

Repeated Games

Olivier Gossner and Tristan Tomala

December 6, 2007

1 The subject and its importance

Repeated interactions arise in several domains such as Economics, Computer Science, and Biology.

The theory of repeated games models situations in which a group of agents engage in a strategic interaction over and over. The data of the strategic interaction is fixed over time and is known by all the players. This is in contrast with stochastic games, for which the data of the strategic interaction is controlled by player's choices, and repeated games with incomplete information, where the stage game is not common knowledge among players¹. Early studies of repeated games include Luce and Raiffa [LR57] and Aumann [Aum60]. In the context of production games, Friedman [Fri71] shows that, while the competitive outcome is the only one compatible with individual profit maximization under a static interaction, collusion is sustainable at an equilibrium when the interaction is repeated.

Generally, repeated games provide a framework in which individual utility maximization by selfish agents is compatible with welfare maximization (common good), while this is known to fail for many classes of static interactions.

1.1 Motivating Example

The discussion of an example shows the importance of repeated games and introduces the questions studied.

Consider the following game referred to as the Prisoner's Dilemma:

¹The reader is referred to the corresponding chapters of this Encyclopedia.

	<i>C</i>	<i>D</i>
<i>C</i>	3, 3	-1, 4
<i>D</i>	4, -1	0, 0

The prisoner's dilemma

Player 1 chooses the row, player 2 chooses the column, and the pair of numbers in the corresponding cell are the payoffs to players 1 and 2 respectively.

In a one-shot interaction, the only outcome consistent with game theory predictions is (D, D) . In fact, each player is better off playing D whatever the other player does.

On the other hand, if players engage in a repeated Prisoner's Dilemma, if they value sufficiently future payoffs compared to present ones, and if past actions are observable, then (C, C) is a sustainable outcome. Indeed, if each player plays C as long as the other one has always done so in the past, and plays D otherwise, both player have an incentive to always play C , since the short term gain that can be obtained by playing D is more than offset by the future losses entailed by the opponent playing D at all future stages.

Hence, a game theoretical analysis predicts significantly different outcomes from a repeated game than from static interaction. In particular, in the Prisoner's Dilemma, the cooperative outcome (C, C) can be sustained in the repeated game, while only the non-cooperative outcome (D, D) can be sustained in one-shot interactions.

In general, what are the equilibrium payoffs of a repeated game and how can they be computed from the data of the static game? Is there a significant difference between games repeated a finite number of times and infinitely repeated ones? What is the role played by the degree of impatience of players? Do the conclusions obtained for the Prisoner's Dilemma game and for other games rely crucially on the assumption that each player perfectly observes other player's past choices, or would imperfect observation be sufficient? The theory of repeated games aims at answering these questions, and many more.

2 Games with observable actions

This section focuses on repeated games with perfect monitoring in which, after each period of the repeated game, all strategic choices of all the players are publicly revealed.

2.1 Data of the game, strategies, payoffs

2.1.1 Data of the stage game

There is a finite set I of players. A stage game is repeated over and over. Each player i 's action set in this stage game is denoted A_i , and $S_i = \Delta(A_i)$ is the set of player i 's mixed actions². Each degenerate lottery in S_i (which puts probability 1 to one particular action in A_i) is associated to the corresponding element in A_i . A choice of action for each player i determines an outcome $a \in \prod_i A_i$. The payoff function of the stage game is $g: A \rightarrow \mathbb{R}^I$. Payoffs are naturally associated to profiles of mixed actions $s \in S = \prod_i S_i$ using the expectation: $g(s) = \mathbf{E}_s g(a)$.

2.1.2 Repeated Game

After each repetition of the stage game, the action profile previously chosen by the players is publicly revealed. After the t first repetitions of the game, player's information consists of the publicly known history at stage t , which is an element of $H_t = A^t$ (by convention, we set $H_0 = \{\emptyset\}$). A strategy in the repeated game specifies the choice of a mixed action at every stage, as a function of the past observed history. More specifically, a *behavioral* strategy for player i is of the form $\sigma_i: \cup_t H_t \rightarrow S_i$. When all the strategy choices belong to A_i ($\sigma_i: \cup_t H_t \rightarrow A_i$), σ_i is called a *pure* strategy.

Other strategy specifications A behavioral strategy allows the player to randomize each action depending on past history. If, at the start of the repeated game, the player was to randomize over the set of behavioral strategies, the result would be equivalent to a particular behavioral strategy choice. This result is a consequence of Kuhn's theorem ([Kuh53], [Aum64]). Furthermore, behavioral strategies are also equivalent to randomizations over the set of pure strategies.

Induced plays Every choice of pure strategies $\sigma = (\sigma_i)_i$ by all the players induces a play $h = (a_1, a_2, \dots) \in A^\infty$ in the repeated game, defined inductively by $a_1 = (\sigma_{i,0}(\emptyset))$ and $a_t = (\sigma_{i,t-1}(a_1, \dots, a_{t-1}))$. A profile of behavioral strategies σ defines a probability distribution P_σ over plays.

Preferences To complete the definition of the repeated game, it remains to define player's preferences over plays. The literature commonly distinguishes infinitely repeated with or without discounting, and finitely repeated games.

²for any finite set X , $\Delta(X)$ denotes the set of probabilities over X .

In infinitely repeated games with no discounting, the players care about their long-run stream of stage payoffs. In particular, the payoff in the repeated game associated to a play $h = (a_1, a_2, \dots) \in A^\infty$ coincides with the limit of the Cesaro means of stage payoffs when this limit exists. When this limit does not exist, the most common evaluation of the stream of payoffs is defined through a Banach limit of the Cesaro means (a Banach limit is a linear form on the set of bounded sequences that lies always between the liminf and the limsup).

In infinitely repeated games with discounting, a discount factor $0 < \delta < 1$ characterizes the player's degree of impatience. A payoff of 1 at stage t is equivalent to a payoff of δ at stage $t + 1$. Player i 's payoff in the repeated game for the play $h = (a_1, a_2, \dots) \in A^\infty$ is the normalized sum of discounted payoffs: $(1 - \delta) \sum_{t \geq 1} \delta^{t-1} g_i(a_t)$.

In finitely repeated games, the game ends after some stage T . Payoffs induced by the play after stage T are irrelevant (and a strategy needs not specify choices after stage T). The payoff for a player is the average of the stage payoffs during stages 1 up to T .

Equilibrium notions What plays can be expected to be observed in repeated interactions of players who observe each other's choices? Non-cooperative Game Theory focuses mainly on the idea of stable convention, i.e. of strategy profiles from which no player has incentives to deviate, knowing the strategies adopted by the other players.

A strategy profile forms a Nash Equilibrium (Nash [Nas51]) when no player can improve his payoff by choosing an alternative strategy, as long as other players follow the prescribed strategies.

In some cases, the observation of past play may not be consistent with the prescribed strategies. When, for every possible history, each player's strategy maximizes the continuation stream of payoffs, assuming that other players abide with their prescribed strategies at all future stages, the strategy profile forms a subgame perfect equilibrium (Selten [Sel65]).

Perfect equilibrium is a more robust and often considered a more satisfactory solution concept than Nash equilibrium. The construction of perfect equilibria is in general also more demanding than the construction of Nash equilibria.

The main objective of the theory of repeated games is to characterize the set of payoff vectors that can be sustained by some Nash or perfect equilibrium of the repeated game.

2.2 Necessary conditions on equilibrium payoffs

Some properties are common to all equilibrium payoffs. First, under the common assumption that all players evaluate the payoff associated to a play in the same way³, the resulting payoff vector in the repeated game is a convex combination of stage payoffs⁴. That is, the payoff vector in the repeated game is an element of the convex closure of $g(A)$, called the set of *feasible payoffs* and denoted F .

Now consider a strategy profile σ , and let τ_i be a strategy of player i that plays after each history (a_1, \dots, a_t) a best response to the profile of mixed actions chosen by the other players in the next stage. At any stage of the repeated game, the expected payoff for player i using τ_i is no less than⁵

$$v_i = \min_{s_{-i} \in S_{-i}} \max_{a_i \in A_i} g_i(s_{-i}, a_i) \quad (1)$$

The payoff v_i is referred to as player i 's min max payoff. A payoff vector that provides each player i with at least [resp. strictly more than] v_i is called *individually rational* [resp. strictly individually rational], and IR [resp. IR^*] denotes the set of such payoff vectors. Since for any strategy profile, there exists a strategy of player i that yields a payoff of no less than v_i , all equilibrium payoffs have to be individually rational.

Also note that players $j \neq i$ collectively have a strategy profile in the repeated game that forces player i 's payoff down to v_i : they play repeatedly a mixed strategy profile that achieves the minimum in the definition of v_i . Such a strategy profile in the one-shot game is referred to as punishing strategy, or min max strategy against player i .

For the Prisoner's Dilemma game, F is the convex hull of $(0, 0)$, $(4, -1)$, $(-1, 4)$ and $(3, 3)$. Both player's min max levels are equal to 0, so that IR is the positive orthant. Figure 1 illustrates the set of feasible and individually rational payoff vectors (hatched area):

The set of feasible and individually rational payoffs can be easily computed from the stage game data.

³A notable exception is the work of Lehrer and Pauzner [LP99] who study repeated games where players has heterogenous impatience levels.

⁴The payoff vector resulting from a play does not necessarily belong to F if players have different evaluations of payoff streams. For instance, in a repetition of the Prisoner's Dilemma, if player 1 cares only about the payoff in stage 1 and player 2 cares only about the payoff in stage 2, it is possible for both players to obtain a payoff of 4 in the repeated game.

⁵If $(E_i)_{i \in I}$ is a collection of sets, e_{-i} denotes an element of $E_{-i} = \prod_{j \neq i} E_j$. A profile $e \in \prod_j E_j$ is denoted $e = (e_i, e_{-i})$ when the i -th component is stressed.




Figure 1: F and IR for the Prisoner's Dilemma

2.3 Infinitely patient players

The following result has been part of the folklore of Game Theory at least since the mid 1960's. Its authorship is obscure (see the introduction of Aumann [Aum81]). For this reason, it is commonly referred to as the "Folk Theorem". By extension, characterization of sets of equilibrium payoffs in repeated games are also referred to as "Folk Theorems".

Theorem 1 *The set of equilibrium payoffs of the repeated game with no discounting coincides with the set of feasible and individually rational payoffs.*

Aumann and Shapley [AS76], [AS94] and Rubinstein [Rub77] show that restricting attention to perfect equilibria does not narrow down the set of equilibrium payoffs. They prove that:

Theorem 2 *The set of perfect equilibrium payoffs of the repeated game with no discounting coincides with the set of feasible and individually rational payoffs.*

We outline a proof of Theorem 2. It is established that any equilibrium payoff is in $F \cap IR$, we need only to prove that each element of $F \cap IR$ is a subgame perfect equilibrium payoff. Let $x \in F \cap IR$, and let $h = a_1, \dots, a_t, \dots$ be a play inducing x . Consider the strategies that play a_t in stage t ; if player i does not respect this prescription at stage t_0 , the other players punish player i for t_0 stages by repeatedly playing the min max strategy profile against player i . After the punishment phase is over, players revert to the play of h , hence playing a_{t_0+1}, \dots

Now we explain why these strategies form a subgame perfect equilibrium. After any history, consider any strategy of player i . The induced play by this strategy for player i and by other player's prescribed strategies is, up to a subset of stages of null density, defined by the sequence h with interweaved periods of punishment for player i . Hence the induced long-run payoff for player i is a convex combination of his punishment payoff and of the payoff induced by h . The result follows since the payoff for player i induced by h is no worse than the punishment payoff.

2.4 Impatient players

The strategies constructed in the proof of the Folk Theorem for repeated games with infinitely patient players (Theorem 1) do not necessarily constitute a subgame perfect equilibrium if players are impatient. Indeed, during a punishment phase, the punishing players may be receiving bad stage payoffs, and these stage payoffs matter in the evaluation of their stream of payoffs. When constructing subgame perfect equilibria of discounted games, one must make sure that after a deviation of player i , players $j \neq i$ have incentives to implement player i 's punishment.

Nash Reversion Friedman [Fri71] shows that every feasible payoff that Pareto dominates a Nash equilibrium payoff of the static game is a subgame perfect equilibrium payoff of the repeated game provided that players are patient enough. In Friedman's proof, punishments take the simple form of reversion to the repeated play of the static Nash equilibrium forever. In the Prisoner's Dilemma, (D, D) is the only static Nash equilibrium payoff, and thus $(3, 3)$ is a subgame perfect Nash equilibrium payoff of the repeated game if players are patient enough. Note however that in some games, the set of payoffs that Pareto dominate some equilibrium payoff may be empty. Also, Friedman's result constitutes a partial Folk Theorem only in that it does not characterize the full set of equilibrium payoffs.

The recursive structure Repeated games with discounting possess a structure similar to dynamic programming models. At any stage in time, players choose actions that maximize the sum of the current payoff and the payoff at the subsequent stages. When strategies form a subgame perfect equilibrium, the payoff vector at subsequent stages must be an equilibrium payoff, and players must have incentives to follow the prescribed strategies at the current stage. This implies that subgame perfect equilibrium payoffs have a recursive structure, first studied by Abreu [Abr88]. Section 3.3.1

presents the recursive structure in more details for the more general model of games with public monitoring.

The Folk Theorem for Discounted Games Relying on Abreu’s recursive results, Fudenberg and Maskin [FM86] prove the following Folk Theorem for subgame perfect equilibria with discounting:

Theorem 3 *If the number of players is 2 or if the set feasible payoff vectors has non-empty interior, then any payoff vector that is feasible and strictly individually rational is a subgame perfect equilibrium of the discounted repeated game, provided that players are sufficiently patient.*

Forges, Mertens and Neyman [FMN86] provide an example for which a payoff which is individually rational but not strictly individually rational is not an equilibrium payoff of the discounted game.

Abreu, Dutta and Smith [ADS94] show that the non-empty interior condition of the theorem can be replaced by a weaker condition of “non equivalent utilities”: no pair of players have same preferences over outcomes. Wen [Wen94] shows that a Folk Theorem still holds when the condition of non equivalent utilities fails if one replaces the minmax level defining individually rational payoffs by some “effective minmax” payoffs.

2.5 Finitely repeated games

Strikingly, equilibrium payoffs in finitely repeated games and in infinitely repeated games can be drastically different. This effect can be best exemplified in repetitions of the Prisoner’s Dilemma.

The Prisoner’s Dilemma Recall that in an infinitely repeated Prisoner’s Dilemma, cooperation at all stages is achieved at a subgame perfect equilibrium if players are patient enough. By contrast, at every Nash equilibrium of any finite repetition of the Prisoner’s Dilemma, both players play D at every stage with probability 1.

Now we present a short proof of this result. Consider any Nash equilibrium of the Prisoner’s Dilemma repeated T times. Let a_1, \dots, a_T be a sequence of action profiles played with positive probability at the Nash equilibrium. Since each player can play D at the last stage of the repetition, and D is a dominating action, $a_T = (D, D)$. We now prove by induction on τ that for any such τ , $a_{T-\tau}, \dots, a_T = (D, D), \dots, (D, D)$. Assume the induction hypothesis valid for $\tau - 1$. Consider a strategy for player i that follows the equilibrium strategy up to stage $T - \tau - 1$, then plays D from stage $T - \tau$

on. This strategy obtains the same payoff as the equilibrium strategy at stages $1, \dots, T - \tau - 1$, and at least as much as the equilibrium strategy at stages $T - \tau + 1, \dots, T - \tau$. Hence, this strategy cannot obtain more than the equilibrium strategy at stage $T - \tau$, and therefore the equilibrium strategy plays D at stage $T - \tau$ with probability 1 as well.

Sorin [Sor86] proves the more general result:

Theorem 4 *Assume that in every Nash equilibrium of G , all players are receiving their individually rational levels. Then, at each Nash equilibrium of any finitely repeated version of G , all players are receiving their individually rational levels.*

The proof of Theorem 4 relies on a backwards induction type of argument, but it is striking that the result applies for *all* Nash equilibria and not only for *subgame perfect* Nash equilibria. This result shows that, unless some additional assumptions are made on the one-shot game, a Folk Theorem cannot obtain for finitely repeated games.

Games with unique Nash payoff Using a backwards induction argument, Benoît and Krishna [BK85] obtain the following result.

Theorem 5 *Assume that G admits x as unique Nash equilibrium payoff. Then every finite repetition of G admits x as unique subgame perfect equilibrium payoff.*

Theorems 4 and 5 rely on the assumption that the last stage of repetition, T , is common knowledge between players. Neyman [Ney99] shows that a Folk Theorem obtains for the finitely repeated Prisoner's Dilemma (and for other games) if there is lack of common knowledge on the last stage of repetition.

Folk Theorems for finitely repeated games A Folk Theorem can be obtained when there are two Nash equilibrium payoffs for each player. The following result is due to Benoît and Krishna [BK85] and Gossner [Gos95].

Theorem 6 *Assume that each player has two distinct Nash equilibrium payoffs in G and that the set of feasible payoffs has non-empty interior. Then, the set of subgame perfect equilibrium payoffs of the T times repetition of G converges to the set of feasible and individually rational payoffs as T goes to infinity.*

Hence, with at least two equilibrium payoffs per player, the sets of equilibrium payoffs of finitely repeated games and infinitely repeated games are asymptotically the same.

The condition that each player has two distinct Nash equilibrium payoffs in the stage game can be weakened, see Smith [Smi95]. Assume for simplicity that *one* player has two distinct Nash payoffs. By playing one of the two Nash equilibria in the last stages of the repeated game, it is possible to provide incentives for this player to play actions that are not part of Nash equilibria of the one-shot game in previous stages. If this construction leads to perfect equilibria in which a player $j \neq i$ has distinct payoffs, we can now provide incentives for both players i and j . If successive iterations of this procedure yield distinct subgame perfect equilibrium payoffs for all players, a Folk Theorem applies.

3 Games with non observable actions

For infinitely repeated games with perfect monitoring, a complete and simple characterization of the set of equilibrium payoffs is obtained: feasible and individually rational payoff vectors. In particular, cooperation can be sustained at equilibrium. How equilibrium payoffs of the repeated game depend on the quality of player's monitoring of each other's actions is the subject of a very active area of research.

Repeated games with imperfect monitoring, in which players observe imperfectly other player's action choices, were first motivated by economic applications. In Stigler [Sti64], two firms are repeatedly engaged in price competition over market shares. Each firm observes its own sales, but not the price set by the rival. While it is in the best interest for both firms to set a collusive price, each firm has incentives to secretly undercut the rival's price. Upon observing plunging sales, should a firm deduce that the rival firm is undercutting prices, and retaliate by setting lower prices, or should lower sales be interpreted as a result of an exogenous shock on market demand? Whether collusive behavior is sustainable or not at equilibrium is one of the motivating questions in the theory of repeated games with imperfect monitoring.

It is interesting to compare repeated games with imperfect monitoring with their perfect monitoring counterparts.

The structure of equilibria used to prove the Folk Theorem with perfect monitoring and no discounting is rather simple: if a player deviates from the prescribed strategies, the deviation is detected and the deviating player is identified, all other players can then punish the deviator. With imper-

fect monitoring, not all deviations are detectable, and when a deviation is detected, deviators are not necessarily identifiable. The notions of detection and identification allow fairly general Folk Theorems for undiscounted games. We present these results in section 3.2.

With discounting, repeated games with perfect monitoring possess a recursive structure that facilitates their study. Discounted games with public monitoring also possess a recursive structure. We review the major results of this branch of the literature in section 3.3.

Almost-perfect monitoring is the natural framework to study the effect of small departures from the perfect or public monitoring assumptions. We review this literature in section 3.4.

Little is known about general discounted games with imperfect monitoring. We present the main known results in section 3.5.

With perfect monitoring, the worst equilibrium payoff for a player is given by the min max of the one-shot game, where punishing (minimizing) players choose an independent profile of mixed strategies. With imperfect monitoring, correlation past signals for the punishing players may lead to more efficient punishments. We present results on punishment levels in 3.6.

3.1 Model

In this section we define repeated games with imperfect monitoring, and describe several classes of monitoring structures of particular interest.

3.1.1 Data of the game

Recall that the one-shot strategic interaction is described by a finite set I of players, a finite action set A_i for each player i , and a payoff function $g: A \rightarrow \mathbb{R}^I$. Player's observation of each other's actions is described by a *monitoring structure* given by a finite set of signals Y_i for each player i and by a transition probability $Q: A \rightarrow \Delta(Y)$ (with $A = \prod_{i \in I} A_i$ and $Y = \prod_{i \in I} Y_i$). When the action profile chosen is $a = (a_i)_{i \in I}$, a profile of signals $y = (y_i)_{i \in I}$ is drawn with probability $Q(y|a)$ and y_i is observed by player i .

Perfect monitoring Perfect monitoring is the particular case in which each player observes the action profile chosen: for each player i , $Y_i = A$ and $Q((y_i)_{i \in I} | a) = \mathbb{1}_{\{\forall i, y_i = a\}}$.

Almost perfect monitoring The monitoring structure is ε -perfect (see Mailath and Morris [MM02]) when each player can identify the other player's

action with a probability of error less than ε . This is the case if there exist functions $f_i: A_i \times Y_i \rightarrow A_{-i}$ for all i such that, for all $a \in A$:

$$Q(\forall i, f_i(a_i, y_i) = a_{-i} | a) \geq 1 - \varepsilon$$

Almost-perfect monitoring refers to ε -perfect monitoring for small values of ε .

Canonical structure The monitoring structure is *canonical* when each player's observation corresponds to an action profile of the opponents, that is, when $Y_i = A_{-i}$.

Public and almost public signals Signals are *public* when all the players observe the same signal, i.e., $Q(\forall i, j y_i = y_j | a) = 1$, for each action profile a . For instance, in Green and Porter [GP84], firms compete over quantities, and the public signal is the realization of the price. Firms can then make inferences on rival's quantities based on their own quantity and market price.

The case in which $Q(\forall i, j y_i = y_j | a)$ is close to 1 for every a is referred to as *almost public monitoring* (see Mailath and Morris [MM02]).

Deterministic signals Signals are *deterministic* when the signal profile is uniquely determined by the action profile. When a is played, the signal profile y is given by $y = f(a)$, where f is called the *signalling function*.

Observable payoffs Payoffs are *observable* when each player i can deduce his payoff from his action and his signal. This is the case if there exists a mappings $\varphi: A_i \times Y_i \rightarrow \mathbb{R}$ such that for each action profile a , $Q(\forall i g_i(a) = \varphi(a_i, y_i) | a) = 1$.

3.1.2 The Repeated Game

The game is played repeatedly and after each stage t , the profile of signals y_t received by the players is drawn according to the distribution $Q(y_t | a_t)$, where a_t is the profile of action chosen at stage t . A player's information consists of his past actions and signals. We let $H_{i,t} = (A_i \times Y_i)^t$ be the set of *player i 's histories* of length t . A strategy for player i now consists of a mapping $\sigma_i: \cup_{t \geq 0} H_{i,t} \rightarrow S_i$. The set of complete histories of the game after t stages is $H_t = (A \times Y)^t$, it describes chosen actions and received signals for all the players at all past stages. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ induces a probability distribution P_σ on the set of plays $H_\infty = (A \times Y)^\infty$.

3.1.3 Equilibrium notions

Nash equilibria Player's preferences over game plays are defined according to the same criteria as for perfect monitoring. We focus on infinitely repeated games, both discounted and undiscounted. A choice of players' preferences defines a set of Nash equilibrium payoffs in the repeated game.

Sequential equilibria The most commonly used refinement of Nash equilibrium for repeated games with imperfect monitoring is the *sequential equilibrium* concept (Kreps and Wilson, [KW82]), which we recall here.

A *belief assessment* is a sequence $\mu = (\mu_{i,t})_{t \geq 1, i \in I}$ with $\mu_{i,t} : H_{i,t} \rightarrow \Delta(H_t)$, i.e., given the private history h_i of player i , $\mu_{i,t}(h_i)$ is the probability distribution representing the belief that player i holds on the full history.

A sequential equilibrium of the repeated game is a pair (σ, μ) where σ is a strategy profile and μ is a belief assessment such that: 1) for each player i and each history h_i , σ_i is a best reply in the continuation game, given the strategies of the other players and the belief that player i holds regarding the past; 2) the beliefs must be consistent in the sense that (σ, μ) is the pointwise limit of a sequence (σ^n, μ^n) where for each n , σ^n is a completely mixed strategy (it assigns positive probability to every action after every history) and μ^n is the unique belief derived from Bayes' law.

Sequential equilibria are defined both on the discounted game and the undiscounted versions of the repeated game.

For undiscounted games, the set of Nash equilibrium payoffs and sequential equilibrium payoffs coincide. The two notions also coincide for discounted games when the monitoring has full support (i.e. under every action profile, all signal profiles have positive probability). The results presented in this survey all hold for sequential equilibria, both for discounted and undiscounted games.

3.1.4 Extensions of the repeated game

When player receive correlated inputs or may communicate between stages of the repeated game, the relevant concepts are correlated and communication equilibria.

Correlated equilibria A correlated equilibrium (Aumann [Aum74]) of the repeated game is an equilibrium of an extended game in which: at a preliminary stage, a mediator chooses a profile of correlated random inputs and informs each player of his own input; then the repeated game is played.

A characterization of the set of correlated equilibrium payoffs for two-player games is obtained by Lehrer [Leh92a].

Correlation arises endogenously in repeated games with imperfect monitoring, as the signals received by the players can serve as correlated inputs that influence player’s continuation strategies. This phenomenon is called *internal correlation*, and was studied by Lehrer, [Leh91], Gossner and Tomala, [GT06], [GT07].

Communication equilibria An (extensive form) communication equilibrium (Myerson [Mye82], Forges [For86]) of a repeated game is an equilibrium of an extension of the repeated game in which after every stage, players send messages to a mediator, and the mediator sends back private outputs to the players. Characterizations of the set of communication equilibrium payoffs are obtained under weak conditions on the monitoring structure, see e.g. Kandori and Matsushima [KM98], Compte [Com98], and Renault and Tomala [RT04].

3.2 Detection and identification

3.2.1 Equivalent actions

A deviation from a player is detectable when it induces a different distribution of signals for other players. When two actions induce the same distribution of actions for other players, they are called equivalent (Lehrer [Leh90], [Leh91], [Leh92a], [Leh92b]):

Definition 1 *Two actions a_i and b_i of player i are equivalent, and we note $a_i \sim b_i$, if they induce the same distribution of other players’ signals:*

$$Q(y_{-i}|a_i, a_{-i}) = Q(y_{-i}|b_i, a_{-i}), \forall a_{-i}$$

Example 1

Consider the two-player repeated Prisoner’s Dilemma where player 2 receives no information about the actions of player 1 (e.g. Y_2 is a singleton). The two actions of player 1 are thus equivalent. The actions of player 2 are independent of the actions of player 1: player 1 has no impact on the behavior of player 2. Player 2 has no power to threaten player 1 and in any equilibrium, player 1 defects at every stage. Player 2 also defects at each stage: since player 1 always defects, he also loses his threatening power. The only equilibrium payoff in this repeated game is thus $(0, 0)$.

Example 1 suggests that between two equivalent actions, a player chooses at equilibrium the one that yields the highest stage payoff. This is indeed the case when the information received by a player does not depend on his own action. Lehrer [Leh90] studies particular monitoring structures satisfying this requirement. Recall from Lehrer [Leh90] the definition of *semi-standard* monitoring structures: each action set A_i is endowed with a partition \bar{A}_i , when player i plays a_i , the corresponding partition cell \bar{a}_i is publicly announced. In the semi-standard case, two actions are equivalent if and only if they belong to the same cell: $a_i \sim b_i \iff \bar{a}_i = \bar{b}_i$ and the information received by a player on other player's action does not depend on his own action.

If player i deviates from a_i to b_i , the deviation is undetected if and only if $a_i \sim b_i$. Otherwise it is detected by all other players. A profile of mixed actions is called *immune to undetectable deviations* if no player can profit by a unilateral deviation that maintains the same distribution of other players' signals. The following result, due to Lehrer [Leh90], characterizes equilibrium payoffs for undiscounted games with semi-standard signals:

Theorem 7 *In a undiscounted repeated game with semi-standard signals, the equilibrium payoffs are the individually rational payoffs that belongs to the convex hull of payoffs generated by mixed action profiles that are immune to undetectable deviations.*

3.2.2 More informative actions

When the information of player i depends on his own action, some deviations may be detected in the course of the repeated game even though they are undetectable in the stage game.

Example 2

Consider the following modification of the Prisoner's dilemma. The action set of player 1 is $A_1 = \{C_1, D_1\} \times \{C_2, D_2\}$ and the action set of player 2 is $\{C_2, D_2\}$. An action for player 1 is thus a pair $a_1 = (\tilde{a}_1, \tilde{a}_2)$. When the action profile $(\tilde{a}_1, \tilde{a}_2, a_2)$ is played, the payoff to player i is $g_i(\tilde{a}_1, a_2)$. We can interpret the component \tilde{a}_1 as a *real* action (it impacts payoffs) and the component \tilde{a}_2 as a *message* sent to player 2 (it does not impact payoffs). The monitoring structure is as follows:

- player 2 only observes the message component \tilde{a}_2 of the action of player 1 and,

- player 1 perfectly observes the action of player 2 if he chooses the cooperative real action ($\tilde{a}_1 = C_1$), and gets no information on player 2's action if he defects ($\tilde{a}_1 = D_1$).

Note that the actions (C_1, C_2) and (D_1, C_2) of player 1 are equivalent, and so are the actions (C_1, D_2) and (D_1, D_2) . However, it is possible to construct an equilibrium that implements the cooperative payoff along the following lines:

i) Using his message component, player 1 reports at each stage $t > 1$ the previous action of player 2. Player 1 is punished in case of a non matching report.

ii) Player 2 randomizes between both actions, so that player 1 needs to play the cooperative action in order to report player 2's action accurately. The weight on the defective action of player 2 goes to 0 as t goes to infinity to ensure efficiency.

Player 2 has incentives to play C_2 most of the time, since player 1 can statistically detect if player 2 uses the action D_2 more frequently than prescribed. Player 1 also has incentives to play the real action C_1 , as this is the only way to observe player 2's action, which need to be reported later on.

The key point in the example above is that the two real actions C_1 and D_1 of player 1 are equivalent but D_1 is *less informative* than C_1 for player 1. For general monitoring structures an action a_i is *more informative* than an action b_i if: *whenever player 1 plays a_i , can reconstitue the signal he would have observed, had he played b_i .* The precise definition of the *more informative* relation relies Blackwell's ordering of stochastic experiments [Bla51]:

Definition 2 *The action a_i of player i is more informative than the action b_i if there exists a transition probability $f : Y_i \rightarrow \Delta(Y_i)$ such that for each a_{-i} and each profile of signals y ,*

$$\sum_{y_i} f(y'_i|y_i)Q(y_i, y_{-i}|a_i, a_{-i}) = Q(y'_i, y_{-i}|b_i, a_{-i})$$

We denote $a_i \succeq b_i$ if $a_i \sim b_i$ and a_i is more informative than b_i .

Assume that prescribed strategies require player i to play b_i at stage t , and let $a_i \succeq b_i$. Consider the following deviation from player i : play a_i at stage t , and reconstruct a signal at stage t that could have arisen from the play of b_i . In all subsequent stages, play as if no deviation took place at stage t , and as if the reconstructed signal had been observed at stage t . Not only such a deviation would be undetectable at stage t , since $a_i \sim b_i$, but it would also be undetectable at all subsequent stages, as it would induce the same probability distribution over plays as under the prescribed strategy. This

argument shows that, if an equilibrium strategy specifies that player i plays a_i , there is no $b_i \succeq a_i$ that yields a higher expected stage payoff than a_i .

Definition 3 *A distribution of actions profiles $p \in \Delta(A)$ is immune to undetectable deviations if for each player i and pair of actions a_i, b_i such that $b_i \succeq a_i$:*

$$\sum_{a_{-i}} p(a_i, a_{-i}) g_i(a_i, a_{-i}) \geq \sum_{a_{-i}} p(a_i, a_{-i}) g_i(b_i, a_{-i})$$

If p is immune to undetectable deviations, and if player i is supposed to play a_i , any alternative action b_i that yields a greater expected payoff can not be such that $b_i \succeq a_i$.

The following proposition gives a necessary condition on equilibrium payoffs that holds both in the discounted and in the undiscounted cases.

Proposition 1 *Every equilibrium payoff of the repeated game is induced by a distribution that is immune to undetectable deviations.*

The condition of Proposition 1 is tight for some specific classes of games, all of them assuming two players and no discounting.

Following Lehrer [Leh92a], signals are *non trivial* if, for each player i , there exist an action a_i for player i and two actions a_j, b_j for i 's opponent such that the signal for player i is different under (a_i, a_j) and (a_i, b_j) . Lehrer [Leh92a] proves:

Theorem 8 *The set of correlated equilibrium payoffs of the undiscounted game with deterministic and non trivial signals is the set of individually rational payoffs induced by distributions that are immune to undetectable deviations.*

Lehrer [Leh92b] assumes that payoffs are observable, and obtains the following result:

Theorem 9 *In a two-player repeated game with no discounting, non-trivial signals and observable payoffs, the set of equilibrium payoffs is the set of individually rational payoffs induced by distributions that are immune to undetectable deviations.*

Finally, Lehrer [Leh91] shows that in some cases, one may dispense with the correlation device of Theorem 8, as all necessary correlation can be generated endogenously through the signals of the repeated game:

Proposition 2 *In two-player games with non-trivial signals such that either the action profile is publicly announced or a blank signal is publicly announced, the set of equilibrium payoffs coincides with the set of correlated equilibrium payoffs.*

3.2.3 Identification of deviators

A deviation is identifiable when every player can infer the identity of the deviating player from his observations. For instance, in a game with public signals, if separate deviations from players i and j induce the same distribution of public signals, these deviations from i or j are not identifiable. In order to be able to punish the deviating player, it is sometimes necessary to know his identity. Detectability and identifiability are two separate issues, as shown by the following example.

Example 3

Consider the following 3-player game where player 1 chooses the row, player 2 chooses the column and player 3 chooses the matrix.

	L	R
T	1, 1, 1	4, 4, 0
B	4, 4, 0	4, 4, 0
	W	

	L	R
0, 3, 0	0, 3, 0	
0, 3, 0	0, 3, 0	
	M	

	L	R
3, 0, 0	3, 0, 0	
3, 0, 0	3, 0, 0	
	E	

Consider the monitoring structure in which actions are not observable and the payoff vector is publicly announced.

The payoff $(1, 1, 1)$ is feasible and individually rational. The associated action profile (T, L, W) is immune to undetectable deviations since any individual deviation from (T, L, W) changes the payoff.

However, $(1, 1, 1)$ is not an equilibrium payoff. The reason is that, player 3, who has the power to punish either player 1 or player 2, cannot punish both players simultaneously: punishing player 1 rewards player 2 and vice-versa. More precisely, whatever weights player 3 puts on the action M and E , the sum of player 1 and player 2's payoffs is greater than 3. Any equilibrium payoff vector $v = (v_1, v_2, v_3)$ must thus satisfy $v_1 + v_2 \geq 3$. In fact, it is possible to prove that the set of equilibrium payoffs of this repeated game is the set of feasible and individually rational payoffs that satisfy this constraint.

Approachability When the deviating player cannot be identified, it may be necessary to punish a group of suspects altogether. The notion of payoff that is enforceable under group punishments is captured by the definition of approachable payoffs:

Definition 1 A payoff vector v is approachable if there exists a strategy profile σ such that, for every player i and unilateral deviation τ_i of player i ,

the average payoff of player i under (τ_i, σ_{-i}) is asymptotically less than or equal to v_i .

Blackwell's [Bla56] approachability theorem and its generalization by Kohlberg [Koh75] provide simple geometric characterizations of approachable payoffs. It is straightforward that approachability is a necessary condition on equilibrium payoffs:

Proposition 3 *Every equilibrium payoff of the repeated game is approachable.*

Renault and Tomala [RT04] show that the conditions of Proposition 1 and 3 are tight for communication equilibria:

Theorem 10 *For every game with imperfect monitoring, the set of communication equilibrium payoffs of the repeated game with no discounting is the set of approachable payoffs induced by distributions which are immune to undetectable deviations.*

Tomala [Tom98] shows that pure strategy equilibrium payoffs of undiscounted repeated games with public signals are also characterized through identifiability and approachability conditions (the approachability definition then uses pure strategies). Tomala [Tom99] provides a similar characterization in mixed strategies for a restricted class of public signals.

Identification through endogenous communication A deviation may be identified in the repeated game even though it cannot be identified in stage game. In a *network game*, players are located at nodes of a graph, and each player monitors his neighbors' actions. Each player can use his actions as messages that are broadcasted to all the neighbors in the graph. The graph is called *2-connected* if no single node deletion disconnects the graph. Renault and Tomala [RT98] show that when the graph is 2-connected, there exists a communication protocol among the players that ensures that the identity of any deviating player becomes common knowledge among all players in finite time. In this case, identification takes place through communication over the graph.

3.3 Public Equilibria

In a seminal paper, Green and Porter [GP84] introduce a model in which firms are engaged in a production game and publicly observe market prices,

which depend both on quantities produced and on non-observable exogenous market shocks. Can collusion be sustained at equilibrium even prices convey imperfect information on quantities produced? This motivates the study of *public equilibria* for which sharp characterizations of equilibrium payoffs are obtained.

Signals are public when all sets of signals are identical, i.e. $Y_i = Y_{pub}$ for each i and $Q(\forall i, j \ y_i = y_j | a) = 1$ for every a . A *public history* of length t is a record of t public signals, i.e. an element of $H_{pub,t} = (Y_{pub})^t$. A strategy σ_i for player i is a *public strategy* if it depends on the public history only: if $h_i = (a_{i,1}, y_1, \dots, a_{i,t}, y_t)$ and $h'_i = (a'_{i,1}, y'_1, \dots, a'_{i,t}, y'_t)$ are two histories for player i such that $y_1 = y'_1, \dots, y_t = y'_t$, then $\sigma_i(h_i) = \sigma_i(h'_i)$.

Definition 4 *A perfect public equilibrium is a profile of public strategies such that after every public history, each player's continuation strategy is a best reply to the opponents' continuation strategy profile.*

The repetition of a Nash equilibrium of the stage game is a perfect public equilibrium, so that perfect public equilibria exist. Every perfect public equilibrium is a sequential equilibrium: any consistent belief assigns probability one to the realized public history and thus correctly forecasts future opponents' choices.

3.3.1 The recursive structure

A perfect public equilibrium (PPE henceforth) is a profile of public strategies that forms an equilibrium of the repeated game and such that, after every public history, the continuation strategy profile is also a PPE. The set of PPEs and the payoffs it induces therefore possesses a recursive structure, as shown by Abreu, Pearce and Stachetti [APS90]. The argument is based on a dynamic programming principle. To state the main result, we first introduce some definitions.

Given a mapping $f : Y_{pub} \rightarrow \mathbb{R}^I$, $G(\delta, f)$ represents the one-shot game where each player i choose actions in A_i and where payoffs are given by:

$$(1 - \delta)g_i(a) + \delta \sum_{y \in Y_{pub}} Q(y|a)f_i(y)$$

In $G(\delta, f)$, the stage game is played, and players receive $f(y)$ as an additional payoff if y is the realized public signal. The weights $1 - \delta$ and δ are the relative weights of present payoffs versus all future payoffs in the repeated game.

Definition 5 *A payoff vector $v \in \mathbb{R}^I$ is decomposable with respect to the set $W \subset \mathbb{R}^I$ if there exists a mapping $f : Y_{pub} \rightarrow W$ such that v is a Nash*

equilibrium payoff of $G(\delta, f)$. $F_\delta(W)$ denotes the set of payoff vectors which are decomposable with respect to W .

Let $E(\delta)$ be the set of perfect public equilibrium payoffs of the repeated game discounted at the rate δ . The following result is due to Abreu *et al.* [APS90]:

Theorem 11 $E(\delta)$ is the largest bounded set W such that $W \subset F_\delta(W)$.

Fudenberg and Levine [FL94] derive an asymptotic characterization of the set of PPE payoffs when the discount factor goes to 1 as follows. Given a vector $\lambda \in \mathbb{R}^I$, define the *score* in the direction λ as:

$$k(\lambda) = \sup \langle \lambda, v \rangle$$

where the supremum is taken over the set of payoff vectors v that are Nash equilibrium payoffs of $G(\delta, f)$, where f is any mapping such that,

$$\langle \lambda, v \rangle \geq \langle \lambda, f(y) \rangle, \quad \forall y \in Y_{pub}$$

Scores are independent of the discount factor. The following theorem is due to Fudenberg and Levine [FL94]:

Theorem 12 Let C be the set of vectors v such that for each $\lambda \in \mathbb{R}^I$, $\langle \lambda, v \rangle \leq k(\lambda)$. If the interior of C is non-empty, $E(\delta)$ converges to C (for the Hausdorff topology) as δ goes to 1.

Fudenberg, Levine and Takahashi [FLT07] relax the non-empty interior assumption. They provide an algorithm for computing the affine hull of $\lim_{\delta \rightarrow 1} E(\delta)$ and provide a corresponding characterization of the set C with continuation payoffs belonging to this affine hull.

3.3.2 Folk Theorems for public equilibria

The recursive structure of Theorem 11 and the asymptotic characterization of PPE payoffs given by Theorem 12 are essential tools for finding sufficient conditions under which every feasible and individually rational payoff is an equilibrium payoff, i.e. conditions under which a Folk Theorem holds.

The two conditions under which a Folk Theorem in PPEs holds are a 1) a condition of detectability of deviations and 2) a condition of identifiability of deviating players.

Definition 2 A profile of mixed actions $s = (s_i, s_{-i})$ has individual full rank if for each player i , the probability vectors (in the vector space $\mathbb{R}^{Y_{pub}}$)

$$\{Q(\cdot|a_i, s_{-i}) : a_i \in A_i\}$$

are linearly independent.

If s has individual full rank, no player can change the distribution of his actions without affecting the distribution of public signals. Individual full rank is thus a condition on detectability of deviations.

Definition 3 A profile of mixed actions s has pairwise full rank if for every pair of players $i \neq j$, the family of probability vectors

$$\{Q(\cdot|a_i, s_{-i}) : a_i \in A_i\} \cup \{Q(\cdot|a_j, s_{-j}) : a_j \in A_j\}$$

has rank $|A_i| + |A_j| - 1$.

Under the condition of pairwise full rank, deviations from two distinct players induce distinct distributions of public signals. Pairwise full rank is therefore a condition of identifiability of deviating players.

Fudenberg *et al.* [FLM94] prove the following theorem:

Theorem 13 Assume the set of feasible and individually rational payoff vectors F has non-empty interior. If every pure action profile has individual full rank and if there exists a mixed action profile with pairwise full rank, then every convex and compact subset of the interior of F is a subset of $E(\delta)$ for δ large enough.

In particular, under the conditions of the theorem, every feasible and individually rational payoff vector is arbitrarily close to a PPE payoff for large enough discount factors. Variations of this result can be found in [FLM94] and [FL94].

3.3.3 Extensions

The public part of a signal The definition of perfect public equilibria extend to the case in which each player's signals consists of two components: a public component and a private component. The public components of all players' signals are the same with probability one. A public strategy is then a strategy that depends only on the public components of past signals, and all the analysis carries through.

Public communication Consider the case of general monitoring structures. In the public communication extension of the repeated game, players make public announcements between any two stages of the repeated game. The profile of public announcements then forms a public signal, and equilibrium characterizations follow from recursive analysis. Ben Porath and Kahneman [BPK96], Kandori and Matsushima [KM98], and Compte [Com98] prove Folk Theorems in games with private signals and public communication.

Private strategies in games with public equilibria PPE payoffs do not cover the full set of sequential equilibrium payoffs, even when signals are public, as some equilibria may rely on players using *private strategies*, i.e. strategies that depend on past chosen actions and past private signals. See [MMS02] and [KO06] for examples. In a *minority games*, there is an odd number of players, each player chooses between actions A and B . Players choosing the least chosen (minority) action get a payoffs of 1, other players get 0. The public signal is the minority action. Renault *et al.* [RSS05], [RSS06] show that, for minority games, a Folk Theorem holds in private strategies but not in public strategies. Only few results are known concerning the set of sequential equilibrium payoffs in private strategies of games with public monitoring, see Mailath and Samuelson [MS06] for a survey.

Almost public monitoring Some PPEs are robust to small perturbations of public signals. Considering strategies with finite memory, Mailath and Morris [MM02] identify a class of public strategies which are sequential equilibria of the repeated game with imperfect private monitoring, provided that the monitoring structure is close enough to a public one. They derive a Folk Theorem for games with almost public and almost perfect monitoring. Hörner and Olszewski [HO07] strengthen this result and prove a Folk Theorem for games with almost public monitoring. Under the detectability and identifiability conditions, they prove that feasible and individually rational payoffs can be achieved by sequential equilibria with finite memory.

3.4 Almost perfect monitoring

Monitoring is almost perfect when each player can identify the action profile of his opponents with near certainty. Almost perfect monitoring is the natural framework to study the robustness of the Folk Theorem to small departures from the assumption that actions are perfectly observed.

The first results were obtained for the Prisoner's Dilemma. Sekiguchi

[Sek97] shows that the efficient payoff can be approximated at equilibrium when players are sufficiently patient and monitoring is almost perfect. Under the same assumptions, Bhaskar and Obara [BO02], Piccione [Pic02] and Ely and Valimaki [EV02] show that a Folk Theorem obtains.

Piccione [Pic02] and Ely and Valimaki [EV02] study a particular class of equilibria called *belief free*. Strategies form a *belief free* equilibrium if, whatever player i 's belief on the opponent's private history, the action prescribed by i 's strategy is a best response to the opponent's continuation strategy.

Ely, Hörner and Olszewski [EHO05] extend the belief free approach to general games. However, they show that, in general, belief free strategies are not enough to reconstruct a Folk Theorem, even when monitoring is almost perfect.

For general games and with any number of players, Hörner and Olszewski [HO06] prove a Folk Theorem with almost perfect monitoring. The strategies that implement the equilibrium payoffs are defined on successive blocks of a fixed length, and are block-belief-free in the sense that, at the beginning of each block, each player is indifferent between several continuation strategies, independently on his belief as to which continuation strategies are used by the opponents. This result closes the almost perfect monitoring case by showing that equilibrium payoffs in the Folk Theorem are robust to a small amount of imperfect monitoring.

3.5 General Stochastic Signals

Besides the case of public (or almost public) monitoring, little is known about equilibrium payoffs of repeated games with discounting and imperfect signals.

The Prisoner's Dilemma game is particularly important for economic applications. Remarkably, it captures the features of collusion with the possibility of secret price cutting, as in Stigler [Sti64].

When signals are imperfect, but independent conditionally on the pair of actions chosen (a condition called conditional independence), Matsushima [Mat04] shows that the efficient outcome of the repeated Prisoner's Dilemma game is an equilibrium outcome if players are sufficiently patient. In the equilibrium construction, each player's action is constant in each block. The conditional independence assumption is crucial in that it implies that, during each block, a player has no feedback as to what signals the other player has received. The conditional independence assumption is non-generic: it holds for a set of monitoring structures of empty interior.

Fong, Gossner, Hörner, and Sannikov [FGHS07] prove that efficiency can be obtained at equilibrium without conditional independence. Their main assumption is that there exists a sufficiently informative signal, but this signal

needs not be almost perfectly informative. Their result holds for a family of monitoring structures of non empty interior. It is the first result that establishes cooperation in the Prisoner’s Dilemma with impatient players for truly imperfect, private and correlated signals.

3.6 Punishment levels

Individual rationality is a key concept for Folk Theorems and equilibrium payoff characterizations. Given a repeated game, define the individually rational (IR) level of player i as the lowest payoff down to which this player may be punished in the repeated game.

Definition 6 *The individual rational level of player i is:*

$$\lim_{\delta \rightarrow 1} \min_{\sigma_{-i}} \max_{\sigma_i} \mathbf{E}_{\sigma_i, \sigma_{-i}} \left[\sum_t (1 - \delta) \delta^{t-1} g_{i,t} \right]$$

where the min runs over profiles of behavior strategies for player $-i$, and the max over behavior strategies of player i .

That is, the individually rational level is the limit (as the discount factor goes to one) of the min max value of the discounted game⁶.

Comparison of the IR level with the min max With perfect monitoring, the IR level of player i is player i ’s min max in the one-shot game, as defined by equation (1). With imperfect monitoring, the IR level for player is never larger than v_i since player i ’s opponents can force down player i to v_{-i} by repeatedly playing the min max strategy against player i .

With two players, it is a consequence of von-Neumann’s min max theorem [vN28] that v_i is the IR level for player i .

For any any number of players, Gossner and Hörner [GH06] show that v_i is equal to the min max in the one-shot game whenever there exists a garbling of player i ’s signal such that, conditionally on i ’s garbled signal, the signals of i ’s opponents are independent. Furthermore, the condition in [GH06] is also a necessary condition in games with information structures. A continuity result in the IR level also applies for monitoring structure close to the ones satisfying the conditional independence condition.

The following example shows that, in general, the IR level can be lower than v_i :

⁶Other approaches, through undiscounted games or limits of finitely repeated games, yield equivalent definitions, see [GT06].

Example 4

Consider the following three-player game. Player 1 chooses the row, player 2 the column and player 3 the matrix. Players 1 and 2 perfectly observe the action profile while player 3 observes player 2's action only. As we deal with the IR level of player 3, we specify the payoff for this player only.

	L	R	
T	0	0	
B	0	-1	
	W		

	L	R	
T	-1	0	
B	0	0	
	E		

A simple computation shows that $v_3 = -\frac{1}{4}$ and that the min max strategies of players 1 and 2 are uniform. Consider the following strategies of players 1 and 2 in the repeated game: randomize uniformly at odd stages, play (T, L) or (B, R) depending on player 1's previous action at even stages. Against these strategies, player 3 cannot obtain better than $-\frac{1}{4}$ at odd stages and $-\frac{1}{2}$ at even stages, resulting in an average payoff of $-\frac{3}{8}$.

Entropy characterizations The exact computation of the IR level in games with imperfect monitoring requires to analyze the optimal trade-off for punishing players between the production of correlated and private signals and the use of these signals for effective punishment. Gossner and Vieille, [GV02] and Gossner and Tomala [GT06] develop tools based on information theory to analyze this trade-off. At any stage, the amount of correlation generated (or spent) by the punishing players is measured using the entropy function. Gossner and Tomala [GT07] derive a characterization of the IR level for some classes of monitoring structures. Gossner, Laraki, and Tomala [GLT06] provide methods explicit computations of the IR level. In particular, for the above example, the IR level computed and is about $-.401$. Explicit computations of IR levels for other games are derived by Goldberg [Gol07].

Primary literature

- [Abr88] D. Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56:383–396, 1988.
- [ADS94] D. Abreu, P. Dutta, and L. Smith. The folk theorem for repeated games: a NEU condition. *Econometrica*, 62:939–948, 1994.

- [APS90] D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58:1041–1063, 1990.
- [AS76] R. J. Aumann and L. S. Shapley. Long-term competition – a game theoretic analysis. preprint, 1976.
- [AS94] R. J. Aumann and L. S. Shapley. Long-term competition—A game theoretic analysis. In N. Megiddo, editor, *Essays on game theory*, pages 1–15. Springer-Verlag, New York, 1994.
- [Aum60] R.J. Aumann. Acceptable points in games of perfect information. *Pacific Journal of Mathematics*, 10:381–417, 1960.
- [Aum64] R.J. Aumann. Mixed and behavior strategies in infinite extensive games. In M. Dresder, L.S. Shapley, and A.W. Tucker, editors, *Advances in Game Theory*, pages 627–650, Princeton, New Jersey, 1964. Princeton University Press.
- [Aum74] R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–95, 1974.
- [BK85] J.-P. Benoît and V. Krishna. Finitely repeated games. *Econometrica*, 53(4):905–922, 1985.
- [Bla51] D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.
- [Bla56] D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956.
- [BO02] V. Bhaskar and I. Obara. Belief-based equilibria in the repeated prisoners’ dilemma with private monitoring. *Journal of Economic Theory*, 102:40–70, 2002.
- [BPK96] E. Ben Porath and M. Kahneman. Communication in repeated games with private monitoring. *Journal of Economic Theory*, 70(2):281–297, 1996.
- [Com98] O. Compte. Communication in repeated games with imperfect private monitoring. *Econometrica*, 66:597–626, 1998.
- [EHO05] J.C. Ely, J. Hörner, and W. Olszewski. Belief-free equilibria in repeated games. *Econometrica*, 73:377–415, 2005.

- [EV02] J.C. Ely and J. Välimäki. A robust folk theorem for the prisoner's dilemma. *Journal of Economic Theory*, 102:84–106, 2002.
- [FGHS07] K. Fong, O. Gossner, J. Hörner, and Y. Sannikov. Efficiency in a repeated prisoner's dilemma with imperfect private monitoring. mimeo, 2007.
- [FL94] D. Fudenberg and D. K. Levine. Efficiency and observability with long-run and short-run players. *Journal of Economic Theory*, 62:103–135, 1994.
- [FLM94] D. Fudenberg, D. K. Levine, and E. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, 1994.
- [FLT07] D. Fudenberg, D. K. Levine, and S. Takahashi. Perfect public equilibrium when players are patient. *Games and Economic Behavior*, 61:27–49, 2007.
- [FM86] D. Fudenberg and E. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54:533–554, 1986.
- [FMN86] F. Forges, J.-F. Mertens, and A. Neyman. A counterexample to the folk theorem with discounting. *Economics Letters*, 20:7–7, 1986.
- [For86] F. Forges. An approach to communication equilibria. *Econometrica*, 54:1375–1385, 1986.
- [Fri71] J. Friedman. A noncooperative equilibrium for supergames. *Review of Economic Studies*, 38:1–12, 1971.
- [GH06] O. Gossner and J. Hörner. When is the individually rational payoff in a repeated game equal to the minmax payoff? DP 1440, CMS-EMS, 2006.
- [GLT06] O. Gossner, R. Laraki, and T. Tomala. On the optimal use of coordination. Cahier du Ceremade 0463, to appear in *Mathematical Programming B*, 2006.
- [Gol07] Y. Goldberg. Secret correlation in repeated games with imperfect monitoring: The need for nonstationary strategies. *Mathematics of Operations Research*, 32:425–435, 2007.

- [Gos95] O. Gossner. The folk theorem for finitely repeated games with mixed strategies. *International Journal of Game Theory*, 24:95–107, 1995.
- [GP84] E. J. Green and R. H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52:87–100, 1984.
- [GT06] O. Gossner and T. Tomala. Empirical distributions of beliefs under imperfect observation. *Mathematics of Operations Research*, 31(1):13–30, 2006.
- [GT07] O. Gossner and T. Tomala. Secret correlation in repeated games with signals. *Mathematics of Operations Research*, 32:413–424, 2007.
- [GV02] O. Gossner and N. Vieille. How to play with a biased coin? *Games and Economic Behavior*, 41:206–226, 2002.
- [HO06] J. Hörner and W. Olszewski. The folk theorem with private almost-perfect monitoring. *Econometrica*, 74(6):1499–1544, 2006.
- [HO07] J. Hörner and W. Olszewski. How robust is the folk theorem with imperfect public monitoring? mimeo, 2007.
- [KM98] M. Kandori and H. Matsushima. Private observation, communication and collusion. *Review of Economic Studies*, 66:627–652, 1998.
- [KO06] M. Kandori and I. Obara. Efficiency in repeated games revisited: The role of private strategies. *Econometrica*, 74:499–519, 2006.
- [Koh75] E. Kohlberg. Optimal strategies in repeated games with incomplete information. *International Journal of Game Theory*, 4:7–24, 1975.
- [Kuh53] H. W. Kuhn. Extensive games and the problem of information. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games, Volume II*, number 28 in Annals of Mathematical Studies, pages 193–216, Princeton, New Jersey, 1953. Princeton University Press.
- [KW82] D.M. Kreps and R.B. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

- [Leh90] E. Lehrer. Nash equilibria of n -player repeated games with semi-standard information. *International Journal of Game Theory*, 19:191–217, 1990.
- [Leh91] E. Lehrer. Internal correlation in repeated games. *International Journal of Game Theory*, 19:431–456, 1991.
- [Leh92a] E. Lehrer. Correlated equilibria in two-player repeated games with nonobservable actions. *Mathematics of Operations Research*, 17:175–199, 1992.
- [Leh92b] E. Lehrer. Two players repeated games with non observable actions and observable payoffs. *Mathematics of Operations Research*, 17:200–224, 1992.
- [LP99] E. Lehrer and A. Pauzner. Repeated games with differential time preferences. *Econometrica*, 67:393–412, 1999.
- [LR57] R.D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. Wiley & Sons, New York, 1957.
- [Mat04] H. Matsushima. Repeated games with private monitoring: Two players. *Econometrica*, 72:823–852, 2004.
- [MM02] G. Mailath and S. Morris. Repeated games with almost-public monitoring. *Journal of Economic Theory*, 102:189–229, 2002.
- [MMS02] G. J. Mailath, S. A. Matthews, and T. Sekiguchi. Private strategies in finitely repeated games with imperfect public monitoring. *The B.E. Journal in Theoretical Economics*, 2, 2002.
- [Mye82] R. B. Myerson. Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, 10:67–81, 1982.
- [Nas51] J. F. Nash. Noncooperative games. *Annals of Mathematics*, 54:289–295, 1951.
- [Ney99] A. Neyman. Cooperation in repeated games when the number of stages is not commonly known. *Econometrica*, 67:45–64, 1999.
- [Pic02] M. Piccione. The repeated prisoner’s dilemma with imperfect private monitoring. *Journal of Economic Theory*, 102:70–84, 2002.

- [RSS05] J. Renault, M. Scarlatti, and M. Scarsini. A folk theorem for minority games. *Games and Economic Behavior*, 53:208–230, 2005.
- [RSS06] J. Renault, M. Scarsini, and S. Scarlatti. Discounted and finitely repeated minority games with public signals. Technical Report 23, Cahiers du CEREMADE, 2006.
- [RT98] J. Renault and T. Tomala. Repeated proximity games. *International Journal of Game Theory*, 27:539–559, 1998.
- [RT04] J. Renault and T. Tomala. Communication equilibrium payoffs of repeated games with imperfect monitoring. *Games and Economic Behavior*, 49:313–344, 2004.
- [Rub77] A. Rubinstein. Equilibrium in supergames. Center for Research in Mathematical Economics and Game Theory, Research Memorandum 25, 1977.
- [Sek97] T. Sekiguchi. Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory*, 76:345–361, 1997.
- [Sel65] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfragentragheit. *Zeitschrift fur die gesamte Staatswissenschaft*, 12:201–324, 1965.
- [Smi95] L. Smith. Necessary and sufficient conditions for the perfect finite horizon folk theorem. *Econometrica*, 63:425–430, 1995.
- [Sor86] S. Sorin. On repeated games with complete information. *Mathematics of Operations Research*, 11:147–160, 1986.
- [Sti64] G. Stigler. A theory of oligopoly. *Journal of Political Economy*, 72:44–61, 1964.
- [Tom98] T. Tomala. Pure equilibria of repeated games with public observation. *International Journal of Game Theory*, 27:93–109, 1998.
- [Tom99] T. Tomala. Nash equilibria of repeated games with observable payoff vector. *Games and Economic Behavior*, 28:310–324, 1999.
- [vN28] J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.
- [Wen94] Q. Wen. The “folk theorem” for repeated games with complete information. *Econometrica*, 62:949–954, 1994.

Reviews

- [Aum81] R.J. Aumann. Survey of repeated games. In R.J. Aumann, editor, *Essays in game theory and mathematical economics in honor of Oskar Morgenstern*, pages 11–42. Wissenschaftsverlag, Bibliographisches Institut, Mannheim, Wien, Zurich, 1981.
- [Mer86] J.-F. Mertens. Repeated games. In *Proceedings of the international congress of Mathematicians*, pages 1528–1577. Berkeley, California, 1986.
- [MS06] G.J. Mailath and L. Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, 2006.
- [MSZ94] J.-F. Mertens, S. Sorin, and S. Zamir. Repeated games. CORE discussion paper 9420-9422, 1994.