

Repeated Games with Perfect Monitoring

Niccolò Lomys*

March 12, 2020

These lecture notes heavily draw upon [Mailath \(2019\)](#), [Hörner \(2015\)](#), [Ali \(2011\)](#), [Mailath and Samuelson \(2006\)](#), [Kandori \(2006\)](#), and [Fudenberg and Tirole \(1991\)](#). All errors are my own. Please bring any error, including typos, to my attention.

1 Introduction

In real life, most games are played within a larger context, and actions in a given situation affect not only the present situation but also the future situations that may arise. When a player acts in a given situation, he takes into account not only the implications of his actions for the current situation but also their implications for the future. If players are patient and current actions have significant implications for the future, then considerations about the future may take over. This may lead to a rich set of behavior that may seem to be irrational when one considers the current situation alone. For instance, opportunistic behavior may be deterred using future rewards and punishments. However, care needs to be taken to make sure that future rewards and punishments are self-enforcing. Such ideas are captured in the repeated games' setup, in which a stage game is played repeatedly over time.

1.1 An Example

Consider the *prisoner's dilemma* for which the *reward* matrix is given by

		Player 2	
		<i>C</i>	<i>D</i>
Player 1	<i>C</i>	1, 1	$-L, 1 + G$
	<i>D</i>	$1 + G, -L$	0, 0

Here, $G, L \in \mathbb{R}_{++}$ measure respectively the additional gain of one player when he defects (i.e. plays *D*) while his opponent cooperates (i.e. plays *C*), and the loss of the latter player. Assume that $G - L < 1$, so that the sum of rewards when both players

*Toulouse School of Economics, University of Toulouse Capitole; niccolo.lomys@tse-fr.eu.

cooperate (equal to 2) exceeds the sum of rewards when one defects (equal to $1 + G - L$). The prisoner's dilemma is the normal form of a two-player, simultaneous-move game. Although mutual cooperation maximizes the two players' joint surplus, defecting is strictly dominant for each player, and so in an isolated interaction, both players defect. Thus, the only Nash equilibrium of the prisoner's dilemma is inefficient.

What happens if players play the prisoner's dilemma over an infinite time *horizon* $t = 0, 1, 2, \dots$ and after each play, both players can observe what the opponent has done? This is an *infinitely repeated game with perfect monitoring*, with the prisoner's dilemma as *stage game*. Consider the following *repeated game strategy*: (i) play (C, C) if $t = 0$ or $t \geq 1$ and (C, C) was played in every prior period; (ii) play (D, D) if some player played D in some prior period. This can be interpreted as an explicit or implicit agreement of the two players: cooperate, or any deviation triggers perpetual mutual defection. Let us now check if each player has any incentive to deviate from this strategy.

- If no one deviates from playing (C, C) , each player enjoys reward 1 in every period. A deviating player enjoys reward $1 + G$ in the period of the deviation but his future reward drops from 1 to 0 in each and every future period. Now assume that the players discount future rewards by the *discount factor* $\delta \in (0, 1)$. The discounted future loss is $\sum_{t=1}^{\infty} \delta^t = \delta/(1 - \delta)$. If this is larger than the short-term gain from defection (equal to G), no one wants to deviate from cooperation. The condition is $\delta/(1 - \delta) \geq G$ or, equivalently, $\delta \geq G/(1 + G)$. That is, players need to be sufficiently *patient*.
- Next, let us check if the players have an incentive to carry out the threat (perpetual defection). Since (D, D) is the Nash equilibrium of the stage game, defecting in each period is a best reply if the opponent always does so. Hence, the players are choosing mutual best replies. In this sense, perpetual defection is a *self-enforcing (or credible) threat*.

To sum up, under the strategy defined above, players are choosing mutual best replies *after any history*, as long as $\delta \geq 1/(1 + G)$. That is, the strategy constitutes a *subgame-perfect Nash equilibrium* of the repeated game. Similarly, in a general infinitely repeated game with perfect monitoring, any outcome which Pareto dominates the Nash equilibrium can be sustained by a strategy that reverts to the Nash equilibrium after a deviation. Such a strategy is called a *grim-trigger strategy*.

Multiple Equilibria. The grim-trigger strategy profile is not the only subgame-perfect Nash equilibrium of the repeated game. The repetition of the stage-game Nash equilibrium (D, D) is also a subgame-perfect Nash equilibrium. Are there any other equilibria? Can we characterize all equilibria in a repeated game? The latter question appears to be formidable at first sight, because there is an infinite number of repeated game strategies, and they can be quite complex. We do have, however, some complete characterizations

of all equilibria of a repeated game, such as *folk theorems* and *self-generation* conditions, as we will discuss.

Credibility of Threats and Renegotiation. One may question the credibility of the threat of defecting forever. In the above example, credibility was formalized as the subgame-perfect Nash equilibrium condition. According to this criterion, the threat is credible because a unilateral deviation by a single player is never profitable. The threat, however, may be upset by *renegotiation*. When players are called upon to carry out this grim threat after a deviation, they may well get together and agree to “let bygones be bygones”. After all, when there is a better equilibrium in the repeated game (for example, the grim-trigger strategy equilibrium), why do we expect the players to stick to the inefficient perpetual mutual defection? This is the problem of *renegotiation proofness* in repeated games. The problem is trickier than it appears. You may get a sense of the difficulty from the following observation. Suppose the players have successfully renegotiated away perpetual mutual defection to play the grim trigger strategy equilibrium again. This is self-defeating, however, because the players now have an incentive to deviate, as they may well anticipate that the threat of perpetual mutual defection will be again subject to renegotiation and will not be carried out.

Finite versus Infinite Horizon. Suppose the prisoner’s dilemma is played *finitely* many times. In the last period, players just play the stage game, and so they play the stage game Nash equilibrium (D, D) . In the penultimate period, they rationally anticipate that they will play (D, D) irrespective of their current play. Hence they are effectively playing the stage game in the penultimate day, and again they play (D, D) . By induction, the only equilibrium of the finitely repeated prisoner’s dilemma is to play (D, D) in every period. This conclusion is driven by *end-game effects*: play in the final period determines play in earlier periods. The impossibility of cooperation holds no matter how long the finite time horizon is. This is in sharp contrast to the infinite horizon case.

If players believe that end-game effects are important, then the finite-horizon model is a better description of their interaction. Although one may argue that players do not live infinitely long, however, there are some good reasons to consider the infinite horizon model. First, even though the time horizon is finite, if players do not know exactly when the game ends, the situation can be formulated as an infinitely repeated game. Suppose that, with probability $r > 0$, the game ends at the end of any given period. This implies that, with probability one, the game ends in a finite horizon. However, the expected *discounted payoff* is equal to $\sum_{t=0}^{\infty} \delta^t (1-r)^t \pi(t)$, where $\pi(t)$ is the reward in period t . This is identical to the payoff in an infinitely repeated game with discount factor $\delta' := \delta(1-r)$. Second, the drastic “discontinuity” between the finite and infinite horizon cases in the repeated prisoner’s dilemma example hinges on the uniqueness of equilibrium in the stage game. [Benoît and Krishna \(1985\)](#) show that, if each player has multiple equilibrium pay-

offs in the stage game, the long but finite-horizon case enjoys the same scope for cooperation as the infinite horizon case (that is, the Folk Theorem, discussed below for infinitely repeated games, approximately holds for the T -period repeated game, when $T \rightarrow \infty$).

Repeated Games and Theories of Efficiency. Economics recognizes (at least) three general ways to achieve efficiency: (i) *competition*, (ii) *contracts*, and (iii) *long-term relationships*. For standardized goods and services, with a large number of traders, promoting market competition is an effective way to achieve efficiency (see the First and Second Welfare Theorems in the *general equilibrium theory*). There are, however, other important resource allocation problems that do not involve standardized goods and services (e.g., resource allocation within a firm or an organization). In such a case, aligning individual incentives with social goals is essential for efficiency, and this can be achieved through incentive schemes (penalties or rewards). The incentive schemes, in turn, can be provided in two distinct ways: by a formal contract or by a long-term relationship. The penalties and rewards specified by a formal contract are enforced by the court, while in a long-term relationship, the value of future interaction serves as the rewards and penalties to discipline the agents' current behavior. The *theory of contracts and mechanism design* concerns the former case; the *theory of repeated games* deals with the latter.

2 Model

2.1 Stage Game

The building block of a repeated game is the simultaneous-move game corresponding to each single interaction, which is referred to as the *stage game*. Here, we consider a finite stage game. To distinguish strategies in the repeated game from those in the stage game, we shall refer to the choices in the stage game as *actions*. A stage game is a triple $G := (N, A, u)$, where: (i) $N := \{1, \dots, n\}$ is the finite set of players, where we denote by i a typical player; (ii) $A := \times_{i=1, \dots, n} A_i$ is the Cartesian product of the finite action sets A_i of each player i , where we denote by a_i a typical action of player i and by a a typical action profile; (iii) $u: A \rightarrow \mathbb{R}^n$ is the utility vector defined pointwise as $u(a) := (u_1(a), \dots, u_n(a))$ that specifies the utility for each player for any given (pure) action profile $a \in A$. A mixed action for player i is an element of $\Delta(A_i)$, denoted by α_i . We write $\alpha_i(a_i)$ for the probability assigned by the mixed action α_i to the pure action a_i . We shall consider the mixed extension of the function u . That is, we enlarge its domain to the set of mixed action profiles $\alpha \in \Delta(A)$ by setting, for each $i \in N$,

$$u_i(\alpha) := \sum_{a \in A} u_i(a) \alpha(a),$$

where $\alpha(a)$ denotes the probability assigned to a by $\alpha \in \Delta(A)$. (Expected) utilities are also referred to as *rewards*, instead of payoffs, for reasons that will become clear.

Remark 1. Since the stage game G is finite, by Nash (1951)'s existence theorem, it has a (possibly mixed) Nash equilibrium. ■

Recall that the *convex hull* of a set $A \subseteq \mathbb{R}^n$, denoted by $\text{co}(A)$, is the smallest convex set containing A .

Definition 1 (Feasible Rewards). *Given a stage game G , the set of feasible rewards, denoted by \mathcal{F} , is the convex hull of the set of stage-game rewards generated by the pure action profiles in A . That is,*

$$\mathcal{F} := \text{co}(\{v \in \mathbb{R}^n : v = u(a) \text{ for some } a \in A\}).$$

That is, the set of feasible rewards is the set of rewards that can be achieved by convex combinations of rewards from pure action profiles. Plainly, we have

$$\mathcal{F} = \{v \in \mathbb{R}^n : v = u(\alpha) \text{ for some } \alpha \in \Delta(A)\}.$$

Some rewards in \mathcal{F} might require players to play correlated actions, because in general the set of independent mixed actions is a strict subset of the mixed action profiles; that is, $\times_{i=1,\dots,n} \Delta(A_i) \subsetneq \Delta(A)$. Therefore, it is customary to assume that players have access to a *public correlating device*, which allows them to replicate the play of correlated action profiles, without modeling it (too) explicitly. We shall assume so whenever convenient. All results that are stated in these lecture notes can be proved without reference to a public correlating device, but the proofs become somewhat more complex.

Exercise 1. Consider the stage game for which the reward matrix is given by

		Player 2	
		L	R
Player 1	U	2, 2	1, 5
	D	5, 1	0, 0

- (a) Represent in a figure the set \mathcal{F} of feasible rewards for this game.
- (b) Given an example of a reward $v \in \mathcal{F}$ that cannot be obtained by any independent mixed action profile.

Solution.

- (a) The set of feasible rewards for this stage game is

$$\mathcal{F} = \text{co}(\{(2, 2), (5, 1), (1, 5), (0, 0)\}),$$

which corresponds to the triangle in the v_1 - v_2 -plane with vertices $(0, 0)$, $(5, 1)$, and $(1, 5)$ (including its boundary).

- (b) With a public correlating device, the players can attach probability $1/2$ to each of the outcomes (L, D) and (R, U) , giving rewards $(3, 3) \in \mathcal{F}$. Such rewards cannot be obtained by any independent mixed action profile. To see this, note that: (i) no pure action profile achieves such rewards; (ii) no independent non-degenerate randomization can achieve such rewards, because any such randomization must attach positive probability to either (U, L) (with rewards $(2, 2)$) or (R, D) (with rewards $(0, 0)$), or both, ensuring that the sum of the two players' average rewards falls below 6. ■

In much of the work on repeated games, payoffs consistent with equilibrium behavior in the repeated game are supported through the use (or threat) of punishments. As such, it is typically important to assess just how much a player i can be punished. This role is played by the minmax.

Definition 2 (Minmax Reward and Minmax Profile). *Given a stage game G , player i 's minmax reward, denoted by \underline{v}_i , is*

$$\underline{v}_i := \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

A minmax profile for player i is a profile $\underline{\alpha}^i := (\underline{\alpha}_i^i, \underline{\alpha}_i^{-i})$ with the properties that $\underline{\alpha}_i^i$ is a stage-game best response for i to $\underline{\alpha}_i^{-i}$ and $\underline{v}_i = u_i(\underline{\alpha}_i^i, \underline{\alpha}_i^{-i})$.

Player i 's minmax reward is the lowest reward player i 's opponents can hold him down to by any independent choice of actions α_j , provided that player i correctly foresees α_{-i} and plays a best-reply to it. Player i always has a best-reply in the set of pure strategies, and therefore restricting him to pure actions does not affect his minmax reward. In general, player i 's opponents are not choosing best responses in profile $\underline{\alpha}^i$; hence, $\underline{\alpha}^i$ need not be a Nash equilibrium of the stage game. Moreover, observe that player i 's opponents may have several actions to minmax player i .¹

Definition 3 (Individually Rational and Strictly Individually Rational Rewards). *A reward vector $v := (v_1, \dots, v_n) \in \mathbb{R}^n$ is: (i) individually rational if $v_i \geq \underline{v}_i$ for all $i \in N$; (ii) strictly individually rational if $v_i > \underline{v}_i$ for all $i \in N$.*

¹An exception is the prisoner's dilemma, where (D, D) is both the unique Nash equilibrium of the stage game and the minmax action profile for both players. Many of the special properties of the infinitely repeated prisoner's dilemma arise out of this coincidence.

Definition 4 (Feasible and Strictly Individually Rational Rewards). *Given a stage game G , the set of feasible and strictly individually rational rewards, denoted by \mathcal{F}^* , is*

$$\mathcal{F}^* := \{v \in \mathcal{F} : v_i > \underline{v}_i \text{ for all } i \in N\}.$$

Remark 2. Note that \mathcal{F}^* is a convex subset of \mathbb{R}^n . ■

[Figure: \mathcal{F} and \mathcal{F}^* for the infinitely repeated prisoner's dilemma in Section 1.1.]

Exercise 2. Solve the following problems.

1. Given a stage game G , player i 's *pure action minmax reward*, denoted by \underline{v}_i^p , is

$$\underline{v}_i^p := \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i}).$$

A *pure action minmax profile* for player i is a profile $\underline{a}^i := (\underline{a}_i^i, \underline{a}_i^{-i})$ with the properties that \underline{a}_i^i is a stage-game best response for i to \underline{a}_i^{-i} and $\underline{v}_i^p = u_i(\underline{a}_i^i, \underline{a}_i^{-i})$.

- (a) Contrast $\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$ from $\max_{a_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})$. Which is (weakly) higher and why? Prove it.
 - (b) Find a game in which the minmax reward is strictly lower than the pure action minmax reward.
2. The definition of minmax reward assumes that player i 's opponents randomize independently of each other. Define the notions of minmax reward and minmax profile in which a player i 's opponents randomize in a correlated manner. Is player i 's *correlated action minmax reward* (weakly) lower, the same, or (weakly) higher than his minmax reward? Prove your assertion.

Solution.

1. (a) *Claim.*

$$\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i}) \geq \max_{a_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}).$$

Proof. For all $\tilde{a}_i \in A_i$ and $\tilde{a}_{-i} \in A_{-i}$ we have

$$u_i(\tilde{a}_i, \tilde{a}_{-i}) \geq \min_{a_{-i} \in A_{-i}} u_i(\tilde{a}_i, a_{-i}).$$

The claim follows by noting that

$$\begin{aligned} u_i(\tilde{a}_i, \tilde{a}_{-i}) &\geq \min_{a_{-i} \in A_{-i}} u_i(\tilde{a}_i, a_{-i}) \quad \text{for all } \tilde{a}_i \in A_i, \tilde{a}_{-i} \in A_{-i} \\ \implies \max_{a_i \in A_i} u_i(a_i, \tilde{a}_{-i}) &\geq \max_{a_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \quad \text{for all } \tilde{a}_{-i} \in A_{-i} \\ \implies \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i}) &\geq \max_{a_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}). \end{aligned}$$

Contrast. Player i 's minmax reward is the lowest reward player i 's opponents can be sure to hold him down to, without knowing player i ' action. In contrast, player i 's maxmin reward is the largest reward the player can be sure to get without knowing his opponents' actions.

(b) Consider *matching pennies* for which the reward matrix is given by

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

Player 1's pure action minmax reward is 1, because for any of player 2's pure actions, player 1 has a best response giving a payoff of 1. Pure action minmax profiles are given by (H, H) and (T, T) . In contrast, player 1's minmax reward is 0, implied by player 2's mixed action of assigning probability 1/2 to H and probability 1/2 to T . An analogous argument applies to player 2. ■

2. Given a stage game G , player i 's *correlated minmax reward*, denoted by \underline{v}_i^c , is

$$\underline{v}_i^c := \min_{\alpha_{-i} \in \Delta(A_{-i})} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}),$$

where $A_{-i} := \times_{j \neq i} A_j$. A *correlated minmax profile* for player i is a profile $\underline{\alpha}^{i,c} := (\underline{\alpha}_i^{i,c}, \underline{\alpha}_{-i}^{-i,c})$ with the properties that $\underline{\alpha}_i^{i,c}$ is a stage-game best response for i to $\underline{\alpha}_{-i}^{-i,c} \in \Delta(A_{-i})$ and $\underline{v}_i^c = u_i(\underline{\alpha}_i^{i,c}, \underline{\alpha}_{-i}^{-i,c})$.

Claim. $\underline{v}_i \geq \underline{v}_i^c$.

Proof. Recall that for any function $f: X \rightarrow \mathbb{R}$ and any two nonempty sets $V, W \subseteq X$, we have

$$V \subseteq W \implies \inf_{x \in V} f(x) \geq \inf_{x \in W} f(x).$$

The desired result then follows by observing that $\times_{j \neq i} \Delta(A_j) \subseteq \Delta(A_{-i})$. ■

Exercise 3. Consider the stage game for which the reward matrix is given by

		Player 2	
		L	R
Player 1	U	-2, 2	1, -2
	M	1, -2	-2, 2
	D	0, 1	0, 1

What are the two players' minmax rewards?

Solution. For this game, we have

$$v_1 = v_2 = 0.$$

To compute player 1's minmax reward, we first compute his rewards to U , M , and D as a function of the probability q that player 2 assigns to L . In the obvious (and somewhat abused) notation, these rewards are $u_U(q) = -3q + 1$, $u_M(q) = 3q - 2$, and $u_D(q) = 0$. Since player 1 can always attain a reward of 0 by playing D , his minmax reward is at least this large; the question is whether player 2 can hold player 1's maximized reward to 0 by some choice of q . Since q does not enter into u_D , we can pick q to minimize the maximum of u_U and u_M , which occurs at the point where the two expressions are equal, i.e. $q = 1/2$. Since $u_U(1/2) = u_M(1/2) = -1/2$, player 1's minmax reward is the 0 reward he can achieve by playing D . Note that $\max\{u_U(q), u_M(q)\} \leq 0$ for any $q \in [1/3, 2/3]$, so we can take player 2's minmax action against player 1 to be any mixed action assigning probability $q \in [1/3, 2/3]$ to L .

Similarly, to find player 2's minmax reward, we first express player 2's reward to L and R as a function of the probabilities p_U and p_M that player 1 assigns to U and M :

$$u_L(p_U, p_M) = 2(p_U - p_M) + (1 - p_U - p_M), \quad (1)$$

$$u_R(p_U, p_M) = -2(p_U - p_M) + (1 - p_U - p_M). \quad (2)$$

Player 2's minmax reward is then determined by

$$\min_{p_U, p_M} \max\{2(p_U - p_M) + (1 - p_U - p_M), -2(p_U - p_M) + (1 - p_U - p_M)\}.$$

By inspection, or by plotting (1) and (2), we see that player 2's minmax reward is 0, which is attained with $p_U = p_M = 1/2$ and $p_D = 0$. Here, unlike the minmax against player 1, the minmax action against player 2 is uniquely determined: if $p_U > p_M$, the reward to L is positive; if $p_M > p_U$, the reward to R is positive; and if $p_U = p_M < 1/2$, then both L and R have positive rewards. ■

2.2 Repeated Game

The repeated game (also sometimes referred to as the *supergame*) is a repetition of the stage game G for each $t = 0, 1, \dots, T$. The parameter T is called the *horizon* of the game; it could be finite, in which case the game is said to be a *finitely repeated game*, or infinite, in which case the game is said to be an *infinitely repeated game*. Since we let time start at 0, if $T < \infty$, the game is actually repeated $T + 1$ times.

Repeated games with *perfect monitoring* are repeated games in which all players observe all realized actions at the end of the period (Note: this means that, if one player

uses a mixed action α_i , his opponents will not observe the lottery itself, but only the realized action a_i). We write a^t for the action profile that is realized in period t . That is, $a^t := (a_1^t, \dots, a_n^t)$ are the actions actually played in period t .

A player's *information set* at the beginning of period t is a vector $(a^0, a^1, \dots, a^{t-1}) \in A^t$ for $t \geq 1$. We define the set of *histories of length t* as the set $H^t := A^t$ for $t \geq 1$, and denote a typical element by h^t . This does not quite address the initial information set, and so, by convention, we set $H^0 := \{\emptyset\}$, and we interpret its single element $h^0 = \emptyset$ as the initial information set. The set of all possible *histories* is $H := \cup_{t=0}^T H^t$, with typical element $h \in H$.

A *pure strategy* for player i is a function $s_i: H \rightarrow A_i$ that specifies, for each history, what action to play. The set of player i 's pure strategies is denoted by S_i . A *mixed strategy* is a mixture over the set of all pure strategies. A *behavior strategy* is a function $\sigma_i: H \rightarrow \Delta(A_i)$, and the set of all player i 's behavior strategies is denoted by Σ_i . We set, as usual, $S := \times_{i=1, \dots, n} S_i$, $\Sigma := \times_{i=1, \dots, n} \Sigma_i$ and write s , resp. σ , for a typical pure, resp. behavior, strategy profile. Kuhn's theorem, establishing the realization-equivalence between mixed and behavior strategies, applies to repeated games as well.²

Example 1. The grim-trigger strategy profile for the infinitely repeated prisoner's dilemma in Section 1.1 is an example of a strongly symmetric strategy profile. Formally, we say the following.

- A *grim-trigger strategy profile* is a strategy profile where a deviation triggers Nash reversion (hence *trigger*) and the Nash equilibrium minmaxes the players (hence *grim*).
- Suppose $A_i = A_j$ for all $i, j \in N$. A strategy profile is *strongly symmetric* if $\sigma_i(h) = \sigma_j(h)$ for all $h \in H$ and all $i, j \in N$ (that is, all players use the same strategy after every history, including asymmetric histories). ■

A strategy profile $\sigma \in \Sigma$ generates a distribution over terminal nodes, that is, over histories H^{T+1} . We call such terminal nodes *outcomes*. When T is finite, defining the probability space is simple: let H^{T+1} be the finite set of outcomes (here, $H^{T+1} := A^{T+1}$ with typical element $h^{T+1} := (a^0, \dots, a^T)$); then the set of events is the set $\mathcal{P}(H^{T+1})$ of all subsets of H^{T+1} . If $T = \infty$, the set of outcomes is the set of infinite sequences $(a^0, a^1, \dots) \in A^\infty$, and so it is infinite as well, and defining the set of events introduces technical details which we ignore in these lecture notes. Hereafter, let $H^\infty := A^\infty$, with typical element $h^\infty := (a^0, a^1, \dots)$.

Note that every period of play begins a proper subgame, and since actions are simultaneous in the stage game, these are the only proper subgames, a fact that we must keep in mind when applying subgame-perfection.

²Two strategies for a player i are realization equivalent if, fixing the strategies of the other players, the two strategies of player i induce the same distribution over outcomes.

In a repeated game, for any non-terminal history $h \in H$, the *continuation game* is the subgame that begins following history h . Observe that a subgame of a finitely repeated game is a finitely repeated game with a shorter horizon, while a subgame of the infinitely repeated game is an infinitely repeated game itself. For a strategy profile σ , player i 's *continuation strategy induced by h* , denoted by $\sigma_i|_h$, is the restriction of $\sigma_i \in \Sigma_i$ to the subgame beginning at h . We represent the *continuation strategy profile induced by h* as $\sigma|_h := (\sigma_1|_h, \dots, \sigma_n|_h)$. In an infinitely repeated game, $\sigma_i|_h$ is a strategy in the original repeated game; therefore, the continuation game associated with each history is a subgame that is strategically identical to the original game. In other words, infinitely repeated games have a convenient *recursive structure*, and this plays an important role in their study.

When T is finite, we shall evaluate outcomes according to the average of the sum of rewards. That is, we define the function $v_i: H^{T+1} \rightarrow \mathbb{R}$ by setting, for all $h^{T+1} := (a^0, \dots, a^T) \in H^{T+1}$,

$$v_i(h^{T+1}) := (T+1)^{-1} \sum_{t=0}^T u_i(a^t).$$

Since a strategy profile generates a probability distribution, we can also extend the domain of the function v_i to all strategies $\sigma \in \Sigma$, by setting

$$v_i(\sigma) := (T+1)^{-1} \mathbb{E}_\sigma \left[\sum_{t=0}^T u_i(a^t) \right],$$

where the operator \mathbb{E}_σ refers to the expectations under the probability distribution over the set of outcomes H^{T+1} generated by σ . This is player i 's *average payoff* in the finitely repeated game.

There are several alternative ways of defining payoffs in the infinitely repeated game. We will focus on the case in which players discount future rewards using a common *discount factor* $\delta \in [0, 1)$. In this case, we define the function $v_i: H^\infty \rightarrow \mathbb{R}$ by setting, for all $h^\infty := (a^0, a^1, \dots) \in H^\infty$,

$$v_i(h^\infty) := (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t).$$

The normalization constant $(1 - \delta)$ that appears in front is a way to make payoffs in the repeated game comparable to rewards in the stage game. Indeed, if a player receives a reward of 1 in every period, the unnormalized discounted sum is equal to $1 + \delta + \delta^2 + \dots = 1/(1 - \delta)$. Once it is normalized then, it is equal to 1 as well. Therefore, when considering the normalized discounted sum rather than the unnormalized one, the set of payoffs that are feasible in the repeated game becomes the same as the set of feasible rewards in the stage game, allowing for meaningful comparisons.

Since a strategy profile generates a probability distribution over outcomes, we extend

here as well the domain of the function v_i to all strategies $\sigma \in \Sigma$, by setting

$$v_i(\sigma) := (1 - \delta) \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \delta^t u_i(a^t) \right],$$

where the operator \mathbb{E}_σ refers to the expectations under the probability distribution over the set of outcomes H^∞ that is generated by σ . This is player i 's *normalized payoff* in the infinitely repeated game.

There are two features of discounting to note. First, assuming that players discount ($\delta < 1$) treats periods asymmetrically: a player cares more about the rewards in an early period than in a later one. In a finite horizon game, there is no challenge in setting $\delta = 1$, but in an infinite horizon game, total payoffs can be unbounded with $\delta = 1$. Second, discounting can represent both time preferences, as we have done so implicitly here, or uncertainty about when the game will end, e.g., suppose that conditional on reaching a period t , the game continues to the next period with a probability $\delta < 1$.

The finitely repeated game with horizon $T < \infty$ is denoted by G^T , while the infinitely repeated game with discount factor δ is denoted by G^δ .

Exercise 4. Consider the stage game for which the reward matrix is given by

		Fabio	
		<i>DE</i>	<i>DT</i>
Eliana	<i>DE</i>	0, 0	3, 1
	<i>DT</i>	1, 3	0, 0

Here, *DT* stands for “Decision Theory” and *DE* stands for “Development Economics”. This stage game is often referred to as the *battle of the sexes*. Answer the following questions.

- (a) Identify the set of feasible and strictly individually rational rewards in the stage game. What are the two players’ minmax rewards? What are their pure action minmax rewards?
- (b) Suppose both players have discount factor δ . If players decide to alternate topics in each period, with Eliana working on Development Economics and Fabio working on Decision Theory at $t = 0$, what are their payoffs?
- (c) Now suppose that Eliana has discount factor $\delta_E \in (0, 1)$, and Fabio has discount factor δ_F with $0 < \delta_F < \delta_E$. Can you suggest a strategy profile that obtains payoffs outside the set of feasible rewards of the stage game?

Solution.

- (a) Let E stand for Eliana and F for Fabio. We have $\underline{v}_E^p = \underline{v}_F^p = 1$ and $\underline{v}_E = \underline{v}_F = 3/4$. The set of feasible and strictly individually rational rewards is

$$\mathcal{F}^* = \text{co}(\{(0, 0), (3, 1), (1, 3)\}) \cap \{(v_E, v_F) \in \mathbb{R}^2 : v_E > 3/4 \text{ and } v_F > 3/4\}.$$

- (b) Eliana's payoff is

$$(1 - \delta) \left(\sum_{t=0}^{\infty} 3\delta^t + \sum_{t=0}^{\infty} \delta^{(2t+1)} \right) = (1 - \delta) \left(\frac{3}{1 - \delta^2} + \frac{\delta}{1 - \delta^2} \right) = \frac{3 + \delta}{1 + \delta}.$$

Similarly, Fabio's payoff is

$$(1 - \delta) \left(\sum_{t=0}^{\infty} \delta^t + \sum_{t=0}^{\infty} 3\delta^{(2t+1)} \right) = (1 - \delta) \left(\frac{1}{1 - \delta^2} + \frac{3\delta}{1 - \delta^2} \right) = \frac{1 + 3\delta}{1 + \delta}.$$

- (c) The repeated game strategy profile that calls for (DT, DE) to be played in periods $t = 0, \dots, T - 1$ for some $T \geq 1$ and (DE, DT) to be played in all subsequent periods yields a repeated game payoff vector outside $\text{co}(\{(0, 0), (3, 1), (1, 3)\})$, being in particular above the line segment joining payoffs $(3, 1)$ and $(1, 3)$. First, note that Eliana's payoff from this strategy profile is $(1 - \delta_E^T) + 3\delta_E^T$ and Fabio's payoff from this strategy profile is $3(1 - \delta_F^T) + \delta_F^T$. By way of contradiction, assume that this pair of payoffs is below, or on, the line segment joining payoffs $(3, 1)$ and $(1, 3)$. Then, for some $\lambda \in [0, 1]$,

$$\begin{cases} 1 + 2\delta_E^T \leq 1 + 2\lambda \\ 3 - 2\delta_F^T \leq 3 - 2\lambda \end{cases} \iff \begin{cases} \delta_E^T \leq \lambda \\ \lambda \leq \delta_F^T \end{cases} \implies \delta_E \leq \delta_F,$$

which contradicts the assumption that $\delta_F < \delta_E$. ■

2.3 Equilibrium Notions for the Repeated Game

Definition 5 (Nash Equilibrium). *A strategy profile $\sigma \in \Sigma$ is a Nash equilibrium of the repeated game if, for all $i \in N$ and $\sigma'_i \in \Sigma_i$, $v_i(\sigma) \geq v_i(\sigma'_i, \sigma_{-i})$.*

The challenge of Nash equilibria in dynamic games is that they permit too much, including non-rational behavior off the equilibrium path. Accordingly, it is more appropriate to restrict attention to subgame-perfect Nash equilibria.

Definition 6 (Subgame-Perfect Nash Equilibrium). *A strategy profile $\sigma \in \Sigma$ is a subgame-perfect Nash equilibrium (hereafter, SPNE) of the repeated game if, for all histories $h \in H$, $\sigma|_h$ is a Nash equilibrium of the subgame beginning at h .*

Existence is immediate: the stage game has a Nash equilibrium; then, any profile of strategies that induces the *same* Nash equilibrium of the stage game after every history of the repeated game is a SPNE of the latter. In principle, checking for subgame perfection involves checking whether an infinite number of strategy profiles are Nash equilibria—the set H of histories is countably infinite even if the stage-game action spaces are finite. Moreover, checking whether a profile σ is a Nash equilibrium involves checking that player i 's strategy σ_i is no worse than an infinite number of potential deviations (because player i could deviate in any period, or indeed in any combination of periods). Fortunately, we can simplify this task immensely, first by limiting the number of alternative strategies that must be examined (this is the topic of Section 3.1), then by organizing the subgames that must be checked for Nash equilibria into equivalence classes (this is the topic of Section 3.2), and finally by identifying a simple constructive method for characterizing equilibrium payoffs (see Section 2.4 in [Mailath and Samuelson \(2006\)](#)).

Theorem 1. *Let G be a stage game that admits a unique Nash equilibrium α . Then, for any finite T , the unique SPNE σ of G^T is such that, for all $i \in N$ and $h \in H$, $\sigma_i(h) = \alpha_i$.*

Proof. The result follows by using backward induction. ■

The next two exercises show that when the stage game G has more than one Nash equilibrium, then in G^T we may have some SPNE where, in some periods, players play some actions that are not played in any Nash equilibrium of the stage game.

Exercise 5. Consider the stage game G for which the reward matrix is given by

		Player 2		
		A	B	C
Player 1	A	3, 3	−1, 4	0, 0
	B	4, −1	0, 0	0, −1
	C	0, 0	−1, 0	x, x

Suppose G is played thrice with no discounting or time-averaging of rewards. Show that if $x \geq 1/2$, then there is a SPNE of $G^{T=3}$ in which (A, A) is played in the first period.

Solution. Suppose $x \geq 1/2$. The strategy profiles (B, B) and (C, C) are NEs of the stage game G . Now consider the following strategy for the repeated game $G^{T=3}$:

- Play (A, A) at $t = 1$.
- If play at $t = 1$ is (A, A) , then play (C, C) at $t = 2$; otherwise, play (B, B) at $t = 2$.
- If play at $t = 1$ is (A, A) and play at $t = 2$ is (C, C) , then play (C, C) at $t = 3$; otherwise, play (B, B) at $t = 3$.

To show that this strategy profile is a SPNE of $G^{T=3}$, we need to show that the strategy profile is a NE in all subgames of $G^{T=3}$. By construction, this strategy profile is a NE of the subgames starting at $t = 2$ and at $t = 3$ after any history of play. For the strategy profile to be a NE of $G^{T=3}$, it must be that (noting that playing B is player i 's most profitable deviation when player $-i$ plays A)

$$3 + 2x \geq 4 \quad \text{or, equivalently,} \quad 2x \geq 1,$$

which is always satisfied for $x \geq 1/2$. ■

Exercise 6. Consider the stage game G for which the reward matrix is given by

		Player 2	
		S	F
Player 1	S	2, 2	0, 3
	F	3, 0	-1, -1

Suppose G is played twice with no discounting or time-averaging of rewards. Show that there is a SPNE of $G^{T=2}$ (possibly involving public correlation) in which (S, S) is played in the first period.

Solution. The strategy profiles (S, F) , (F, S) are NEs of the stage game G . Now consider the following strategy for the repeated game $G^{T=2}$:

- At $t = 1$: play (S, S) .
- At $t = 2$: if play at $t = 1$ is (S, S) or (F, F) , then play (S, F) with probability $1/2$ and (F, S) with probability $1/2$ —this is done by using a public correlating device; if play at $t = 1$ is (S, F) , then play (F, S) , if play at $t = 1$ is (F, S) , then play (S, F) .

To show that this strategy profile is a SPNE of $G^{T=2}$, we need to show that the strategy profile is a NE in all subgames of $G^{T=2}$. Clearly, a NE is played at $t = 2$ after any history of play. For the strategy profile to be a NE of $G^{T=2}$, it must be that

$$2 + \frac{1}{2}(3 + 0) \geq 3 + 0,$$

which is always satisfied. ■

Player i 's minmax reward is often referred to as player i 's *reservation utility*. The reason for this name is the following result.

Proposition 1. *Player i 's reward is at least \underline{v}_i in any Nash equilibrium of the stage game G . Player i 's payoff is at least \underline{v}_i in any Nash (and therefore, also in any SPNE) equilibrium of G^T for any $T < \infty$ and G^δ for any $\delta \in [0, 1)$.*

Proof. Let $\hat{\alpha}$ be a Nash equilibrium of G . Thus, $\hat{\alpha}_i$ is a best response to $\hat{\alpha}_{-i}$. By definition of minmax reward, it follows that $u_i(\hat{\alpha}_i, \hat{\alpha}_{-i}) \geq \underline{v}_i$, which establishes the first statement.

Now suppose that G is played repeatedly over time. Player i can always use the following strategy $\sigma_i \in \Sigma_i$ in the repeated game:

$$\sigma_i(h) \in \arg \max_{a_i \in A_i} u_i(a_i, \sigma_{-i}(h)) \quad \text{for all } h \in H.$$

That is, the strategy that player i picks in every period is some best-reply to the action profile played by players $-i$. This strategy may not be optimal, since it ignores the fact that the future play of i 's opponents may depend on how player i plays today. However, because all players have the same information at the beginning of each period t , the probability distribution over the actions of players $-i$ in period t , conditional on player i 's information, corresponds to independent randomizations by player i 's opponents. Thus, in every period, player i guarantees at least the minimum over his opponents independent randomizations of the maximum over his best-replies, i.e. he guarantees at least \underline{v}_i in every period, and so secures

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t \underline{v}_i = \underline{v}_i$$

in the infinitely repeated game G^δ , and also

$$(T + 1)^{-1} \sum_{t=0}^T \underline{v}_i = \underline{v}_i$$

in the finitely repeated game G^T . This establishes the second statement. ■

3 Preliminary Results

3.1 One-Shot Deviation Principle

Definition 7 (One-Shot Deviation). *A one-shot deviation for player i from strategy $\sigma_i \in \Sigma_i$ is a strategy $\sigma'_i \in \Sigma_i$ such that $\sigma'_i(h') = \sigma_i(h')$ for all $h' \in H$ with $h' \neq h$, and $\sigma'_i(h) \neq \sigma_i(h)$.*

A one-shot deviation thus agrees with the original strategy everywhere except at one history where the one-shot deviation occurs.

Definition 8 (Profitable One-Shot Deviation). *Fix a strategy profile $\sigma \in \Sigma$. A one-shot deviation σ'_i from strategy σ_i is profitable for player i if, at the history h for which $\sigma'_i(h) \neq \sigma_i(h)$,*

$$v_i(\sigma'_i|_h, \sigma_{-i}|_h) > v_i(\sigma|_h).$$

The profitability of a one-shot deviation is evaluated *ex interim*, i.e. based on payoffs realized once histories in which the deviation differs are realized, and not evaluated *ex ante* at the beginning of time. Thus, the history that makes a deviation profitable may be off the path of play. This should make you reflect on how the one-shot deviation principle relates to Nash equilibria and SPNE: a Nash equilibrium can have profitable one-shot deviations if these deviations occur off the equilibrium path; in contrast, because a SPNE must be immune to deviations on and off the equilibrium path, there can be no profitable one-shot deviations. In other words, the absence of profitable one-shot deviations is necessary for a profile to be SPNE; it is also sufficient.

Theorem 2 (One-Shot Deviation Principle). *A strategy profile $\sigma \in \Sigma$ is a SPNE of G^δ if and only if no player has a profitable one-shot deviation.*

To confirm that a strategy profile σ is a SPNE, we thus need only consider alternative strategies that deviate from the action proposed by σ once and then return to the prescriptions of the equilibrium strategy. This does not imply that the path of generated actions will differ from the equilibrium strategies in only one period. The deviation prompts a different history than does the equilibrium, and the equilibrium strategies may respond to this history by making different subsequent prescriptions.

The importance of the one-shot deviation principle lies in the implied reduction in the space of deviations that need to be considered. In particular, we do not have to worry about alternative strategies that might deviate from the equilibrium strategy in period t , and then again in period $t' > t$, and again in period $t'' > t'$, and so on.

Proof. Necessity (“only if”) is immediate: if σ is a SPNE, then there are no profitable deviations, whether one-shot or not. Sufficiency (“if”) is more subtle. We need to show that, whenever a profitable deviation exists, we can define a profitable one-shot deviation. Suppose then that $\sigma \in \Sigma$ is not a SPNE. Then, there exists $i \in N$, $h^t \in H$, and σ'_i such that

$$v_i(\sigma'_i|_{h^t}, \sigma_{-i}|_{h^t}) - v_i(\sigma|_{h^t}) > 0.$$

Because payoffs are discounted, it follows that there exists $T \in \mathbb{N}$ such that, defining $\sigma''_i \in \Sigma_i$ by $\sigma''_i(h^\tau) := \sigma_i(h^\tau)$ for all $h^\tau \in H$ with $\tau \geq t + T$ and $\sigma''_i(h^\tau) := \sigma'_i(h^\tau)$ otherwise, we have

$$v_i(\sigma''_i|_{h^t}, \sigma_{-i}|_{h^t}) - v_i(\sigma|_{h^t}) > 0.$$

That is, σ''_i is a profitable deviation that only differs from σ_i at finitely many histories. We now proceed by induction. Consider all histories $h \in H^{t+T-1}$. We have two possibilities:

- Either $v_i(\sigma''_i|_h, \sigma_{-i}|_h) - v_i(\sigma|_h) > 0$ for some $h \in H^{t+T-1}$. In this case, define $\sigma'''_i \in \Sigma_i$ as the one-shot deviation from σ_i at h , with $\sigma'''_i(h) := \sigma''_i(h)$. This is a one-shot deviation, and so we are done.

- Or $v_i(\sigma_i''|_{h^t}, \sigma_{-i}|_{h^t}) - v_i(\sigma|_{h^t}) \leq 0$. In this case, define $\sigma_i''' \in \Sigma_i$ by $\sigma_i'''(h^\tau) = \sigma_i(h^\tau)$ for all $h^\tau \in H$ with $\tau \geq t + T - 1$ and $\sigma_i'''(h^\tau) := \sigma_i'(h^\tau)$ otherwise. Again, σ_i''' is a profitable deviation that only differs from σ_i at finitely many histories.

Repeat the above process iteratively until the difference is positive. ■

The one-shot deviation principle is due to [Blackwell \(1965\)](#). The essence of the proof is that, because payoffs are discounted, any strategy that is a profitable deviation to the proposed profile σ must be profitable within a finite number of periods. Once we can restrict attention to finite steps, backward induction proves the existence of a profitable one-shot deviation. Discounting is not necessary for [Theorem 2](#). [Fudenberg and Tirole \(1991\)](#) show that the one-shot deviation principle holds for any extensive-form game with perfect monitoring in which payoffs are *continuous at infinity*—a condition that essentially requires that actions in the far future have a negligible impact on current payoffs. In the context of infinitely repeated games, continuity at infinite is guaranteed by discounting.

Naturally, you may be wondering about the relationship between the one-shot deviation principle and Nash equilibria that are not subgame perfect. Since Nash equilibria do not demand incentives off the equilibrium path, profitable (one or multi-shot) deviations off the path of play do not prevent a strategy profile from being a Nash equilibrium. A more reasonable hope is that restricting attention to histories that arise on the equilibrium path would relate Nash equilibria to one-shot deviations. That is, you may conjecture the following.

Conjecture 1. A strategy profile $\sigma \in \Sigma$ is a Nash Equilibrium if and only if there are no profitable one-shot deviations from histories that arise on the path of play.

The next exercise helps you evaluate the validity of this conjecture.

Exercise 7. Consider the prisoner’s dilemma for which the reward matrix is given by

		Player 2	
		<i>C</i>	<i>D</i>
Player 1	<i>C</i>	3, 3	−1, 4
	<i>D</i>	4, −1	1, 1

Answer the following questions.

- What is the set of feasible and individually rational rewards?
- If the game is repeated infinitely often, what conditions on δ would have to be satisfied so that both choose to cooperate in each period in a SPNE? Describe the SPNE strategy profile.

- (c) Consider the *tit-for-tat strategy*: play C at $t = 0$; at each $t \geq 1$, play whatever the opponent played at $t - 1$.
- (i) Suppose that Player 1 defects forever while Player 2 follows tit-for-tat. What is Player 1's expected payoff from this "infinite-shot" deviation?
 - (ii) Consider the following one-shot deviation: Player 1 chooses to deviate only at $t = 0$, but then follows tit-for-tat forever after. What are the two players' expected payoffs? Under what condition on δ is this not a profitable one-shot deviation?
 - (iii) Under what condition on δ is tit-for-tat a Nash equilibrium profile? Compare the condition with that above to evaluate Conjecture 1.
 - (iv) Prove that tit-for-tat is not a SPNE, regardless of the value of δ .

Solution.

- (a) We have $\underline{v}_1 = \underline{v}_2 = 1$. Thus, the set of feasible and strictly individually rational rewards is

$$\mathcal{F}^* = \text{co}(\{(3, 3), (-1, 4), (4, -1), (1, 1)\}) \cap \{(v_1, v_2) \in \mathbb{R}^2 : v_1 > 1 \text{ and } v_2 > 1\}.$$

- (b) Consider the following "grim-trigger" strategy profile:

- Play (C, C) if: (i) $t = 0$ or (ii) $t \geq 1$ and (C, C) was played in every prior period.
- Play (D, D) if some player played D in some prior period.

We apply the one-shot deviation principle to show that this is a SPNE of the repeated game for $\delta \geq 1/3$. There are two cases to consider:

- Suppose (C, C) was played in every prior period. Player i 's incentive condition to play C in the current period is

$$(1 - \delta)(4) + \delta(1) \leq 3,$$

which is satisfied for any $\delta \geq 1/3$.

- Suppose D was played by some player in some prior period. Player i 's incentive condition to play D in the current period is satisfied for any $\delta \in [0, 1)$.

Thus, for $\delta \geq 1/3$, this is a SPNE of G^δ in which players cooperate in each period.

- (c) (i) Player 1's expected payoff from this "infinite-shot" deviation is

$$(1 - \delta)(4) + \delta(1) = 4 - 3\delta.$$

(ii) Player 1's payoff is

$$(1 - \delta) \left(\sum_{t=0}^{\infty} \delta^t (4) + \sum_{t=0}^{\infty} \delta^{(2t+1)} (-1) \right) = \frac{4 - \delta}{1 + \delta}.$$

Similarly, player 2's payoff is

$$(1 - \delta) \left(\sum_{t=0}^{\infty} (-1) \delta^t + \sum_{t=0}^{\infty} \delta^{(2t+1)} (4) \right) = \frac{4\delta - 1}{1 + \delta}.$$

This deviation is not a profitable one-shot-deviation for player 1 if and only

$$3 \geq \frac{4 - \delta}{1 + \delta},$$

or, equivalently,

$$\delta \geq \frac{1}{4}.$$

(iii) Tit-for-tat is a NE of the game only if a deviation to always defecting is unprofitable, which requires

$$3 \geq 4 - 3\delta,$$

or, equivalently,

$$\delta \geq 1/3.$$

This, together with part (c)-(ii) shows that Conjecture 1 is false since for any $\delta \in [1/4, 1/3)$ there is no profitable one shot deviation from the path of play of tit-for-tat, but tit-for-tat is not a NE.

(iv) Consider a period- t history where (D, C) has been played in period $t - 1$. Player 1 has no profitable one-shot deviation if (make sure you understand the inequality, in the exam you would have to briefly explain it)

$$\frac{4\delta - 1}{1 + \delta} \geq 1,$$

or, equivalently,

$$\delta \geq \frac{2}{3}. \tag{3}$$

Player 2 has no profitable one-shot deviation if (make sure you understand the inequality, in the exam you would have to briefly explain it)

$$\frac{4 - \delta}{1 + \delta} \geq 3,$$

or, equivalently,

$$\delta \leq \frac{1}{4}. \tag{4}$$

Since (3) and (4) cannot be simultaneously satisfied, the desired result follows by the one-shot deviation principle. ■

3.2 Automaton Representation of Strategy Profiles

To further simplify the task of finding SPNE of the repeated game, we can group histories into equivalence classes, where each member of an equivalence class induces an identical continuation strategy. We achieve this grouping by representing repeated-game strategies as automata, where the states of the automata represent equivalence classes of histories.

An *automaton* is a quadruple (W, w^0, f, τ) where:

- W is a set of states,
- $w^0 \in W$ is an initial state,
- $f: W \rightarrow \times_i \Delta(A_i)$ is an output function,
- $\tau: W \times A \rightarrow W$ is a transition function.

The transition function identifies the next state of the automaton, given its current state and the realized stage-game pure action profile. If f specifies a pure output at state w , we write $f(w)$ for the resulting action profile. If a mixed action is specified by f at w , $f^w(a)$ denotes the probability attached to profile a . Note: even if two automata only differ in their initial state, they nonetheless are different automata.

An automaton (W, w^0, f, τ) with f specifying a pure action at every state induces an outcome (a^0, a^1, \dots) as follows:

$$\begin{aligned} a^0 &= f(w^0), \\ a^1 &= f(\tau(w^0, a^0)), \\ a^2 &= f(\tau(\tau(w^0, a^0), a^1)), \\ &\vdots \end{aligned}$$

We extend this to identify the strategy induced by an automaton. First, extend the transition function from the domain $W \times A$ to the domain $W \times H \setminus \{\emptyset\}$ by recursively defining

$$\tau(w, h^t) = \tau(\tau(w, h^{t-1}), a^{t-1}).$$

With this definition, we have the strategy σ described by $\sigma(\emptyset) = f(w^0)$ and

$$\sigma(h^t) = f(\tau(w^0, h^t)).$$

Similarly, an automaton for which f sometimes specifies mixed actions induces a path of play and a strategy.

Conversely, any strategy profile can be represented by an automaton. Take the set of histories H as the set of states, the null history \emptyset as the initial state, $f(h^t) := \sigma(h^t)$, and $\tau(h^t, a) := h^{t+1}$, where $h^{t+1} := (h^t, a)$ is the concatenation of the history h^t with the action profile a .

This representation leaves us in the position of working with the full set of histories H . However, strategy profiles can often be represented by automata with finite sets W . The set W is then a partition on H , grouping together those histories that prompt identical continuation strategies. The advantage of the automaton representation is most obvious when W can be chosen finite.

Example 2. The grim-trigger strategy profile for the infinitely repeated prisoner's dilemma in Section 1.1 has as its automaton representation (W, w^0, f, τ) , where:

- $W = \{w_{CC}, w_{DD}\}$;
- $w^0 = w_{CC}$;
- $f(w_{CC}) = (C, C)$ and $f(w_{DD}) = (D, D)$;

•

$$\tau(w, a) := \begin{cases} w_{CC} & \text{if } w = w_{CC} \text{ and } a = (C, C) \\ w_{DD} & \text{otherwise} \end{cases} .$$

The next figure is the automaton representation of the grim-trigger strategy profile for the infinitely repeated prisoner's dilemma in Section 1.1.



Circles are states and arrows transitions, labeled by the profiles leading to the transitions. The subscript on a state indicates the action profile to be take at the state. ■

We say that a state $w' \in W$ is accessible from another state $w \in W$ if there exists sequence of action profiles such that beginning at w , the automaton eventually reaches w' More formally, there exists h^t such that $w' = \tau(w, h^t)$. Accessibility is not symmetric. Consequently, in an automaton (W, w^0, f, τ) , even if every state in W is accessible from the initial state w^0 , this may not be true if some other state replaced w^0 as the initial state (e.g., see the last example).

Remark 3. The complexity of a strategy is sometimes defined by the number of states of the automaton with the smallest number of states (hereafter, the smallest automaton) that implements it. ■

Remark 4. When a strategy profile σ is described by the automaton (W, w^0, f, τ) , the continuation strategy profile after the history h^t , $\sigma|_{h^t}$, is described by the automaton obtained by using $\tau(w^0, h^t)$ as the initial state, that is, $(W, \tau(w^0, h^t), f, \tau)$. If every state in W is accessible from w^0 , then the collection of all continuation strategy profiles is described by the collection of automata $\{(W, w^0, f, \tau)\}_{w \in W}$. ■

The advantage of the automaton representation is that we need only verify the strategy profile induced by (W, w^0, f, τ) is a Nash equilibrium, for all $w \in W$, to confirm that the strategy profile induced by (W, w^0, f, τ) is a subgame-perfect equilibrium. The following result is immediate from the last remark.

Proposition 2. *The strategy profile with representing automaton (W, w^0, f, τ) is a SPNE if and only if, for all $w \in W$ accessible from w^0 , the strategy profile induced by (W, w, f, τ) is a Nash equilibrium of the repeated game.*

Each state of the automaton identifies an equivalence class of histories after which the strategies prescribe identical continuation play. The requirement that the strategy profile induced by each state of the automaton (i.e. by taking that state to be the initial state) corresponds to a Nash equilibrium is then equivalent to the requirement that we have Nash equilibrium continuation play after every history.

This result simplifies matters by transferring our concern from the set of histories H to the set of states of the automaton representation of a strategy. If $W = H$, little has been gained. However, often W is considerably smaller than the set of histories, with many histories associated with each state of (W, w, f, τ) , as the last example shows is the case with the grim-trigger strategy profile. Verifying that each state induces a Nash equilibrium is then much simpler than checking every history.

Exercise 8. Consider the infinitely repeated prisoner's dilemma in Section 1.1.

- (a) Formally write and represent with a figure the smallest automaton representation of the tit-for-tat strategy profile: both players play C at $t = 0$; at each $t \geq 1$, each player plays whatever his opponent played at $t - 1$.
- (b) Formally write and represent with a figure the smallest automaton representation of the following strategy profile: player 1 alternates between C and D (beginning with C), player 2 always plays C , and any deviation results in perpetual mutual defection.

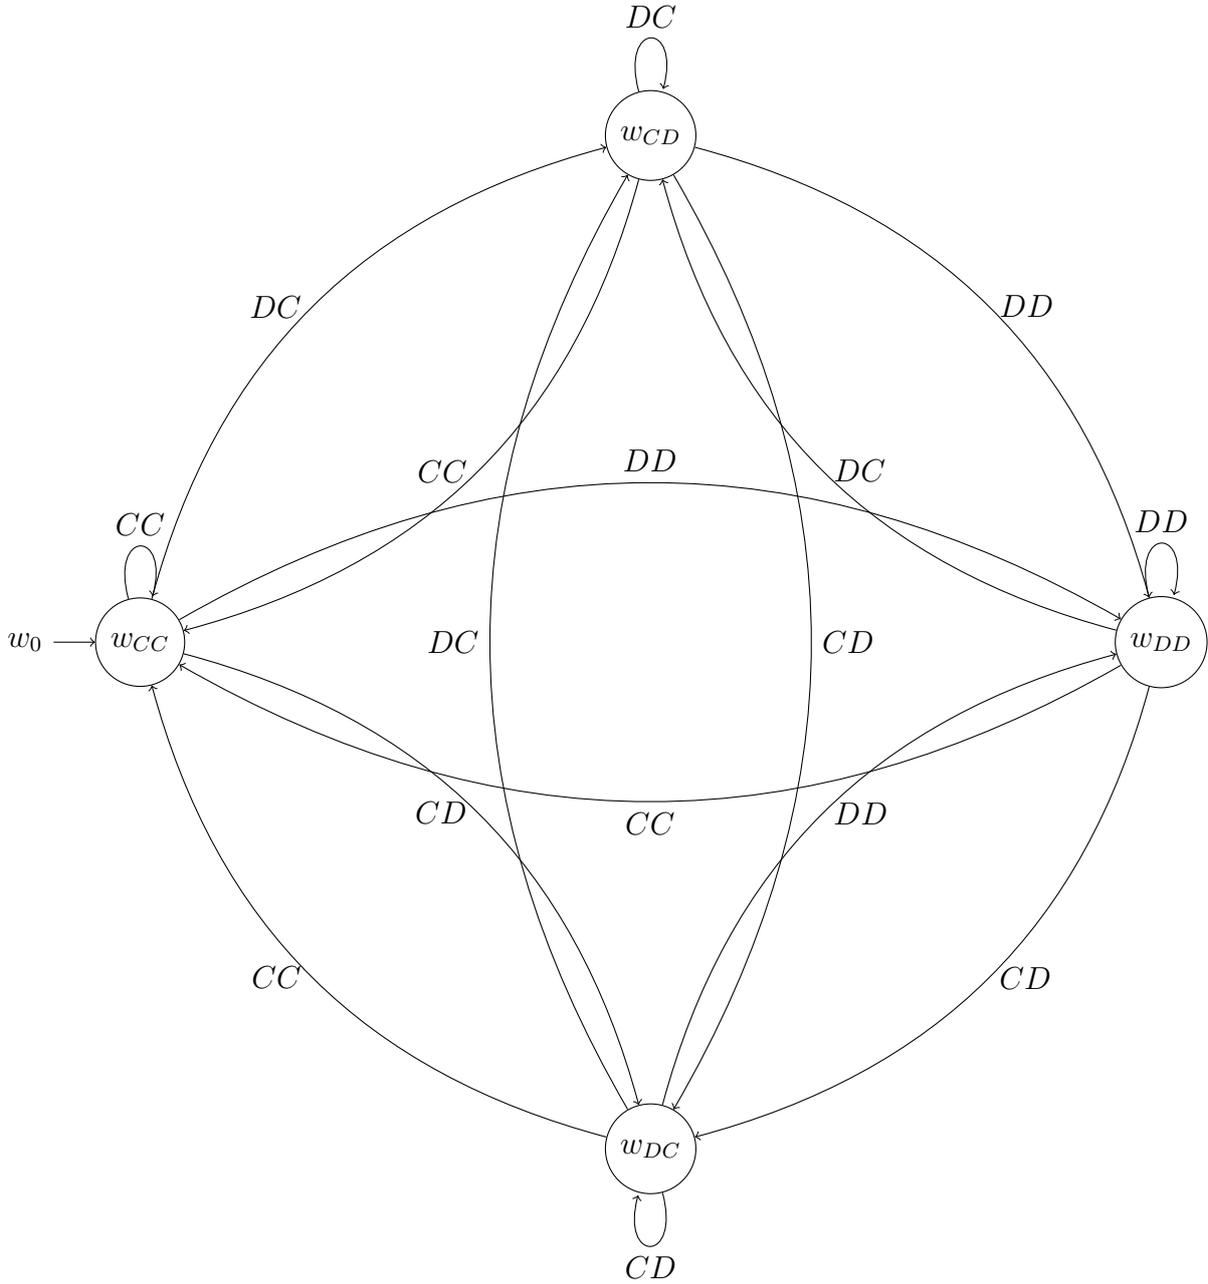
Solution.

- (a) Formally, the tit-for-tat strategy profile has as its smallest automaton representation (W, w^0, f, τ) , where:

$$- W = \{w_{CC}, w_{CD}, w_{DC}, w_{DD}\};$$

- $w^0 = w_{CC}$;
- $f(w_{a_1 a_2}) = (a_1, a_2)$;
- $\tau(w_{a_1 a_2}, (a'_1, a'_2)) = w_{a'_2, a'_1}$.

In a figure, the automaton of the tit-for-tat strategy profile is as follows:



(b) This strategy profile has as its smallest automaton representation (W, w^0, f, τ) , where:

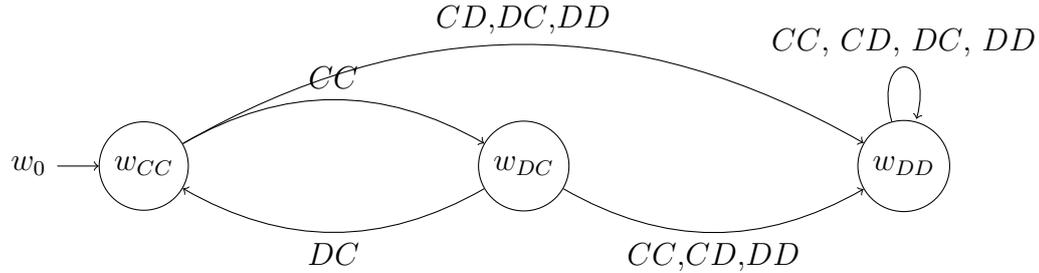
- $W = \{w_{CC}, w_{DC}, w_{DD}\}$;
- $w^0 = w_{CC}$,

$$- f(w_{a_1 a_2}) = (a_1, a_2);$$

-

$$\tau(w, a) := \begin{cases} w_{DC} & \text{if } w = w_{CC} \text{ and } a = (C, C) \\ w_{CC} & \text{if } w = w_{DC} \text{ and } a = (D, C) \\ w_{DD} & \text{otherwise} \end{cases} .$$

In a figure, the automaton of this strategy profile is as follows:



4 Folk Theorems with Perfect Monitoring

4.1 Introduction

Folk Theorems are staples in repeated games: one must consume or be aware of them in thinking about strategic interactions that persist over time. A Folk Theorem asserts that every feasible and strictly individually rational payoff can be associated with some SPNE so long as players are sufficiently patient.³ Thus, in the limit of extreme patience, repeated play allows virtually any payoff to be an equilibrium outcome. The idea of a Folk Theorem is to use variations in payoffs to create incentives towards particular kinds of behavior; the variations in any single stage game may be small, but are amplified when they exist in each period as $\delta \rightarrow 1$.⁴ We will consider three different versions of this result.

1. *Nash Threats Folk Theorem.* A weaker version of Folk Theorem in which payoffs above some stage-game Nash rewards are supportable as SPNE using stage-game Nash equilibria as punishments. The set of payoffs supportable here is generally smaller than the set of feasible and strictly individually rational payoffs. This Folk Theorem is established in [Friedman \(1971\)](#).
2. *Folk Theorem for Nash Equilibrium.* This Folk Theorem supports payoffs that are feasible and strictly individually rational, but only as Nash equilibrium.

³The term “folk” arose because its earliest versions were part of the informal tradition of the game theory community for some time before appearing in a publication.

⁴For Folk Theorems in repeated games with undiscounted payoff criteria, see [Aumann and Shapley \(1994\)](#) and [Rubinstein \(1994\)](#).

3. *Folk Theorem for SPNE.* This is the strongest version (in these lecture notes) that supports payoffs that are feasible and strictly individually rational as SPNE. This Folk Theorem is established in [Fudenberg and Maskin \(1986\)](#).

We return to questions of how to interpret such results after presenting them. Now, let us be a little more formal about public correlating devices. Let $\{\omega^0, \dots, \omega^t\}$ be a sequence of independent draws from, without loss of generality, a uniform distribution on $[0, 1]$ and assume that players observe ω^t at the beginning of period t . With public correlation, the relevant history in period t is $h^t := (\omega^0, a^0, \dots, \omega^{t-1}, a^{t-1}, \omega^t)$, so it includes all prior actions, all prior realizations of the public random variable, and the current realization of it. A behavior strategy for player i is now a function that specifies for each history of this kind which mixed action in $\Delta(A_i)$ to pick. Importantly, with correlating devices, the profitability of a deviation is evaluated *ex post*, that is, conditional on the realization of ω , because the realization of the correlating device is public. As already mentioned, the following results can be established without reference to a public correlating device, but the proofs become more complex (see [Sorin \(1986\)](#) and [Fudenberg and Maskin \(1991\)](#) for the details). Hereafter, denote by $\tilde{\omega}$ a typical realization of ω .

4.2 Nash Threats Folk Theorem

This result is often also called the Nash Reversion Folk Theorem, and is most commonly used in applications: the idea is that for each player i , the punishment from not engaging in equilibrium behavior is a repetition of the stage game Nash equilibrium that offers the lowest payoff to him. By reverting to a stage-game Nash equilibrium, once players are in a Punishment Phase, dynamic incentives are no longer necessary to ensure that all players cooperate in the punishment.

Theorem 3 (Nash Threats Folk Theorem). *Let $\alpha \in \Delta(A)$ be a Nash equilibrium of the stage game G with expected rewards v^{NE} . Then, for all $v \in \mathcal{F}^*$ with $v_i > v_i^{NE}$ for all $i \in N$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a SPNE of G^δ with payoffs v .*

Proof. Fix $v \in \mathcal{F}^*$ with $v_i > v_i^{NE}$ for all $i \in N$. First, suppose that there is a pure action profile $a \in A$ such that $u(a) = v$, and consider the following strategy profile:

- *Cooperation Phase:* If $t = 0$ or $t \geq 1$ and a was played in every prior period, play a .
- *Punishment Phase:* If any other action profile is played in any prior period, then play α for every subsequent period.

Let us argue that this is a SPNE: once a deviation occurs, players are repeating a stage-game Nash equilibrium regardless of the history; since this is a Nash equilibrium of the repeated game at every history, there is nothing more to be established about histories off the equilibrium path. So it suffices to argue that no player has a unilateral profitable

one-shot deviation in the Cooperation Phase. By the one-shot deviation principle, the relevant incentive condition for this to happen is

$$(1 - \delta) \max_{\tilde{a}_i \in A_i} u_i(\tilde{a}_i, a_{-i}) + \delta v_i^{NE} < v_i. \quad (5)$$

Since (5) holds at $\delta = 1$ (as $v_i^{NE} < v_i$) and its the left-hand side is continuous in δ , there must be some $\underline{\delta} < 1$ such that (5) is satisfied for any $\delta \in (\underline{\delta}, 1)$, as desired.

Next, suppose that there is no pure action profile $a \in A$ such that $u(a) = v$. The argument is a bit trickier, but not overly so: replace the pure action profile $a \in A$ with a public randomization stage-game action $a(\omega)$ that yields expected payoff v (that is, such that $\mathbb{E}_\omega[u(a(\omega))] = v$; this is always possible, as $v \in \mathcal{F}^*$, hence, $v \in \mathcal{F}$). Modify the Cooperation Phase above in the obvious way. The Punishment Phase incentives are unaffected, but the Cooperation Phase incentives change in the following sense: depending on the realization $\tilde{\omega}$ of ω , player i may not receive exactly v in the current period⁵, and his deviation payoff also depends on the realization of the public correlating device. However, by the one-shot deviation principle, the Cooperation Phase incentive condition is satisfied so long as

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta v_i^{NE} < (1 - \delta) u_i(a(\tilde{\omega})) + \delta v_i \quad \text{for all } \tilde{\omega} \in [0, 1]. \quad (6)$$

As before, since (6) holds at $\delta = 1$, by continuity in δ , there must be some $\underline{\delta} < 1$ such that (6) is satisfied for any $\delta \in (\underline{\delta}, 1)$, as desired. ■

The following exercise highlights how focusing on Nash reversion can exclude certain repeated game payoffs.

Exercise 9. Give an example of a stage game that has a strictly dominant action for each player (and so a unique Nash equilibrium), but in which there is a SPNE of the repeated game in which each player obtains a strictly lower payoff than the reward in the Nash equilibrium of the stage game. Prove that the strategy profile you construct is a SPNE.

Solution. Consider the stage game G for which the reward matrix is given by

		Player 2		
		A	B	C
Player 1	A	4, 4	3, 0	1, 0
	B	0, 3	2, 2	0, 0
	C	0, 1	0, 0	0, 0

For each of the two players, A is a strictly dominant action and (A, A) is the unique NE. Each player's minmax reward is 1. In the unique NE, on the other hand, each

⁵The payoff v is only the expected ex ante payoff in each period in the Cooperation Phase but may not be the ex post payoff.

player's reward is 4. In the repeated game, payoffs between 1 and 4 cannot be achieved by strategies that react to deviations by choosing A , since one player's choosing A allows the other to obtain a payoff of 4 (by choosing A also), which exceeds his payoff if he does not deviate.

Nevertheless, such payoffs can be achieved in SPNEs. The punishments built into the players' strategies in these equilibria need to be carefully designed. Specifically, consider the following strategy profile, described in two phases as follows:

- *Phase I*: Play (B, B) at every date.
- *Phase II*: Play (C, C) for two periods, then return to Phase *I*.

Play starts with Phase *I*. If there is any deviation, start up Phase *II*. If there is any deviation from that, start Phase *II* again.

We show that if δ is close enough to 1, this strategy profile is a SPNE. Consider player 1 (the game and the strategy profile are symmetric, and so the same reasoning applies to player 2).

- Phase *I*. Player 1's payoffs from following the strategy in the next three periods are $(2, 2, 2)$, whereas his payoffs from a one-shot deviation are at most $(3, 0, 0)$; in both cases, his payoff is subsequently 2 in each period. Thus, following the strategy is optimal if $2 + 2\delta + 2\delta^2 \geq 3$, or, equivalently,

$$\delta \geq \frac{\sqrt{3} - 1}{2}.$$

- Phase *II*, first period. Player 1's payoffs from following the strategy in the next three periods are $(0, 0, 2)$, whereas his payoffs from a one-shot deviation are at most $(1, 0, 0)$; in both cases, his payoff is subsequently 2 in each period. Thus, following the strategy is optimal if $2\delta^2 \geq 1$, or, equivalently,

$$\delta \geq \frac{\sqrt{2}}{2}.$$

- Phase *II*, second period. Player 1's payoffs from following the strategy in the next three periods are $(0, 2, 2)$, whereas his payoffs from a one-shot deviation are at most $(1, 0, 0)$; in both cases, his payoff is subsequently 2 in each period. Thus, following the strategy is optimal if $2\delta + 2\delta^2 \geq 1$, or certainly so if $2\delta^2 \geq 1$, as required by the previous case.

We conclude by the one-shot deviation principle that this strategy profile forms a SPNE if

$$\delta \geq \underline{\delta} := \max \left\{ \frac{\sqrt{3} - 1}{2}, \frac{\sqrt{2}}{2} \right\} = \frac{\sqrt{2}}{2}.$$

Such SPNE yields the two players a payoff of 2, which is strictly lower than the reward in the unique NE of the stage game. ■

4.3 Folk Theorem for Nash Equilibrium

Theorem 4 (Folk Theorem for Nash Equilibrium). *For all $v \in \mathcal{F}^*$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a Nash equilibrium of G^δ with payoffs v .*

The set of payoffs that are supportable is now larger than that of the Nash Threats Folk Theorem, but the solution-concept is more permissive. The proof is straightforward and mirrors the previous proof.

Proof. Fix $v \in \mathcal{F}^*$. First, suppose that there is a pure action profile $a \in A$ such that $u(a) = v$, and consider the following strategy profile:

- *Cooperation Phase:* Play a if: (i) $t = 0$, or (ii) $t \geq 1$ and a was played in every prior period, or (iii) $t \geq 1$ and the realized action profile differs from a in two or more components.
- *Punishment Phase:* If player i was the only one to not follow profile a , then in each period, each player $j \neq i$ plays her component of a mixed strategy that makes player i attain his minmax payoffs.

It may seem odd that when two or more players deviate, no one is punished for it; this does not pose an issue to equilibrium because players compute the expected payoffs of deviating assuming that no one else deviates. Moreover, note that punishment-phase incentives are trivial as in a Nash equilibrium there is nothing to be established about histories off the equilibrium path.

In the Cooperation Phase only behavior that corresponds to (i) and (ii) needs to be incentivized; the other histories are off the equilibrium. Incentives are satisfied if

$$(1 - \delta) \max_{\tilde{a}_i \in A_i} u_i(\tilde{a}_i, a_{-i}) + \delta \underline{v}_i < v_i. \quad (7)$$

As before, since (5) holds at $\delta = 1$ and the left-hand side is continuous in δ , there must be some $\underline{\delta} < 1$ such that (5) is satisfied for any $\delta \in (\underline{\delta}, 1)$, as desired.

Next, suppose that there is no pure action profile $a \in A$ such that $u(a) = v$. Replace the pure action profile $a \in A$ with a public randomization stage-game action $a(\omega)$ that yields expected payoff v . Modify the Cooperation Phases above in the obvious way. Similarly to the proof of Theorem 3 the Cooperation Phase incentive condition is satisfied so long as

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta \underline{v}_i < (1 - \delta) u_i(a(\tilde{\omega})) + \delta v_i \quad \text{for all } \tilde{\omega} \in [0, 1]. \quad (8)$$

As before, since (8) holds at $\delta = 1$, by continuity in δ , there must be some $\underline{\delta} < 1$ such that (8) is satisfied for any $\delta \in (\underline{\delta}, 1)$, as desired. ■

The Folk Theorem for Nash Equilibrium is seldom applied because the punishments used may be implausible. That is, punishments may be very costly for the punisher to carry out and so they represent non-credible threats. In other words, the strategies used may not be subgame perfect. For example, consider the “*destroy the world*” game for which the reward matrix is given by

		Player 2	
		<i>L</i>	<i>R</i>
Player 1	<i>U</i>	6, 6	0, -100
	<i>D</i>	7, 1	0, -100

The minmax rewards are $\underline{v}_1 = 0$ and $\underline{v}_2 = 1$. Note also that Player 2 has to play *R* to minmax Player 1. Theorem 4 informs us that (6, 6) is possible as a Nash equilibrium payoff of the repeated game, but the strategies suggested in the proof require Player 2 to play *R* in every period following a deviation. While this will hurt Player 1, it will hurt Player 2 a lot more. Thus, it seems unreasonable to expect Player 2 to carry out the threat.

The lack of incentives at the punishment phase motivates the search for a Folk Theorem that supports the set of feasible and strictly individually rational payoffs as SPNE of the repeated game.

4.4 Folk Theorem for SPNE

Given some nonempty convex set $B \subseteq \mathbb{R}^n$, the *dimension* of B , denoted by $\dim(B)$, is the dimension of the smallest affine set containing it. A set $A \subseteq \mathbb{R}^n$ is *affine* if for all $x, y \in A$ and $r \in \mathbb{R}$, we have $rx + (1 - r)y \in A$ (and so affine sets are translation of subspaces). The dimension of an affine set is the dimension of the vector subspace of its translations. For a convex set $B \subseteq \mathbb{R}^n$, the condition that $\dim(B) = n$ is topologically equivalent to the condition that B has non-empty interior. The *interior* of a set $B \subseteq \mathbb{R}^n$, denoted by $\text{Int}(B)$, is the largest open set contained in B .

Theorem 5 (Folk Theorem for SPNE). *Suppose $\dim(\mathcal{F}^*) = n$. For all $v \in \mathcal{F}^*$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a SPNE of G^δ with payoffs v .*⁶

Proof. Fix $v \in \mathcal{F}^*$. Suppose for simplicity that:

- There exists a pure action profile $a \in A$ such that $u(a) = v$;
- For each player i , assume that a profile \underline{a}^i that minmaxes player i is in pure actions.

⁶For $n = 2$, one can relax the full dimensionality assumption $\dim(\mathcal{F}^*) = 2$. See Fudenberg and Maskin (1986) for the details.

Both of the above ensure that deviations are detectable, whereas with mixed action profiles, deviations are not detectable with probability 1. Tackling mixed action requires more care. We will briefly return to this point after the proof.

Pick some $\varepsilon > 0$ and some $v' \in \text{Int}(\mathcal{F}^*)$ such that, for all $i \in N$,

$$\underline{v}_i < v'_i < v_i,$$

and the vector

$$v'(i) := (v'_1 + \varepsilon, \dots, v'_{i-1} + \varepsilon, v'_i, v'_{i+1} + \varepsilon, \dots, v'_n + \varepsilon)$$

is in \mathcal{F}^* . The full-dimensionality assumption ensures that $v'(i)$ exists for some v' and ε .⁷ The profile $v'(i)$ is an ε -reward for other players relative to v' , but not for player i . Again, to avoid the details of public randomizations, assume that there exists an action profile $a(i)$ such that $u(a(i)) = v'(i)$. Choose $T \in \mathbb{N}$ such that for all $i \in N$,

$$\max_{\tilde{a} \in A} u_i(\tilde{a}) + T\underline{v}_i < \min_{\tilde{a} \in A} u_i(\tilde{a}) + Tv'_i. \quad (9)$$

In other words, the punishment length T is sufficiently long that player i is worse off from deviating and then being minmaxed for T periods than obtaining the lowest possible payoff once and then T periods of v'_i . Finally, consider the following strategy profile:

- *Phase I.* Play begins in Phase I. Play a so long as no player deviates or at least two players deviate. If a single player i deviates from a , go to Phase II_i .
- *Phase II_i .* Play \underline{a}^i for T periods so long as no one deviates or two or more players deviate. Switch to Phase III_i after T successive periods in Phase II_i . If player j alone deviates, go to Phase II_j .
- *Phase III_i .* Play $a(i)$ so long as no one deviates or two or more players deviate. If a player j alone deviates, go to Phase II_j .

Notice that the above construction makes explicit what actions players should follow when two or more players deviate; such a specification is necessary because strategies are complete contingent plans. Again, it may seem odd that when two or more players deviate, no one is punished for it; this does not pose an issue to equilibrium because players compute the expected payoffs of deviating assuming that no one else deviates.

We use the one-shot deviation principle to check that this is a SPNE.

Phase I

- Player i 's payoff from following the strategy: v_i .

⁷Since $\mathcal{F}^* \neq \emptyset$, we can always find an element of $\text{Int}(\mathcal{F}^*)$ in any open neighborhood of any element of \mathcal{F}^* .

- Player i 's payoff from deviating once: at most $(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i$.
- The relevant incentive condition is

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < v_i. \quad (10)$$

Since (10) holds at $\delta = 1$ (as $v'_i < v_i$) and its left-hand side is continuous in δ , the deviation is unprofitable for δ sufficiently large.

Phase II_i (suppose there are $T' \leq T$ periods left in this phase)

- Player i 's payoff from following the strategy: $(1 - \delta^{T'}) \underline{v}_i + \delta^{T'} v'_i$.
- Player i 's payoff from deviating once (recall that i is being minmaxed): at most $(1 - \delta) \underline{v}_i + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i$.
- The relevant incentive condition is

$$(1 - \delta) \underline{v}_i + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < (1 - \delta^{T'}) \underline{v}_i + \delta^{T'} v'_i,$$

which is satisfied (independently of the value of δ) because the deviation offers no short-term reward and increases the length of punishment so it is not profitable.

Phase II_j ($j \neq i$; suppose there are $T' \leq T$ periods left in this phase)

- Player i 's payoff from following the strategy: $(1 - \delta^{T'}) u_i(\underline{a}^j) + \delta^{T'} (v'_i + \varepsilon)$.
- Player i 's payoff from deviating once: at most $(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i$.
- The relevant incentive condition is

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < (1 - \delta^{T'}) u_i(\underline{a}^j) + \delta^{T'} (v'_i + \varepsilon). \quad (11)$$

Since (11) holds at $\delta = 1$ (as $v'_i < v'_i + \varepsilon$), by continuity in δ , the deviation is unprofitable for δ sufficiently large. That is, for δ close to 1, the ε difference between what i gets in Phase II_j vs. what he gets in Phase II_i makes the deviation unprofitable. This is a place in which rewarding people for punishing others helps with incentives.

Phase III_i

- Player i 's payoff from following the strategy: v'_i .
- Player i 's payoff from deviating once: at most $(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i$.

- The relevant incentive condition is

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < v'_i.$$

The inequality in (9) ensures that the deviation is unprofitable for δ sufficiently close to 1.

Phase III_j ($j \neq i$)

- Player i 's payoff from following the strategy: $v'_i + \varepsilon$.
- Player i 's payoff from deviating once: at most $(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i$.
- The relevant incentive condition is

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < v'_i + \varepsilon. \quad (12)$$

Since (12) holds at $\delta = 1$ (as $v'_i < v'_i + \varepsilon$) and its left-hand side is continuous in δ , the deviation is unprofitable for δ sufficiently large.

Thus, the strategy profile is a SPNE. ■

Remark 5. Throughout the proof of Theorem 5, we have assumed that all payoffs are attained using pure action profiles, in which case deviations are immediately detectable. Matters are more subtle when the minmax requires randomization: then a player must receive the same payoff for each action in the support of his mixed action, although he may realize different payoffs in the stage game. Thus, continuation values must be adjusted in subtle ways based on the realized actions of the punishers. See [Fudenberg and Maskin \(1986\)](#) for more details. ■

Remark 6. The equilibrium constructed in the proof of Theorem 5 involves both the “stick” (Phase II) and the “carrot” (Phase III). Often, however, only the stick is necessary. The carrot phase is needed only if the parties punishing in Phase II get less than their min-max rewards. ■

Remark 7 (Patience and Frequency of Interactions). The crucial condition in Folk Theorems is a high discount factor. Although the discount factor may not be directly observable, it should be high when one period is short. An empirically testable implication is that players who have daily interaction have a better scope for cooperation than those who interact only once a year. An important message of Folk Theorems is that a *high frequency of interaction* is essential for the success of a long term relationship. ■

4.4.1 Dimension and Interiority

The assumption that $\dim(\mathcal{F}^*) = n$ ensures that players need not be simultaneously minmaxed. More specifically, it ensures that player-specific punishments are possible and that, at the same time, punishers can be rewarded for punishing, as the proof of Theorem 5 makes clear. When the assumption fails, punishing a deviator may also punish someone else (i.e. a punisher). In such a circumstance, it would be difficult to incentivize someone else to punish the deviator. The following example (due to [Fudenberg and Maskin \(1986\)](#)) illustrates this point.

Consider the three-player stage game G for which the reward matrix is given by

		Player 2				Player 2		
		L	R			L	R	
Player 1	U	1, 1, 1	0, 0, 0	Player 1	U	0, 0, 0	0, 0, 0	
	D	0, 0, 0	0, 0, 0		D	0, 0, 0	1, 1, 1	
		A				B		
				Player 3				

In the above game, Player 1 chooses rows (U or D), Player 2 chooses columns (L or R), and Player 3 chooses matrices (A or B). We will now show that, whereas each player's minmax reward in this game is 0, the minimum payoff attainable in any SPNE for any player and any discount factor δ is at least $1/4$. Thus, Theorem 5 fails because there are feasible strictly individually rational payoff vectors, namely, those giving some player a payoff in $(0, 1/4)$, that cannot be obtained as the result of any SPNE of the repeated game, for any discount factor.

Let \underline{v} be the minimum payoff attainable in any SPNE of the repeated game, for any player and any discount factor, in this three-player game. Given the symmetry of the game, if one player achieves a given payoff, then all do, obviating the need to consider player-specific minimum payoffs. Let α_i denote the probability with which players use their first action (i.e. U , L , or A) in the first period. Note that each player i can secure as a reward at least

$$\max\{\alpha_j\alpha_k, (1 - \alpha_j)(1 - \alpha_k)\}$$

where i, j, k are distinct, in the first period. Now, without loss, assume that $\alpha_1 \leq \alpha_2 \leq \alpha_3$. If $\alpha_2 \leq 1/2$, then $(1 - \alpha_1)(1 - \alpha_2) \geq 1/4$ and Player 3 can secure a reward of $1/4$. If instead $\alpha_2 \geq 1/2$, then $\alpha_2\alpha_3 \geq 1/4$ and the same holds from Player 1's point of view. That is, at least one player can secure $1/4$ as a reward in the first period. His continuation payoff being at least \underline{v} , it follows that

$$\underline{v} \geq (1 - \delta)\frac{1}{4} + \delta\underline{v},$$

and so $\underline{v} \geq 1/4$.

Clearly, full dimensionality fails in the previous example: all three players have the same payoff function and so the same preferences over action profiles. Therefore, it is impossible to increase or decrease one player's payoff without doing so to all other players.

How much stronger is the sufficient condition $\dim(\mathcal{F}^*) = n$ than the conditions necessary for the Folk Theorem for SPNE? [Abreu, Dutta and Smith \(1994\)](#) show that it suffices for the existence of player-specific punishments (and so for the Folk Theorem) that no two players have identical preferences over action profiles in the stage game. Recognizing that the payoffs in a game are measured in terms of utilities and that affine transformations of utility functions preserve preferences, they offer the following formulation.

Theorem 6 (NEU Folk Theorem for SPNE). *Suppose the game G satisfies the NEU (non-equivalent utilities) condition; that is, there are no two distinct players $i, j \in N$ and real constants c and $d > 0$ such that $u_i(a) = c + du_j(a)$ for all $a \in A$. Then, for all $v \in \mathcal{F}^*$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a SPNE of G^δ with payoffs v .*

Proof. Omitted. ■

When NEU is violated, standard minmax payoff is not the right reference point. [Wen \(1994\)](#) introduced the notion of *effective minmax* to general stage games. Player i 's effective minmax is defined as

$$\underline{v}_i := \min_{a \in A} \max_{j \in J_i} \max_{a_j \in A_j} u_i(a_j, a_{-j}),$$

where J_i is the set of players whose utilities are equivalent to player i 's. Plainly, the effective minmax coincides with the standard minmax when NEU is satisfied. [Wen \(1994\)](#) shows that, if we use the effective minmax payoff instead of the minmax payoff, we can drop the dimensionality assumption in the statement of the Folk Theorem. In the above example, each player's effective minmax is $1/4$ assuming observable mix strategies (thus all mixed strategies are essentially pure strategies). Therefore the SPNE payoff set coincides with $\{(v, v, v) \in \mathbb{R}^3 : v \in (1/4, 1]\}$ in this example (in the limit as $\delta \rightarrow 1$).

Exercise 10. Consider the three-player stage game G for which the reward matrix is given by

		Player 2				Player 2		
		L	R			L	R	
Player 1	U	1, 1, -1	0, 0, 0		Player 1	U	0, 0, 0	0, 0, 0
	D	0, 0, 0	0, 0, 0			D	0, 0, 0	1, -1, 1
		A				B		
				Player 3				

As before, Player 1 chooses rows (U or D), Player 2 chooses columns (L or R), and Player 3 chooses matrices (A or B). Answer the following questions.

- (a) What is the set of feasible and individually rational rewards? What is its dimension? What is its interior? What is the set of feasible and strictly individually rational rewards?
- (b) Does game G satisfies the NEU condition?
- (c) Fix $\delta \in [0, 1)$ and find the set of payoffs that can be supported as Nash equilibrium of G^δ . Does this set depends on δ ?

Solution. Note: this example is due to [Forges, Mertens and Neyman \(1986\)](#).

- (a) Players 2 and 3 can each secure 0 and the sum of their payoffs is also 0. Thus, the set of feasible and individually rational rewards is $\{(v, 0, 0) \in \mathbb{R}^3 : v \in [0, 1]\}$. It has dimension 1 and empty interior. The set of feasible and strictly individually rational rewards is empty.
- (b) Yes, players have non-equivalent utilities.
- (c) The set of payoffs that can be supported as Nash equilibrium of G^δ , no matter the value of $\delta \in [0, 1)$, is $\{(0, 0, 0)\}$. Otherwise, there is a first stage at which players play (U, L, A) or (D, R, B) with positive probability. Without loss, suppose it is (U, L, A) . Then player 2 can secure a positive payoff by playing the equilibrium before and including this stage, and L afterwards, giving him an expected payoff that is strictly positive, a contradiction. ■

5 Repeated Games with Fixed $\delta < 1$

The Folk Theorem for SPNE shows that many payoffs are possible as SPNE of the infinitely repeated game. However, the construction of strategies in the proof is fairly complicated, since we need to have punishments and then rewards for punishers to induce them not to deviate. In general, an equilibrium may be supported by an elaborate hierarchy of punishments, and punishments of deviations from the prescribed punishments, and so on. Also, Folk Theorems are concerned with limits as $\delta \rightarrow 1$, whereas we may be interested in the set of equilibria for a particular value of $\delta < 1$.

We will now approach the question of identifying equilibrium payoffs for a given $\delta < 1$. In repeated games with perfect information, it turns out that an insight of [Abreu \(1988\)](#) will simplify the analysis greatly: equilibrium strategies can be enforced by using a worst possible punishment for any deviator.

5.1 Constructing Equilibria: Simple Strategies and Penal Codes

Hereafter, we restrict attention to pure strategies. The analysis can be extended to mixed strategies, but at a significant cost of increased notation. Let E_δ^p be the set of pure-strategy SPNE payoffs when players have discount factor δ . As we will show in the next set of lecture notes, E_δ^p is a compact set.

Definition 9 (Simple Strategies). *Given $(n + 1)$ outcomes $\{h_0^\infty, h_1^\infty, \dots, h_n^\infty\}$, the associated simple strategy profile, denoted by $\sigma(h_0^\infty, h_1^\infty, \dots, h_n^\infty)$, consists of a prescribed outcome h_0^∞ and a “punishment” outcome h_i^∞ for every player i . Under the profile, play continues according to the outcome h_0^∞ . Players respond to any deviation by player i with a switch to player i ’s punishment outcome h_i^∞ . If player i deviates from outcome h_i^∞ , then h_i^∞ starts again from the beginning. If some other player j deviates, then a switch is made to player j ’s punishment outcome h_j^∞ .*

A critical feature of simple strategy profiles is that the punishment for a deviation by player i is independent of when the deviation occurs and of the nature of the deviation, that is, it is independent of the “crime”. For instance, the profiles used to prove the Folk Theorems for infinitely repeated games with perfect monitoring are simple.

Let

$$v_i^t(h^\infty) := (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_i(a^\tau)$$

be the payoff to player i from the outcome (a^t, a^{t+1}, \dots) . Moreover, let $a^t(j)$ be the action profile in period t according to outcome h_j^∞ . We can use the one-shot deviation principle to identify necessary and sufficient conditions for a simple strategy profile to be a SPNE.

Lemma 1. *The simple strategy profile $\sigma(h_0^\infty, h_1^\infty, \dots, h_n^\infty)$ is a SPNE of G^δ if and only if*

$$v_i^t(h_j^\infty) \geq \max_{a_i \in A_i \setminus \{a_i^t(j)\}} (1 - \delta) u_i(a_i, a_{-i}^t(j)) + \delta v_i^0(h_i^\infty) \quad (13)$$

for all $i \in N$, $j = 0, 1, \dots, n$, and $t = 0, 1, \dots$.

Proof. The right side of (13) is the payoff to player i from deviating from outcome h_j^∞ in the t -th period, and the left side is the payoff from continuing with the outcome. If this condition holds, then no player will find it profitable to deviate from any of the outcomes $\{h_0^\infty, h_1^\infty, \dots, h_n^\infty\}$. Condition (13) thus suffices for subgame perfection

Because a player might be called on (in a suitable out-of-equilibrium event) to play any period- t of any outcome h_j^∞ in a simple strategy profile, condition (13) is also necessary for subgame perfection. ■

A simple strategy profile specifies an equilibrium path h_0^∞ and a *penal code* $\{h_1^\infty, \dots, h_n^\infty\}$ describing responses to deviations from equilibrium play. We are interested in optimal

penal codes, embodying the most severe such punishments. Let

$$v_i^{min} := \min\{v_i \in \mathbb{R} : v \in E_\delta^p\}$$

be the smallest pure-strategy SPNE payoff for player i (which is well defined by the compactness of E_δ^p).

Definition 10 (Optimal Penal Codes). *Let $\{h_i^\infty\}_{i \in N}$ be n outcomes satisfying*

$$v_i^0(h_i^\infty) = v_i^{min} \tag{14}$$

for all $i \in N$. The collection of n simple strategy profiles $\{\sigma(i)\}_{i \in N}$,

$$\sigma(i) := \sigma(h_i^\infty, h_1^\infty, \dots, h_n^\infty)$$

is an optimal penal code if, for all $i \in N$, the strategy profile $\sigma(i)$ is a SPNE.

Do optimal penal codes exist? Compactness of E_δ^p yields the subgame-perfect outcomes h_i^∞ satisfying (14). The remaining question is whether the associated simple strategy profiles constitute equilibria. The first statement of the following proposition shows that optimal penal codes exist. The second, reproducing the key result of [Abreu \(1988\)](#), is the punchline of the characterization of SPNE: simple strategies suffice to achieve any feasible SPNE payoff.

Proposition 3. *1. Let $\{h_i^\infty\}_{i \in N}$ be n outcomes of pure-strategy SPNE $\{\sigma(i)\}_{i \in N}$ satisfying $v_i(\sigma(i)) = v_i^{min}$ for all $i \in N$. Then, the simple strategy profile $\sigma(i) := \sigma(h_i^\infty, h_1^\infty, \dots, h_n^\infty)$ is a pure-strategy SPNE, for all $i \in N$, and hence $\{\sigma(i)\}_{i \in N}$ is an optimal penal code.*

2. The pure outcome h_0^∞ can be supported as an outcome of a pure-strategy SPNE if and only if there exist pure outcomes $\{h_1^\infty, \dots, h_n^\infty\}$ such that the simple strategy profile $\sigma(h_0^\infty, h_1^\infty, \dots, h_n^\infty)$ is a SPNE.

Hence, anything that can be accomplished with a SPNE in terms of payoffs can be accomplished with simple strategies: any SPNE payoff that can be achieved can be done so by specifying $n + 1$ outcomes, one that must be followed as long as no player deviates, and one for each player in case of deviation by this player. As a result, we need never consider complex hierarchies of punishments when constructing SPNE, nor do we need to tailor punishments to the deviations that prompted them (beyond the identity of the deviator). It suffices to associate one punishment with each player, to be applied whenever needed.

Proof. The “if” direction of statement 2 is immediate.

To prove statement 1 and the “only if” direction of statement 2, let h_0^∞ be the outcome of a pure strategy SPNE denoted by $\hat{\sigma}(0)$ (note: $\hat{\sigma}(0)$ is not necessarily a simple strategy profile). Let $\{h_i^\infty\}_{i \in N}$ be outcomes of pure-strategy SPNE $\{\hat{\sigma}(i)\}_{i \in N}$, with $v_i(\hat{\sigma}(i)) = v_i^{\min}$ for all $i \in N$. Now consider the simple strategy profile given by $\sigma(h_0^\infty, h_1^\infty, \dots, h_n^\infty)$. We claim that this strategy profile constitutes a SPNE. Considering arbitrary h_0^∞ , this argument establishes statement 2. For $h_0^\infty \in \{h_1^\infty, \dots, h_n^\infty\}$, it establishes statement 1.

By Lemma 1, it suffices to fix a player i , an index $j \in \{0, 1, \dots, n\}$, a time t , and action $a_i \in A_i \setminus \{a_i^t(j)\}$, and show

$$v_i^t(h_j^\infty) \geq (1 - \delta)u_i(a_i, a_{-i}^t(j)) + \delta v_i^0(h_i^\infty). \quad (15)$$

Now, by construction, h_j^∞ is the outcome of a SPNE—the outcome h_0^∞ is by assumption produced by a SPNE, whereas each of $h_1^\infty, \dots, h_n^\infty$ is part of an optimal penal code. This ensures that for any t and $a_i \in A_i \setminus \{a_i^t(j)\}$,

$$v_i^t(h_j^\infty) \geq (1 - \delta)u_i(a_i, a_{-i}^t(j)) + \delta v_i^d(h_i^\infty, t, a_i), \quad (16)$$

where $v_i^d(h_i^\infty, t, a_i)$ is the continuation payoff received by player i in equilibrium $\hat{\sigma}(j)$ after the deviation to a_i in period t . Since $\hat{\sigma}(j)$ is a SPNE, the payoff $v_i^d(h_i^\infty, t, a_i)$ must itself be a SPNE payoff. Hence,

$$v_i^d(h_i^\infty, t, a_i) \geq v_i^{\min} = v_i^0(h_i^\infty),$$

which with (16) implies (15), giving the result. ■

It is an immediate corollary that not only can we restrict attention to simple strategies but we can also take the penal codes involved in these strategies to be optimal.

Corollary 1. *Suppose h_0^∞ is the outcome of some SPNE. Then, the simple strategy $\sigma(h_0^\infty, h_1^\infty, \dots, h_n^\infty)$, where each h_i^∞ yields the lowest possible SPNE payoff v_i^{\min} to player i , is a SPNE.*

For a nice and classic illustration of the use of simple strategies and penal codes, see [Abreu \(1986\)](#).

6 Concluding Remarks

These lecture notes are only a first introduction to the theory of repeated games with perfect monitoring. If you wish to learn more, an excellent starting point is Part I in [Mailath and Samuelson \(2006\)](#). Among others, we do not cover (or only briefly touch upon, either in the lecture notes or as suggested readings for the report) the following important topics (with some references to seminal work).

Folk Theorem with Long-Lived and Short-Lived Players

- Sections 2.7 and 3.6 in [Mailath and Samuelson \(2006\)](#).
- [Fudenberg, Kreps and Maskin \(1990\)](#).

Folk Theorem with Overlapping Generations of Players

- [Smith \(1992\)](#).
- [Kandori \(1992a\)](#).

Folk Theorem with Random Matching

- Section 5.1 in [Mailath and Samuelson \(2006\)](#).
- [Kandori \(1992b\)](#).
- [Ellison \(1994\)](#).
- [Okuno-Fujiwara and Postlewaite \(1995\)](#).
- [Ghosh and Ray \(1996\)](#) (Adverse selection, not on folk theorem).

Finitely Repeated Games

- Section 4.4 in [Mailath and Samuelson \(2006\)](#).
- [Friedman \(1985\)](#) (Nash threats Folk Theorem).
- [Benoît and Krishna \(1985\)](#) (Folk Theorem with observable mixed strategy).
- [Smith \(1995\)](#) (Necessary and sufficient condition for Folk Theorem).
- [Gossner \(1995\)](#) (Folk Theorem without observable mixed strategies).

Renegotiation in Repeated Games

- Section 4.6 in [Mailath and Samuelson \(2006\)](#) and Section 4 in [Pearce \(1992\)](#).
- [Benoît and Krishna \(1993\)](#).
- [Farrell and Maskin \(1989\)](#).
- [Bernheim and Ray \(1989\)](#).
- [Abreu, Pearce and Stacchetti \(1993\)](#).

Repeated Extensive Forms

- Section 5.4 in [Mailath and Samuelson \(2006\)](#).

References

- Abreu, Dilip (1986), “Extremal Equilibria of Oligopolistic Supergames.” *Journal of Economic Theory*, 39, 191–225.
- (1988), “On the Theory of Infinitely Repeated Games with Discounting.” *Econometrica*, 56, 383–396.
- Abreu, Dilip, Prajit K. Dutta, and Lones Smith (1994), “The Folk Theorem for Repeated Games: A NEU Condition.” *Econometrica*, 62, 939–948.
- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1993), “Renegotiation and Symmetry in Repeated Games.” *Journal of Economic Theory*, 60, 217–140.
- Ali, S. Nageeb (2011), “Notes for Economics 200c.” *Lecture Notes*.
- Aumann, Robert J. and Lloyd S. Shapley (1994), “Long-Term Competition—A Game-Theoretic Analysis.” Nimrod Megiddo ed. *Essays in Game Theory (In Honor of Michael Maschler)*. Springer, 1–15.
- Benoît, Jean-Pierre and Vijay Krishna (1985), “Finitely Repeated Games.” *Econometrica*, 53, 905–922.
- (1993), “Renegotiation in Finitely Repeated Games.” *Econometrica*, 61, 303–323.
- Bernheim, Douglas B. and Debraj Ray (1989), “Collective Dynamic Consistency in Repeated Games.” *Games and Economic Behavior*, 1, 296–326.
- Blackwell, David (1965), “Discounted Dynamic Programming.” *The Annals of Mathematical Statistics*, 36, 226–235.
- Ellison, Glenn (1994), “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching.” *The Review of Economic Studies*, 61, 567–588.
- Farrell, Joseph and Eric S. Maskin (1989), “Renegotiation in Repeated Games.” *Games and Economic Behavior*, 1, 327–360.
- Forges, Françoise, Jean-François Mertens, and Abraham Neyman (1986), “A Counterexample to the Folk Theorem with Discounting.” *Economics Letters*, 20, p. 7.
- Friedman, James W. (1971), “A Non-cooperative Equilibrium for Supergames.” *The Review of Economic Studies*, 38, 1–12.
- (1985), “Cooperative Equilibria in Finite Horizon Non-Cooperative Supergames.” *Journal of Economic Theory*, 52, 390–398.
- Fudenberg, Drew, David M. Kreps, and Eric S. Maskin (1990), “Repeated Games with Long-Run and Short-Run Players.” *The Review of Economic Studies*, 57, 555–573.
- Fudenberg, Drew and Eric S. Maskin (1986), “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information.” *Econometrica*, 54, 533–554.

- (1991), “On the Dispensability of Public Randomization in Discounted Repeated Games.” *Journal of Economic Theory*, 53, 428–438.
- Fudenberg, Drew and Jean Tirole (1991), “Game Theory.” *MIT Press*.
- Ghosh, Parikshit and Debraj Ray (1996), “Cooperation in Community Interaction without Information Flows.” *The Review of Economic Studies*, 63, 491–519.
- Gossner, Olivier’ (1995), “The Folk Theorem for Finitely Repeated Games with Mixed Strategies.” *International Journal of Game Theory*, 24, 95–107.
- Hörner, Johannes (2015), “Repeated Games, Part I: Perfect Monitoring.” *Lecture Notes*.
- Kandori, Michihiro (1992a), “Repeated Games Played by Overlapping Generations of Players.” *The Review of Economic Studies*, 59, 81–92.
- (1992b), “Social Norms and Community Enforcement.” *The Review of Economic Studies*, 59, 63–80.
- (2006), “Repeated Games.” *The New Palgrave Dictionary of Economics*.
- Mailath, George J. (2019), “Modeling Strategic Behavior: A Graduate Introduction to Game Theory and Mechanism Design.” *World Scientific*.
- Mailath, George J. and Larry Samuelson (2006), “Repeated Games and Reputations: Long-Run Relationships.” *Oxford University Press*.
- Nash, John Forbes (1951), “Non-Cooperative Games.” *Annals of Mathematics*, 54, 286–295.
- Okuno-Fujiwara, Masahiro and Andrew Postlewaite (1995), “Social Norms and Random Matching Games.” *Games and Economic Behavior*, 9, 79–109.
- Pearce, David (1992), “Repeated Games: Cooperation and Rationality.” Jean-Jacques Laffont ed. *Advances in Economic Theory: Sixth World Congress of the Econometric Society, Volume 1*.
- Rubinstein, Ariel (1994), “Equilibrium in Supergames.” Nimrod Megiddo ed. *Essays in Game Theory (In Honor of Michael Maschler)*. Springer, 17–27.
- Smith, Lones (1992), “Folk theorems in Overlapping Generations Games.” *Games and Economic Behavior*, 4, 426–449.
- (1995), “Necessary and Sufficient Conditions for the Perfect Finite Horizon Folk Theorem.” *Econometrica*, 63, 425–430.
- Sorin, Sylvain (1986), “On Repeated Games with Complete Information.” *Mathematics of Operations Research*, 11, 147–160.
- Wen, Quan (1994), “The “Folk Theorem” for Repeated Games with Complete Information.” *Econometrica*, 62, 949–954.