Lecture 7

Today's agenda

- The Extensive Form Representation of a Game
- Solving an Extensive Form Game
- Reminder: A Static Cournot Duopoly
- Infinitely Repeated Games
- The Folk Theorem
- The Rotemberg and Saloner Model

Industrial Economics (EC5020), Spring 2010, Sotiris Georganas, March 1, 2010

Aims

- be familiar with the concept of a "super game" or a "repeated game" and be able to solve for an equilibrium of such a game.
- understand the requirements for self-enforcing collusion in repeated games.
- understand why price wars are more likely to be observed in good times

Tirole, Ch. 6, pp. 239-253 (except 6.2, "Static Approaches to Dynamic Price Competition" and 6.4, "Price Rigidities").

The Extensive Form Representation of a Game I

Information provided by an extensive form game:

1. The number of players.
2. When each player can take an action.
3. What actions are available for a player when it's her turn to move.
4. Each player's payoff for all possible outcomes of the game.

The Extensive Form Representation of a Game II

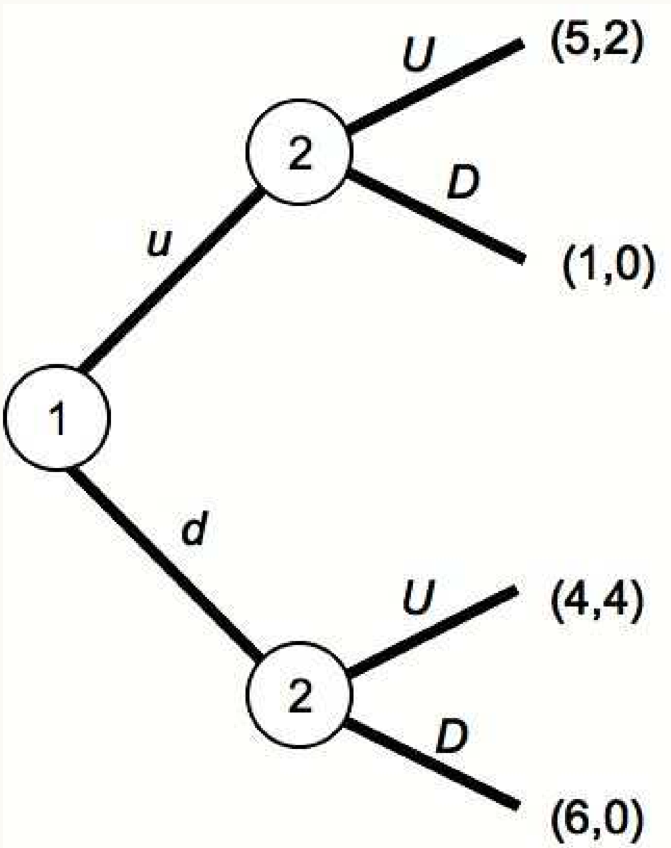

The Extensive Form Representation of a Game III

- Example — see also figure.
  - Two players: 1 and 2.
  - Player 1 chooses "u" or "d".
  - Player 2: observes 1's choice; and then chooses "U" or "D".
  - The players get payoffs as indicated in the figure.
    - Convention: Player 1's payoff is the first number, and Player 2's is the second.
- Player 1's strategy set: $S_1 = \{u, d\}$.
- Player 2's strategy set: $S_2 = \{UU, UD, DU, DD\}$.
  - 2's strategy is a *function*: what to do if 1 chose u, and what to do if 1 chose d.
  - So, for example, UD means: "choose U if 1 chose u, and choose D if 1 chose d".

Solving an Extensive Form Game I

- Identify the (pure strategy) NE!
  - (u,UU) is one.
  - (d,DU) is another.
- How can (d,DU) be a NE?
  - Effectively, by playing the strategy DU, Player 2 threatens Player 1: "If you play u, then I'll play D, in which case you will get the payoff 1 — the worst payoff you can get in this game."
  - Player 1 believes that 2 would carry out this threat, and accordingly chooses d.
- Problem: perhaps not a plausible equilibrium.
  - Player 2's threat is *not credible*: if Player 1 actually played u, Player 2 would be better off playing U rather than D.
  - If Player 1 acknowledged that the threat is not credible, she should expect the response U if playing u, in which case playing u would indeed be optimal — the equilibrium breaks down.

Solving an Extensive Form Game II

- Put differently:
  - In a NE, no player has an incentive to deviate unilaterally *along the equilibrium path*.
  - But, the NE concept does not require the players to behave optimally *off the equilibrium path*.
- How to rule out NE that rely on noncredible threats: **subgame perfect Nash equilibrium (SPNE)**.
- First: a **subgame** is defined as a single node and all the branches and nodes and payoffs that flow from that node.
  - The full game counts as a subgame.

We can now define a subgame perfect Nash equilibrium (Selten, 1965):

- **Subgame perfect Nash equilibrium (SPNE)**: a strategy profile that is a NE in every subgame.

How to identify the SPNE in a finite horizon game:

Solving an Extensive Form Game III

- In a game with a finite horizon, we can find the subgame perfect Nash equilibria by **backward induction**:
  1. Identify the smallest possible subgames.
  2. Ask: for each of these subgames, what is the NE (or, if relevant, the optimal choice of the single player) in that subgame?
  3. Replace these subgames with the implied payoffs, making them terminal nodes of the new reduced-form game. Then go back to 1.

Reminder: A Static Cournot Duopoly I

- Two firms in a market.
- They compete in quantities.
  - Firm 1's quantity: q_1.
  - Firm 2's quantity: q_2.
- The firms' profits:

$$\Pi_1(q_1, q_2) = D(q_1 + q_2)\, q_1 - C_1(q_1),$$

$$\Pi_2(q_1, q_2) = D(q_1 + q_2)\, q_2 - C_2(q_2).$$

- A **Cournot-Nash equilibrium**:
  - (q_1^c, q_2^c) is a Cournot-Nash equilibrium if neither firm can increase its profits by deviating unilaterally:

$$\Pi_1(q_1^c, q_2^c) \geq \Pi_1(q_1, q_2^c) \qquad \text{for every } q_1,$$

$$\Pi_2(q_1^c, q_2^c) \geq \Pi_2(q_1^c, q_2) \qquad \text{for every } q_2.$$

Reminder: A Static Cournot Duopoly II

- In contrast, the **collusive outcome**, (q_1^m, q_2^m), maximizes total profits:

$$\max_{q_1, q_2} \Pi_1(q_1, q_2) + \Pi_2(q_1, q_2).$$

- The collusive outcome cannot be sustained as a Cournot-Nash equilibrium in a one-shot game.
  - Suppose Firm 1 expects Firm 2 to produce its share of the optimal cartel output, q_2^m. Then Firm 1's best response would be the solution to

$$\max_{q_1} \Pi_1(q_1, q_2^m).$$

    - Denote the solution to this problem by q_1^r.
    - And denote Firm 1's profit if deviating to q_1^r by Π_1^r [$=\Pi_1^r(q_1^r, q_2^m)$].
  - We have $q_1^r > q_1^m$ and $\pi_1^r > \pi_1^m > \pi_1^c$.

Infinitely Repeated Games I

- "Infinitely repeated games" are also called "supergames."
- Imagine that there is an infinite sequence of time periods: $t = 1, 2, 3, \ldots$
- In each period, the duopolists simultaneously choose their quantities q_1^t and q_2^t.
- They then get the profits $\Pi_1(q_1^t, q_2^t)$ and $\Pi_2(q_1^t, q_2^t)$.
- This is repeated in every period, and both players know all previously chosen quantities.
- The list of all chosen quantities prior to period t is called the **period t history of the game** [=everything that has happened previously in the game].

Infinitely Repeated Games II

- Each firm maximizes the **discounted sum** of all its future profits. For Firm 1:

$$\begin{aligned} V_1 &= \Pi_1\left(q_1^1, q_2^1\right) + \delta \Pi_1\left(q_1^2, q_2^2\right) \\ &\quad + \delta_1^2 \Pi_1\left(q_1^3, q_2^3\right) + \delta_1^3 \Pi_1\left(q_1^4, q_2^4\right) + \cdots \\ &= \sum_{t=1}^{\infty} \delta_1^{t-1} \Pi_1\left(q_1^t, q_2^t\right), \end{aligned}$$

  where δ is a **discount factor** [recall that $\delta^0 = 1$].
- Assumption: $0 < \delta < 1$.
- Interpretation:
  - $\delta = \frac{1}{1+r}$, where r is an interest rate.
  - δ could reflect the possibility that, with some probability, the game ends after the current period.
- A strategy in a repeated game is a function.
  - This function specifies, for any possible history of the game, which quantity a player chooses.

Infinitely Repeated Games III

Grim trigger strategies

- Consider the following "grim trigger strategy" for Firm 1 in period t [see also figure on whiteboard]:
  - If both firms have played the collusive output (q_i^m) in all previous periods, play the collusive output in this period too.
  - If at least one firm did not play the collusive output (some $q_i \neq q_i^m$) in at least one previous period, play the Cournot-Nash output (q_i^c).
- Check if it is a SPNE for the firms to use this strategy:
  - First we check that no firm has an incentive to deviate unilaterally along the equilibrium path. (Requirement for having a Nash equilibrium.)
  - Then we check the same thing off the equilibrium path. (Requirement for subgame perfection.)

Infinitely Repeated Games IV

- Reminder:
  - The sum of a finite geometric series:

$$1 + \delta + \delta^2 + \delta^3 + \cdots + \delta^{T-1} = \frac{1 - \delta^T}{1 - \delta}$$

  - The sum of an infinite geometric series:

$$1 + \delta + \delta^2 + \delta^3 + \cdots = \frac{1}{1 - \delta}$$

- Player 1's payoff if both players play the grim trigger strategy (from $t = 1$ onwards):

$$\begin{aligned} V_1^e &= \sum_{t=1}^{\infty} \delta^{t-1} \Pi_1^m = \Pi_1^m \sum_{t=1}^{\infty} \delta^{t-1} \\ &= \Pi_1^m \left(1 + \delta + \delta^2 + \delta^3 + \cdots\right) \\ &= \frac{\Pi_1^m}{1 - \delta}. \end{aligned}$$

Infinitely Repeated Games V

- Payoff if deviating (from the equilibrium path) at $t = 1$:

$$\begin{aligned} V_1^d &= \Pi_1^r + \sum_{t=2}^{\infty} \delta^{t-1} \Pi_1^c = \Pi_1^r + \Pi_1^c \sum_{t=2}^{\infty} \delta^{t-1} \\ &= \Pi_1^r + \Pi_1^c \left(\delta + \delta^2 + \delta^3 + \cdots\right) \\ &= \Pi_1^r + \delta \Pi_1^c \left(1 + \delta + \delta^2 + \cdots\right) \\ &= \Pi_1^r + \frac{\delta \Pi_1^c}{1 - \delta}. \end{aligned}$$

- That is, no incentive to deviate if

$$V_1^e \geq V_1^d \Leftrightarrow \frac{\Pi_1^m}{1-\delta} \geq \Pi_1^r + \frac{\delta \Pi_1^c}{1-\delta}.$$

- Solving this expression for δ yields

$$\delta \geq \frac{\Pi_1^r - \Pi_1^m}{\Pi_1^r - \Pi_1^c}.$$

Infinitely Repeated Games VI

- Interpretation:
  - By deviating, you make a short-term gain but get a lower profit in all future periods. So if you're patient enough (sufficiently large δ), then you resist the temptation to deviate.
- Checking subgame perfection:
  - Any deviation is effectively punished by the competitor.
  - Is carrying out this punishment credible?
  - Imagine that we are in a subgame where at least one firm has previously chosen some quantity differing from the collusive output (some $q_i \neq q_i^m$).
  - The grim trigger strategy prescribes that then each firm should choose the Cournot output ($q_i = q_i^c$).
  - We must verify that this is a Nash equilibrium. Clearly it is!

Infinitely Repeated Games VII

- Conclusion:
  We can sustain the outcome (q_1^m, q_2^m) (in every period) as an SPNE of the infinitely repeated game if the players care sufficiently much about the future (or, the interest rate r low enough):
  $$\delta \geq \frac{\Pi_1^r - \Pi_1^m}{\Pi_1^r - \Pi_1^c}.$$
- Are there other equilibria?
  - Yes! For example, always playing the Cournot-Nash quantity is also a SPNE.
  - Multiplicity of equilibria a problem with this theory — no obvious prediction.
  - The typical approach among IO economists:
    - Assume the firms are able to coordinate on a collusive equilibrium whenever such an equilibrium exists.

The Folk Theorem I

- The result above that cooperation is possible for large enough values of δ is a special case of a more general result called the **Folk Theorem**.
- Let Π be a vector of per-period payoffs in a repeated game.
  - Π is **feasible** if there exists some strategy profile that gives rise to this payoff profile.
    - In our duopoly example, $\Pi = (\Pi_1, \Pi_2)$ is feasible if $\Pi_1 + \Pi_2 \leq \Pi^m$ (the sum of profits cannot exceed the monopoly profit).
  - Π_i is **individually rational** if it exceeds player i's reservation payoff (the highest payoff player i can guarantee himself).
    - In our duopoly example, Π_i is individually rational if $\Pi_i > 0$ (a firm can always guarantee itself a zero profit by producing nothing).

The Folk Theorem II

- The Folk Theorem says that in an infinitely repeated game with observable actions and with δ sufficiently close to unity, if Π is a vector of per-period payoffs that is feasible and if each Π_i is individually rational, then there exists an SPNE in which the components of Π are the per-period equilibrium profits.
  - Put differently: In a repeated game in which the players are sufficiently patient, every payoff above maximin can be achieved in some equilibrium.
- The Folk Theorem is in a way a problem for the theory: we can explain too much!
  - The approach taken by economists:
    - Assume the players can coordinate their behavior on some "focal" equilibrium. For example, in a symmetric game, the players coordinate on a symmetric equilibrium, and this equilibrium is Pareto efficient from the point of view of these players (e.g., the firms).

The Rotemberg and Saloner Model I

- Model:
  - A duopoly market with two identical firms.
  - Constant MC=c.
  - An infinite sequence of time periods: $t = 1, 2, 3, \ldots$
  - In each period, the duopolists simultaneously choose their prices p_1^t and p_2^t.
  - Demand is stochastic:
    - With probability $\frac{1}{2}$: demand is low, $q = D_L(p)$.
    - With probability $\frac{1}{2}$: demand is high, $q = D_H(p)$.
    - For all p, $D_H(p) > D_L(p)$.
  - The demand shock is identically and independently distributed across time periods.
  - The firms learn the current state before choosing their prices simultaneously.

The Rotemberg and Saloner Model II

- Let's look for a pair of prices (p_L, p_H) such that:
  (a) Both firms charge p_s when the state is s.
  (b) p_L and p_H are part of an SPNE.
  (c) The expected present discounted profits of each firm along the equilibrium path

$$\begin{aligned} & \frac{1}{2}\sum_{t=0}^{\infty}\delta^t\left[\frac{D_L(p_L)(p_L-c)}{2}+\frac{D_H(p_H)(p_H-c)}{2}\right] \\ = & \frac{1}{(1-\delta)2}\left[\frac{D_L(p_L)}{2}(p_L-c)+\frac{D_H(p_H)}{2}(p_H-c)\right] \end{aligned}$$

  $\equiv V$ is not Pareto dominated by other equilibrium payoffs.

- In particular, look for a fully collusive outcome (monopoly prices $(p_L, p_H) = (p_L^m, p_H^m)$ is each state).
  - In a fully collusive outcome, monopoly profits is

$$\Pi_s^m = D_s(p_s^m)(p_s^m - c).$$

The Rotemberg and Saloner Model III

- With such behavior,

$$V = \frac{\Pi_L^m + \Pi_H^m}{4(1-\delta)}. \tag{1}$$

- Assume grim trigger strategies where a deviation by anyone leads to marginal cost pricing (Bertrand eq) for the rest of the game (the harshest possible deviation).
  - If following equilibrium when state is s, a firm's overall payoff is

$$\frac{1}{2}\Pi_s^m + \delta V.$$

  - If deviating (just undercutting the rival's price), the firm can get (almost)

$$\Pi_s^m + 0$$

  (from next period onwards the firm gets zero profit).

The Rotemberg and Saloner Model IV

- That is, no incentive to deviate if

$$\frac{1}{2}\Pi_s^m + \delta V \geq \Pi_s^m \Leftrightarrow \delta V \geq \frac{1}{2}\Pi_s^m. \qquad (2)$$

- Eq. (**??**) must hold both for $s = L$ and $s = H$. This will be the case if and only if it holds for $s = H$ (since $\Pi_H^m > \Pi_L$):

$$\delta V \geq \frac{1}{2}\Pi_H^m. \qquad (3)$$

- Plugging (**??**) into (**??**) and simplifying yield

$$\delta \geq \delta_0 \equiv \frac{2\Pi_H^m}{3\Pi_H^m + \Pi_L^m}.$$

- Because $\Pi_H^m > \Pi_L$,

$$\delta_0 \in \left(\frac{1}{2}, \frac{2}{3}\right)$$

The Rotemberg and Saloner Model V

- The main insight can be seen already here: the condition for collusion is more stringent in a high-demand state.
  - Tirole goes on and derives the Pareto optimal prices when the fully collusive outcome is not obtainable.
  - We can stop here, and note that for $\delta \in \left[\frac{1}{2}, \delta_0\right)$, collusion can be sustained with deterministic demand. However with stochastic demand and a high-demand state, this is not possible.
  - Rotemberg and Saloner interpret this as demonstrating the existence of a price war during booms — there is less collusion in good times.
  - They also discuss empirical evidence of this.