

1 Repeated Games

We are involved in the same interactions again and again and again, often with the same people.

Our behavior often does not satisfy the predictions of our static analysis.

Prisoner's Dilemma:

	C	D	
C(cooperate)	3; 3	0; 4	(1)
D(defect or cheat)	4; 0	2; 2*	

Restaurant Quality Game

	B	N	
H	2; 2	-1; 0	(2)
L	3; -1	0; 0*	

Why do we see this divergence between predictions and what actually happens?

When you were small, did you have a curfew (time you had to be home)? Did you follow it? If you did, why would you?

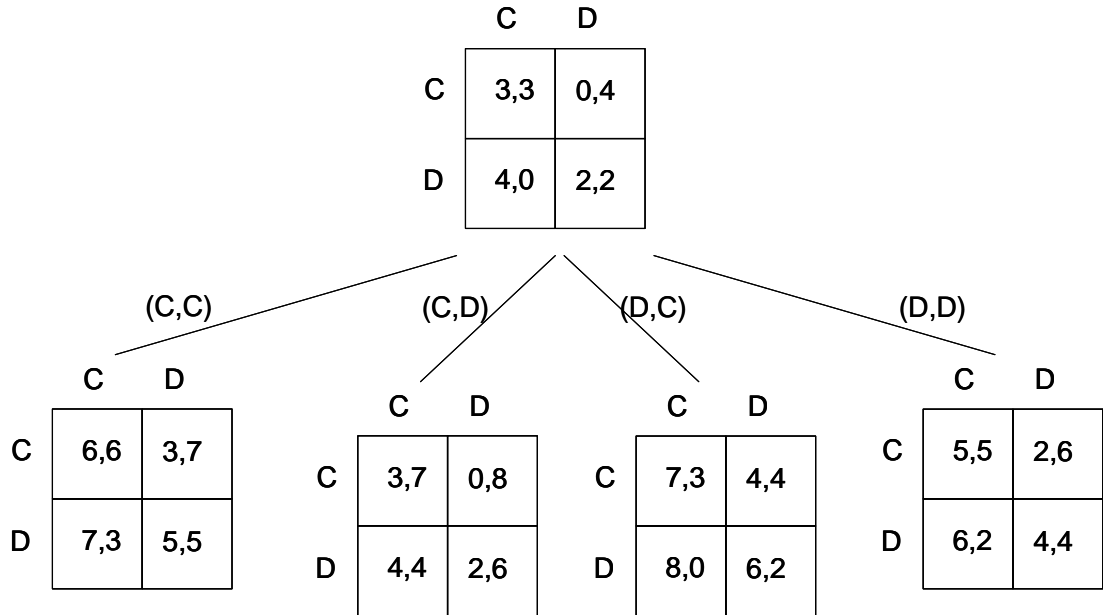
Because you expected punishment. If you stayed out too late they would punish you.

Reward and Punishment. We use them all of the time, and yet how do we analyze them in interactions? Can they be a formal part of equilibrium? Yes they can in a repeated game.

1.1 The Twice (Finite) Repeated Game

In a finitely repeated game we will get the value of a sequence of payoffs just by summing them. Furthermore, while it may not seem obvious, the analysis for the twice repeated game is essentially identical for the $T < \infty$ period repeated game, so throughout this section we will focus completely on the twice repeated game.

How do we represent the twice repeated Prisoner's Dilemma? Well we can do it as a sequential game where each "decision node" is a normal form game, and the "actions" at that decision node are the pairs of strategies that players might play. In other words we can think of it as something like the graph below:



(3)

We derive the payoffs in the second period by adding the payoffs from the first period action pair to every cell in the second period. Notice that these additions do not affect the payoff differences, only their absolute level. Decisions will be based on the payoff differences, thus they will be identical to the decisions in the simultaneous game. Thus we can see that the second decision by both players is going to be always D .

Given these payoffs the total payoff in the first period can be written as:

	C	D
C	5; 5	2; 6
D	6; 2	4; 4

and again the absolute payoffs don't affect incentives, all that matters is the payoff differences which are unchanged. Thus both players have a dominant strategy of playing D . So the only subgame perfect equilibrium is always playing D .

More generally we have the following proposition:

Proposition 1 *If the stage game has only one Nash equilibrium, then in the twice (finitely) repeated game there is only one subgame perfect equilibrium, which is to repeat that equilibrium every time.*

Proof. In the second period history will only add a constant to payoffs, which does not affect incentives and thus a NE must be played no matter what

occured in the first period. Since there is a unique NE that equilibrium must be played.

Now we know that in the first period no matter what happens today the future is fixed—the unique NE will be played in the second period. This means the future is fixed, only adding a constant to payoffs, which does not affect incentives and the unique NE must be played in period one.

To extend this to $T > 2$ just begin in period T , and then the working backward (starting at $t = T - 1$) the second step is applied iteratively (noting that both history and the future only add a constant to payoffs) until you reach $t = 1$.

Remark 2 *To go into more detail, in the final period ($t = T$) the utility of the strategy pair (X, Y) will be $u_1(X, Y) + h_1^T$ and $u_2(X, Y) + h_2^T$ where h_1^T is the amount player 1 earned up to this period, h_2^T is the same for player 2. These constants will not affect incentives, thus whatever is done in this period must be a Nash equilibrium.*

This is always true, what is unique about this game is that there is one NE, thus no matter what (h_1^T, h_2^T) we must always play this (X^, Y^*) .*

Going to $t = T - 1$ since we know that we will always play (X^, Y^*) tomorrow the future is fixed, no matter what is done today it will always be the same. This means our utilities today can be written as $u_1(X, Y) + h_1^{T-1} + f_1^{T-1}$ $u_2(X, Y) + h_2^{T-1} + f_2^{T-1}$ where (h_1^{T-1}, h_2^{T-1}) are the earnings up to the current period and $(f_1^{T-1}, f_2^{T-1}) = U(X^*, Y^*)$. The key point is that since (h_1^{T-1}, h_2^{T-1}) and (f_1^{T-1}, f_2^{T-1}) are both fixed, they just add a constant to payoffs and thus we must play a NE today—which must be the unique NE.*

In $t = T - 2$ we now know that $(f_1^{T-2}, f_2^{T-2}) = 2U(X^, Y^*)$ and all the rest of the analysis goes through, thus in this manner we can work back to $t = 1$, at which point we will have established that we must play (X^*, Y^*) in every period.*

■

How does this change if there are multiple equilibria? Consider the Battle of the Sexes, or as I usually like to think of it the Flirting game.

	S	M	
S	4; 3	1; 1	(4)
M	0; 0	3; 4	

Now the equilibria in the second period are (S, S) or (M, M) . The equilibria in the first period? Obviously it could be either (S, S) or (M, M) again, and you can take it from me there are no other equilibria. The full set of equilibria are:

- 1 : (S, S) , 2 : (S, S)
- 1 : (S, S) , 2 : (M, M)
- 1 : (M, M) , 2 : (S, S)
- 1 : (M, M) , 2 : (M, M)

This probably sounds like the equilibrium you played with that "special someone." Occasionally you would switch coffee shops—just to see if he would

follow you. Of course in equilibrium there still won't be the confusion you had in reality—sometimes you went to different coffee shops, etceteras. Hopefully you can see how this will generalize, if there are $N \geq 2$ Nash equilibria of the stage game then there will be at least N^2 Nash equilibria of the twice repeated game, or in the T period repeated game N^T . Thus if there are multiple stage game equilibria, we find radically different results from if there is only one.

Now let's be honest. When we say that a game has a unique equilibrium it's usually by construction. We basically pare down reality until there is only one sensible strategy for both players. This is useful and a good thing to do when analyzing a static interaction. You want to predict what can happen, not say "it depends." However now we have found an absolutely peculiar result. If we did construct the game so that there is only one equilibrium then we get the weird fact that there will only be one equilibrium in the finitely repeated game. Would this result change dramatically if we did not pare down reality so severely? Yes. So perhaps in the repeated game we should always allow for at least two different equilibria.

How do we get these? Well think about most of your interactions, don't you always have the option of just not getting into the interaction at all? We all do, usually, and this will always result in a Nash equilibrium because if you expect me to avoid you then you can sensibly avoid me as well—in fact often trying to interact and being rebuffed has negative consequences so this is the only optimal strategy if I choose not to interact with you.

When we have multiple equilibria—especially if they are Pareto ranked (one is strictly better for both parties)—we can use a reward and punishment scheme in the last period. If you do the right thing in the first period I reward you with the Pareto Dominant equilibrium, if you do the wrong thing I punish you with the Pareto dominated equilibrium. Using these rewards and punishments we can then get behavior in the first period that is "cooperative." In other words people will do things that are not Nash Equilibria of the static game. Consider the following prisoner's dilemma.

	C	D	N	
C	3; 3	0; 4 ²	0; 0 ¹	(5)
D	4; 0 ¹	2; 2 ^{12*}	0; 0 ¹	
N (not interact)	0; 0 ²	0; 0 ²	0; 0 ^{12*}	

N is the option of no interaction. Usually in these types of situations that will be even worse than (D, D) . Why? Usually (D, D) only comes up sometimes—when you are caught doing something wrong. So usually you get the benefit of the crime without ever paying the cost.

Now in the second period we have two NE, (D, D) and (N, N) and $u_1(D, D) > u_1(N, N)$ plus $u_2(D, D) > u_2(N, N)$ —these equilibria are Pareto ranked.

Now, how can we get you to cooperate? Consider the strategy:

If we do the right thing in the first period we play (D, D) in the second period (reward)

If we do not, then we play (N, N) in the second period (punishment).

If this is the strategy in the second period, would we be willing to play (C, C) in the first period?

$$\begin{aligned} v_1^*((C, C), (D, D)) &= u_1(C, C) + u_1(D, D) = 3 + 2 = 5 \\ \hat{v}_1((D, C), (N, N)) &= u_1(D, C) + u_1(D, D) = 4 + 0 = 4 \end{aligned}$$

yes we would!

The main problem with these types of equilibria is that formal analysis is basically too complicated. Thus to find them you basically have to guess and verify. These equilibria, in general, might be very complicated, but in this class we will always use ones where "if you do the right thing we play the Pareto Dominant NE, if you do the wrong thing we play the other one."

The second Problem with these types of equilibria is that in general we can get really strange payoffs, like playing (N, N) almost all of the time. I.e. while we can support Pareto Efficient payoffs we can also support payoffs that are worse than the best NE.

You don't think this happens in reality? You think all equilibria are Pareto Efficient? How about the social norm that women must work at home? In the old days it made sense (for the poor and middle class) because of the amount of work needed to maintain a house. Imagine doing laundry without a washing machine, and in the same day cooking dinner on a wood stove. It was not easy. But then why where rich landlords' wives not allowed to pursue careers? What about brilliant women who could have significantly advanced science? They existed, the classic example is Marie Curie. She was born in 1867 and the first person to win two Nobel prizes (1903 and 1911). She is famous for her work on radioactivity. Not all repeated game equilibria are Pareto Efficient just like not all static Nash equilibria are Pareto Efficient.

1.2 The Infinitely Repeated Game

Another case that is simple (just my little joke, "simple" indeed) to analyze is the infinitely repeated game. Why? Because now there is no last period, no backward induction. The value will still be the sum of payoffs, but we have to add a discount factor δ , $0 < \delta < 1$. In other words the payoffs t periods in the future will be discounted by δ^{t-1} . Otherwise almost all values are infinite, and we don't want to compare infinities.

What does δ mean? It can be interpreted in several ways. $\delta \rightarrow 1$ can either mean people are more patient or that they interact more often. Notice that $\delta = \frac{1}{1+\rho}$, where $\rho > 0$ is the discount rate. Loosely speaking ρ is their personal interest rate, and should be approximately the same as the real interest rate. (This equivalence is precise when we are analyzing firms, in fact r should be the opportunity cost of capital). Notice that if interaction is more frequent obviously the relevant r will decrease (daily interest rate versus annual), and thus so should ρ . We can also interpret $1 - \delta$ as the probability of interaction breaking down. In other words in each period with probability δ the players interact that period (and potentially in the future) however with probability

$1 - \delta$ players stop interacting forever. If we use this interpretation we have to rescale payoffs, but this has no impact on analysis.

Then the value of a constant stream of x is easy to calculate:

$$\begin{aligned} V(x) &= x + \delta x + \delta^2 x + \delta^3 x + \dots \\ (1 - \delta)V(x) &= x + \delta x + \delta^2 x + \delta^3 x + \dots \\ &\quad 0 - \delta x - \delta^2 x - \delta^3 x - \dots \\ (1 - \delta)V(x) &= x \\ V(x) &= \frac{1}{1 - \delta}x \end{aligned}$$

Frankly this math magic is one reason we use this model of discounting, it is so easy to calculate the value of constant streams of payoffs. Like before we will only consider a very simple class of strategies. (Too much complexity can be a real pain in the neck here. Almost suicidal.) These strategies will all be: (A, B) today if we played (A, B) yesterday. (C, D) today otherwise. In this strategy (A, B) is the "reward." and (C, D) is the "punishment." Notice that these strategies are especially simple because cooperation is its own reward. These are the *Grimm* (after the Brother's Grimm Fairy Tales) or the *Trigger* strategies. The term trigger strategies is particularly descriptive because one mistake leads to an infinite amount of punishment. It's like the person you're playing with pulled a trigger and shot you. These strategies are the harshest, and harsher than most of us will use in our daily life, but we analyze them because they are simple.

These strategies will always be *Subgame Perfect equilibria* for high enough δ if:

1. (C, D) is a NE of the stage game.
2. $u_1(A, B) > u_1(C, D)$, $u_2(A, B) > u_2(C, D)$ (people are always doing better at (A, B)).

Let me give you some examples:

Example 3 *Prisoner's dilemma with the Payoffs Re-normalized.*

	C	D	
C(cooperate)	1; 1	-2; 2	(6)
D(defect or cheat)	2; -2	0; 0	

The strategy we are interested in is (C, C) today if we played (C, C) yesterday. (D, D) today otherwise. The value of following the equilibrium strategy given you played (C, C) yesterday is:

$$\begin{aligned} V|(C, C) &= 1 + \delta + \delta^2 + \delta^3 \dots \\ &= \frac{1}{1 - \delta} \end{aligned}$$

the value you get from deviating is:

$$\begin{aligned}\hat{V}|(C, C) &= 2 + \delta * 0 + \delta^2 * 0 + \delta^3 * 0 + \delta^4 * 0 \dots \\ &= 2\end{aligned}$$

$$\begin{aligned}V|(C, C) &\geq \hat{V}|(C, C) \\ \frac{1}{1-\delta} &\geq 2 \\ 1 &\geq 2(1-\delta) \\ 1 &\geq 2-2\delta \\ 2\delta &\geq 1 \\ \delta &\geq \frac{1}{2}\end{aligned}$$

so we conclude that if $\delta \geq \frac{1}{2}$ then this is an equilibrium.

Example 4 *Restaurant Quality Game*

Lets look at another game, which is not symmetric but is always easy to solve.

	B	N
H	2; 2	-1; 0
L	3; -1	0; 0

Consider the strategy (H, B) today if we played (H, B) yesterday. (N, N) today otherwise.

$$\begin{aligned}V_1|(H, B) &= 2 + 2\delta + 2\delta^2 + \dots \\ &= \frac{2}{1-\delta} \\ \hat{V}_1|(H, B) &= 3 + \delta * 0 + \delta^2 * 0 + \delta^3 * 0 + \delta^4 * 0 \dots \\ &= 3\end{aligned}$$

$$\begin{aligned}V_1|(H, B) &\geq \hat{V}_1|(H, B) \\ \frac{2}{1-\delta} &\geq 3 \\ 2 &\geq 3-3\delta \\ \delta &\geq \frac{1}{3}\end{aligned}$$

$$\begin{aligned}
V_2|(H, B) &= 2 + 2\delta + 2\delta^2 + \dots \\
&= \frac{2}{1 - \delta} \\
\hat{V}_2|(H, B) &= 0 + \delta * 0 + \delta^2 * 0 + \delta^3 * 0 + \delta^4 * 0 \dots \\
&= 0 \\
V_2|(H, B) &\geq \hat{V}_2|(H, B) \\
\frac{2}{1 - \delta} &\geq 0
\end{aligned}$$

which is always true. In fact I didn't have to check the customer at all, all I really needed to say was notice that he is always best responding to his beliefs in the static game.

Example 5 *An Alternative Strategy in the Restaurant Game.*

Grimm or Trigger strategies are unreasonable harsh. Most people never use them, they avoid a restaurant for a while and then start going to it again. This is the normal strategy. Will it work as an equilibrium? Yes, as long as they avoid the restaurant for long enough. In my example one period is enough. Thus the strategy we are interested in is: (H, B) today unless they were supposed to play (H, B) yesterday and did not. (N, N) today otherwise. Notice that this last case means: only if we were supposed to play (H, B) yesterday and we did not.

$$\begin{aligned}
V_1|(H, B) &= \frac{2}{1 - \delta} \\
\hat{V}_1|(H, B) &= 3 + \delta * 0 + 2\delta^2 + 2\delta^3 + 2\delta^4 + \dots \\
&= 3 + 2\delta^2 (1 + \delta + \delta^2 + \dots) \\
&= 3 + \frac{2\delta^2}{1 - \delta}
\end{aligned}$$

$$\begin{aligned}
\frac{2}{1 - \delta} &\geq 3 + \frac{2\delta^2}{1 - \delta} \\
\frac{2}{1 - \delta} - \left(\frac{2\delta^2}{1 - \delta} \right) &\geq 3 \\
\frac{2(1 - \delta^2)}{1 - \delta} &\geq 3
\end{aligned}$$

Now notice that $(1 - \delta^2) = (1 - \delta)(1 + \delta)$, thus this is the same as:

$$\begin{aligned}
\frac{2(1 - \delta)(1 + \delta)}{1 - \delta} &\geq 3 \\
2(1 + \delta) &\geq 3 \\
2\delta &\geq 1 \\
\delta &\geq \frac{1}{2}
\end{aligned}$$

Will this work for you? Well if we interpret δ as "the frequency of interaction" then if you go to a restaurant every day then δ will be high. If you go to the restaurant only once in a great deal of time then δ will be very low. The basic intuition is that the more frequently you go to the restaurant the fewer number of times you have to punish them. The less frequently you go to a restaurant the more you have to punish them (in terms of number of times you don't go back.) To understand this intuition notice that the one period punishment works if $\delta \geq \frac{1}{2}$, the Grimm strategy works if $\delta \geq \frac{1}{3}$. This means that any time the one period strategy will work so will the Grimm strategy, but if δ is low enough your only option is the Grimm strategy.

1.2.1 An Inefficient example.

	H	C	D	
H	2; 2	0; 0	-2; x^2	(7)
C	0; 0	1; 1	-2; 2^2	
D	x ; -2^1	2 ; -2^1	0 ; 0^{12}	

$$x > 2$$

Notice that in this game, just like in the Prisoner's Dilemma above, (C, C) is an equilibrium if $\delta \geq \frac{1}{2}$. However now we have a new, Pareto Dominant, option of playing (H, H) forever. Consider our standard strategy of (H, H) today if we played (H, H) yesterday. (D, D) today otherwise.

$$\begin{aligned} V|(H, H) &= \frac{2}{1-\delta} \\ \hat{V}|(H, H) &= x + \delta * 0 + \delta^2 * 0 + \delta^3 * 0 + \delta^4 * 0 \dots \\ &= x \end{aligned}$$

$$\begin{aligned} V|(H, H) &\geq \hat{V}|(H, H) \\ \frac{2}{1-\delta} &\geq x \\ 2 &\geq x(1-\delta) = x - x\delta \\ x\delta &\geq x - 2 \\ \delta &\geq \frac{x-2}{x} \end{aligned}$$

$\frac{x-2}{x} < 1$ so they will cooperate, but will it require a higher δ ? A lower δ ? Who really can say? And what do we think the equilibrium will be when both strategies are equilibrium? We can not say, hopefully it will be the Pareto Dominant equilibrium, but it may not be. It depends on society, and what society expects.