

Repeated Games with Perfect Monitoring

Niccolò Lomys*

February 9, 2019

Logistics

Topics.

1. Repeated Games with Perfect Monitoring
2. Repeated Games with Imperfect Public Monitoring
3. Signaling Games
4. Reputations in Repeated Games

References. For each topic, lecture notes will be available on my website. I am not planning to discuss all of the lecture notes in class, but I expect that you understand them and solve the exercises they contain for the exam. You should think of my lectures as the “highlights” of the material, and that a thorough learning shall come from reading these notes, solving problem sets, and discussing the material with your classmates. The lecture notes also reference the relevant textbooks and papers in the literature. Reading textbooks and papers, though not mandatory, is a good way to sharpen your understanding of the topics we cover.

Grading. Your final grade will depend on two problem sets (30%) and a final exam (70%). The first problem set is on topics 1 and 2; the second problem set is on topics 3 and 4. Each problem set consists in solving the exercises that are contained in the relevant parts of the lecture notes. You have to solve problem sets in groups of two or three students (no less, no more), and hand in a solution sheet for each group.

Office Hours. By appointment (just send me an email).

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

Preamble

- These notes heavily draw upon [Hörner \(2015\)](#), [Ali \(2011\)](#), [Mailath and Samuelson \(2006\)](#), and [Fudenberg and Tirole \(1991\)](#). All errors are my own. Please bring any error, including typos, to my attention.
- These notes are only a first introduction to the theory of repeated games with perfect monitoring. If you want (or need) to learn more, an excellent starting point is Part I in [Mailath and Samuelson \(2006\)](#).

1 A Very Short Introduction

Most interactions involve a dynamic element. Let us consider for now the *prisoners' dilemma* for which the payoff matrix is given by

| | | Player 2 | |
|----------|----------|---------------------------|---------------------------|
| | | <i>C</i> | <i>D</i> |
| Player 1 | <i>C</i> | 1, 1 | − <i>L</i> , 1 + <i>G</i> |
| | <i>D</i> | 1 + <i>G</i> , − <i>L</i> | 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. It is also customary to assume that $G - L < 1$, so that the sum of payoffs when both players cooperate (2) exceeds the sum when one defects ($1 + G - L$).

The prisoners' dilemma is the normal form of a two-player, simultaneous-action game. Defecting is strictly dominant for each player, and so in an isolated interaction both players defect. What happens if this simultaneous-move game is repeated twice or, more generally, T times, and after each play, both players can observe what the opponent has done? How about if it is repeated infinitely often? How about if we consider other simultaneous-action games? This is the topic of *repeated games with perfect monitoring*, introduced in these notes.

2 Preliminaries

2.1 Model

Stage Game

The building block of a repeated game is the normal form game corresponding to each single interaction, which is referred to as the *stage game*. Here, we consider a finite stage game, and 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 generic 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 generic action of player i and by a a generic 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 . A mixed action for player i is an element of $\Delta(A_i)$, denoted 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 $\alpha(a)$ is the probability assigned to a by $\alpha \in \Delta(A)$. This (expected) utility is also referred to as the *reward*, instead of payoff, 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.

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

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

That is, 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 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 | |
| | | U | D |
| Player 1 | L | 2, 2 | 1, 5 |
| | R | 5, 1 | 0, 0 |

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

¹The convex hull of a set $A \subseteq \mathbb{R}^n$, denoted $\text{co}(A)$, is the smallest convex set containing A .

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 \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, α_{-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 .²

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 $i \in N$.*

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

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

Exercise 2. Solve the following problems.

1. Given a stage game G , player i 's pure action minmax reward, denoted \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 conceptually $\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_i, a_{-i})$ 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.

²An exception is the prisoners' 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 prisoners' dilemma arise out of this coincidence.

2. The definition of minmax reward assumes that player i 's opponents randomize independently of each other. Define a notion 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.

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?

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 , called the *horizon* of the game, 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*. Observe that, since we let time start at 0, if $T < \infty$, the game is actually repeated $T + 1$ times.

In this first set of notes, we study repeated games with *perfect monitoring*; that is, we assume that 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 , thus, 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 its elements 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 as the initial information set. The set of *histories* is defined as $H := \cup_{t=0}^T H^t$, with generic element $h \in H$.

A *pure strategy* for player i , then, 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 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 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 pure, resp. behavior, strategy profile. The statement of Kuhn's theorem, establishing the realization-equivalence between mixed and behavior strategies, applies to repeated games as well.³

A strategy profile $\sigma \in \Sigma$ generates a distribution over terminal nodes, that is, over histories H^{T+1} . Again, it is clear what is meant when T is finite: if Z is finite, defining the probability

³Two 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.

space is simple. Let Z be the set of outcomes, and the set of events is the set $\mathcal{P}(Z)$ of all subsets of Z . If $T = \infty$, the set of outcomes Z is the set of infinite sequences $(a^0, a^1, \dots) \in A^{\mathbb{N}_0}$, and so it is infinite as well, and defining the set of events introduces technical details, which it is best to ignore. Hereafter, we denote by \mathbf{a} a generic outcome path.

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

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 $\sigma_i|_h$, 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 repeated games have a convenient *recursive structure*, and this play 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 terminal histories 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 game, player i 's payoff given some infinite history $h^\infty = (a^0, a^1, \dots)$, is

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

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 infinite histories, 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 infinite histories 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 $\delta < 1$.

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

1. 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?
2. 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?
3. 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?

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 irrational 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 2.2), then by organizing the subgames that must be checked for Nash equilibria into equivalence classes (see Section 2.3 in [Mailath and Samuelson \(2006\)](#)), 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 $h \in H$, $\sigma_i(h) = \alpha_i$.*

Exercise 5. Prove Theorem 1.

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 subgame-perfect Nash equilibria where, in some periods, players play some actions that are not played in any Nash equilibrium of the stage game.

Exercise 6. 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.

Exercise 7. 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.

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, any subgame-perfect Nash) equilibrium of G^T for any $T < \infty$ and G^δ for any $\delta \in [0, 1)$.*

Exercise 8. Prove Proposition 1.

2.2 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 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 $v_i(\sigma''_i|_{h^t}, \sigma_{-i}|_{h^t}) - v_i(\sigma|_{h^t}) > 0$ 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: i.e., 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 exercise below helps you evaluate the validity of this conjecture.

Exercise 9. Consider the prisoners' 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.

1. What is the set of feasible and individually rational rewards?
2. 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.
3. Consider a “tit-for-tat” strategy profile in which both players cooperate at time $t = 0$, and for $t \geq 1$, each player mimics the action that the opponent picked in the previous period.
 - (a) Suppose that Player 1 shirks forever while Player 2 follows tit-for-tat. What is Player 1's expected payoff from this “infinite-shot” deviation?
 - (b) Consider the following one-shot deviation from the equilibrium path: 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?
 - (c) Under what condition on δ is tit-for-tat a Nash equilibrium profile? Compare the condition with that above to evaluate Conjecture 1.
 - (d) Prove that tit-for-tat is not a SPNE, regardless of the value of δ .

3 Folk Theorems for Infinitely Repeated Games

3.1 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.⁴ 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 $\delta \rightarrow 1$. We will consider three different versions of this result.

⁴The 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.

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.
3. *Folk Theorem for SPNE.* This is the strongest version (in these 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 from $[0, 1]$ and assume that players observe ω^t at the beginning of period t . With public correlation, the history relevant 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).

3.2 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 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. First, suppose that there is a pure action profile 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. The relevant incentive condition for this to happen is

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

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

Next, suppose that there is no pure action profile a such that $u(a) = v$. The argument is a bit trickier, but not overly so: replace the pure action profile a with a public randomization stage-game action $a(\omega)$ that yields expected payoff v (this is always possible, as $v \in \mathcal{F}^*$, hence, $v \in \mathcal{F}$). Modify the Cooperation Phases above in the obvious way. The Punishment Phase incentives are unaffected, but the Cooperation Phase incentives change in the following sense: depending on 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, the Cooperation Phase incentive condition is satisfied so long as

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta v_i^{NE} < (1 - \delta) u_i(a(\omega)) + \delta v_i. \quad (2)$$

As before, since (2) holds at $\delta = 1$ and its left-hand side is continuous in δ , there must be some $\underline{\delta} < 1$ such that (2) 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 10. Give an example of a stage game that has a strictly dominant strategy for each player (and hence 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.

3.3 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. First, suppose that there is a pure action profile a such that $u(a) = v$, and consider the following strategy profile:

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.

- *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

$$(1 - \delta) \max_{\tilde{a}_i \in A_i} u_i(\tilde{a}_i, a_{-i}) + \delta \underline{v}_i < v_i. \quad (3)$$

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

Next, suppose that there is no pure action profile a such that $u(a) = v$. Replace the pure action profile 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

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta \underline{v}_i < (1 - \delta) u_i(a(\omega)) + \delta v_i. \quad (4)$$

As before, since (4) holds at $\delta = 1$ and its left-hand side is continuous in δ , there must be some $\underline{\delta} < 1$ such that (4) 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 | |
| | | L | R |
| Player 1 | U | 6, 6 | 0, -100 |
| | D | 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 stage motivates the search for a Folk Theorem that supports the set of feasible and strictly individually rational payoffs as SPNE of the repeated game.

3.4 Folk Theorem for SPNE

Given some convex set $B \in \mathbb{R}^n$, the *dimension* of B , denoted $\dim(B)$, is the maximum number of linearly independent vectors in B . Topologically, the condition that $\dim(B) = n$ is equivalent to the condition that B has non-empty interior.⁶

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 .*

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

- There exists a pure action profile 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.

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 ε and v' .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 $i \in N$,

$$\max_{\tilde{a} \in A} u_i(\tilde{a}) + T\underline{v}_i < \min_{\tilde{a} \in A} u_i(\tilde{a}) + Tv'_i. \quad (5)$$

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 .

⁶The interior of a set $B \subseteq \mathbb{R}^n$, denoted $\text{Int}(B)$, is the largest open set contained in B .

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}^* .

- *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 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 .
- 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

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < v_i. \quad (6)$$

Since (6) 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

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < (1 - \delta^{T'}) u_i(\underline{a}^j) + \delta^{T'} (v'_i + \varepsilon). \quad (7)$$

Since (7) holds at $\delta = 1$ (as $v'_i + \varepsilon > v'_i$) and its left-hand side is continuous 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 (5) 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

$$(1 - \delta) \max_{\tilde{a} \in A} u_i(\tilde{a}) + \delta(1 - \delta^T) \underline{v}_i + \delta^{T+1} v'_i < v'_i + \varepsilon. \quad (8)$$

Since (8) holds at $\delta = 1$ (as $v'_i + \varepsilon > v'_i$) 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 2. 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 3. 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.

3.4.1 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 | |

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 (weakly) greater than $1/4$. 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 11. Consider the three-player stage game G for which the reward matrix is given by

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

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.

1. 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?
2. Does game G satisfies the NEU condition?
3. 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 ?

4 Generating Equilibria

Many arguments in repeated games are based on constructive proofs. We are often interested in equilibria with certain properties, such as equilibria featuring the highest or lowest payoffs

for some players, and the argument proceeds by exhibiting an equilibrium with the desired properties. When reasoning in this way, we are faced with the prospect of searching through a prohibitively immense sets of possible strategies and equilibria. In the next two sections, we describe two complementary approaches for finding equilibria.

Disclaimer. For expositional clarity, throughout Section 4 we restrict attention to pure actions and strategies.

4.1 Dynamic Programming and Self-Generation: A First Look

Abreu, Pearce and Stacchetti (1986, 1990) recognized that one can decompose repeated games into straightforward dynamic programming problems in which behavior today is implemented with self-enforcing “utility payments” tomorrow; this problem is analogous to a principal-agent setting in which the agent is rewarded or punished for certain actions.⁸ The caveat is that these promises of utility have to be “self-enforcing”, or in other words, generated from equilibria in the repeated game. As we will now show, this decomposition of SPNE payoffs into flow payoffs today and promised utilities tomorrow simplifies the study of repeated interactions.

Fix $\delta \in [0, 1)$. Consider functions of the form $\gamma: A \rightarrow W$, where W is a non-empty subset of \mathbb{R}^n . Denote with W^A the set of all such functions. Here, you should think of $\gamma_i(a)$ as the payment promises that player i obtains when the action profile is a . That the range of γ is W means that only payments in W can be used. These payments do not happen immediately but in the future; their value determines the strength of the incentives that they create today.

Definition 9 (Enforceability). *A pure action profile a is enforceable on W if there exists some specification of (not necessarily credible) continuation promises $\gamma \in W^A$ such that, for all $i \in N$ and $\tilde{a}_i \in A_i$,*

$$(1 - \delta)u_i(a) + \delta\gamma_i(a) \geq (1 - \delta)u_i(\tilde{a}_i, a_{-i}) + \delta\gamma_i(\tilde{a}_i, a_{-i}),$$

In this case, we say that γ enforces a on W .

In other words, when the other players play their part of an enforceable profile a , the continuation promises γ_i make (enforce) the choice of a_i optimal (incentive compatible) for i .

We now iterate the logic: we have a set of promised utilities W , and actions that are enforceable given those promised utilities; why not think about payoffs that emerge from combining actions that are enforceable on W and the enforcing continuation promises?

Definition 10 (Decomposability). *A payoff vector $v \in \mathbb{R}^n$ is pure-action decomposable on W if there exists a pure action profile a enforceable on W such that, for all $i \in N$,*

$$v_i = (1 - \delta)u_i(a) + \delta\gamma_i(a),$$

where γ is a function enforcing a . In this case, we say that v is pure-action decomposed by (a, γ) on W .

⁸Abreu et al. (1986) developed the main ideas behind self-generation, though similar ideas were introduced in Mertens and Parthasarathy (2003). The function B_δ , defined below, was introduced by Shapley (1953).

Exercise 12. Consider a payoff v such that, for some $i \in N$, $v_i < \underline{v}_i$. Is v pure-action decomposable on some set $W \subseteq \mathbb{R}^n$ for some $\delta \in [0, 1)$? Motivate your answer.

Let $\mathcal{P}(\mathbb{R}^n)$ denote the set of all subsets of \mathbb{R}^n . We define the function

$$B_\delta^p: \mathcal{P}(\mathbb{R}^n) \rightarrow \mathcal{P}(\mathbb{R}^n), \quad W \mapsto B_\delta^p(W),$$

where

$$B_\delta^p(W) := \{v \in \mathbb{R}^n : v \text{ is pure-action decomposable on } W \text{ with discount } \delta\}.$$

These are all the payoffs that can be obtained by pure-action decomposition, by using continuation promises from W only when the discount factor is δ .

Proposition 2 (Properties of B_δ^p). *The following statements hold true.*

1. B_δ^p is a monotone function; that is, $W \subseteq W' \implies B_\delta^p(W) \subseteq B_\delta^p(W')$.
2. B_δ^p maps compact sets into compact sets; that is, W compact $\implies B_\delta^p(W)$ compact.
3. $B_\delta^p(\mathcal{F}) \subseteq \mathcal{F}$.

Exercise 13. Prove Proposition 2.

Example 1. To make matters more concrete, consider the prisoners' dilemma for which the reward matrix is given by

| | | | |
|----------|-----|----------|-------|
| | | Player 2 | |
| | | C | D |
| Player 1 | C | 3, 3 | -1, 4 |
| | D | 4, -1 | 1, 1 |

Set $W := \{(3, 3), (1, 1)\}$. Let us see whether the action profile (C, C) is enforceable on W for some $\delta \in [0, 1)$. Consider $\gamma \in W^A$ defined pointwise as $\gamma(C, C) := (3, 3)$ and $\gamma(C, D) = \gamma(D, C) = \gamma(D, D) := (1, 1)$. Then, (C, C) is enforceable on W if

$$(1 - \delta)(3) + \delta(3) \geq (1 - \delta)(4) + \delta. \tag{9}$$

Note that (9) is satisfied for any $\delta \geq 1/3$. Therefore, for any $\delta \geq 1/3$, the action profile (C, C) is enforceable on W and the payoff vector $(3, 3)$ is pure-action decomposable on W . Moreover, $(1, 1)$, being the payoff from the stage game Nash equilibrium, is also pure-action decomposable on W for any δ (use a constant γ mapping each strategy profile into $(1, 1)$). To sum up, W is pure-action decomposable on itself for any $\delta \geq 1/3$. This is our first example of a pure-action *self-generating set* of payoffs.

Definition 11 (Self-Generation). *The set W is pure-action self-generating if every payoff in W is pure-action decomposable on W , that is, if $W \subseteq B_\delta^p(W)$.*

Notice that W may be pure-action self-generating, but $B_\delta^p(W)$ may contain payoffs not in W . That is, more payoffs are pure-action decomposable on W than those in W . This should naturally motivate you to wonder about the largest pure-action self-generating set. We will answer this question soon.

One reason to be interested in self-generating sets is their tight connection to SPNE payoffs. Let E_δ^p be the set of pure-strategy SPNE payoffs when players have discount factor δ . Then, we have the following.

Theorem 7 (Self-Generation and PPE Payoffs). *If the set $W \subseteq \mathbb{R}^n$ is bounded and pure action self-generating, then $B_\delta^p(W) \subseteq E_\delta^p$ (and hence $W \subseteq E_\delta^p$).*

We will not present the formal argument here (we will do so in the context of repeated games with imperfect public monitoring). However, the intuition should be clear. If a payoff vector is pure-action decomposable on W , then it is “one-period credible” with respect to promises in the set W . If those promises are themselves pure-action decomposable on W , then the original payoff vector is “two-period credible”. One can iterate this argument to show that, if a set of payoffs W is pure-action decomposable on itself, then each such payoff has “infinite-period credibility”; thus, it is a pure-strategy SPNE payoff of the repeated game.

Proposition 3 (Factorization). *The following statements hold true.*

1. $E_\delta^p = B_\delta^p(E_\delta^p)$, that is, E_δ^p is a fixed point of the function B_δ^p .⁹
2. E_δ^p is the largest pure-action self-generating set.

Proof. Omitted. See lecture notes on repeated games with imperfect public monitoring. ■

Proposition 4. *The set E_δ^p is compact.*

Proof. Since E_δ^p is a bounded subset of \mathbb{R}^n ($E_\delta^p \subseteq \mathcal{F}$ and \mathcal{F} is bounded), it suffices to show that E_δ^p is closed. In turn, by part 3 of Proposition 3, to show that E_δ^p is closed it suffices to show that its closure, denoted $\text{cl}(E_\delta^p)$, is pure-action self-generating. Let $v \in \text{cl}(E_\delta^p)$ and suppose $\{v^\ell\}_{\ell \in \mathbb{N}}$ is a sequence of pure-strategy SPNE payoffs, converging to v , with each v^ℓ decomposable via the action profile a^ℓ and continuation promises γ^ℓ . Since every A_i is finite, $\{(a^\ell, v^\ell)\}_{\ell \in \mathbb{N}}$ lies in the compact set $A \times \text{cl}(E_\delta^p)^A$, where $\text{cl}(E_\delta^p)^A$ denotes the set of all maps from A to $\text{cl}(E_\delta^p)$. Thus, there is a subsequence converging to $(a^\infty, \gamma^\infty)$, with $a^\infty \in A$ and $\gamma^\infty(a) \in \text{cl}(E_\delta^p)$ for all $a \in A$. Moreover, it is immediate that $(a^\infty, \gamma^\infty)$ decomposes v on $\text{cl}(E_\delta^p)$. ■

Exercise 14. Consider an infinitely repeated version of the stage game G for which the reward matrix is given by

| | | Player 2 | | |
|----------|-----|----------|-------|--------|
| | | U | M | D |
| Player 1 | L | 1, 1 | 3, 0 | −2, 0 |
| | M | 0, 3 | 2, 2 | −2, 0 |
| | R | 0, −2 | 0, −2 | −4, −4 |

⁹We say that a set of payoffs can be *factorized* if it is a fixed point of B_δ^p . Hence the name of Proposition ??.

1. Suppose $\delta \geq 1/2$. Show that the set $\{(1, 1), (2, 2)\}$ is self-generating.
2. Suppose $\delta = 1/3$. Find a self-generating set that contains the set $\{(2, 2)\}$.
3. Show that there is no self-generating set that contains the set $\{(-3, 3)\}$.

4.2 Simple Strategies and Penal Codes

By part 2 of Proposition 3, any action profile appearing on a pure-strategy SPNE path is decomposable on the set E_δ^p of pure-strategy SPNE payoffs. By compactness of E_δ^p , there is a collection of pure-strategy SPNE profiles $\{\sigma^1, \dots, \sigma^n\}$, with σ^i yielding the lowest possible pure-strategy SPNE payoff for player i . In this section, we will show that out-of-equilibrium behavior in any pure-strategy SPNE can be decomposed on the set $\{v(\sigma^1), \dots, v(\sigma^n)\}$. This in turn leads to a simple recipe for constructing equilibria.

Definition 12 (Simple Strategies). *Given $(n + 1)$ outcomes $\{\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n)\}$, the associated simple strategy profile $\sigma(\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n))$ consists of a prescribed outcome $\mathbf{a}(0)$ and a “punishment” outcome $\mathbf{a}(i)$ for every player i . Under the profile, play continues according to the outcome $\mathbf{a}(0)$. Players respond to any deviation by player i with a switch to the player i punishment outcome path $\mathbf{a}(i)$. If player i deviates from the path $\mathbf{a}(i)$, then $\mathbf{a}(i)$ starts again from the beginning. If some other player j deviates, then a switch is made to the player j punishment outcome $\mathbf{a}(j)$.*

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 perfect-monitoring repeated games are simple.

We can use the one-shot deviation principle to identify necessary and sufficient conditions for a simple strategy profile to be a SPNE. Let

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

be the payoff to player i from the outcome path (a^t, a^{t+1}, \dots) .

Lemma 1. *The simple strategy profile $\sigma(\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n))$ is a SPNE of G^δ if and only if*

$$v_i^t(\mathbf{a}(j)) \geq \max_{a_i \in A_i \setminus \{a_i^t(j)\}} (1 - \delta) u_i(a_i, a_{-i}^t(j)) + \delta v_i^0(\mathbf{a}(i)) \quad (10)$$

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

Exercise 15. Prove Lemma 1.

A simple strategy profile specifies an equilibrium path $\mathbf{a}(0)$ and a *penal code* $\{\mathbf{a}(1), \dots, \mathbf{a}(n)\}$ 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 13 (Optimal Penal Codes). *Let $\{\mathbf{a}(i)\}_{i \in N}$ be n outcome paths satisfying*

$$v_i^0(\mathbf{a}(i)) = v_i^{min} \tag{11}$$

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

$$\sigma(i) := \sigma(\mathbf{a}(i), \mathbf{a}(1), \dots, \mathbf{a}(n))$$

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 outcome paths $\mathbf{a}(i)$ satisfying (11). 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 subgame-perfect Nash equilibria: simple strategies suffice to achieve any feasible SPNE payoff.

Proposition 5. *1. Let $\{\mathbf{a}(i)\}_{i \in N}$ be n outcome paths of pure-strategy subgame-perfect Nash equilibria $\{\hat{\sigma}(i)\}_{i \in N}$ satisfying $v_i(\hat{\sigma}(i)) = v_i^{min}$ for all $i \in N$. The simple strategy profile $\sigma(i) := \sigma(\mathbf{a}(i), \mathbf{a}(1), \dots, \mathbf{a}(n))$ 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 path $\mathbf{a}(0)$ can be supported as an outcome of a pure-strategy SPNE if and only if there exist pure outcome paths $\{\mathbf{a}(1), \dots, \mathbf{a}(n)\}$ such that the simple strategy profile $\sigma(\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n))$ is a SPNE.

Hence, anything that can be accomplished with a SPNE in terms of payoffs can be accomplished with simple strategies. As a result, we need never consider complex hierarchies of punishments when constructing subgame-perfect Nash equilibria, 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 $\mathbf{a}(0)$ be the outcome of a SPNE. Let $(\mathbf{a}(1), \dots, \mathbf{a}(n))$ be outcomes of SPNE $(\hat{\sigma}(1), \dots, \hat{\sigma}(n))$, with $v_i(\hat{\sigma}(i)) = v_i^{min}$. Now consider the simple strategy profile given by $\sigma(\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n))$. We claim that this strategy profile constitutes a SPNE. Considering arbitrary $\mathbf{a}(0)$, this argument establishes statement 2. For $\mathbf{a}(0) \in \{\mathbf{a}(1), \dots, \mathbf{a}(n)\}$, 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

$$v_i^t(\mathbf{a}(j)) \geq (1 - \delta)u_i(a_i, a_{-i}^t(j)) + \delta v_i^0(\mathbf{a}(i)). \quad (12)$$

Now, by construction, $\mathbf{a}(j)$ is the outcome of a SPNE—the outcome $\mathbf{a}(0)$ is by assumption produced by a SPNE, whereas each of $\mathbf{a}(1), \dots, \mathbf{a}(n)$ is part of an optimal penal code. This ensures that for any t and $a_i \in A_i \setminus \{a_i^t(j)\}$,

$$v_i^t(\mathbf{a}(j)) \geq (1 - \delta)u_i(a_i, a_{-i}^t(j)) + \delta v_i^d(\mathbf{a}(i), t, a_i), \quad (13)$$

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

$$v_i^d(\mathbf{a}(i), t, a_i) \geq v_i^0(\mathbf{a}(i)) = v_i^{\min},$$

which with (13) implies (12), 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 $\mathbf{a}(0)$ is the outcome path of some SPNE. Then, the simple strategy $\sigma(\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(n))$, where each $\mathbf{a}(i)$ yields the lowest possible SPNE payoff v_i^{\min} to player i , is a SPNE.*

For a nice illustration of the use of simple strategies and penal codes, see Section 2.6.2 in [Mailath and Samuelson \(2006\)](#) (based, in turn, on [Abreu \(1986\)](#)).

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 (1986), “Optimal Cartel Equilibria with Imperfect Monitoring.” *Journal of Economic Theory*, 39, 251–269.
- (1990), “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 58, 1041–1063.
- Ali, S. Nageeb (2011), “Notes for Economics 200c.” *Lecture Notes*.

- Blackwell, David (1965), “Discounted Dynamic Programming.” *The Annals of Mathematical Statistics*, 36, 226–235.
- Friedman, James W. (1971), “A Non-cooperative Equilibrium for Supergames.” *The Review of Economic Studies*, 38, 1–12.
- 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*.
- Hörner, Johannes (2015), “Repeated Games, Part I: Perfect Monitoring.” *Lecture Notes*.
- Mailath, George J. and Larry Samuelson (2006), “Repeated Games and Reputations: Long-Run Relationships.” *Oxford University Press*.
- Mertens, Jean-François and T. Parthasarathy (2003), “Equilibria for Discounted Stochastic Games.” Sylvain Sorin and Abraham Neyman eds., *Stochastic Games and Applications. NATO Science Series (Series C: Mathematical and Physical Sciences)*. Springer, Dordrecht.
- Nash, John Forbes (1951), “Non-Cooperative Games.” *Annals of Mathematics*, 54, 286–295.
- Shapley, Lloyd Stowell (1953), “Stochastic Games.” *Proceedings of National Academy of Science*, 39, 1095–1100.
- 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.