuitive criterion is violated.

Formalizing the Intuitive Criterion

- After hearing a message, $m \in M$, R's beliefs are $\mu(t|m)$.

- Denote the set of best responses of the R as:

$$BR(\mu, m) \equiv argmax_{a \in A} \sum_{t \in T} \mu(t|m) U^r(t, m, a).$$

- Define a subset, $I \subset T$, and let $BR(I, m)$ be the set of best responses for R to beliefs concentrated on I:

$$BR(I, m) \equiv \bigcup_{\{\mu : \mu(I)=1\}} BR(\mu, m).$$

- Then given strategies $\pi^s(m|t)$ and $\pi^r(a|m)$, the equilibrium payoff to a sender of type t is:

$$U^*(t) \equiv \sum_{a \in A} \sum_{m \in M} \pi^r(a|m) \pi^s(m|t) U^s(t, m, a).$$

- Then a (sequential) equilibrium FAILS to satisfy the intuitive criterion if there exist:
  - (1) an unsent message, $m' \in M$ (so this eliminates considering the criterion for separating equilibria)
  - (2) a subset $J \in T$ of types of senders where the deviation is a bad thing (always)
  - (3) a type $t' \in \neg J$ such that $\forall\ a \in BR(\neg J, m')$, $U^*(t') < U^s(t', m', a)$.

  So the t' guy is like the surly sender in the pooling on quiche example above. The second item is important because the receiver must have an inclination that the out of equilibrium signal did not come from types in J.

Job Market Signaling

- Spence.
- Workers have privately observed productivity, $t \in \{H, L\}$, and choose a level of education, $e \in [0, \infty)$.
- The market observes e and then offers a wage, $w \in [0, \infty)$.
- Worker's preferences:

$$U^s(t, e, w) = w - c(t, e).$$

  $c(t, e)$ is the psychic cost for a worker of type t to acquire education e.

- There will be a lot of equilibria in the this problem but we'll eliminate some via sequential rationality and others via elimination of weakly dominated strategies. We'll finally get a unique solution by imposing the intuitive criterion.
- More next time.

20 Lecture 20: November 21, 2006

20.1 More on Job Market Signaling (Spence)

- We assume that since the L type worker has a higher marginal cost of education, the indifference curves of the H and L types satisfy the single cross property (SCP) as in G-20.1.
- We'll show that this is all we need to get a unique separating equilibrium (where H types successfully signal and get a high wage) if we impose some equilibrium refinements.
- Let a worker of ability t and education e produce output, $y(t, e)$, such that:

$$y(H, e) > y(L, e) \ \forall \ e$$

$$\frac{\partial y(t, e)}{\partial e} > 0 \ \forall \ t, e.$$

- Spence argues (assumes) that with a competitive market, the wage will satisfy: $w(e) = y(t, e)$, ie zero profits!
- We have to consider three types of equilibria: 1) Separating - two types choose two different levels of education; 2) Pooling - two types choose the same level of education; and 3) Hybrid - at least one type randomizes between pooling with the other type and distinguishing himself.
- See G-20.2 for the separating case. L type gets e_L education and H type gets e_H education. Neither type wants to deviate (ie, choose the other's education) so we have an (infinite number of) NE.
- See G-20.3 for a pooling equilibrium. Firms pay the ex-ante expectation of productivity:

$$w_{pool} = py(H, e) + (1 - p)y(L, e).$$

  Note the wage schedule just has to be sufficiently kinked to get this equilibrium. Again there are an infinite number of them but the wage always falls on the dotted line.
- See G-20.4 for the "envy case". This is where the low type would like to pretend to be a high type. If there is no envy, we'll just get a nice separating equilibrium.
- So where to go next? We have all these NE! We'll now use sequential rationality, elimination of WDS, and the intuitive criterion to arrive at a unique (separating) equilibrium.
- In a SE, the market must have beliefs, $\mu(t, e)$, and since the wage is the expected productivity of these beliefs:

$$y(L, e) \leq w(e) \leq y(H, e).$$

- We can also apply sequential rationality to show that the worst the low and high types can do are the indifference curves illustrated in G-20.5. So in any equilibrium, they must get at least that much utility. Start with the productivity curves, draw the tangent indifference curve for the low type and then draw in the indifference curve of the high type intersecting I_L and $y(H, e)$ simultaneously.

- After finally applying the intuitive criterion, we are left with a UNIQUE SEPARATING EQUILIBRIUM, where the low type chooses $\underline{e}$ and the high type chooses $\overline{e}$. Wages for L and H are thus $y(L, \underline{e})$ and $y(H, \overline{e})$ respectively.

20.2 Signaling Games and the Intuitive Criterion

- Recall the definition of the intuitive criterion. Given an unsent message, m', and set of types, J, REJECT any SE satisfying the following conditions

  - (1) $\forall\ t \in J$, and $\forall\ a \in BR(T, m')$, $U^*(t) > U^s(t, m', a)$
  - (2) $\exists\ t' \in T\neg J$ such that $\forall\ a \in BR(T\neg J, m'), U^*(t') < U^s(t', m', a)$.

  So we have that the receiver places zero weight on those types in J having played m'. All the weight should be on the types in $T\neg J$. We also need some type that definitely would like to deviate given the concentrated beliefs and given ALL best responses by the receiver when his beliefs are concentrated on $T\neg J$.

- What if we reverse the order on the second part of the definition: Given an unsent message, m', and set of types, J, REJECT any SE satisfying the following conditions

  - (1) $\forall\ t \in J$, and $\forall\ a \in BR(T, m')$, $U^*(t) > U^s(t, m', a)$
  - (2) $\forall\ a \in BR(T\neg J, m')$, $\exists\ t' \in T\neg J$ such that: $U^*(t') < U^s(t', m', a)$.

  This is called "extended intuitive criterion" or "introspective consistency." We now throw out an equilibria if there are ANY beliefs that would make t' want to deviate. So this is a bit weaker than the intuitive criterion in that it will throw out more equilibria. Is it realistic? - maybe not. It's just a way to judge the strength of a candidate equilibrium.

- An equilibrium that satisfies the extended intuitive criterion is a forward induction equilibrium. See PS 7 for a SE that satisfies the intuitive criterion but is rejected by the extended version.

- **Proposition:** A strategy profile, π, is a forward induction equilibrium iff π is a SE satisfying the extended intuitive criterion.

- See G-20.6 for an example of using the extended intuitive criterion. Note there is a SE with player 1 playing A and player 2 playing r (if player 2 has high enough beliefs that Player 1 played R). However, if player 2 observes player 1 not playing A (ie, player 2 gets to move), he sees that player 1 passed up a payoff of 1 by playing A. Thus player 1 must have choosen L and hopes to get 2 versus having chosen R and getting 0. So any beliefs that put positive weight on R do not satisfy the extended intuitive criterion.

- So is the extended intuitive criterion enough? Sometimes there are NO bad deviations and all SE are forward induction equilibria. We must resort to ad hoc arguments to capture forward induction.

20.3 Other Refinements

- **Definition:** Dominance: eliminate a type t if m is sent and m is dominated by m' for type t.
- (Equilibrium) Dominance turns out to be too weak of a requirement. So we have other refinements involving the sets of best responses by the receiver such that the sender will want to (sometimes strictly) deviate from the equilibrium. These are somewhat strange and completely theoretical constructs. See notes for the D1 refinement, the D2 refinement, the Divinity refinement, the Universal Divinity refinement, and finally the Never Weak Response refinement.
- The nice thing about applying the D1 refinement to the Spence signaling game is that with more than two types, we can get an unique equilibrium only after refining with D1. The intuitive criterion is enough to get uniqueness in the two type case.
- So what's the order of the refinements. We have:
  - For general games:

$$NE \supset SPNE \supset SE \supset PE \supset ProperE \supset IC \supset EIC$$

  - And for signaling games:

$$EIC \supset Div \supset D1 \supset D2 \supset UniDiv \supset NWR \supset Stable \neq \emptyset.$$

  So it's nice that no matter how many refinements we make, we still end up with at least one equilibrium.

20.4 Cheap Talk Games

- Cheap talk games involve the sender sending a message that does not affect his payoff. (Education signaling is costly so does not qualify as cheap talk).
- We'll show that there always exists some equilibrium where the signal is just ignored and the receiver always does the same (pooling) action.
- However, depending on how aligned the sender and receiver's preferences are, the message can be useful.
- Timing:
  - Suppose S observes the state of the world, $t \in [0, 1]$.

- S sends a message, $m \in M$, to R.
- R takes an action, $a \in (-\infty, \infty)$.
- R does not observe t but knows $F(t) \sim [0,1]$.
- Payoffs are $U^s(a,t,b)$ and $U^r(a,t)$, where b is a common knowledge measure of how nearly the agents' interests coincide. Note the payoffs do NOT depend on m, hence cheap talk.

- See G-20.7 for an example of payoff functions. If $b=0$, their interests are perfectly aligned so the sender should just reveal the state of the world, ie $m=t$, and both will be at their optimum.

- All we assume is that the payoff functions are concave with a unique maximum.

- A strategy is a complete plan of action so the sender's strategy is a mapping from types to messages. Let $q(m|t)$ be the density of sender's choice of message when the state is t. Let $a(m)$ specify the action of the receiver given the message, m.

- The receiver updates his beliefs using Bayes rule:

$$Pr(t|m) = \frac{q(m|t) * f(t)}{\int_0^1 q(m|x)f(x)dx}.$$

- So strategies, $\{q(m|t), a(m)\}$, form a BE if the usual Nash conditions hold:

$$m^* = argmax_{m \in M} \ \ U^s(a(m), t, b),$$

$$a(m) = argmax_{a \in A} \ \ \int_0^1 U^r(a,t)Pr(t|m)dt.$$

- **Lemma:** Suppose b is such that no $t \in [0,1]$ satisifes $a^s(t,b) = a^r(t)$ (so the preferences are not aligned). Then $\exists \ \epsilon > 0$ such that if u and v are actions induced in equilibrium, then:

$$|u - v| \geq \epsilon.$$

  Further, the set of actions induced in equlibrium is finite.

- So we have a partioned equilibrium where a finite number of messages are sent and a finite number of actions are then taken. If b is large, ie preferences are dissimilar, then we'll get fewer and few equilibria with the limit being the unique pooling equilibrium where the receiver ignores the message and always chooses the same action (based on the prior probabilities). If $b \to 0$, we get more and more equilibria (still finite) because preferences are more aligned. There is a graph in the notes which might become relevant.

- Next we'll move to auction theory but there is more in Cramton's notes on signaling games.

21 Lecture 21: November 28, 2006

21.1 Single Item Auctions

- Follows McAfee and McMillan (JEL 1987) and Milgrom and Webber (Econometrica, 1982).
- Auctions are a nice application of game theory because the rules of the game are made explicit and they provide a nice (sometimes efficient) mechanism for allocating goods and services for some agreed upon price.
- Nice statistic: 450% of the world's GNP is traded each year by auction.
- Four common types of auctions:
  - (1) English: ascending dynamic auction where highest bidder wins and pays his bid.
  - (2) Dutch: descending dynamic auction where the price falls continuously until someone jumps in and accepts. Winner pays stopping price.
  - (3) Second Price: sealed bid auction where bids are submitted secretly and highest bidder wins and pays second highest bid.
  - (4) First Price: sealed bid auction where bids are submitted secretly and the highest bidder wins and pays his bid.
- The Dynamic (Oral) auctions are becoming more popular with the growth of the internet auction sites.
- Which auction is best for allocating a good depends crucially on the setting.
- Modeling Information. There are three common assumptions on information:
  - (1) Independent Private Values: $v_i \sim F_i$ independently of v_j for $j \neq i$. Appropriate if bidders are heterogeneous in their preferences over the object's attributes.
  - (2) Common Values: $e_i = v + \epsilon_i$, $\epsilon_i \sim F_i$ with mean 0. Appropriate if bidders have homogeneous preferences but have different signals about the true value.
  - (3) Affiliated Values: $u_i(v_i, e_i)$, so my value depends on a private value as well as a common component (eg state of the world). Appropriate if a bidder has private information (v_i) that is positively correlated with the bidder's value of the object.

Some results in common value auctions

- Unlike IPV auctions where bidders bid truthfully in SPSB and shade their bids in FPSB formats, when we have common values, bidders need to worry about the "winner's curse." Winning may provide you the satisfaction of attaining the object but the event also reveals to the winner that they bid more than everyone else and probably bid above the true common value.

- The key thing in common value auctions is to bid optimally *conditional on winning the auction!* We want to avoid the winners curse. Your value of the object conditional on having won the auction is much less than your initial estimate of the object's worth.

- The winner's estimate in a common value auction is always substantially positively biased.

- So what should bidder's optimally do?

  - (1) SPSB: shade your bid below your value depending on your estimate of how noisy your signal of the common value is.
  - (2) FPSB: shade as in IPV setting but then shade some more to avoid the winner's curse.

- How does the competitive effect influence all of this? Usually as we increase the number of bidders, in a FPSB auction for example, you have to shade less because there is less of a chance that you can win and capture a small surplus. With common values and more than 3 bidders, the winner's curse effect dominates the competitive effect so as the number of bidder's increases, you should shade more.

Benchmark Model: IPV, Symmetric, Risk Neutral Bidders

- Assume buyers have values, $v_1, \ldots, v_n \sim F$ on $[0, \infty)$.

- Seller has commonly known value, v_0.

- In a SPSB or English setting, bid your value or up to your true value. The winner gets $v_{(1)} - v_{(2)}$ ex-post and expects in the interim state to get:

$$E_{v_{(1)}}[v_{(1)} - v_{(2)}] = E\left(\frac{1 - F(v_{(1)})}{f(v_{(1)})}\right).$$

  Verify this.

- In the FPSB (and equivalently the Dutch Auction), we can write the bidder's expected profit as:

$$\pi(v, b(v)) = (v - b(v)) * Pr(Win|b(v)).$$

- By the envelope theorem:

$$\begin{aligned}\frac{d\pi}{dv} &= \underbrace{\frac{\partial\pi}{\partial b}}_{=0}\frac{\partial b}{\partial v} + \frac{\partial\pi}{\partial v} \\ &= \frac{\partial\pi}{\partial v} \\ &= Pr(Win|b(v)) \\ &= Pr(b(v) \text{ is the highest bid}) \\ &= Pr(v \text{ is the highest value}) \\ &= [F(v)]^{n-1}\end{aligned}$$

  Note in this derivation we used the symmetry of the bid functions and assumed that $b'(v) \geq 0$.

- So since $\frac{d\pi}{dv} = [F(v)]^{n-1}$, by the FTC,

$$\pi(v) = \underbrace{\pi(0)}_{const\ of\ integration=0} + \int_0^v F(x)^{n-1}dx = \int_0^v F(x)^{n-1}dx. \quad (*)$$

- Substituting (*) into the expected profit:

$$\begin{aligned}\pi(v, b(v)) &= (v - b(v)) * Pr(Win|b(v)) \\ \int_0^v F(x)^{n-1}dx &= (v - b(v)) * Pr(Win|b(v)) \\ \int_0^v F(x)^{n-1}dx &= (v - b(v)) * F(v)^{n-1} \\ F(v)^{-(n-1)}\int_0^v F(x)^{n-1}dx &= v - b(v) \\ b(v) &= v - F(v)^{-(n-1)}\int_0^v F(x)^{n-1}dx\end{aligned}$$

- Note that this is, in general, a non-linear bidding function (which is increasing in v).

If $v \sim U[0,1]$, then $F(v) = v$, so:

$$\begin{aligned} b(v) &= v - F(v)^{-(n-1)} \int_0^v F(x)^{n-1} dx \\ &= v - v^{-(n-1)} \int_0^v x^{n-1} dx \\ &= v - v^{-(n-1)} \left[\frac{x^n}{n} \right]_0^v \\ &= v - v^{-(n-1)} \frac{v^n}{n} \\ &= v - \frac{1}{n} v \\ &= \frac{n-1}{n} v \end{aligned}$$

And as $n \to \infty$, $b(v) \to v$, so in the limit, the seller is able to extract the full surplus.

- So in words, "the bidder bids the expected value of the second highest value GIVEN that the bidder has the highest value."

22 Lecture 22: November 30, 2006

22.1 More on Single Item Auctions

- **Theorem:** Revenue Equivalence. The seller's expected revenue in an English, Dutch, FPSB, and SPSB auction are all the SAME if we assume: IPV, symmetric bidding functions, and risk neutrality. This follows because in all these auctions,

$$\frac{d\pi}{dv} = Pr(win),\ \pi(0) = 0.$$

  So each auction assigns the item to the same (highest value) bidder. Thus, if different auction formats assign the good to the same bidder, they must yield the same expected revenue to the seller.

- However, usually these assumptions do NOT hold! Running the different formats on the same item will often lead to different revenue for the seller because at least one of the assumptions is violated.

- Other assumptions, such as the format not affecting the bidder's participation in the auction and no-resale of the item, are also often violated.

Royalty Schemes

- What if in addition to the winning bid, the seller also received a portion of the (ex-post) value of the item. For example, a linear royalty scheme:

$$p = b + r\tilde{v}.$$

  So the price the seller receives is the bid, b, plus a portion, r, of the realized value, $\tilde{v}$.

- These are common in textbook contracts and oil rights leases by governments. For example, the US auctions off lease rights to explore areas for oil fields. Bidders submit simultaneous sealed bids and then also must pay a 1/6 royalty on the oil revenues to the government once the oil is extracted. This is just a distortionary tax and may effect the incentives by the firm to extract efficiently.

- Other countries avoid the taxation distortion by auctioning off a joint ownership contract. Ie, the government pays half the cost and receives half the profits from any successful extraction. See Venezuela and Libya.

- One problem with these types of contracts is the fear of expropriation. After the oil field is explored and, say, deemed profitable, the government can come in and instate some large profit windfall tax on the firms. The US does this often and can affect the bidding strategies of firms in the initial auction.

Risk Aversion

- If sellers are not risk neutral, then in an ascending English auction, they continue to bid the same way (stay in up to their valuation). However, in a FPSB auction where they would normally shade their bids by some amount, they now shade less because doing so increases their probability of winning and reduces the amount that they win if they are the highest bidder. This reduces the risk so a risk averse bidder prefers less shading.

- So you might conclude that if bidders are risk averse (as they often are), then you should always prefer (from a seller's perspective) a FPSB to an English auction. But we are assuming IPV ! With a common value component, bidders want to shade MORE in a FPSB auction to avoid the winner's curse. So now it's not so clear which one will yield a higher expected revenue.

Collusion Among Bidders

- What if, as a seller, you suspect that bidders are colluding and essentially submitting a single bid. Then the seller should set a reserve price. If we assume IPV and symmetric bidders, the reserve price, r, satisfies:

$$v_0 = r - \frac{1 - F(r)^n}{nF(r)^{n-1}f(r)},$$

where v_0 is the seller's valuation.

Many Items

- Suppose the seller can produce an unlimited amount of a good at marginal cost, c. The optimal selling mechanism would NOT be an auction. Instead the seller should just offer essentially the monopoly price as a take it or leave it offer. The problem with this is the Coase Conjecture which says that the seller needs to be able to commit to NEVER lower his price at any time in the future. This is difficult in practice though most sellers can do it up to a point. If the seller cannot commit at all, the only equilibrium is marginal cost pricing.

All Pay Auctions

- Assuming IPV and $b'(v) > 0$, the highest value guy still wins the object so the expected revenue to the seller must be the same. Bidders all shade substantially. Same assignment of the object $\Longrightarrow$ Same expected revenue to the seller.

Costly to Enter a Bid

- In auctions where it is costly to enter a bid, ie a firm has to do some research to determine its private valuation, it might be optimal (in a sequential setting) for one firm to make the investment and then enter a fairly high bid in order to discourage others from making the same investment. This is particularly relevant to firm takeover games where many potential buyers are competing for a hostile takeover of a firm.

Correlated Values

- What if instead of the IPV assumption, we assume bidder's valuations, though private, are correlated.
- Milgrom and Webber (Econometrica 1982), show that the reason English auctions are so common is that they reveal the most information in the bidding process when bidder's have affiliated values, ie their estimates of the object's worth are positively correlated.
- The main result when bidder's have affiliated information about the value of an object is that the expected revenue to the seller satisfies:

$$English \geq SPSB \geq FPSB = Dutch.$$

- The intuition is that "the equilibrium bid function depends on everyon's information. The more affiliated information you condition on, the higher the bid." If your information is really good (zero variance of your signal say), then you can just bid your valuation.

23 Lecture 23: December 5, 2006

23.1 Multi Item Auctions

- Examples of multi-item auctions: treasury bills, spectrum and electric power auctions.
- There are several ways to auction many similar items. We highlight several of them below.

Sealed Bid: Pay-as-Bid Auction

- This is also called a discriminatory auction. The treasury used to use this format to auction t-bills.
- Bidders submit their demand schedule and all bids above, p_0, the clearing price win and pay their bid.
- See G-23.1. Note that the demand is the aggregate demand of all bidders.

Sealed Bid: Uniform-Price Auction

- Same idea as pay-as-bid, but all bidders with a bid above p_0 win, but they all pay p_0.
- See G-23.2.

Sealed Bid: Vickrey Auction

- See G-23.3. Bidders still submit demand schedules but now price is determined using the bidder's individual demand schedule along with the "Residual Supply" of all other bidders, ie $Q_s - \sum_{j \neq i} Q_j(p)$.
- The residual supply curve represents the (social) opportunity cost to everyone else of allocating the item to bidder i. Bidder i pays this opportunity cost up to the clearing price.
- The nice thing about vickrey pricing is that it is ALWAYS efficient because it gives bidder i 100% of the incremental gains from trade that bidder i brings to the table. This results in ideal incentives so each bidder has a weakly dominate strategy of bidding truthfully.
- For example, consider one item and two bidders with valuations $v_1 = 90$ and $v_2 = 100$. Bidder 2 will win the auction and should pay $p = 90$, the opportunity cost (to society) of NOT allocating the item to bidder 1.
- Consider a simple example of 3 bidders and 3 items in a vickrey format auction. Marginal values are:

	Bidder A	Bidder B	Bidder C
First Item	10	8	4
Second Item	6	7	2
Third Item	3	5	1

The first thing to notice is that since Vickrey is always efficient, the 3 items will go to the three bidders that value them the most. So bidder A will get 1, and bidder B will get 2. What should they pay? Use the following rule: "The winning bidder pays the k^{th} highest LOSING bid of the OTHER bidders on the k^{th} item won." So for bidder A, the highest marginal value of either bidder B or C that is losing is 5. Thus bidder A pays 5 for his one item. The two highest losing bids by bidders A and C are 6 and 4, so bidder B pays $6 + 4 = 10$ for his two items.

- One result that some people don't like about vickrey auctions is that they are NOT anonymous! Two bidders winning the exact same thing may pay different prices.

Bidding Strategies in Sealed Bid Multi-item Auctions

- See G-23.4. Recall that bidders optimally bid truthfully in a vickrey auction. What about the others?
- Pay-as-bid auctions involve bidders trying to "guess the clearing price". They want to submit all their demand at the clearing price to avoid paying a higher amount for the some of the units. In practice this is exactly what happens. Most bids are submitted very close to the ex-post clearing price. This allowed for a "short-squeeze" in treasury bill auctions and prompted the government to switch to a uniform price auction for t-bills in 1998.
- In uniform-price auctions, bidders bid truthfully for the first unit and then shade more and more for each additional unit.

Ascending Bid: Standard Auction

- See G-23.5. We start the clock price very low and slowly raise the price if there is excess demand. Stop when the market clears. All bidders pay their bids.

Ascending Bid: Ausubel Auction

- See G-23.6. Same clock mechanism, but compute the residual supply as in Vickrey and determine clinching prices off the residual supply curve.

Other Types and Remarks

- There are of course other types of mechanisms for multi-item auctions. The FCC uses simultaneous ascending bids for the spectrum auctions and a sequential dutch auction are used for tulips in the Holland.
- How do the different multi-item auction formats compare in terms of efficiency and revenue generation ?
- The revenue results are ambiguous and depend crucially on the specifics of the environment.
- As for efficiency, we have the following:
  - (1) Uniform-price and standard ascending-bid auctions: inefficient due to demand reduction.
  - (2) Pay-as-bid: inefficient due to different shading.
  - (3) Vickrey: Efficient in private value settings, but not with affiliated information.
  - (4) Ausubel: Efficient in private value and affiliated value settings.

So Larry wins.

24 Lecture 24: December 7, 2006

24.1 More on Multi Item Auctions - Identical Objects Model

- Suppose a seller has 1 DIVISIBLE good that he values at zero.
- There are n bidders that each can consume $q_i \in [0, \lambda_i]$.
- t_i is bidder i's type which is private information and $t_i \sim F_i$.
- Types are independent and the marginal value to bidder i from winning a proportion q_i of the good is:
$$v_i(t, q_i).$$
Note that v_i depends on EVERYONE's type. So there is a common value (affiliated value) element to the model.
- Bidder i's payoff if he gets q_i and pays x_i is:
$$\int_0^{q_i} v_i(t, s)ds - x_i.$$
- We have some assumptions on the marginal value:
  - (1) Value Monotonicity: non-negative, increasing in t_i, weakly increasing in t_j, and weakly decreasing q_i.
  - (2) Value Regularity: for all $i, j, q_i, q_j, t_{-i}, t_i' > t_i$:
  $$v_i(t_i, t_{-i}, q_i) > v_j(t_i, t_{-i}, q_j) \Rightarrow v_i(t_i', t_{-i}, q_i) > v_j(t_i', t_{-i}, q_j).$$
  So this means that if we have two levels of private information for bidder i with t_i' being the better signal, then if bidder i does better than bidder j at t_i, bidder i also does better than j at t_i'.
- Define bidder i's marginal revenue, ie the marginal revenue the SELLER gets from awarding the additional quantity to bidder i:
$$MR_i(t, q_i) = v_i(t, q_i) - \frac{1 - F_i(t_i)}{f_i(t_i)} \frac{\partial v_i(t, q_i)}{\partial t_i}.$$
Note in the IPV setting, $v_i = t_i$, so the last term was just 1. We called this the "virtual valuation" in the Myerson paper.
- **Theorem:** Revenue Equivalence (Again). In any equilibrium of any auction game in which the lowest-type bidders receive an expected payoff of zero, the seller's expected revenue equals:
$$E[Rev] = E_t \left[\sum_{i=1}^{n} \int_0^{q_i(t)} MR_i(t, s)ds \right].$$

So once again, the revenue to the seller ONLY depends on the assignment of the good. This is a pretty strong result only making some fairly reasonable assumptions above. And this holds even with a common value auction.

24.2 Practical Auction Design

- In most realistic settings, we have an auction for many goods and there may be complementarities between the goods.
- If there are say 20 items to be auctioned, then there are $2^{20} = 1,048,576$ potential allocations of the items among the bidders. Define the set of potential allocation as S.
- Let $v_i(s_i)$ be bidder i's valuation with his assignment, s_i.
- What would we do in a Vickrey auction?
  - (1) First find the efficiency allocation. Thus solve:

$$s^* \in arg\ max_{s \in S} \sum_i v_i(s_i).$$

    This alone may be hard to do.
  - (2) Find prices for each winner to pay based on the social opportunity cost (SOC) of their winnings. Ie, bidder i pays:

$$SOC_i = max_{s_{-i} \in S_{-i}} \underbrace{\sum_{j \neq i} v_j(s_{-i})}_{what\ others\ could\ get} - \underbrace{\sum_{j \neq i} v_j(s_j^*)}_{what\ others\ actually\ get}.$$

So the first term is the value that the other bidders would get if they received bidder i's winnings instead of him. The second term is what the other bidders are actually getting in the s^* allocation.

Example 1

- Consider the following marginal valuation table for an auction with two bidders and two items.

	Item A	Item B	Bundle AB
Bidder 1	0	0	10
Bidder 2	6	6	6

So bidder 1 only want both items (complementarities) and bidder 2 only cares about getting at least one.

- The vickrey outcome would be for bidder 1 to get both items since he places the highest marginal value on them together. No other allocation achieves as high of a value. Bidder 1 should pay the SOC, which in this case is 6. (Not 12).

Example 2

- Consider the following marginal valuation table for an auction with two bidders and two items.

	Item A	Item B	Bundle AB
Bidder 1	8	8	12
Bidder 2	5	5	10

So bidder 2 has a constant marginal valuation.

- The vickrey outcome would be for bidder 1 to get one item and bidder 2 to get the other. This results in total value of 13.
- Bidder 1 should pay $10 - 5 = 5$ for his item applying the formula above. Bidder 2 should pay $12 - 8 = 4$ for his item.
- The problem with all of this is that bidders often do not know their marginal valuations so an economic advisor to a bidder must try to recover it.

24.3 Strategic Issues in Bidding Strategies

- It is well known that bidders should seek to avoid the winner's curse in common value auctions. What are some of the other pitfalls?
- Consider the following marginal valuation table:

	Item A	Item B	Bundle AB
Bidder 1	100	100	200
Bidder 2	90	90	180

Both bidders have constant marginal valuation but bidder 1 is stronger.

- Consider an ascending style format where the price keeps rising until there is no excess demand. If both bidders stay in until the price hits their valuation, we end up at a price of 90 where bidder 2 drops out and bidder 1 gets both items at 90, so makes 20 in profit. However, consider what happens if both bidders reduce their demand to 1 immediately, ie when price is zero. The auction ends, each attains one item and make profits of 100 and 90 respectively. Clearly much better for everyone involved (besides the seller).

- So this illustrates demand reduction. Bidders may have an incentive to NOT be greedy because doing so only shrinks the size of the pie and bidders end up with LESS than they would get by reducing their demand earlier on in the auction.
- We'll continue with other strategic issues in bidding behavior next time.

25 Lecture 25: December 12, 2006

25.1 More on Practical Auction Design

Demand Reduction

- To counter this behavior by bidders, a seller can either introduce a reserve price or attempt to increase the intensity of competition.

Inefficiency Theorem

- In any equilibrium of a uniform-price auction, with positive probability, the objects are won by bidders other than those with the highest values.
- So the winning bidder influences the price with positive probability.
- This creates an incentive to shade below your value, and this incentive grows with more and more units.
- Differential shading implies inefficiency!
- The exceptions to the theorem include pure common value auctions (where everyone has the same value so all allocations are efficient) and single unit auctions (where bidders bid truthfully for the first unit).

Exposure Problem

- This reflects the idea that a bidder may have demand complementarities. Eg, in an auction for two items, the bidder may only value them together and gets no value from just attaining one.
- If the bidder cannot bid on a bundle of items, you might end up being exposed to winning only one object which is a pure loss.
- Thus the equilibrium in auctions like this, bidders either "go for it" and bid high to guarantee they win both items, or they simply stay out. The outcome is often inefficient.
- We can eliminate this problem with package bids.

Asymmetries

- If one bidder has even a small advantage (ie it is common knowledge that they will benefit more from winning), then that bidder can always win the auction.

Other Considerations

- We also have several other strategic issues in designing a practical auction:
  - Collusion: we want to stifle tacit cooperation between bidders
  - Complexity: we want to minimize the amount of "gaming" that bidders are required to do when choosing their bids
  - Determining objectives: we need to determine what the seller wants to accomplish (revenues versus efficiency) and what the bidder wants to do (what is his objective function?). How important is competition in the post-auction market?
- See Cramton slides for more issues in practical auction design.