9

Repeated Games

Tomorrow, and tomorrow, and tomorrow,
Creeps in this petty pace from day to day
To the last syllable of recorded time

Shakespeare

When a game $\mathcal{G}$ is repeated an indefinite number of times by the same players, many of the anomalies associated with finitely repeated games (§4.2) disappear. Nash equilibria of the repeated game arise that are not Nash equilibria of $\mathcal{G}$. The exact nature of these equilibria is the subject of the folk theorem (§9.10). We have encountered many games $\mathcal{G}$ in which most, or all Nash equilibria are Pareto-inefficient. Indeed, all the generic two-player games with two strategies, the prisoner's dilemma (§3.11), the battle of the sexes (§3.9), and the hawk-dove game (§3.10) are of this type.[1] The folk theorem asserts that if the signals of defection (that is, the signals that a player deviated from the behavior specified by the Nash equilibrium) are of sufficiently high quality, and if players have sufficiently long time horizons, the repeated game based on $\mathcal{G}$ can attain Pareto-efficiency, or at least approximate Pareto-efficiency as closely as desired.

The folk theorem requires that a defection on the part of a player carry a signal that is conveyed to other players. We say a signal is *public* if all players receive the same signal. We say the signal is *perfect* if it accurately reports whether or not the player in question defected. The first general folk theorem that does not rely on incredible threats was proved by Fudenberg and Maskin (1986) for the case of perfect public signals (§9.10).

Repeated game models provide elegant and compelling explanations of many on-going strategic interactions. The folk theorem, however, is highly overrated. The folk theorem achieves its magic through the unrealistic assumption that the defection signal can be made arbitrarily close to public and arbitrarily accurate. For instance, as the number of players increases,

[1] We say a finite normal form game is *generic* if all the entries in the game matrix are distinct. More generically, we say that a property of a set of equations is generic if, when it holds for a particular set of parameters, it holds in a sufficiently small neighborhood of these parameters.

the folk theorem continues to hold only if the signal becomes more accurate and closer to public. In fact, of course, as group size increases, the signal will, under plausible conditions, become both less public and noisier. A cogent analysis of cooperation in repeated games with self-regarding players under empirically plausible conditions remains to be developed.

9.1 Death and Discount Rates in Repeated Games

Suppose an agent plays a repeated game in which the payoff at the end of each period is π, the agent's discount rate is ρ, and the probability that the agent dies at the end of each period is $\sigma > 0$. We then can write the present value v as

$$v = \frac{\pi + (1-\sigma)v}{1+\rho},$$

because the agent receives π and, unless he dies, plays the lottery in which he wins v again, both at the end of the current period. Solving for v,we get

$$v = \frac{\pi}{\sigma + \rho}.$$

This result gives us some information as to what plausible discount rates are for humans in cases where we cannot arbitrage our gains and losses by banking them at the current interest rate. Given human life expectancy, this argument suggest a discount rate of about 2% to 3% per year.

9.2 Big Fish and Little Fish

Many species of fish are attacked by parasites that attach to their gills and inner mouth parts. Often such fish will form a symbiotic relationship with a smaller species of fish for whom the parasite is a food source. Mutual trust is involved, however, because the larger fish must avoid the temptation of eating the smaller fish, and the smaller fish must avoid the temptation of taking a chunk out of the larger fish, thereby obtaining a meal with much less work than picking around for the tiny parasites. This scenario, which is doubtless even more common and more important for humans than for fish, is explored in the following problem.

Suppose Big Fish and Little Fish play the prisoner's dilemma shown in the diagram. Of course, in the one-shot game there is only one Nash equilibrium, which dictates that both parties defect. However, suppose the same players play the game at times $t = 0, 1, 2, \ldots$ This is then a new game, called a *repeated game*, in which the payoff to each is the sum of the payoffs over all periods, weighted by a *discount factor* δ, with $0 < \delta < 1$. Note that a discount factor δ relates to a *discount rate* ρ by the formula $\delta = 1/(1 + \rho)$. We call the game played in each period the *stage game* of a *repeated game* in which at each period the players can condition their moves on the complete previous history of the various stages. A strategy that dictates following one course of action until a certain condition is met, and then following a different strategy for the rest of the game is called a *trigger strategy*.

	C	*D*
C	5,5	-3,8
D	8,-3	0,0

THEOREM 9.1 *The cooperative solution (5,5) can be achieved as a subgame perfect Nash equilibrium of the repeated game if δ is sufficiently close to unity, and each player uses the trigger strategy of cooperating as long as the other player cooperates, and defecting forever if the other player defects on one round.*

PROOF: We use the fact that for any discount factor δ with $0 < \delta < 1$,

$$1 + \delta + \delta^2 + \ldots = \frac{1}{1 - \delta}.$$

To see this, write

$$\begin{aligned} x &= 1 + \delta + \delta^2 + \ldots \\ &= 1 + \delta(1 + \delta + \delta^2 + \ldots) = 1 + \delta x, \end{aligned}$$

from which the result follows.

By the way, there is a faster way of arriving at the same result. Consider a repeated game that pays 1 now and in each future period, and the discount factor is δ. Let x be the value of the game to the player. The player receives 1 now, and then gets to play exactly the same game in the next period. Because the value of the game in the next period is x, its present value is δx. Thus $x = 1 + \delta x$, so $x = 1/(1 - \delta)$.

Now suppose both agents play the trigger strategy. Then, the payoff to each is $5/(1 - \delta)$. Suppose a player uses another strategy. This must involve cooperating for a number (possibly zero) of periods, then defecting forever; for once the player defects, his opponent will defect forever, the best response to which is to defect forever. Consider the game from the time t at which the first player defects. We can call this $t = 0$ without loss of generality. A fish that defects receives 8 immediately and nothing thereafter. Thus the cooperate strategy is Nash if and only if $5/(1-\delta) \geq 8$, or $\delta \geq 3/8$. When δ satisfies this inequality, the pair of trigger strategies is also subgame perfect, because the situation in which both parties defect forever is Nash subgame perfect. ■

There are lots of other subgame perfect Nash equilibria to this game. For instance,

THEOREM 9.2 *Consider the following trigger strategy for Little Fish: alternate $C, D, C, \ldots$ as long as Big Fish alternates $D, C, D, \ldots$. If Big Fish deviates from this pattern, defect forever. Suppose Big Fish plays the complementary strategy: alternate $D, C, D, \ldots$ as long as Little Fish alternates $C, D, C, \ldots$. If Little Fish deviates from this pattern, defect forever. These two strategies form a subgame perfect Nash equilibrium for δ sufficiently close to unity.*

PROOF: The payoffs are now $-3, 8, -3, 8, \ldots$ for Little Fish and $8, -3, 8, -3, \ldots$ for Big Fish. Let x be the payoffs to Little Fish. Little Fish gets -3 today, 8 in the next period, and then gets to play the game all over again starting two periods from today. Thus, $x = -3 + 8\delta + \delta^2 x$. Solving this, we get $x = (8\delta - 3)/(1 - \delta^2)$. The alternative is for Little Fish to defect at some point, the most advantageous time being when it is his turn to get -3. He then gets zero in that and all future periods. Thus, cooperating is Nash if and only if $x \geq 0$, which is equivalent to $8\delta - 3 \geq 0$, or $\delta \geq 3/8$. ■

9.3 Alice and Bob Cooperate

Alice and Bob play the game in the figure to the right an indefinite number of times. They use trigger strategies but do not discount the future. Show that if the probability p of continuing the game in each period is sufficiently large, then it is Nash for both Alice and Bob to cooperate (play C) in each period. What is the smallest value of p for which this is true?

Alice \ Bob	C	D
C	3,3	0,5
D	5,0	1,1

To answer this, let v be the present value of cooperating forever for Alice. Then $v = 3 + pv$, because cooperation pays 3, plus with probability p Alice gets the present value v again in the next period. Solving, we get $v = 3/(1-p)$. If Alice defects, she gets 5 now and then 1 forever, starting in the next period. The value of getting 1 forever is $v_1 = 1 + p \cdot v_1$, so $v_1 = p/(1-p)$. Thus Alice's total return to defecting is $5 + p/(1-p)$. Cooperating beats defecting for Alice if $3/(1-p) > 5+p/(1-p)$. Solving, we find Alice should cooperate as long as $p > 50\%$.

9.4 The Strategy of an Oil Cartel

	Low	High
Low	184,168	104,176
High	208,88	128,96

Suppose there are two oil-producing countries, Iran and Iraq. Both can operate at either of two production levels: 2 or 4 million barrels a day. Depending on their decisions, the total output on the world market will be 4, 6, or 8 million barrels a day, and the price per barrel in these three cases is \$100, \$60, and \$40, respectively. Costs of production are \$8 per barrel for Iran and \$16 per barrel for Iraq. The normal form of the game is shown in the diagram. It is clear a prisoner's dilemma.

Suppose this game is repeated every day, and both countries agree to cooperate by producing the low output, each one threatening the other with a *trigger strategy:* "If you produce high output, even once, I will produce high output forever." Show that cooperation is now a Nash equilibrium if the discount rate is sufficiently low. What is the maximum discount rate that will sustain cooperation?

9.5 Reputational Equilibrium

Consider a firm that can produce a good at any quality level $q \in [0,1]$. If consumers anticipate quality q_a, their demand x is given by

$$x = 4 + 6q_a - p.$$

Suppose the firm knows this demand curve, and takes q_a as given but can set the quality q supplied. The firm has no fixed costs, and the cost of producing one unit of the good of quality q is $2 + 6q^2$.

In each period $t = 1, 2, \ldots$ the firm chooses a quality level q and a price p. Consumers see the price but do not know the quality until they buy the good. Consumers follow a trigger strategy, in which they buy the good in each period in which $q \geq q_a$, but if $q < q_a$ in some period, they never buy from the firm again.

Suppose the firm uses discount factor $\delta = 0.9$. Define a *reputational equilibrium* as one in which quality q_a is supplied in each period. What are the conditions for a reputational equilibrium?

9.6 Tacit Collusion

Consider a duopoly operating over an infinite number of periods $t = 1, 2, \ldots$. Suppose the duopolists are price setters, so each pure strategy for firms 1 and 2 in period t takes the form of setting prices $p_1^t, p_2^t \geq 0$, respectively, conditioned on the history of prices in previous time periods, and a pure strategy for the whole game is a sequence of strategies, one for each period t. Suppose the profits in period t are given by $\pi_1(p_1^t, p_2^t)$ for firm 1 and $\pi_2(p_1^t, p_2^t)$ for firm 2. The payoffs to the firms for the whole game are then

$$\pi_1 = \sum_{t=1}^{\infty} \delta^t \pi_1(p_1^t, p_2^t), \qquad \pi_2 = \sum_{t=1}^{\infty} \delta^t \pi_2(p_1^t, p_2^t),$$

where δ is the common discount factor for the firms.

To specify the function $\pi_i(p_1^t, p_2^t)$, suppose the two firms have no fixed cost and constant marginal cost c, the firms face a downward-sloping demand curve, and the lowest price producer gets the whole market. Also, if the two producers have the same price, they share the market equally.

THEOREM 9.3 *There is a subgame perfect Nash equilibrium of this game in which* $p_1^t = p_2^t = c$ *for* $t = 1, 2, \ldots$.

Note that this is the "competitive" equilibrium in which profits are zero and price equals marginal cost. The existence of this Nash equilibrium is called Bertrand's paradox because it seems implausible (though hardly paradoxical!) that two firms in a duopoly actually behave in this manner.

THEOREM 9.4 *Suppose* $\delta > 50\%$, p^m *is the monopoly price (that is, the price that maximizes the profits of a monopolist) and* $c \leq p \leq p^m$. *Let* s *be a strategy profile which firm 1 sets* $p_1^t = p$, *firm 2 sets* $p_2^t = p$, $t = 1, 2, \ldots$, *and each firm responds to a deviation from this behavior on* *the part of the other firm by setting price equal to c forever. Then s is a subgame perfect Nash equilibrium.*

We call a Nash equilibrium of this type *tacit collusion*.

PROOF: Choose p satisfying the conditions of the theorem. Let $\pi(p)$ be total industry profits if price p is charged by both firms, so $\pi(p) \geq 0$. Suppose firm 2 follows the specified strategy. The payoff to firm 1 for following this strategy is

$$\pi^1(p) = \frac{\pi(p)}{2}\left(1 + \delta + \delta^2 + \ldots\right) = \frac{\pi(p)}{2(1-\delta)}. \tag{9.1}$$

The payoff to firm 1 for defecting on the first round by charging an amount $\epsilon > 0$ less than p is $\pi^1(p - \epsilon)$. Thus, $2(1-\delta) < 1$ is sufficient for Nash. Clearly, the strategy is subgame perfect, because the Bertrand solution in which each firm charges marginal cost is Nash. ■

Intuition tells us that tacit collusion is more difficult to sustain when there are many firms. The following theorem, the proof of which we leave to the reader (just replace the "2" in the denominators of equation (9.1) by "n") shows that this is correct.

THEOREM 9.5 *Suppose there are $n > 1$ firms in the industry, but the conditions of supply and demand remain as before, the set of firms with the lowest price in a given time period sharing the market equally. Then, if $\delta > 1 - 1/n$ and $c \leq p \leq p^m$, the trigger strategies in which each firm sets a price equal to p in each period, and each firm responds to a deviation from this strategy on the part of another firm by setting price equal to c forever, form a subgame perfect Nash equilibrium.*

Another market condition that reduces the likelihood that tacit collusion can be sustained is incomplete knowledge on the part of the colluding firms. For instance, we have the following.

THEOREM 9.6 *Suppose there is an n-firm oligopoly, as described previously, but a firm that has been defected upon cannot implement the trigger strategy until $k > 1$ periods have passed. Then, tacit collusion can hold only if $\delta^{k+1} > 1 - 1/n$.*

This theorem shows that *tacit collusion is less likely to hold the more easily a firm can "hide" its defection*. This leads to the counterintuitive, but quite correct, conclusion that making contractual conditions public knowledge ("putting all your cards on the table") may *reduce* rather than *increase* the degree of competition in an industry.

PROOF: If it takes k periods after a defection to retaliate, the gain from defection is

$$\pi(p)(1+\delta+\delta^2+\ldots+\delta^k) = \frac{\pi(p)}{1-\delta}\left(1-\delta^{k+1}\right),$$

from which the result immediately follows, because the present value of cooperating forever is $\pi(p)/n(1-\delta)$. ∎

9.7 The One-Stage Deviation Principle

Suppose s is a strategy profile for players in a repeated game, so s specifies what move each player makes at each stage of the game, depending on the prior history of moves in the game. We say s satisfies the *one-stage deviation principle* if no player can gain by deviating from s, either on or off the equilibrium path of the game tree, in a single stage, otherwise conforming to s. The following theorem is often very useful in analyzing repeated games, because it says that a strategy profile is subgame perfect if it satisfies the one stage deviation principle.

THEOREM 9.7 The One-Stage Deviation Principle. *A strategy profile s for a repeated game with positive discount rates, based on a finite stage game, is a subgame perfect Nash equilibrium if and only if it satisfies the one-stage deviation principle.*

For a formal proof of this theorem (it's not difficult) and some extensions see Fudenberg and Tirole (1991). Here is an informal proof. Obviously, subgame perfection implies the one-stage deviation principle. Suppose s is not subgame perfect. Then for some player i, there is an alternative strategy profile $\tilde{s}_i$ that offers i a higher payoff against s_{-i}, starting at one of i's information sets ν_i. If $\tilde{s}_i$ differs from s_i in only a finite number of places, we can assume $\tilde{s}_i$ has the fewest deviations from s_i among such alternatives for i. If $\tilde{s}_i$ has only one deviation, we are done. If there are more than one, then the final deviation must make i better off and again we are done. If $\tilde{s}_i$ has an infinite number of deviations, because the payoffs are bounded, and the discount rate is strictly positive, we can cut the game off at some period

t such that any possible gain to i after time t are as small as we want. This shows that there is also a finite deviation (cutting off $\tilde{s}_i$ after time t) that improves i's payoff starting at ν_i.

9.8 Tit for Tat

Consider a repeated prisoner's dilemma with two players. The payoffs are (r, r) for mutual cooperation, (p, p) for mutual defection, (t, s) when player 1 defects, and player 2 cooperates, and (s, t) when player 1 cooperates, and player 2 defects, where $t > r > p > s$ and $2r > s + t$. A tit for tat player cooperates on the first round of a repeated prisoner's dilemma, and thereafter does what the other player did on the previous round.

It is clear that tit for tat is not Nash when played against itself if and only if player 1 gains by defecting on the first move. Moreover, defecting on the first round, and going back to tit for tat when you partner plays tit for tat increases your payoff by $t + \delta s - r(1 + \delta)$, which is negative for $\delta_1 > (t - r)/(r - s)$. Note that $0 < \delta_1 < 1$.

After the first defect, the gain from defecting for k more rounds before returning to cooperate is $\delta^{k+1}(p - s(1 - \delta) - \delta r)$. If this is negative, then δ_1 is the minimum discount factor for which tit for tat is Nash against itself. If this is positive, however, then defecting forever increases payoff by $t - r - (r - p)\delta/(1 - \delta)$, which is negative for a discount factor greater than $\delta_2 = (t - r)/(t - p)$. Because $0 < \delta_2 < 1$, we find that there is always a discount factor less than unity, above which tit for tat is a Nash equilibrium.

However, tit for tat is not subgame perfect, because if a player defects on a particular round, tit for tat specifies that the two players will exchange cooperate and defect forever, which has a lower payoff than cooperating forever. Fortunately, we can revise tit for tat to make it subgame perfect. We define *contrite tit for tat* (Boyd 1989) as follows. A player is in either *good standing* or *bad standing*. In period 1 both players are in good standing. A player is in good standing in period $t > 1$ only if in period $t - 1$ (i) he cooperated and his partner was in good standing; (ii) he cooperated and he was in bad standing; or (iii) he defects and he was in good standing, while his partner was in bad standing. Otherwise the player is in bad standing in period $t > 1$. Contrite tit for tat says to cooperate unless you are in good standing and your partner is in bad standing. Using the one-stage deviation principle (§9.7), it is then easy to show that contrite tit for tat is subgame perfect.

9.9 I'd Rather Switch Than Fight

Consider a firm that produces a *quality good*, which is a good whose quality is costly to produce, can be verified only through consumer use, and cannot be specified contractually. In a single-period model, the firm would have no incentive to produce high quality. We develop a repeated game between firm and consumer, in which the consumer pays a price *greater* than the firm's marginal cost, using the threat of switching to another supplier (a trigger strategy) to induce a high level of quality on the part of the firm. The result is a nonclearing product market, with firms enjoying price greater than marginal cost. Thus, they are quantity constrained in equilibrium.

This model solves a major problem in our understanding of market economies: markets in quality goods do not clear, and the success of a firm hinges critically on its ability to sell a sufficient quantity of it product, something that is assured in a clearing market. Thus, the problem of competition in quality-goods markets is quite different from Walrasian general equilibrium models, in which the only problem is to produce at minimum cost.

Every Monday, families in Pleasant Valley wash their clothes. To ensure brightness, they all use bleach. Low-quality bleach can, with low but positive probability, ruin clothes, destroy the washing machine's bleach delivery gizmo, and irritate the skin. High-quality bleach is therefore deeply pleasing to Pleasant Valley families. However, high-quality bleach is also costly to produce. Why should firms supply high quality?

Because people have different clothes, washing machines, and susceptibility to skin irritation, buyers cannot depend on a supplier's reputation to ascertain quality. Moreover, a firm could fiendishly build up its reputation for delivering high-quality bleach and then, when it has a large customer base, supply low-quality bleach for one period, and then close up shop (this is called "milking your reputation"). Aggrieved families could of course sue the company if they have been hurt by low-quality bleach but such suits are hard to win and very costly to pursue. So no one does this.

If the quality q of bleach supplied by any particular company can be ascertained only after having purchased the product, and if there is no way to be compensated for being harmed by low-quality bleach, how can high quality be assured?

Suppose the cost to a firm of producing a gallon of the bleach of quality q is $b(q)$, where $b(0) > 0$ and $b'(q) > 0$ for $q \geq 0$. Each consumer