Chapter Five: Repeated Complete Information Games*

5.1 Some General Principles

As mentioned in Chapter One, repeating a game raises two further issues about the players:

1. How do they remember the past?

2. How do they appraise the future?

The "history" $\mathcal{H}$ of a repeated game is simply a record of what all players did at all prior iterations of the game. Technically, if $\mathcal{S}$ denotes the action space of the constituent game then the set of possible histories of any given length n is the Cartesian product of $\mathcal{S}$ n times. The set $\mathcal{H}$ of all possible histories of any length is then the reunion:

$$\mathcal{H} = \bigcup_{n=1}^{\infty} \left(\prod_{i=1}^{n} \mathcal{S} \right) \quad (1)$$

This is an awfully large set of possible pasts for an average player to remember. So, in practice, players only remember finitely many possible developments called the "states" of the (repeated) game. Many histories will therefore belong to the same state when they share some important qualities. A typical definition of a state is what the players did at the very last turn. But one could similarly keep a record of several past turns. More interestingly, the states could distinguish among patterns of behavior, such as a propensity to respect certain norms.

In most cases, there is a finite set of states $\mathcal{F}$ and a "transition rule" $\mathcal{T}$ from state to state that spells out precisely how the current state together with the current play define the next state. Technically:

$$\mathcal{T}: \mathcal{F} \times \mathcal{S} \rightarrow \mathcal{F} \quad (2)$$

The simplest example of such a transition rule is given by the two-memory-states repeated Prisoner's Dilemma of Chapter One: joint cooperation from the Cooperated state leads back to the same Cooperated state. Anything else leads to the Defected state.

The most widely accepted answer to the second question is that players discount future payoffs in geometric fashion: if a player discounts the next turn by a factor d ($0 < d < 1$) then s/he discounts the next-to-next turn by factor d^2, the one after by d^3, and so on… A basic assumption is that the sequence of future turns does not end deterministically, although it can end probabilistically. The reason is that a game that is repeated only a finite number of turns is merely a "non-repeated" game as was studied in Chapter Three.

To understand the relationship between discounting and probabilistic ending, one only needs to look at Figure 5.1. In the upper part of Figure 5.1, after the blue player chooses Play, Nature ends the game in the outcome (10, 20), with probability $p = 0.1$, or continues it with probability $p = 0.9$. At node B, each side therefore appraises the future as the expected payoffs 0.9×(10, 20)=(1, 2), should the game end, and whatever expected payoffs hold at node R, pre-multiplied by probability $p = 0.9$ of reaching it. In the lower part of Figure 5.1, the move Play produces the immediate payoffs (1, 2) and whatever expected payoffs hold at node R, pre-multiplied by the discount factor $d = 0.9$. As far as expected payoffs at node B are concerned, the results in either case are exactly the same.

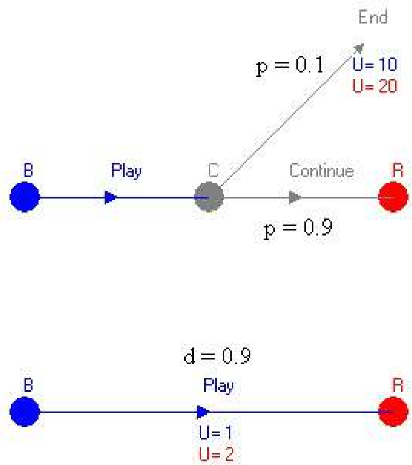

Figure 5.1: Two Equivalent Formulations of Discounting

In repeated games the concept of Nash equilibrium can quickly become unsatisfactory for the same reason as in the example of Figure 1.4 discussed in Chapter One: one can easily design suboptimal strategic plans for the future that result in an optimal plan today based on non-credible threats or pledges (see 5.2.1 below.) The remedy is the concept of sequential rationality: strategies must be optimal at every single turn of the game for the player(s) deciding at that turn. This prevents today's optimal plans to rely on suboptimal plans for tomorrow. The basic solution concept is called a "subgame perfect equilibrium (SPE)." A subgame is any part of the whole game that defines a game all by itself. A SPE is an equilibrium that is Nash in all subgames. And since the repeated game viewed from tomorrow is a subgame, the SPE should always be Nash in tomorrow's repeated game. This precludes today's optimal plan to rely on tomorrow's sub-optimal plans.

In practice, a SPE is based on a description of finitely many memory states of the repeated game. In the example of Figure 1.24 of Chapter One, there are only two memory states: Cooperated and Defected. The Grim Trigger is indeed a SPE but it is more appropriately described as a "Markov perfect equilibrium" (MPE) because of the flavor of Markov chain implied by the memory states and the transitions. In a MPE, play must simply be optimal at every defined state of the game for the deciding player(s). It is easy to show that a MPE is always a SPE. It is also easy to show that a MPE always exists (see homework 5.6.2.)

5.2 Repeating a Normal Form Game

Repeating a normal form game with finitely many memory states is particularly easy using GamePlan. One simply defines the constituent game and duplicates it to create as many memory states as desired. Then, by editing the "upto" of each cell one defines the (common) discount factor as well as the transition rule from state to state. Let us proceed by examples.

5.2.1 Tit for Tat

The strategy "an eye for an eye.." is as old as the Bible. One would expect the wisdom of such an old and venerated text to be grounded in sound Game Theory. To investigate, consider the repetition of the standard two-player Prisoner's Dilemma of Figure 1.23, in Chapter One. Suppose that the two sides are currently cooperating. Tit for Tat (TFT) prescribes that they should then continue doing so at the next turn, and at the next to next turn, and so on, indefinitely. The discounted payoff of doing so is clearly 0 for both sides. Now suppose that one side sneaks in a unilateral Dfct: according to TFT, the other side should retaliate with a Dfct at the next turn while the defector should return to Coop according to the book. So, the initial {Coop, Dfct} is followed by a {Dfct, Coop} which will be followed by {Coop, Dfct}, and so on, indefinitely. The expected value for the side contemplating that initial Dfct is easily calculated:

$$E = 1 + d \times (-2) + d^2 \times (1) + d^3 \times (-2) + \ldots = \frac{1-2d}{1-d^2} < 0 \qquad (3)$$

provided that $d > ½$. So, with enough concern for the future, TFT seems to deter individual defection since it yields a worse discounted payoff than that of always cooperating.

Let us, however, probe that issue one step further. Suppose that, at the very moment when the victim of the initial defection is preparing to retaliate, the initial offender comes to him and makes the following plea: "look, this was all a big mistake. I did not intend to defect on you. I just did it by accident. Please forgive me and skip the retaliation since it would place you in exactly the same situation that I am in now. Indeed, I will cooperate according to the book. But if you retaliate, you will initiate the sequence described above and will get the exact same negative expected payoff E, in (3). So, it is in your interest to forgive me and maintain the cooperation I should not have breached in the first place."

The initial offender is merely pointing out that TFT is not sequentially rational in this game when $d > ½$. The threat of retaliation it relies on turns out to be suboptimal when tested. It is therefore not credible in the first place. In fact, it would only be credible in the case of very shortsighted players with a low discount factor $d < ½$. In that case, the expected payoff E turns positive and better than continued cooperation. Unfortunately, this also means that one side immediately defects at the start of the game and cooperation cannot arise between shortsighted players in that game.

5.2.2 Guilt and the Prisoner's Dilemma

It has already been observed in Chapter One that the repetition of the Prisoner's Dilemma creates a game where rational retaliation becomes possible and can therefore foster cooperation. The simplistic Grim Trigger illustrated in Chapter One fits that bill but seems unsatisfactory: a single deviation by either side compromises cooperation forever! Could there be more sophisticated designs that would not just foster cooperation but promote it and even re-establish it after episodes of defection? The answer is yes and, although there are many such schemes possible, there is a particular instructive one called "Contrite Tit-for-Tat": The idea is to introduce a concept of guilt that leads to the definition of three memory states: One side or the other is guilty (of inappropriate defection) or neither side is. One becomes guilty by defecting unilaterally on a non-guilty side. One remains non-guilty when defecting in retaliation on a guilty side. And one always becomes non-guilty by cooperating. The *GamePlan* model looks like Figure 5.2.

In this example, the (common) discount factor is set to $d = 0.9$ for every transition. In the No-Guilt memory state, bilateral cooperation and bilateral defection create no guilt and lead back to the No-Guilt state. A unilateral defection leads to the corresponding player's guilt state. In either guilt state, the guilty party will remain guilty by defecting and will become non-guilty by cooperating. The non-guilty party will remain non-guilty no matter what it does.

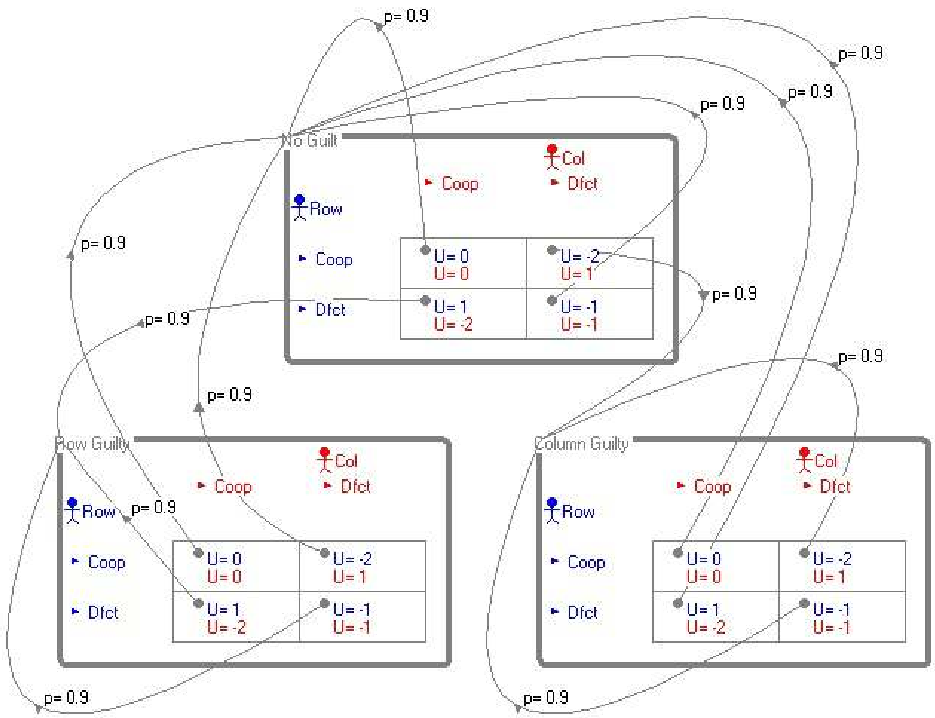

Figure 5.2: A Three-Memory State Repeated Prisoner's Dilemma

Solving the game for pure strategy equilibria yields two solutions: one is Defect all the time in all states of the game. This is a general fact: the repetition of the Nash equilibrium of the constituent game (here Defect & Defect) always yields a Markov perfect equilibrium of the discounted repeated game. The other solution is of great interest: always cooperate in the no-guilt state, always cooperate when guilty and always defect when your opponent is guilty. As a result, any prior sequence of play quickly re-establishes joint cooperation.

5.2.3 An Environmental Treaty

The neighboring states of Megasmog and Pristina have a serious dispute that threatens their longstanding peace: the Blue River that flows from the Megasmog industrial region in the North, along their common border to the South, has become increasingly polluted. Fortunately for Megasmog, it has access to the sources of the river and therefore enjoys a clean water supply. But Pristina's citizens are reduced to filter their water of to buy bottled water produced by the Megasmog Upper River Water Company. Some of Pristina's businesses are pushing to relax its strict anti-pollution laws in order to retaliate. But polluting the environment further would be to the detriment of both sides. Game Theory Associates (GTA), a consulting firm, describes the situation by the normal form game of Figure 5.3.

The situation appears hopelessly disadvantageous for Pristina. But GTA contends that an environmental treaty that would maintain clean policies on both sides is entirely possible. It is only a matter of design. After long negotiations, the two sides agree to consider three "states" of the treaty: Compliance, Megasmog non-compliance, and Pristina non-compliance. The non-compliance state will be reached by the side that is found to unilaterally dirty the environment. The two sides will then remain in that state for a few turns before returning to Compliance. While in non-compliance, the state responsible will clean the environment while the other will be expected to play Dirty. Return to Compliance will be decided by an independent panel with a given probability p. GTA has proposed the model of Figure 5.4.

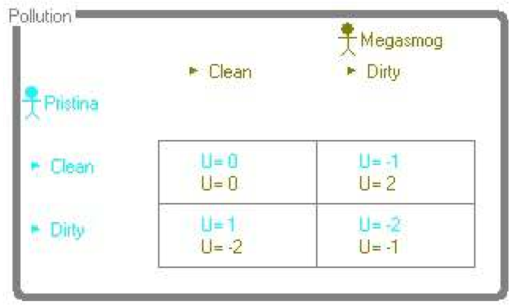

Figure 5.3: The Pollution Game

The Megasmog and Pristina delegations to the talks find that proposition dubious, to say the least. They immediately assail the GTA representative, Dr. Green, with questions: why should Pristina dirty the environment when Megasmog is in non-compliance since their objective is to protect the environment? Why should the panels decide on a probabilistic return to compliance rather than after a fixed period of time? Dr. Green explains that this doesn't make any difference. Even if Pristina is not required to dirty the environment while Megasmog is in non-compliance, it will still do so under the pressure of its business community, since that is allowed by the treaty terms. And if it's not allowed, the treaty has no teeths. And as for a fixed number of turns, it makes no difference since a probability of return to compliance defines an expected number of turns of non-compliance (see Homework 5.6.2.)

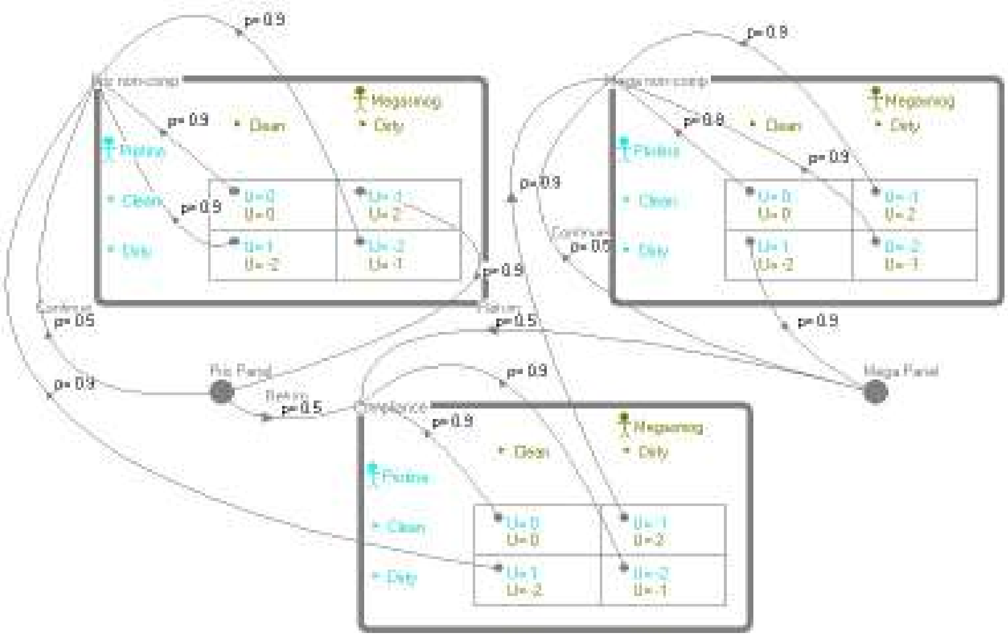

Figure 5.4: An Environmental Treaty

Dr. Green explains that this setup yields a Markov perfect equilibrium (MPE) with Compliance as a stable steady state. You can adjust some of the parameters, she adds, such as the probability of return, but the treaty should succeed (see Homework 5.6.2.)

5.2.4 The Tragedy of the Commons

There are several generalizations of the two-player prisoner's dilemma to three or more players even when symmetry is preserved. It all depends on the effect of accumulating defections. In the game of Figure 3.4, in Chapter Three, a unanimous defection brings the worst possible result for all three players. When played just once, this game has three very symmetrical pure Nash equilibria: one side cooperates while the other two defect. But any of the three sides can be the "victim" and it is therefore hard to predict whom that will be when the game is played for real. Such a situation has been described as the "Tragedy of the Commons", a social dilemma involving a population of self-interested decision makers whose rational individualistic behavior can lead to social catastrophe.

The repetition of that game can easily yield cooperation, depending on how the memory states and the state transitions are defined. For instance, one can create four memory states: Cooperation and one for each possible victim. When all sides cooperate or all simultaneously defect, the state of Cooperation endures. Any deviation by one or two sides lead to a victim state where the victim is expected to defect in retaliation while at least one of the defectors will cooperate. With high enough discount factors this yields a MPE where full cooperation endures. However, the transitions can be engineered in such a way that cooperation will be quickly reestablished rationally (see homework..).

5.2.5 Folk Theorems

In mathematics, folk theorems are well known results whose authorship is unclear. In game theory the term refers to various statements about simple equilibria of repeated games. Perhaps the simplest and most powerful Folk Theorem concerns the Grim Trigger: suppose that a constituent game admits a strategy profile that is not in equilibrium but yields a strictly better outcome than a Nash equilibrium of the same game, for all players. Then, if the players have enough concern for the future (i.e. high enough discount factors), the Grim Trigger that sustains the better outcome through the threat of perpetual reversion to the Nash equilibrium forms a MPE in the repeated game. The typical example is given in Figure 1.24 in Chapter One, but there are numerous other cases.

5.3 The Graph Form

The graph form can accommodate far more diverse game conditions that the simple repetition of a normal form game. In particular, it allows sequential play instead of the implicit simultaneous play of the normal form. The MPE is still the standard solution concept and one must carefully design the graph in order to represent the various possible states of memory.

5.3.1 The Dollar Auction

Professor Gotcha teaches at a state university where he thinks he is badly underpaid for his hard work. In order to supplement his income he devises the following game for his Game Theory class: he will auction a brand new $10 bill. The students will be free to bid up, but only $1 at a time. However, there is a catch in the rules: the highest bidder will indeed get the $10 bill in exchange for his/her bid, but the second highest bidder will also pay his/her bid and will get only the professor's thanks. When he shares his idea with a game-loving colleague, professor Gotcha adds: "At worst, it will cost me about one Dollar." Always up to the challenge, his colleague takes two Dollars out of his pocket and hands them to his friend saying: "Go right ahead then. You will now make a profit if you play the game."

Professor Gotcha teaches Game Theory using GamePlan and devised the model of Figure 5.5. To simplify his analysis, he assumed that two students called Blue and Red would want to play the game and that his only uncertainty is about who will move first, an issue he models by a Chance node with equal probability of either student being first to make up his/her mind about what to do. Then, he carefully distinguishes two possible turns per player. Of course, he considers the possibility of not playing the game altogether (stay) but accounts for his colleague's contribution that he can keep if he goes ahead. The discount factor $d = 0.999$ accounts for the very fast back and forth of a live auction.

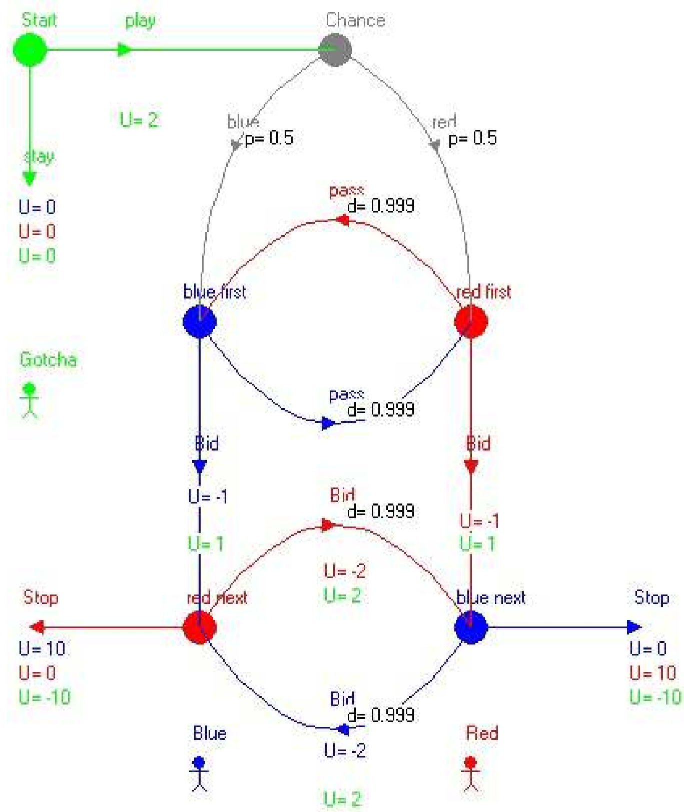

Figure 5.5: The Dollar Auction

Professor Gotcha reported a successful auction to his colleague: he walked away with \$11 net (i.e., the winning bid was \$11) not counting his friend's \$2. In fact he had predicted that he would win at least that much with probability p=10.71%. But Dr. Gotcha had turned a blind eye to a winning strategy for his students: whoever made up his/her mind first should bid \$1 and whoever would go next should abstain from bidding any further. The professor would lose \$9 and be unable to brag. But he was careful not to release his lecture notes before the game.

5.3.2 Repeated Sequential Play

Sometime, repeating a game has very counterintuitive results. The alternate form of the simplest game illustrated in Figure 1.4 in Chapter One highlighted the importance of forward thinking: the threat by Red to move Left in order to deter Blue to choose Continue was dismissed as non-credible by virtue of its non-optimality. But should that simplest game be repeated, the thinking can change drastically. Figure 5.6 shows a version of that repeated game with three memory states representing the three possible plays of the one-shot game. Stop leads to the State 1 node that leads back to the Start node. But Continue followed by either of Red's moves leads to two other possible memory states: now, Continue followed by Left would yield State 2 in which the game would unfold again.

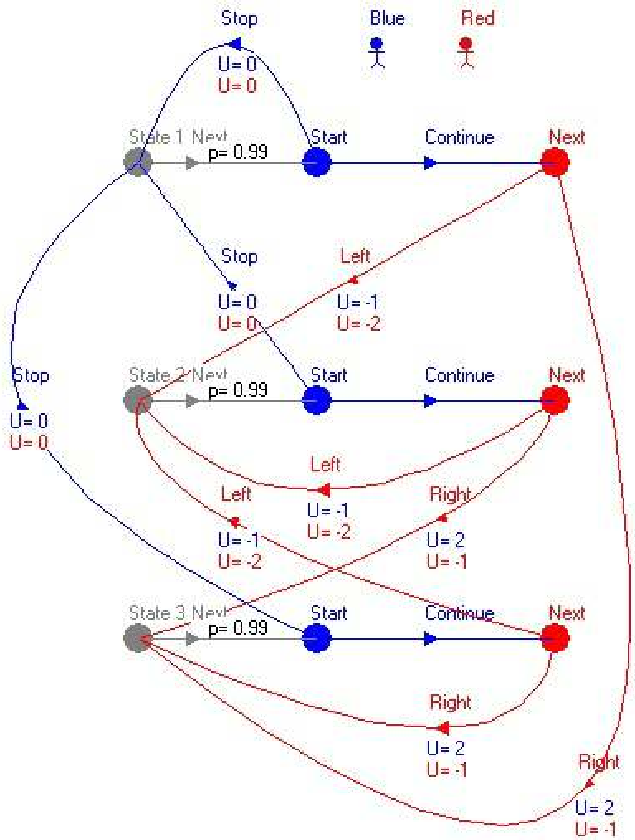

Figure 5.6: A Repeated Simplest Game

The repetition of the constituent game equilibrium {Continue, Right} is of course a MPE of this repeated game. But there are other solutions of interest. In one MPE, Red will always choose Left with probability $p = 2/3$ and Right with probability $p = 1/3$. As a result, and in both States 1 and 2, Blue chooses Stop with certainty. He only chooses Continue with probability $p = 50/99$ in State 3. In other words, Blue is deterred from choosing Continue unless he just observed the sequence {Continue, Right}. And even in that case, he still is deterred with probability $p = 49/99$. The mere repetition can turn what looks like a completely non-credible threat into a perfectly rational one in both simultaneous and sequential play games.

5.4 Repeated Continuous Games

Continuous games such as the duopoly and oligopoly studied in Chapter One and Chapter Three are great candidates for repetition. Indeed, many economic games are by nature repeated in time and should therefore be studied in that perspective. The Cournot duopoly and oligopoly give rise to equilibria that are not as efficient as what could be achieved with some collusion (see homework.) So, they lend themselves to the same improvements using trigger schemes. However, trigger schemes require a clear understanding between the players of what point within an entire continuum will be chosen as the target to maintain as well the exact schedule of retaliation needed to sustain it. This creates some credibility issues as far as applications are concerned.

So, the following question arises: is it possible to design perfect equilibria that sustain an efficient outcome, maintain that outcome dynamically and do not entail anything more that unilateral pledges or threats to formulate? The answer is yes, but it requires the technical developments of the next section.

5.4.1 The Decomposition Theorem

Let $x_i^t \in A_i$ denote Player i's decision in her action space A_i and let $X_{-i}^t \in \bigtimes_{j \neq i} A_j$ denote all other players' decisions in their own action spaces, at turn t. Further let $U_i(x_i^t, X_{-i}^t)$ denote i's constituent game payoff resulting from such decisions at turn t. Player i' objective in the discounted repeated game (with discount factor δ_i for i) is to maximize, at each turn t, the discounted sum of present and future payoffs:

$$E_i(\xi_i^t, \Xi_{-i}^t) = \sum_{s=0}^{\infty} \delta_i^s U_i(x_i^{t+s}, X_{-i}^{t+s}) \tag{4}$$

where $\xi_i^t = \{x_i^{t+s}\}_{s \in \mathbb{N}}$ denotes Player i's present and future choices and $\Xi_{-i}^t = \{X_{-i}^{t+s}\}_{s \in \mathbb{N}}$ denotes all other players' expected present and future choices. The history of the game evolves according to $h^{t+1} = h^t \cup (x_i^t, X_{-i}^t)$.[1] More generally, if the game has memory states and a transition rule $\boldsymbol{\mathcal{T}}$, one has:[2]

$$h^{t+1} = \boldsymbol{\mathcal{T}}(h^t, (x_i^t, X_{-i}^t)) \tag{5}$$

If $h^t \in \boldsymbol{\mathcal{H}}$ denotes the history, or state, of the repeated game at turn t, a strategy for Player i is a map

$$\psi_i : h^t \in \boldsymbol{\mathcal{H}} \to x_i^t \in A_i$$

We also denote by $\Psi = (\psi_i, \Psi_{-i})$ a strategy profile such that $x_i^t = \psi_i(h^t)$, and similarly for all players. One has:[3]

Theorem 5.1: Ψ is a MPE if and only if there exists for each player i two functions $\pi_i \geq 0$ and g_i (of any sign) such that

$$g_i(X_{-i}^t) - \pi_i(\xi_i^t, X_{-i}^t) = U_i(x_i^t, X_{-i}^t) + \delta_i g_i(\Psi_{-i}(h^{t+1})) \tag{6a}$$

with $\quad \pi_i(\xi_i^t, \Psi_{-i}(h_t)) = 0 \quad$ if $x_i^{t+s} = \psi_i(h^{t+s})$ for all $s \geq 0$ (6b)

This is merely a version of the Bellman Equation of Dynamic Programming with a twist that will be extremely useful in applications. But each of the two funtions involved in (6a) finds an interesting interpretation: suppose that one chooses $\pi_i \equiv 0$. Then, whatever Player i could gain by a choice $x_i^t \neq \psi_i(h^t)$ in $U_i(x_i^t, X_{-i}^t)$ would be cancelled *exactly* by the term $\delta_i g_i(\Psi_{-i}(h^{t+1}))$ according to the other players reaction $\Psi_{-i}(h^{t+1})$, since the two terms must add up to $g_i(X_{-i}^t)$ which is *independent* of x_i. For that reason, the g_i term has been called the "countervailing" part of Ψ_{-i}. Equilibria with that property are entirely feasible and have been called countervailing. As shown in examples below, they can arise spontaneously from threats and pledges and need not require any coordination in the strategic choice of the players. When it is not identically nil, the π_i term has the effect of holding the players to a specific strategic choice ψ_i as suggested by (6b). This has been called the "coercive" part of Ψ_{-i}. This may be a desirable feature, for instance if the equilibrium is the result of a treaty design.

5.4.2 Unilateral threats and pledges

[1] If the game begins at time $t = 0$ one has $h_0 = \emptyset$.
[2] $\boldsymbol{\mathcal{T}}(h_t, \Psi(h_t)) = h_t \cup \Psi(h_t)$ can be viewed as a trivial case of transition rule in $\boldsymbol{\mathcal{H}}$.
[3] This theorem first appeared in Langlois & Langlois (1996).

The normal form Prisoner's Dilemma of Figure 1.23 in Chapter One can be reinterpreted as a continuous game with choices $x_i \in [0,1]$ and payoffs

$$U_i(x_i, x_j) = x_i - 2x_j \tag{7}$$

where x_i is interpretable as a "level of defection." The four corners of the resulting (square) action space provide exactly the same payoffs as in Figure 1.23. This continuous Prisoner's Dilemma can serve as a generalization of the discrete version. Its repetition with discount factor δ (common for simplicity) is a typical case where Theorem 5.1 can be applied. It is usually easiest to construct countervailing equilibria by setting $\pi_i \equiv 0$. It is usually quite easy to then modify that equilibrium into a coercive one if need be. In the continuous case is also quite helpful to define the state of the game as simply the last players' choices. In this two player case, this means: $h_t = (x_i^{t-1}, x_j^{t-1})$. Equation (6a) then reduces to:

$$g_i(x_j^t) = x_i^t - 2x_j^t + \delta g_i(\psi_j(x_i^t, x_j^t)) \tag{8}$$

So, if one "knows" g_i and that it is monotonic, it is easy to reconstruct ψ_j by simply solving (8). The question is how g_i could be known? It turns out that g_i can be completely determined by some unilateral threats or pledges made by the players. For instance, suppose that Player j offers to progressively reciprocate i's full cooperation by cutting her defection level in half at each turn. He is pledging:

$$\psi_j(x_i^t = 0, x_j^t) = x_j^t \div 2$$

(9)

But this entirely determines g_i in (8). Indeed, one can write:

$$g_i(x_j^t) = -2x_j^t + \delta g_i(x_j^t \div 2) = \tfrac{-4}{2-\delta} x_j^t \tag{10}$$

Replacing in (8) yields the formula:

$$\psi_j(x_i^t, x_j^t) = \tfrac{(2-\delta)x_i^t + 2\delta x_j^t}{4\delta} \in [0,1] \tag{11}$$

provided $\delta > \frac{2}{3}$. Player i can formulate independently his own pledge or threat and obtain the corresponding strategy in similar fashion. If this also yields a true strategy as in (11), the result is a MPE. And should Player i make the symmetric pledge, he would obtain the symmetric strategy and the two sides would find themselves in an MPE with an elegant property: it sustains and dynamically re-establishes cooperation after any episode of unilateral or bilateral defection.

Instead of making the above pledge of partial reciprocation, Player j may instead make a threat of partial retaliation. For instance, she can threaten progressive retaliation to full defection by cutting in two her current distance to full defection. This means:

$$\psi_j(x_i^t = 1, x_j^t) = 1 - (1 - x_j^t) \div 2 = (1 + x_j^t) \div 2 \tag{12}$$

Again, this determines g_i in (8), at least in the countervailing case:

$$g_i(x_j^t) = 1 - 2x_j^t + \delta g_i((1 + x_j^t) \div 2) = \tfrac{2-3\delta}{(1-\delta)(2-\delta)} - \tfrac{4}{2-\delta} x_j^t \tag{13}$$

Replacing in (8) yields:

$$\psi_j(x_i^t, x_j^t) = \tfrac{(3\delta\text{-}2) + (2-\delta)x_i^t + 2\delta x_j^t}{4\delta} \in [0,1] \tag{14}$$

provided $\delta > \frac{2}{3}$. Again, Player can formulate his own pledge or threat independently. If he adopts the symmetric threat, one again obtains an MPE but with a far less attractive property: it sustains and dynamically re-establishes full defection after any episode of unilateral or bilateral cooperation.

This is not to say that threats are inappropriate for all sorts of games. It only illustrates how unilateral statements coupled with a countervailing assumption can yield interesting MPEs.

5.4.3 Reaction Function Equilibria

In the late 1960's, it was conjectured that Cournot's reaction functions (see Section 1.2.6) could be replaced by a MPE that would promote a more cooperative outcome than the Nash-Cournot equilibrium.[4] The conjecture was proven correct in the early 1990's.[5] The technique is illustrated with the simplest revenue model of Chapter One.

If the two sides of the duopoly were to collude and "fix" prices, they could agree to produce equal levels $q_1 = q_2 = q$ and split the proceeds. They would thus jointly maximize

$$qf(2q) = q(60 - 2q)$$

by choosing $q = 15$ (rather than the Nash-Cournot equilibrium $q = 20$.) They would each enjoy the collusive revenue $qf(2q) = 15 \times 30 = 450$ (instead of $20 \times 20 = 400$.) In order to obtain a MPE that would achieve a collusive outcome, one can define a countervailing g_i function and solve for the reaction function ψ_j

$$g_i(q_j) = q_i(60 - q_i - q_j) + \delta g_i(\psi_j) \tag{15}$$

A simple choice is

$$g_i(q_j) = \lambda - \mu q_j$$

which yields by (15)

$$\lambda - \mu q_j = q_i(60 - q_i - q_j) + \delta(\lambda - \mu\psi_j) \tag{16}$$

or

$$\psi_j(q_i, q_j) = \frac{q_i(60 - q_i - q_j) + \mu q_j - (1 - \delta)\lambda}{\delta\mu} \tag{17}$$

One can now impose on (16) the condition that some $q = \psi_j(q, q)$ is a steady state of the MPE. This yields a relation between λ and μ. It is entirely possible to choose these in such a way that the collusive outcome $q = 15$ be the desired steady state. Unfortunately, this fails to provide dynamic stability to that steady state, a very desirable property. Instead, one can choose an intermediate value such as $q = 16$. Setting $\mu = q = 16$ then yields (with, say δ=0.9) $\lambda = 4,736$ and

$$\psi_j(q_i, q_j) = \frac{q_i(60 - q_i - q_j) + \mu q_j - (1 - \delta)\lambda}{\delta\mu} = \frac{5}{8}\Big(q_i(60 - q_i - q_j) + 16q_j - 473.6\Big)$$

One can verify that ψ_j, together with its symmetric ψ_i have $(q_i, q_j) = (16, 16)$ as a dynamically stable steady state.[6] As a result, they map a neighborhood Ω of that point into itself. Together with the trigger condition $\psi_i = \psi_j = 20$ outside Ω, this pair forms a MPE and exhibits an example of the reaction function equilibria conjectured by Friedman.[7]

---

[4]That conjecture was expressed by James W. Friedman.

[5]Langlois & Sachs (1993) and Friedman & Samuelson () independently achieved the result.

[6]The eigenvalues of the Jacobian matrix $D[\psi_i, \psi_j]$ at $(16, 16)$ are $\lambda_i = \lambda_j = \frac{5}{6}$, less than one in absolute value.

[7]By involving a coercive term π_i, as in (6b), it is possible to construct reaction function equilibria that support the collusive point $(15, 15)$.

5.5 Attrition and Bargaining

An interesting twist on repeated games arises when the players have the capability to end the game by their own choice. The two major examples are wars of attrition and the repeated game model of bargaining.

5.5.1 The War of Attrition

Consider a two-player repeated game where a player's choice at their turn is whether to end the game in a loss for themselves, and a gain for the other, or continue playing the game at a cost, thereby giving the other side the symmetric choice at the next turn. The situation is best pictured as the most basic graph form of Figure..

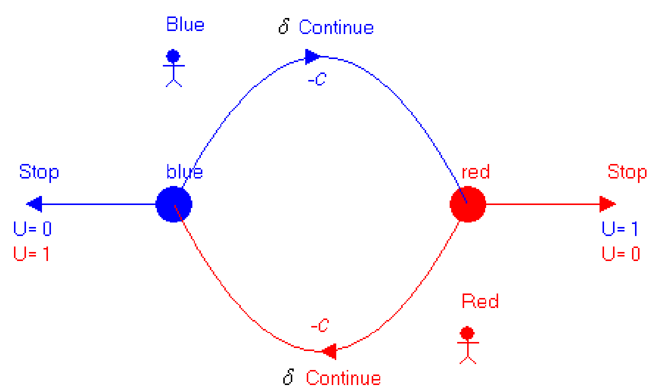

Figure 5.7: The War of Attrition

This game has three MPEs: in two pure equilibria, one side chooses Stop while the other chooses Continue. In a more interesting symmetric MPE, each side continues with a probability that is linked to the game parameters. Here, it is assumed that some prize of value $U = 1$ is at stakes and that the player who chooses Stop hands it to the other (and keeps nothing.) It is easy to solve for the mixed MPE: by symmetry, we may assume a common probability p of continue (and $(1-p)$ of Stop.) At node blue, for instance, player Blue anticipates an expected payoff for Continue:

$$E(@\text{blue}) = -c + \delta E(@\text{red}) = -c + \delta\big(p\delta E(@\text{blue}) + (1-p)\big) = \frac{\delta(1-p)-c}{1-\delta p}$$

In order for the probability p to be rational for Blue at node blue, one must have $E(@\text{blue}) = = 0$ or:

$$p = 1 - \frac{c}{\delta} \qquad (18)$$

So, as long as $c < \delta$, the War of Attrition can continue on with that probability at each turn.

5.5.2 The Rubinstein Bargaining Model

The bargaining problem is as old as Game Theory. In fact, John Nash made his contribution with what is now known as the "Nash Bargaining Solution." This was an axiom-based formula for what bargain should emerge given certain parameters. What became known as the "Nash Program" is the goal of explaining such outcomes through the non-cooperative game theoretic approach. Rubinstein's bargaining model is typical of the Nash Program. The game structure is in fact the same as the above War of Attrition one. The only difference is that the payoff to each side is the result of an offer by the other at the previous turn. The result is shown in Figure 5.8:

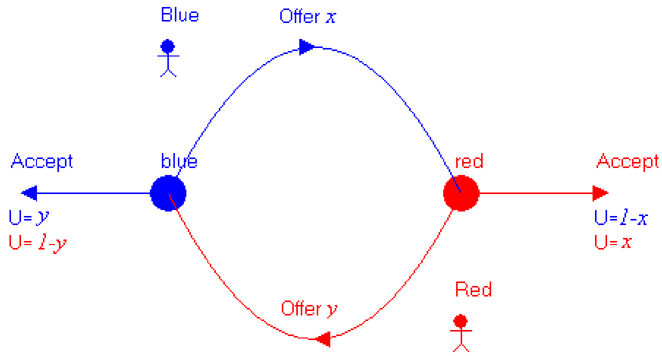

Figure 5.8: The Rubinstein Bargaining Model

Again, the future is discounted by factor δ and player Blue can make the following calculation at node blue: $E(\text{Offer } x) = \delta E(\text{@red})$. But now, one focuses on adjusting x and y to reach the earliest possible bargain since the others will be less valauble as discounted further. This implies an immediate acceptance by each side of the corresponding current offer. This yields:

$$y = E(\text{Offer } x) = \delta(1 - x)$$

With the symmetric condition $x = \delta(1 - y)$, one finds the optimal bargain $x = y = \frac{1}{1+\delta}$.

5.6 Homework

5.6.1 The Repeated Nash Equilibrium as MPE

Argue in your own word why the repetition of a same Nash equilibrium of a repeated constituent game is a MPE of the discouted repeated game. Hint: consider an arbitrary number of memory states and pick any one of them. If the play of the Nash equilibrium is expected in all other memory states, what is best in the picked memory state?

5.6.2 The Best Environmental Treaty

In probability theory, if an event recurs with fixed probability p and ends with probability $(1 - p)$, the expected number ν of event turns is defined by:

$$\nu = (1-p)\sum_{n=1}^{\infty} np^{n-1} = \frac{1}{1-p}$$

(a) What probability p corresponds to an expected number of 10 turns?

(b) What is the maximum probability p that is compatible with the success of the environmental agreement described in §5.2.3?

5.6.3 The Repeated Battle of the Sexes

Construct a repeated game model of the Battle of the Sexes and obtain a MPE that will sustain the alternate and joint choice of Ballet and Fight by the two players.

5.6.4 Collusion in Oligopoly

Generalize the construction of section 5.4.3 to the case of three oligopolists (i, j, k).

(a) Show that the Nash-Cournot equilibrium is at $q_i = q_j = q_k = 15$.

(b) Show that the collusive point is at $q_i = q_j = q_k = 10$.

(c) Let $g_i(q_j, q_k) = \lambda - \mu(q_j + q_k)$, and symmetrically for j and k. Using $\mu = 11$ and $\lambda = 3,212$, solve a system of three equations in three unknowns (ψ_i, ψ_j, ψ_k) (here written for i):

$$g_i(q_j, q_k) = q_i(60 - q_i - q_j - q_k) + \delta g_i(\psi_j, \psi_k)$$

Verify that $q_i = q_j = q_k = 11$ is a steady state for the reaction functions (ψ_i, ψ_j, ψ_k).[8]

5.6.5 The Cuban Missile Crisis

Nuclear crises were described by Herman Kahn as a Game of Chicken. Consider the following continuous game utility functions (and symmetrically, by exchanging i and j):[9]

$$U_i(x_i, x_j) = x_i - x_j - 2x_i x_j$$

This is the continuous extension of the Chicken (normal form) game of Figure 5.9 where $x_i = 0$ means Swerve and $x_i = 1$ means Drive On. In the context of a nuclear crisis, one may interpret x_i as a "level of aggression."

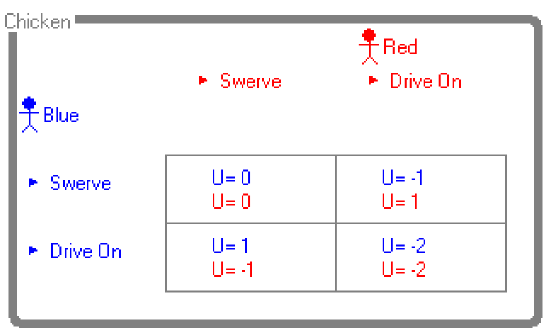

Figure 5.9: The Game of Chicken

Suppose that the future is discounted by $\delta > \frac{1}{2}$. Further assume that the US (as j) offers to reciprocate full cooperation ($x_i = 0$) by the Soviet Union (SU) according to the formula:

$$\psi_j(x_i = 0, x_j) = x_j \div 2\delta$$

In essence, if δ is close enough to 1, this means that US will cut its level of aggression in half at each turn, should SU stick to $x_i = 0$. It is a pledge of incremental reciprocation.

(a) Show that this pledge is equivalent to a countervailing strategy with $g_i(x_j) = -2x_j$. Hint: assume $g_i(x_j) = -\mu x_j$, solve for ψ_j and apply the above condition.

(b) Argue that the best SU can hope for in the long run is to maintain full cooperation. Hint: show that, whatever steady state (x_i, x_j) is ever reached, it will have to satisfy

$$(2\delta - 1)x_j = (1 - 2x_j)x_i$$

So, the best "long term" $g_i(x_j) = -2x_j$ can only occur if $x_i = 0$.

(c) Show that the very same conclusions can be reached if US instead threatens incremental retaliations $\psi_j(x_i = \frac{1}{2}, x_j) = (1 + x_j) \div 4\delta$.

---

[8] It is possible to show that this steady state is dynamically stable under the dynamics defined by the reaction functions.

[9] One justification for such a structure was offered in Langlois (1991.)