MA300.2 Game Theory II, LSE

Lecture 5: Variations on Repeated Games

1. Introduction

The Pandora's Box problem generated by repeated games finds its most extreme expression in the Folk Theorem. In that theorem, "everything better than minmax" is sustainable as a subgame perfect equilibrium outcome, provided players are patient enough.

Does this mean that the techniques we have been studying have no bite at all in analyzing real-world situations?

There are several answers to this question.

First, there is the entire conceptual understanding of the restrictions imposed by credibility (subgame perfection). Notice that we do not get all these equilibria "free of charge". There are constraints on what can be achieved. For instance, if we fix the discount factor (and do not choose the degree of patience as we did in the folk theorem) then we know that various equilibrium structures work for some discount factors and not for others. [Think of the three-action example in the last class. Or think of the Bertrand competition example a couple of lectures ago in which we wrote down the exact discount factor restrictions needed to achieve various degrees of collusion.]

Second, We might simply focus on the "best possible equilibrium" achievable under any coalition structure, given the discount factor. In this way, repeated games would give us a prediction about the maximal degree of cooperation sustainable. Notice, however, that there is a tension here. These cooperative outcomes will generally have to be backed up by very dire punishments. The dire punishments do not have to be used "in equilibrium" but they must be used in the event of a deviation. So we might worry, not about the "individual credibility" of punishment paths (which is guaranteed by subgame perfection), but their "collective credibility". This is the theme of Section 2.

Third, there are other extensive-form games, even ones that could continue indefinitely and look pretty close to repeated games, in which the notio of subgame perfection pins down the equilibrium very tightly irrespective of the degree of patience. Section 3 discusses a celebrated example of this phenomenon, the Rubinstein bargaining model.

2. Renegotiation

A theme that has recurred again and again in class is why would players "deliberately" select on equilibria with bad outcomes if "better" equilibria are available. My answer to this is that *individual rationality* (along with the common knowledge of the game and strategic beliefs) does not take us further than equilibrium behavior. In particular, it does not permit us to choose among equilibria on the basis of *collective* rationality.

At the same time, we do study repeated games in the hope that it will be able to explain "collusive" outcomes, in which the players can get high payoffs. Thus one justification of the bad equilibria is that we are not really interested in them *per se*, but only insofar as they support the good outcomes. Thus we, the players, might have a conversation about how to proceed and then turn off all conversation, so that when a deviation has to be punished we go ahead and punish, on the expectation that others will too.

But what if we cannot turn off the conversation? Then after a deviation the past is sunk. Might we not have a collective incentive to ignore the deviation and simply start cooperating all over again? We might, but in that case cooperation may not be sustainable to start with! This leads us into the realm of equilibria that are "immune to collective rethinking", or *renegotiation-proof.*

The following instance illustrates the main point:

Example. Consider the 3×3 two-player game shown below:

	A_2	B_2	C_2
A_1	$4,4$	$0,5$	$0,0$
B_1	$5,0$	$3,3$	$0,0$
C_1	$0,0$	$0,0$	$1,1$

Think of repeating this game once. We know that if β is high enough there is a SGP equilibrium in which we play (A_1, A_2) today, following it up with (B_1, B_2) in the case of compliance and (C_1, C_2) in the case of noncompliance. This is a subgame perfect equilibrium and its equilibrium payoff is the very best one among all equilibria.

Is that good enough grounds to settle on this equilibrium? Maybe, if players can talk just once at the start of the game and not thereafter? But what if they can talk at the start of the repetition as well? Is it then reasonable, using the same collective rationality postulate that started us off on cooperation in the first place, to assume that (C_1, C_2) will ever be played? Probably not.

But now notice the paradox: if it is not reasonable, then we destroy the even more cooperative outcome (A_1, A_2) in the first period! It cannot be sustained any longer.

Renegotiation-proof equilibria are a way to formalize the idea that players can not just talk and cooperate, they can do so repeatedly as the game wears on. As we have just seen, such equilibria do *not* generally give you the best SGP outcome.

A full definition of the concept for infinitely repeated games is a bit complicated, so we shall restrict ourselves to thinking about finitely repeated games.

Consider a game played over dates $0, 1, \ldots, T$. For any subset of payoff vectors — call it A — denote by $\text{Par}(A)$ its *Pareto frontier*, the set of all efficient payoffs in A:

$$\text{Par}(A) = \{\mathbf{p} \in A | \text{ there is no } p' \in A \text{ such that } p' \gg p\}.$$

The next concept we need is what it means for a set of payoffs in an S-period game to *support* a set of payoffs in an $S+1$-period game. Let W be such a set of payoffs in the S-period game and let p be a feasible payoff vector in the $S+1$-period game. Say that W *supports* p if there are $n+1$ payoff vectors (perhaps not all distinct) — call them $\mathbf{w}; \mathbf{w}^1, \ldots, \mathbf{w}^n$ — and an action profile $\mathbf{a}$ such that

(1) $$\mathbf{p} = (1-\beta)\mathbf{f}(\mathbf{a}) + \beta\mathbf{w},$$

and, for every i,

(2) $$p_i \geq (1-\beta)d_i(\mathbf{a}) + \beta w_i^i.$$

Denote by $\phi(W)$ the set of all payoff vectors supportable by W.

This concept is very useful indeed and has led to some real insights in the theory of repeated games. While we won't go into all that here, notice that this shorthand allows us to write down very quickly the set of SGP equilibrium payoff vectors. Let V^0 be the set of all equilibrium (=subgame perfect equilibrium) payoffs in the one-shot game. Recursively, if V^S is the set of subgame-perfect payoffs for an S-period game, it is easy to see that the corresponding set for $S+1$ is given by

$$V^{S+1} = \phi(V^S),$$

and this way we can "recurse backwards" to find the set of all subgame perfect payoffs at the start of the full repeated game.

It turns out that this construction is also useful in constructing the renegotiation-proof payoffs. Once again, begin with V^0 (call it W^0 for notational symmetry below), but this time look only at the set of payoffs

$$R^0 = \text{Par}(W^0),$$

because we know that no other payoff vector will be available, given collective rationality. So the set of supportable payoffs is then

$$W^1 = \phi(R^0),$$

and we must take the Pareto-frontier of W^1:

$$R^1 = \text{Par}(W^1).$$

In this way, given that R^S is the set of renegotiation-proof payoff vectors for the S-game, R^{S+1} must solve the equation

$$R^{S+1} = \text{Par}\phi(R^S).$$

Proceed in this way all the way to R^T.

This isn't just some abstract definition. It has some intriguing properties. To see this, consider the following

Example. [A lender-borrower game.] A lender lends to a borrower using a fixed rate of interest r. The borrower has access to two projects, but the lender can specify and monitor for which project the funds are to be used. One project has a unit return (net of the rate of interest) of a to the borrower, the other has a return of b, and in fact $a > b > 0$. Both projects can be scaled up or down with these unit returns, but there is an upper limit on the scale of the bad project; call it L.

The lender does not care about the projects. She just wants her money back with interest.

The borrower can default on the loan, in which case an exogenous penalty of monetary value π is imposed on her.

Start at stage 0 (counting from the very last period of the game). There are only two equilibria. Define

$$\ell_0 \equiv \frac{\pi}{1+r};$$

then

$$W^0 = \{(r\ell_0, a\ell_0), (r\ell_0, b\ell_0)\}$$

and

$$R^0 = \text{Par}(W^0) = W^0.$$

At stage 1, define

$$\ell_1 \equiv \ell_0 + \frac{(a-b)\ell_0}{1+r};$$

this is now the maximum loan size which can be sustained without default. Consequently,

$$R^1 = \{(r\ell_1, a\ell_1), (r\ell_1, b\ell_1)\}.$$

Indeed, recursively, as long as $\ell_t \leq L$, define

$$\ell_{t+1} \equiv \ell_t + \frac{(a-b)\ell_t}{1+r};$$

and you will have

$$R^{t+1} = \{(r\ell_{t+1}, a\ell_{t+1}), (r\ell_{t+1}, b\ell_{t+1})\}$$

as long as $\ell_{t+1} \leq L$.

But at some point S ℓ_S will cross L. At this stage the set of renegotiation-proof payoffs will collapse to a single point:

$$R^S = \{(r\ell_S, a\ell_S)\}$$

and in any longer horizon the loan size will need to collapse back to ℓ_0. This is because there is no "punishment power" left in the S-period game. Thus in the renegotiation-proof equilibrium of this game, the loan size will cycle up and down with periodicity S!

Collective rationality applied throughout leads us into paradoxes similar to individual rationality.

3. Rubinstein Bargaining

3.1. **Basics.**

Rubinstein bargaining represents possibly one of the simplest examples of an infinite game. It *looks* "repeated", as you will see in a minute, but it is not formally a repeated game. It is also a beautiful example of the bite created by subgame perfection. Finally, Rubinstein bargaining is of interest in itself and has found many applications in the literature.

Suppose there are two persons, call them 1 and 2. They are dividing a cake of size 1. They take turns in proposing divisions of the cake; at each round, person i proposes a division and person j must accept or reject. If there is an acceptance, the game ends and the proposed division is implemented. If there is a rejection, we move on to the next round, and proposer and responder switch roles.

If a period passes, the next period is discounted. The discount factor of player i is $\beta_i \in (0,1)$. Thus, if a division $(x, 1-x)$ is settled on at date t, the two payoffs are $\beta_1^t x$ and $\beta_2^t(1-x)$, and if no division is ever settled on at all, then the payoffs are zero.

A *history* is a listing of all that has transpired up to any proposal or response stage. A strategy for player i prescribes an initial proposal (if she is the very first proposer) and then conditional on every history, a fresh proposal (if she is the proposer at that history) or an accept-reject decision (if she is the responder at that history).

3.2. **Nash, Without Subgame Perfection.** There are lots of Nash equilibria of this game. Specifically, fix *any* division of the cake; call it $(x_1, x_2) = (y, 1-y)$. Now think of the following strategies. Each proposer proposes this division initially and following any history. Each responder i uses the response rule following any history and proposal: yes if and only if she is given at least x_i.

Why is this a nash equilibrium? Well, given you — player j — are playing the strategy described above it is best for me, player i, to follow the strategy prescribed for me. And vice versa obviously. So these two strategies "lock together", and we have a Nash equilibrium.

However, such an equilibrium is not subgame perfect. Under the strategies described above, nodes in which player i is offered a bit less than x_i will never be visited. But subgame perfection requires that the optimality of the strategy at those nodes, *visited or not*, should be checked. If we do so, we see that the above strategies are not optimal. If person i feels that she will get x_i tomorrow she should certainly be willing to accept $\beta_i x_i + \epsilon$ today. But for small ϵ this number is in fact smaller than x_i, and the going strategy is telling her to refuse such an offer. So the above strategies — the components of them that describe responses — are not "credible".

3.3. **Perfect Equilibrium in the Rubinstein Model.** The amazing thing about the two-person Rubinstein model is that subgame perfection wipes out all but one of these multiple Nash equilibria! Unlike repeated games, there is no folk-theorem-like diversity of equilibria here.

THEOREM **1.** *There is a unique subgame perfect equilibrium payoff vector in the two-person Rubinstein bargaining model.*

Proof. Let M_i be the *maximum* equilibrium payoff and m_i be the *minimum* equilibrium payoff to person i.

Note that when it is i's turn to propose, she can be sure that a proposed payment of $\beta_j M_j + \epsilon$ to j will be accepted by j for every $\epsilon > 0$. This means that i can always get at least $1 - \beta_j M_j$ in equilibrium. This proves that

$$m_i \geq 1 - \beta_j M_j. \tag{3}$$

Now examine M_j. Suppose that j wants to try and clinch an agreement today. She cannot get more than $1 - \beta_i m_i$. On the other hand, if j makes an unacceptable offer, the max she can get from tomorrow (discounted to today) is $\beta_j M_j$. It follows that

$$M_j \leq \max\{1 - \beta_i m_i, \beta_j M_j\}.$$

It is easy to see that $M_j > 0$ (why?). Therefore, the second term on the RHS above cannot be the one that attains the max. Consequently,

$$M_j \leq 1 - \beta_i m_i. \tag{4}$$

Combining (3) and (4), it is easy to see that

$$m_i \geq 1 - \beta_j M_j \geq 1 - \beta_j(1 - \beta_i m_i),$$

or

$$m_i \geq \frac{1 - \beta_j}{1 - \beta_i \beta_j}. \tag{5}$$

Now combining (3) and (4) in a slightly different way and using the same logic,

$$M_j \leq 1 - \beta_i m_i \leq 1 - \beta_i(1 - \beta_j M_j),$$

so that

$$M_j \leq \frac{1 - \beta_i}{1 - \beta_i \beta_j}.$$

Flipping the indices i and j,

$$M_i \leq \frac{1 - \beta_j}{1 - \beta_i \beta_j}. \tag{6}$$

Combining (5) and (6), we conclude that

$$m_i = M_i = \frac{1 - \beta_j}{1 - \beta_i \beta_j}. \tag{7}$$

□

We can now unpack this to figure out supporting strategies. Player 1 as proposer will always propose the division $(x^*, 1 - x^*)$, where

$$x^* = \frac{1 - \beta_2}{1 - \beta_1 \beta_2},$$

and accept any proposal that gives her at least $\beta_1 x^*$. Likewise, Player 1 as proposer will always propose the division $(1 - y^*, y^*)$, where

$$y^* = \frac{1 - \beta_1}{1 - \beta_1 \beta_2},$$

and accept any proposal that gives her at least $\beta_2 y^*$.

If the two are equally patient with common discount factor β, then the proposer picks up $1/(1 + \beta)$ and the responder picks up $\beta/(1 + \beta)$. The difference arises from the first-mover advantage that the proposer has but in any case will wash out as the discount factor converges to one.

3.4. **More Than Two Players.** What if there are three or more players engaged in bargaining? Let's take a natural extension of the two-player model. Initially a proposer is chosen, and she proposes a division of the cake: a vector $\mathbf{x} \in \mathbb{R}^n$ such that $\sum_i x_i \leq 1$. Responders say yes or no; if all accept, the proposed allocation is implemented and the game is over. Otherwise, one unit of time passes (which everyone discounts at the common rate β), and then the rejector makes a fresh proposal.

A *stationary strategy* (sometimes called a Markovian strategy) calls upon the agent to *take the same action* always when it is her turn to propose or respond, provided that the ambient situation is the same.[1] Neither history nor calendar time matters in these strategies.

Focus on responses. With stationary strategy profiles they must look like this: say yes if the amount given to you, person i, is above some threshold m_i, otherwise say no.

Quick Digression. Actually it is a bit more complicated than that because if the proposal gives you more (than your threshold) but gives some responder who responds after you less (than her threshold), you don't want to accept because you want to grab the initiative, at least in the rejector-proposes protocol. So the technically correct description of a response vector is a collection $(m_1^*, \ldots m_n^*)$ such that each responder i says yes if the proposal gives $x_j^* \geq m_j$ to every responder j who responds after i and says no if the proposal gives $x_j^* < m_j$ to some responder j who comes after i.[2]

Now go back to proposals made by i. She will simply get try to

$$z_i \equiv 1 - \sum_{j \neq i} m_j,$$

assuming this is nonnegative, and zero otherwise.

Turn to i's role as responder. By rejecting a proposal, i takes the initiative and gets z_i tomorrow, valued today at βz_i. Therefore

$$m_i = \beta z_i = \beta(1 - \sum_{j \neq i} m_j)$$

Subtract βm_i from both sides; we then get

$$(1-\beta)m_i = \beta(1-M), \tag{8}$$

where M is just the sum of all the m_j's: $M = \sum_j m_j$. Adding the above equation over all i, we see that

$$(1-\beta)M = n\beta(1-M),$$

or

$$M = \frac{n\beta}{1+(n-1)\beta}.$$

[1] In the present context, "the ambient situation is the same" simply means that she is responding to the same proposal.

[2] Strictly speaking, i is perfectly allowed to randomize if $x_i = m_i$ but we simply presume that he will accept then. It is easy to see that if he rejects with some probability in the indifference case then this cannot form part of any stationary equilibrium, because the corresponding proposer will not have a well-defined best response.

Using this information in (8), we see that

$$(1-\beta)m_i = \frac{\beta(1-\beta)}{1+(n-1)\beta},$$

or

$$m_i = m^* = \frac{\beta}{1+(n-1)\beta}, \tag{9}$$

for every i, and indeed, this is what each responder gets. The proposer picks up the remainder, which is easily seen to be

$$\frac{1}{1+(n-1)\beta}.$$

All of this goes to equal division as $\beta \to 1$.

Notice the connection with the two-player model (set $n = 2$ in equation (9) above).

Unfortunately, the beautiful uniqueness result for two-person Rubinstein bargaining no longer survives with three or more players. I will discuss this if we have time, though I doubt we will.