Repeating Yourself

Marco Pagnozzi

`mpagnozzi@unisa.it`

(Based on Notes by David Myatt)

Objective:

- Repeated games with discounting.
- Conditional equilibrium selection.
- Finitely vs. infinitely repeated games.
- Collusion and conditions for collusion.
- The Nash-threats folk theorems.

Repeated Games with Discounting

Repeated simultaneous move games with discounting are a particular class of multi-stage games. They are often used to analyze collusive arrangements.

Players: A finite set $I = \{1, 2, \ldots, n\}$, with member i.

Stages: There are $T + 1$ stages, with $t \in \{0, 1, \ldots, T\}$.

Actions: At each stage t, player i chooses an action $a_i^t \in A_i$. (The (stage) strategy space includes mixtures over A_i.) The choice may be conditioned on the history of action choices in all stages up to $t - 1$.

Payoffs: The (vNM) payoff to player i in period t is $u_i(a^t)$ where $a^t = \times_i a_i^t$. The payoffs for the whole game are:

$$\Pi_i = (1-\delta)\sum_{t=0}^{T} \delta^t u_i(a^t).$$

(The leading constant in the payoffs is a normalisation; it gives the *average* per-period payoffs.)

- In an infinitely repeated game, the payoffs are:

$$\Pi_i = \lim_{T\to\infty}\left[(1-\delta)\sum_{t=0}^{T} \delta^t u_i(a^t)\right].$$

Conditional Equilibrium Selection

- Consider the following symmetric game:

	Left	Middle	Right
Top	0 0	3 4	6 0
Middle	4 3	0 0	0 0
Bottom	0 6	0 0	5 5

- Consider the mixed strategy of $[\mathrm{T}(1-\epsilon)+\mathrm{M}\epsilon]$ versus the pure strategy B, for ϵ very small. Row's payoffs are:

	L	M	R
$\mathrm{T}(1-\epsilon)+\mathrm{M}\epsilon$	4ϵ	$3(1-\epsilon)$	$6(1-\epsilon)$
B	0	0	5

- Bottom (and by symmetry, Right) is strictly dominated by the mixing. Hence, delete them to obtain:

	L (y)	M $(1-y)$	Expected
T (x)	0 0	3 4	$3(1-y)$
M $(1-x)$	4 3	0 0	$4y$
Expected	$3(1-x)$	$4x$	

⇒ There are 3 equilibria: 2 in pure strategies, 1 in mixed strategies.

- Examine the payoffs graph:

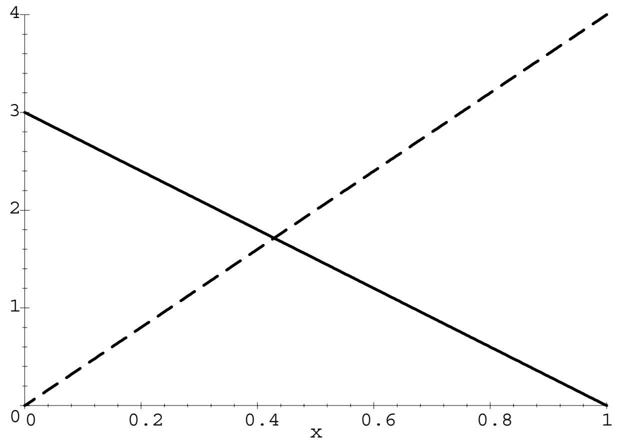

Left (or $y=1$): solid; Middle (or $y=0$): dash

$$\Rightarrow \quad \text{NE} = \begin{cases} \{\mathrm{M},\mathrm{L}\} & \text{with payoff } (4,3) \\ \{\mathrm{T},\mathrm{M}\} & \text{with payoff } (3,4) \\ \left\{\frac{3}{7},\frac{3}{7}\right\} & \text{with payoff } \left(\frac{12}{7},\frac{12}{7}\right) \end{cases}$$

- But all three NE are Pareto dominated by $\{\mathrm{B},\mathrm{R}\}$.

Repeating the Game

	L	M	R
T	0 0	3 4	6 0
M	4 3	0 0	0 0
B	0 6	0 0	5 5

- Can players do any better? Assume they repeat the game twice, with a discount rate δ.
- Consider the following strategy:

(*i*) Play {B,R} in the first period, and {M,L} in the second.

(*ii*) If either player deviates in the first period, play the mixed equilibrium $\left\{\frac{3}{7}, \frac{3}{7}\right\}$ in the second period.

- In equilibrium, the strategy yields:

$$\Pi_i^E \geq 5 + \delta 3.$$

- If a player deviates in the first period he can obtain 6 instead of 5, but this leads to a penalty in the second period.
- By deviating from equilibrium, a player obtains:

$$\Pi_i^D = 6 + \delta \frac{12}{7}.$$

⇒ The strategy described is an equilibrium if a deviation is not profitable, i.e. if:

$$\Pi_i^E > \Pi_i^D \quad \Leftarrow \quad \delta > 7/9.$$

- If players are sufficiently patient, then they can obtain more than the stage game NE payoff.

- Notice that the outcome required:

(*i*) To play a stage game NE in the second period;

(*ii*) The presence of multiple stage game NE.

- Why?

Conditional Equilibrium Selection II

	Left	Middle	Right
Top	2 2	0 0	6 0
Middle	0 0	4 4	0 0
Bottom	0 6	0 0	5 5

- Again, there are 3 Nash Equilibria in this game, but they are all dominated by $\{$B,R$\}$.
- Assume players repeat the game twice, with a discount rate δ.
- Consider the following strategy:

(i) Play $\{$B,R$\}$ in the first period, and $\{$M,M$\}$ in the second.

(ii) If either player deviates in the first period, play the pure strategy equilibrium $\{$T,L$\}$ in the second period.

- If a player deviates in the first period he can obtain 6 instead of 5, but this leads to a penalty of $(4-2)=2$ in the second period.

⇒ The strategy described is an equilibrium for $\delta > 1/2$.

- Again, this required a stage game NE in the second period.

Equilibria of Finitely Repeated Games

- Payoff outcomes from repeated games can at least mimic the stage game:
  - Play a stage game NE at every stage, contingent on the stage but not the history of plays.
  - Then deviating cannot benefit within a stage, and does not change future play.

⇒ We have a subgame perfect Nash Equilibrium.

- But *multiple* stage NE allow a more involved outcomes since future choice may be contingent on current play (e.g. conditional equilibrium selection).
- Then deviation in a stage has two effects: it alters stage payoffs, and may change the equilibrium played in future stages.
- This enables various outcomes, conditional on a high enough discount rate.

- But when there is a *unique* stage game NE, the only equilibrium of the finitely repeated game is to play it in every stage.

Finitely Repeated Prisoners' Dilemma

- Consider the two level pricing game — a Prisoners' Dilemma:

	High	Low
High	3 3	5 0
Low	0 5	1 1

⇒ There is a *unique* dominant strategy NE: {L,L}.

- Consider repeating this stage game T times:
  - In the last stage, it is a dominant strategy for both firms to price low.
  - In the penultimate stage, current behaviour does not influence future payoffs. Hence both firms price low.
  - Iterating back, firms price low at every stage.

⇒ Through the backward induction logic, the unique NE of the finitely repeated game is to price low every period.

- Hence *finitely* repeated games may not yield added stage outcomes.
- The argument depends on the existence of a final period.

Repeated Competition and Collusion

- In a one-shot and in a finitely-repeatedBertrand game, firms price at marginal cost and earn zero profits.
- But infinite repetition yields the possibility of collusion.

- Suppose firms adopt the following (trigger) strategy:
  - Charge the monopoly price in each period (and obtain an equal share of the monopoly profits), unless a player deviates, in which case switch to Bertrand equilibrium forever.

- The payoff from colluding, with n firms in the market, is:

$$\Pi^C = \sum_{t=0}^{\infty} \delta^t \frac{\pi_M}{n} = \frac{1}{1-\delta}\frac{\pi_M}{n}.$$

- A firm can deviate and undercut its rival. This yields π_M for one period, but zero profits forever after (since it destroys collusion):

$$\Pi^D = \pi_M + 0.$$

- Hence a firm will not deviate iff:

$$\Pi_C > \Pi_D \Leftrightarrow \quad \frac{1}{1-\delta}\frac{\pi_M}{n} > \pi_M$$
$$\Leftrightarrow \quad 1 \geq n(1-\delta) \quad \Leftrightarrow \quad \delta \geq 1-\frac{1}{n}.$$

⇒ If players are sufficiently patient, then tacit collusion is possible.

- But there are many equilibria in this game:
  - Consider price $p \in \left[c, p^M\right]$ yielding payoff $\pi\left(p\right)$.
  - The trigger strategy "choose price p in each period; if a player deviates choose c forever" is an equilibrium iff:

$$\frac{1}{1-\delta}\frac{\pi\left(p\right)}{n} > \pi\left(p\right) \quad \Leftrightarrow \quad \delta \geq 1-\frac{1}{n}.$$

⇒ Any price between marginal cost and monopoly price can be sustained in equilibrium as long as $\delta \geq 1-\frac{1}{n}$.

Conditions for Collusion

- The previous argument required an *infinite* horizon. This might seem unrealistic.
- An alternative definition is an *indefinite* horizon. For instance, the discount parameter δ may be the probability that the world continues tomorrow.

- The collusive outcome required a sufficiently high discount rate:
  - Players need to care about the future; they trade the benefit of cheating today against the penalty of punishment tomorrow.
- δ becomes higher as the time periods shorten. Hence regular interaction helps to support collusion.
- If the future is uncertain, this corresponds to a lower δ. Hence collusion is harder.

- The required δ depends on the number of firms.
- The temptation to cheat is the capture of the whole market, but the punishment is the loss of the nth share of the market during collusive phases.

The Nash-Threats Folk Theorem

- Consider the prisoners' dilemma:

	High	Low
High	3 3	5 0
Low	0 5	1 1

- The convex hull of the strategy profile payoff represents the *feasible* payoffs.

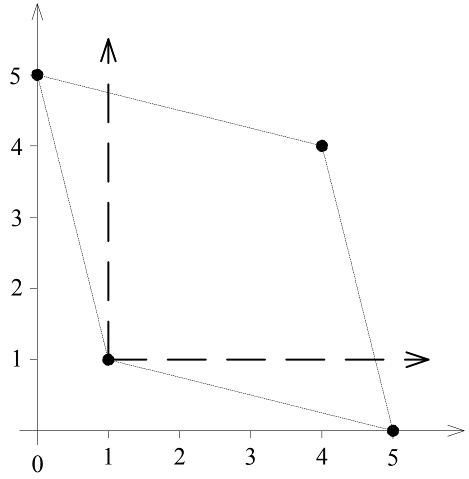

- For sufficiently high δ, any per-period feasible payoff vector higher than the stage-game NE payoff vector may be achieved.

- Such a payoff vector can be supported in the repeated game by a trigger strategy:
  - Play the prescribed strategies to yield the required average payoff.
  - If anyone deviates, revert to the stage NE forever.
- This is the Nash-threats *Folk Theorem*.

Minmax Threats and the Folk Theorem

- **Definition**: Player i's *minmax* payoff (or *reservation utility*) is the lowest payoff that the other players can force upon i:

$$v_i = \min_{a_{-i} \in A_{-i}} \left\{ \max_{a_i \in A_i} u_i(a_i, a_{-i}) \right\}.$$

- This is the worst other players can force on player i, recognizing that player i will do her best in the circumstance.

⇒ Player i cannot get less than v_i in the stage game.

- The set of *individually rational feasible payoffs* is the subset of the feasible payoffs that give each player at least her minmax payoff.

- **Folk Theorem:** For every individually rational feasible payoff vector v, there exists a discount factor $\underline{\delta} < 1$ such that, for all $\delta \in (\underline{\delta}, 1)$, there is a Nash equilibrium of the repeated game with payoff v.

⇒ If players are sufficiently patient then any feasible, individually rational payoff can be obtained in equilibrium.

- Consider the following game:

	C	D
C	$\underline{4}$ $\underline{4}$	0 $\underline{2}$
D	$\underline{2}$ 0	1 1

- C is a dominant strategy ⇒ (C, C) is the unique Nash equilibrium.
- The Nash-threat Folk Theorem only indicates that $(4, 4)$ can be achieved.

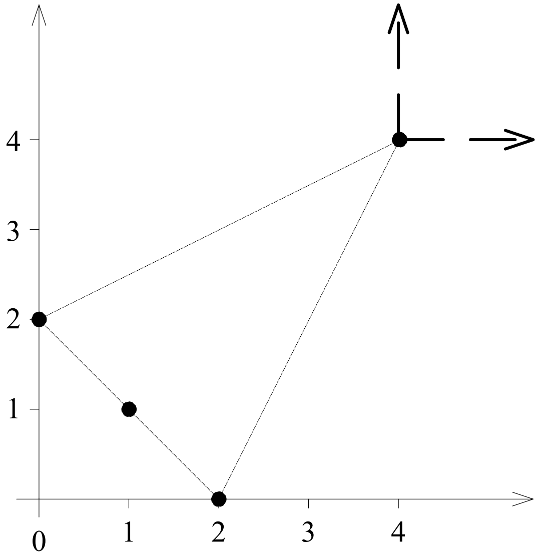

- But consider the *minmax* payoffs:
  - If Column chooses C, Row's best reply is D and Row gets 4
  - If Column chooses D, Row's best reply is D and Row gets 2

$\Rightarrow$ Column can force Row to get 2.

- Minmax payoff for both players is 2. So many other equilibria can be achieved as subgame perfect equilibria:

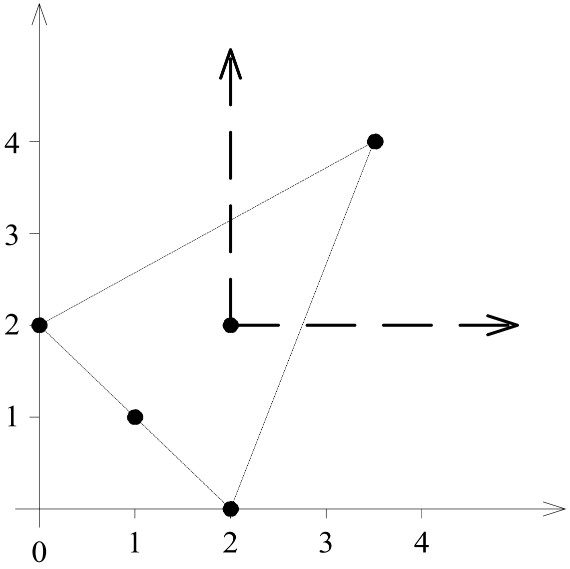

Repeated Competition with Noisy Demand

- The standard supergame-theoretic reasoning assumes that prices are observed.
- If secret-price cutting is available (i.e., a firm cannot observe the price set by its opponent), then the collusion rule must depend on a publicly-observed state variable.
- Consider Tirole's version of the Green and Porter model of collusion with noisy demand:
  - With probability $(1-\alpha)$, demand follows from the standard model.
  - With probability α, market demand falls to zero.
- Then a player can have zero sales for two reasons:

(a) Her opponent has undercut her;

(b) Demand was zero.

Strategies with Noisy Demand

- Consider the following strategy:
  - Play p_M (monopoly price) as split the market as long as demand is positive;
  - Play p_C (marginal cost) for T periods if zero demand is observed by either firm.

$\Rightarrow$ Players enter a price war phase whenever one of them observes zero demand.

- Notice that the event of entering a price war is commonly known to both players:
  - If a player colludes and observes zero demand, she knows there will be a price war;
  - If a player deviates, she knows her opponent observes zero demand and there will be a price war.

Calculating Payoffs and Supporting Collusion

- Suppose that players collude and charge p_M.
- Denote the expected payoffs in the collusive phase as V_C and the expected payoff in the price war phase as V_W:

$$\begin{cases} V_C = \underbrace{(1-\alpha)\left[\frac{\pi_M}{2}+\delta V_C\right]}_{\text{Keep colluding}} + \underbrace{\alpha\delta V_W}_{\text{Enter war}} \\ V_W = \underbrace{\delta^T V_C}_{\text{Collude after } T \text{ Periods}} \end{cases}$$

- A deviating player obtains V_D:

$$V_D = (1-\alpha)\left[\pi_M + \delta V_W\right] + \alpha\delta V_W.$$

- To avoid a deviation by a player we require:

$$\begin{aligned} & V_D \leq V_C \\ \Leftrightarrow \quad & (1-\alpha)\left[\pi_M + \delta V_W\right] \leq (1-\alpha)\left[\frac{\pi_M}{2}+\delta V_C\right] \\ \Leftrightarrow \quad & \pi_M + \delta V_W \leq \frac{\pi_M}{2} + \delta V_C \\ \Leftrightarrow \quad & \frac{\pi_M}{2} \leq \delta(V_C - V_W). \end{aligned}$$

- Solving V_C and V_W simultaneously (from the initial inequalities) yields:

$$\begin{cases} V_C = \dfrac{(1-\alpha)\pi_M/2}{1-(1-\alpha)\delta-\alpha\delta^{T+1}} \\[2ex] V_W = \dfrac{(1-\alpha)\delta^T\pi_M/2}{1-(1-\alpha)\delta-\alpha\delta^{T+1}} \end{cases}$$

- Substituting for V_C and V_W in the condition for δ we obtain:

$$2(1-\alpha)\delta-(1-2\alpha)\delta^{T+1}\geq 1.$$

- This is satisfied for appropriate values of δ and given a large enough T (i.e., for α small, δ large and T large).
- Folk Theorems are available under incomplete information.

- The model has the following features:
  - Agents never cheat in equilibrium, yet price wars still occur.
  - Price wars are required to give players incentive to behave in equilibrium.
  - Given low demand, a player would rather not “pull the trigger”.
  - Players enter a price war phase, since they expect heir opponents to enter too.

- The existence of a public state variable is critical in the Green and Porter model.