Games and Economic Behavior 146 (2024) 137–159

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Folk theorems in repeated games with switching costs ☆

Yevgeny Tsodikovich [a], Xavier Venel [b,*], Anna Zseleva [c]

[a] *Department of Economics, Bar Ilan University, Israel*
[b] *Dipartimento di Economia e Finanza, Luiss university, Italy*
[c] *Department of Quantitative Economics, Maastricht University, Netherlands*

ARTICLE INFO

JEL classification:
C73

Keywords:
Switching costs
Repeated games
Folk theorem
Stochastic games
Non-zero-sum games

ABSTRACT

We study how switching costs affect the subgame perfect equilibria in repeated games. We show that (i) the Folk Theorem holds whenever the players are patient enough; (ii) the set of equilibrium payoffs is obtained by considering the payoffs of a simple one-shot auxiliary game; and (iii) the switching costs have a negative impact on a player in the infinitely undiscounted repeated game but can be beneficial for him in a finitely repeated game or in a discounted game.

1. Introduction

Switching costs appear naturally in many real-life scenarios, as changing an action might incur additional costs compared to maintaining it for an extra period. This can occur due to set-up costs (Akerlof and Yellen, 1985a,b), movement costs (Filar and Schultz, 1986), or costs of time of inactivity (Yavuz and Jeffcoat, 2007). For example, in inspection games (Avenhaus et al., 2002) the inspector typically pays a price for moving between inspected locations and gives an opportunity for undetected violations to occur while he is commuting. This additional cost is taken into account when studying different inspection models, such as environmental protection (Jørgensen et al., 2010), arms race verification (O'Neill, 1994), border protection (Darlington et al., 2022) and the cyber version of border protection (Rass and Rainer, 2014; Rass et al., 2017). In this paper, we study the general effect of switching costs on non-zero-sum scenarios, characterize the set of equilibrium payoffs, and prove the Folk Theorem for different time horizons and payoff accumulation methods. Our analysis relies on results in the Folk Theorem literature, and results specific to games with switching costs, as elaborated below.

The effect of switching costs. The introduction of switching costs to a repeated game has two major effects. First, it causes some of the payoff to dissipate: alternating between two actions yields a lower payoff than their average due to the cost of switching. This impacts both the worst-case payoff a player can defend (i.e., the individually rational level) and the payoffs a player can receive in equilibrium. Second, switching costs serve as a commitment device, as it is costly to change actions between subsequent stages. For

☆ We thank Stephane Gaubert, Chantal Marlats, and Guillaume Vigeral for their highly valuable comments. YT acknowledges the support of The Israel Science Foundation grants #2566/20, #1626/18 & #448/22. XV is a member of the Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni. XV is part of the Italian MIUR PRIN PNRR project "Supply chain disruptions, Financial losses and their preventions" (project P2022XT8C8) and the MIUR PRIN 2022 project "Learning in Markets and Society" (project 2022EKNE5K).

* Corresponding author.
E-mail addresses: yevgets@gmail.com (Y. Tsodikovich), xvenel@luiss.it (X. Venel), anna.zseleva@maastrichtuniversity.nl (A. Zseleva).

https://doi.org/10.1016/j.geb.2024.05.006
Received 15 September 2023

Available online 20 May 2024

example, when the switching costs are significantly larger than the stage payoffs, in equilibrium, actions are changed only finitely many times (Chakrabarti, 1990). The natural questions that follow are which equilibrium payoffs can a player obtain in a repeated game with switching costs, how they depart from the settings without switching costs, and how they change as the switching costs increase.

Our Model. We consider a repeated game and assume that in each time step, in addition to the payoffs of the strategic interaction, players pay some cost if they change their previous action. These switching costs are fixed throughout the game but can depend on the actions being switched. To change the relative weight between the switching costs and the payoffs from the strategic interaction, we use a multiplicative factor that applies to all the switching costs together, without affecting their inner structure. For example, in the Traveling Inspector Model (Filar and Schultz, 1986), this captures the idea that the distance between locations is fixed whereas the cost of the movement depends on the fuel prices which influences all possible routes in the same multiplicative manner (see Tsodikovich et al. (2022) for details). This leads to a dynamic game with a stage payoff defined as a weighted sum of the payoffs of the strategic interaction and the switching cost.

We consider three different variations regarding payoff accumulation and time horizon. In Section 3, we consider an infinitely repeated game with undiscounted payoffs. This serves as a benchmark for the infinitely repeated game with discounted payoffs (Section 4), and for the finitely repeated game (Section 5). In each model, we study the shape of the set of feasible payoffs and the individually rational level of each player, which are combined to provide the set of subgame perfect equilibrium (SPE) payoffs through an adapted version of the Folk Theorem.

In addition, we perform comparative statics in each model to study the effect of an increase in the switching costs on the equilibrium payoffs of the players.

Our Results. Our paper contains three parts. In the first part, we focus on infinitely repeated games with undiscounted payoffs, allowing us to establish a Folk Theorem without the use of a public correlation device and without the Full Dimensionality assumption (Dutta, 1995). The Folk Theorem is not only interesting on its own, but it is also a tool to understand the dependence of the equilibrium payoffs on the weight of the switching costs relative to the single-stage payoffs from the strategic interaction.

We show that when the switching costs are symmetric, the set of equilibrium payoffs is equal to the intersection of the individually rational payoffs with the feasible set of a one-shot auxiliary game, constructed by considering only the first two stages of the repeated game (Section 2.2). Our results provide a *novel and simple method to calculate the feasible set* since we show that it depends solely on the primary parameters of the game. When the switching costs are asymmetric, the feasible set is only a subset of the feasible set of the auxiliary game described above, as we show in Example 2. Although the symmetry assumption seems restrictive at first, in fact, it is quite general as the common assumption in the switching costs literature is that the costs are not only symmetric but also independent of the actions being changed (Chakrabarti, 1990; Lipman and Wang, 2000, 2009).

In the second part, we analyze repeated games with discounted payoffs. We show that the Folk Theorem holds asymptotically when players are patient enough and the Full Dimensionality assumption is satisfied. Moreover, *the limit set is the same as for the undiscounted evaluation.* The Full Dimensionality assumption can be satisfied directly by the stage game, and if not, we show *it is satisfied in the presence of non-zero switching costs* (Proposition 5). The intuition is that adding switching costs to players with equivalent utility functions "breaks" the equivalence and adds the missing dimensions. It is also possible that the Folk Theorem does not hold, as our model does not assume nonzero switching costs.

In the third part, we analyze finitely repeated games. Here too we prove a Folk Theorem which is based on the Folk Theorem for finite games of Marlats (2015), and show that the results from the second part hold here too. This implies that *the set of equilibrium payoffs with discounted payoffs or finite horizon (with appropriate conditions)* is asymptotically equal to the set of equilibrium payoffs with undiscounted payoffs, as discussed in the subsequent sections.

Throughout the paper, we provide *comparative statics* and study how the set of equilibrium payoffs change when the relative weight between the switching costs and the single-stage payoffs change, in a similar fashion to Tsodikovich et al. (2022). We deduce from our characterizations that *whenever one of our Folk Theorems holds, a player cannot do better with higher switching costs (maybe, up to an error of ε).* On the contrary, in cases where the Folk Theorems do not hold, it is possible for players to benefit from an increase in the switching costs, as we exemplify in a finitely repeated game. In these cases, it is possible to sustain a better SPE for a player with high enough switching costs. The intuition is that a high switching cost introduces a commitment power to the player, and forces the rest to react to his action knowing he will not change it. As a consequence, he can obtain a higher payoff with the higher switching costs (effectively, it becomes a sequential game). The positive effect of high switching costs falls in line with other well-established ideas, that when players are prohibited (here: costly prohibited) from changing actions and the players are myopic enough, they can force their preferred equilibrium on the others.

Structure of the paper. This introduction is followed by a short literature review regarding repeated games with switching costs and the relevant Folk Theorems. In Section 2 we present the notation and the model of our paper. This section is divided into five parts, as we actually deal with three sub-models. In the first part of the section, we present the common parts of all the models, which include the one-shot game and the switching costs. Then, we present each of the three sub-models which differ in the way the payoff is accumulated: undiscounted, discounted, and time-average with a finite horizon. Finally, we present the auxiliary one-shot game (Section 2.2) which is used to calculate the feasible set of the repeated game.

In Section 3 we discuss the scenario where the payoffs are undiscounted and present our main results for this case. To our knowledge, we establish for the first time a Folk Theorem for subgame perfect equilibrium with infinite horizon and undiscounted evaluation in the framework of stochastic games. We show in Section 4 that our results hold also in the more interesting case of discounted payoffs. In Section 5 we discuss the applicability of the results to finitely repeated games. In each section, we provide a Folk Theorem and discuss the implications of increased switching costs, in terms of comparative statics. We also provide counterexamples showing our assumptions are indeed necessary for the results to hold. We conclude our work in Section 6. To improve readability the proofs are relegated to the Appendix.

1.1. Literature review

Our work relates to two strands of literature, dealing both with the Folk Theorem and with switching costs in repeated games.

Switching Costs. Switching costs have been studied in the literature in multiple scenarios, mainly considering the switching costs that consumers pay when changing firms (Klemperer, 1987, 1995; Beggs and Klemperer, 1992), or setup costs firms have when setting new prices or starting new projects (Akerlof and Yellen, 1985a,b).

The model of repeated normal-form games with switching costs was proposed by Filar and Schultz (1986) (see also Filar (1985)) as a method to incorporate movement costs and lost time into repeated games. In their framework (namely, the *Traveling Inspector Model*), only one player pays some switching costs (a moving inspector) whereas the rest are stationary inspectees. They used this assumption to aggregate all the inspectees into one player and reduce the repeated game into a two-player stochastic game where only the inspector controls the transitions. This significantly simplifies the problem and allows the utilization of the theory of single-controller stochastic games (Filar, 1980, 1981; Filar and Raghavan, 1984).

The model was later extended by Chakrabarti (1990) and Lipman and Wang (2000, 2009) to scenarios where all players pay switching costs. Lipman and Wang assumed that all the players face the same cost and that the cost is the same for any pair of actions being switched. In their early paper, they focused on the equilibria in finite horizon whereas in the latter they studied the infinitely repeated game (with discounting or time-averaging). Both models are in discrete time but justified by the discretization of a continuous model. This leads to a specific structure of the general payoff, which is different from ours, in which the strategic payoff is normalized whereas the switching costs are not normalized. In particular, in their model, infinite switching leads to an infinitely negative (time-averaged) payoff. As a consequence, we obtain substantially different results. Depending on the application, both models can be interesting. The model of Lipman and Wang (2009) is suitable for games with fixed horizon where the frequency of the decisions of the players is arbitrarily large. On the contrary, our model is suited for games where each stage has a fixed length and the number of stages goes to infinity. We discuss these differences along with our results.

In Lipman and Wang (2000), they showed that the introduction of switching costs has two effects and that both can lead to their result: a switching cost may serve as a commitment to keep the current action (diminishing the profitability of a deviation), but at the same time can make punishment harder (switching to and from the punishment strategy incurs costs). In Lipman and Wang (2009), they showed that the asymptotic behavior of the set of Nash Equilibria in the discounted case depends on the ratio between the cost of switching and the discount factor. They highlighted especially that some payoffs disappear due to the necessity to switch infinitely often.

Repeated games with switching costs can also be seen as a particular case of supergames. Friedman (1971) introduced the notion of supergames as a sequence of one-shot games where the final payoff of each agent is a discounted average of the stage payoff. Friedman (1974) extended his model by introducing time-dependent supergames where there is a connection between the actions played at two consecutive stages. In particular, this includes repeated games with switching costs: the payoffs today depend on the cost of switching from yesterday's actions. Switching costs models provide additional structure compared to supergames, such as time-homogeneity and partial separation of the players' payoffs.[1]

Finally, zero-sum games with switching costs were studied in Tsodikovich et al. (2022). Similarly to Filar and Schultz (1986), the authors focused on models where only one player has switching costs that are paid to the other adversarial player (otherwise, the game is no longer zero-sum). They studied the regularity of the value function as a function of the scale of the switching costs. In the current work, we follow the same line and investigate the set of equilibria in a non-zero-sum framework as a function of this scale.

We note that under some algebraic manipulation, this model (and games with switching costs in general) can describe other situations. For example, Schoenmakers et al. (2008) studied a model of learning by doing, where repeating an action grants a bonus. Up to an affine transformation, this is identical to penalizing players for switching. In this general framework of repeated games with inter-temporal externalities, Tsodikovich et al. (2024) compare optimal strategies to the commonly used time-independent strategies.

Finally, there also exist articles that analyze switching costs in a one-shot game framework. Guney and Richter (2018) focused on the notion of *status quo* in the presence of switching costs whereas Guney and Richter (2022) developed a notion of Nash equilibrium with switching cost for one-shot games.

Folk Theorem. The classical Folk Theorem states that any feasible payoff (can be achieved by some strategy) and individually rational (above what each player can guarantee for himself, his minmax level) is supported in equilibrium in the repeated game with undiscounted payoffs. This proposition, focusing on the Nash equilibrium, can be proven by relying on trigger strategies. The players agree on a joint plan to generate a feasible payoff as long as all players cooperate. If one of the players decides to deviate, the others punish him to his minmax level. Aumann and Shapley (1994) extended this to Subgame Perfect Equilibria (SPE) by using the two following ideas: punish only for a finite time instead of forever, and punish for longer and longer periods each time someone deviates.

There exist versions of the Folk Theorem for different repeated interaction models. The closest to our model are of Friedman (1990), who established a Folk Theorem for supergames with compact action spaces and some separation between the past actions of one player and the present payoff of another, and of Dutta (1995), who established the main Folk Theorem for stochastic games.

[1] The payoff of Player ℓ today does not depend on the actions of the others yesterday.

The latter is a natural generalization of the classical Folk Theorem under some ergodicity assumptions and the Full Dimensionality assumption and is similar to the classical one established by Fudenberg and Maskin (1986) for repeated games. In particular, Dutta showed that three assumptions are required for the Folk Theorem to hold in this case: independence of the asymptotic feasible set of the initial state, independence of the asymptotic maxmin of the initial state, and full dimensionality of the feasible sets. This result was extended to public monitoring signaling simultaneously in Fudenberg and Yamamoto (2011) and in Hörner et al. (2011) under similar assumptions. Finally, Marlats (2015) provided a Folk Theorem for stochastic games with a finite horizon, by adding a richness condition on the limit set of finite SPE payoffs (see Assumption A4).

How does the introduction of switching costs change the story? As pointed out in Lipman and Wang (2009), switching costs play two different roles. First, they change the set of feasible payoffs. When playing a joint plan, it is possible to switch rarely compared to the global payoff such that switching costs are negligible, but it is also possible to use switches to decrease the payoffs, even outside the feasible set of the one-shot game. Second, they affect the minimax payoffs. The presence of costs for Player ℓ constrains him and therefore the other players can punish him more. Notice that when punishing, players suffer first from having bad payoffs but also from the cost of switching actions, hence it is necessary to threaten them to obtain an SPE or to reward them afterward. Still, they established a new Folk Theorem (concerning the time-averaging evaluation) by introducing a suitable notion of individual rationality, while Chakrabarti (1990) provided a complete characterization of the payoffs supported by an equilibrium under the additional assumption that the switching costs out-scale any possible gain in the one-shot game.

From a technical perspective, expressing a repeated game with switching costs as a stochastic game has been suggested in Lipman and Wang (2009) (also in Filar and Schultz (1986), but when only one player pays the switching costs) and used extensively in Tsodikovich et al. (2022). In this representation, the states correspond to the pure actions played in the previous time step, and the payoffs in each state comprise the standard single-stage payoff and the switching costs. These stochastic games have additional structure compared to the general ones studied by Dutta (1995), which facilitates some of the results. For example, the resulting stochastic game is in fact a dynamic game, as the transitions are deterministic and depend only on the actions of the players. Similarly, any state is reachable from any state in a single step, and there are no absorbing states. More importantly, we show that whenever all players have non-zero switching costs, the special structure of the game fulfills the Full Dimensionality assumption (Proposition 5), and allows the use of the Folk Theorem from Dutta (1995).

This relates to the general discontinuity of the set of equilibrium payoffs and strategies when the switching costs go to zero, which was studied by Lipman and Wang (2009). They showed that even a small non-zero switching cost can have a significant effect on the equilibrium structure. For example, small switching costs can create a multiplicity of equilibria in a setting where without switching costs there is a unique one, and vice-versa. Such phenomena are well understood in light of Proposition 5, as small non-zero switching costs are enough for a game to fulfill the Full Dimensionality assumption in cases that the assumption is not fulfilled without switching costs.

2. The switching costs model

We consider an *n-player game* with action-dependent switching costs. It is formally defined by a tuple $\Gamma = (N, (I^\ell)_{\ell\in N}, (u^\ell)_{\ell\in N}, (S^\ell)_{\ell\in N}, c)$ where N is a finite set of n players. For every Player $\ell \in N$, I^ℓ is a finite set of actions of size m^ℓ, u^ℓ is his payoff function from $I = \bigtimes_{\ell\in N} I^\ell$ to $\mathbb{R}$, and $S^\ell = (s^\ell_{ij})$ is an $m^\ell \times m^\ell$ switching costs matrix. We identify the sets of actions with the sets $I^\ell = \{1, \dots, m^\ell\}$. The relative weight of the switching costs compared to the stage payoff is $c \geqslant 0$.

At each time step $t > 1$, each Player ℓ chooses an integer $i^\ell(t) \in \{1, \dots, m^\ell\}$. Denote by $i(t) = (i^\ell(t))_{\ell\in N}$ the profile of actions played at stage t. The stage payoff of Player ℓ is $\widetilde{u}^\ell(i(t-1), i(t)) := u^\ell(i(t)) - c \cdot s^\ell_{i^\ell(t-1)i^\ell(t)}$, so Player ℓ is penalized for switching the previous action $i^\ell(t-1)$ to the action $i^\ell(t)$ by $c \cdot s^\ell_{i^\ell(t-1)i^\ell(t)}$. Naturally, at the first time step, $t = 1$, switching costs are not paid and the payoffs are according to u^ℓ. This creates an asymmetry between the first stage and the rest of the stages. To simplify the notation, and although $i(0)$ is not defined, we set $\widetilde{u}^\ell(i(0), i(1)) = u^\ell(i(1))$. Alternatively, our results hold if we assume that there exists some action profile at time $t = 0$ and switching costs are paid in every stage $t \geqslant 1$ of the game.

We assume that for all i, j and for all $\ell \in N$, $s^\ell_{ij} \geqslant 0$ and $s^\ell_{ii} = 0$. We distinguish between several sets of agents. If $S^\ell = 0$, then Player ℓ has no switching costs. We denote the set of such players by N^0 and the set of players who have switching costs by $N^1 = N \setminus N^0$. If S^ℓ is a matrix where all non-diagonal elements are strictly positive, Player ℓ pays a cost for any change of actions. We say that he has *no free switches* and we denote the set of such players by N^{nf}.

Note that even when there are two players and $u^1 = -u^2$, the addition of switching costs turns the game into a non-zero-sum game, since $\widetilde{u}^1 \neq -\widetilde{u}^2$. Tsodikovich et al. (2022) analyzed a similar two-player framework, but assumed that only Player 2 pays switching costs and that they are paid to Player 1, which makes the game indeed a zero-sum game. Their analysis can serve as a worst-case analysis in our model.

At each time period, the players are also allowed to play mixed actions, where a mixed action of Player ℓ is a probability distribution over I^ℓ. As usual, the set of all probability distributions over some finite set A is denoted by $\Delta(A)$, so the set of mixed actions of Player ℓ is $\Delta(I^\ell)$. Note that at time t, $i^\ell(t-1)$ is already known, even if Player ℓ played a mixed action at time $t-1$. This information is important in our case, as the utilities at stage t are determined by the pure actions played at stage $t-1$. In addition, we assume that players have access to a public randomization device as it is classic in the literature of Folk Theorems in dynamic environments (Dutta, 1995; Marlats, 2015).

2.1. Payoff accumulation and subgame perfect equilibria

We consider three models which differ in the way that the payoff is accumulated and the horizon of the game. In all the versions, the equilibrium notion we are interested in is subgame perfect equilibrium (SPE). Informally, in these equilibria, after each history, the continuation strategies are the best responses to each other and are equilibrium strategies in the subgame that starts at this point. These are all formally defined in the rest of this subsection.

2.1.1. Finite game of length T

The game is played for $T \geqslant 1$ stages and the payoff of each player is the average per-stage payoff. Formally, let $\sigma = (\sigma^{\ell})_{\ell\in N}$ be an n-tuple of strategies for the entire game and $(\sigma^{\ell}(t))_{\ell\in N}$ be the mixed actions played at stage t (given the history). The average per-stage payoff of Player ℓ is therefore

$$\gamma_T^{\ell}(\sigma) = \mathbb{E}_{\sigma}\left(\frac{1}{T}\sum_{t=1}^{T}\widetilde{u}^{\ell}(i(t-1), i(t))\right). \tag{1}$$

For $T = 1$ this is simply the one-shot game without switching costs ($\widetilde{u}^{\ell} = u^{\ell}$), which is hereafter referred to as *the one-shot game u*. This game is relevant, as the possible payoffs in this game are the baseline for the possible payoffs and possible equilibrium payoffs in the repeated game with and without switching costs.

In particular, two important sets are the *feasible set* of the one-shot game and the *individually rational payoffs set* of the one-shot game. The feasible set of the game with utilities $u = (u^{\ell})_{\ell\in N}$ is the convex combination of all the possible vector payoffs $F(u) = Conv(\{u(i), i \in I\})$. When it is clear from the context, we will simply denote this set by F.

Similarly, the individually rational level in the repeated game is related to the individually rational level in the one-shot game, i.e. to the minimax value. The *minimax for Player ℓ* in the one-shot game is defined as the maximal payoff that he can defend when the other players try to minimize his payoff in an uncorrelated manner:

$$\overline{v}^{\ell} = \min_{\sigma^{-\ell}\in\Sigma^{-\ell}} \max_{x^{\ell}\in\Delta(I^{\ell})} u^{\ell}(x^{\ell}, \sigma^{-\ell}),$$

where $\Sigma^{-\ell} = \bigtimes_{j\neq\ell} \Delta(I^j)$. The set of *individually rational payoffs* is the set of all payoffs above the minimax level for all players:

$$IR = \left\{(y^{\ell})_{\ell\in N} \in \mathbb{R}^n, \text{ for all } \ell \in N,\ y^{\ell} \geqslant \overline{v}^{\ell}\right\}.$$

If $n = 2$, then $\overline{v}^{\ell}$ is the value of the zero-sum finite game where Player ℓ has payoff u^{ℓ} whereas Player $-\ell$ has payoff $-u^{\ell}$. If $n > 2$, it is also the value of a zero-sum game but it is no longer a finite game since the set of strategies $\Sigma_{-\ell}$ is not a product state space.

As mentioned above, the notion of equilibrium we study in the paper is SPE, i.e., equilibria in which after each history no player has an incentive to deviate. However, it is not possible to consider the continuation payoffs after each history in the standard manner, since the subgame starting after each history of length $t > 1$ includes an initial state and possible payment of switching costs, while the original game is defined as such that switching costs do not exist in the first stage. We resolve this issue by a slight abuse of notation. Fix $\sigma = (\sigma^{\ell})_{\ell\in N}$ be a profile of strategies and let h_t be a history of length $t < T$, i.e., $h_t \in I^t$. We define $\gamma_T(\sigma, h_t)$ to be the payoff of a game whose first t stages are according to h_t (regardless of the strategy profile) and the rest of the $T - t$ stages are according to the strategy profile σ. A strategy profile is a subgame perfect equilibrium if after every history no player has a profitable deviation, even if this history is not part of the equilibrium path.

Definition 1. A profile of strategies $\sigma_* = (\sigma_*^{\ell})_{\ell\in N}$ is a *Subgame Perfect Equilibrium (SPE)* in the finite game of length T if after every history h_t of length $t < T$, every $\ell \in N$, and for every σ^{ℓ} strategy of Player ℓ,

$$\gamma_T^{\ell}\left((\sigma^{\ell}, (\sigma_*)^{-\ell}), h_t\right) \leqslant \gamma_T^{\ell}(\sigma_*, h_t),$$

where as usual, $(\sigma_*)^{-\ell}$ is the vector σ_* without σ_*^{ℓ}. The vector $(\gamma_T^{\ell}(\sigma_*))_{\ell\in N}$ is then called an SPE Payoff of the finitely repeated game of length T. We denote by $SPE_T(c)$ the set of SPE Payoffs.

2.1.2. δ-discounted payoffs game

The game has infinitely many stages and the payoff of each player is discounted according to $\delta \in (0, 1)$. Formally, let $\sigma = (\sigma^{\ell})_{\ell\in N}$ be an n-tuple of strategies for the entire game and $(\sigma^{\ell}(t))_{\ell\in N}$ be the mixed actions played at stage t (given the history). The δ-discounted payoff of Player ℓ is therefore

$$\gamma_{\delta}^{\ell}(\sigma) = \mathbb{E}_{\sigma}\left((1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\widetilde{u}^{\ell}(i(t-1), i(t))\right). \tag{2}$$

The closer the discount factor to 1, the more patient the players are. We define the SPE for this case in an analogous manner to Definition 1.

Definition 2. A profile of strategies $\sigma_* = (\sigma_*^\ell)_{\ell\in N}$ is a *Subgame Perfect Equilibrium (SPE)* in the δ-discounted payoffs game if after every finite history h, every $\ell \in N$, and for every σ^ℓ strategy of Player ℓ,

$$\gamma_\delta^\ell\left((\sigma^\ell,(\sigma_*)^{-\ell}),h\right) \leqslant \gamma_\delta^\ell(\sigma_*,h).$$

The vector $(\gamma_\delta^\ell(\sigma_*))_{\ell\in N}$ is then called an SPE Payoff of the δ-discounted game. We denote by $SPE_\delta(c)$ the set of SPE Payoffs.

2.1.3. *Undiscounted payoffs game*

The game is played infinitely many stages and the payoff of each player is the limit of the average payoffs. Formally, let $\sigma = (\sigma^\ell)_{\ell\in N}$ be an n-tuple of strategies for the entire game and $(\sigma^\ell(t))_{\ell\in N}$ be the mixed actions played at stage t (given the history). The undiscounted payoff of Player ℓ is therefore

$$\gamma^\ell(\sigma) = \mathbb{E}_\sigma\left(\liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\widetilde{u}^\ell(i(t-1),i(t))\right). \tag{3}$$

We define the SPE for this case in an analogous manner to Definitions 1 and 2. Note that here the continuation payoffs can be considered in the regular manner, as the switching costs at the first stage of the subgame are negligible compared to the long-term average. On the contrary, we impose a stronger assumption on the convergence of the payoffs, as is done in Maschler et al. (2013) (Definition 13.16).

Definition 3. A profile of strategies $\sigma_* = (\sigma_*^\ell)_{\ell\in N}$ is a *Subgame Perfect Equilibrium (SPE)* in the undiscounted game if, with probability 1, the mean-average payoff converges and after every history, no player has an incentive to deviate. Formally, for every $\ell \in N$, the following limit exists with probability 1 under σ_*:

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\widetilde{u}^\ell(i(t-1),i(t))$$

Let h be a finite history of length t, denote by $\sigma_*^\ell(h)$ the continuation strategy after the finite history h of Player ℓ. Then for every $\ell \in N$ and for every σ^ℓ strategy of Player ℓ,

$$\gamma^\ell\left((\sigma^\ell,(\sigma_*(h))^{-\ell})\right) \leqslant \gamma^\ell(\sigma_*(h)). \tag{4}$$

The vector $(\gamma^\ell(\sigma_*))_{\ell\in N}$ is then called an SPE Payoff of the infinitely repeated game. We denote by $SPE_\infty(c)$ the set of SPE Payoffs.

Remark 1. There exist other definitions of the infinite payoff and equilibrium concept in the undiscounted games. For example, one can apply the result in Ashkenazi-Golan et al. (2022) to obtain the existence of the weaker notion of Nash equilibrium in the game with payoffs $(\gamma^\ell)_{\ell\in N}$. This approach does not yield the convergence on-path and does not guarantee subgame perfection.

Another possibility is to replace the inferior limit in Equation (3) with the superior limit yielding a more optimistic evaluation of infinite sequences of payoffs. The results will be the same. Indeed, we first impose that under σ_* there is convergence so both the inferior limit and superior limit coincide. Second, punishment against deviation can be adapted to guarantee that both an optimistic and pessimistic player have no incentives to deviate.

When $c = 0$, the above definitions reduce to the regular finite and infinite games without switching costs, with the vector of utilities $(u^\ell)_{\ell\in N}$. This article aims to study the properties of the sets $SPE_T(c)$, $SPE_\delta(c)$ and $SPE_\infty(c)$ as a function of the cost factor c.

2.2. *The auxiliary one-shot game*

We present a novel method to calculate the feasible set and the set of equilibrium payoffs in the undiscounted infinitely repeated game. We show (Theorem 1) that when all the matrices S^ℓ are symmetric, the feasible set is equal to the feasible set of a one-shot game constructed by merging two stages of the game. This result will extend asymptotically to the feasible sets for the infinitely repeated game with discounted payoffs (Section 4), and for the finitely repeated game (Section 5). In this section, we formally define this one-shot game.

Let $\Gamma = (N,(I^\ell)_{\ell\in N},(u^\ell)_{\ell\in N},(S^\ell)_{\ell\in N},c)$ be a repeated game with switching costs. For each Player ℓ, define the new action spaces $J^\ell = I^\ell \times I^\ell$ and $J = \bigtimes_{\ell\in N} J^\ell$. We define naturally the projections of each $j \in J$ on the first and second coordinates by j_1 and j_2. The payoff of the one-shot game is then defined for every $\ell \in N$ and for every $(j^\ell)_{\ell\in N} \in J$ by

$$g_c^\ell(j) := \frac{1}{2}\left(u^\ell(j_1)+u^\ell(j_2)\right) - c\cdot s^\ell_{j_1^\ell j_2^\ell}. \tag{5}$$

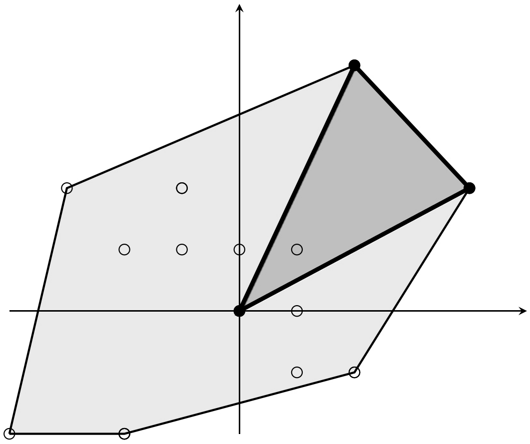

Fig. 1. The feasible set of the standard one-shot "Battle of the Sexes" game (dark gray and filled dots) and the feasible set of the auxiliary one-shot game for $c = 1$ (dark/light gray and empty dots), discussed in Example 1. Dots correspond to payoffs obtained by pure actions in their respective one-shot game.

Let us define by $g_c = (g_c^\ell)_{\ell\in N}$ the payoff function of this one-shot game. By our previous notation, the set of feasible payoffs in this auxiliary game is denoted by $F(g_c)$.

Example 1. *The repeated "Battle of the Sexes" game with switching costs.* Consider the standard two-player "Battle of the Sexes" game augmented with two switching cost matrices

	$A2$	$B2$
$A1$	2, 1	0, 0
$B1$	0, 0	1, 2

, $S^1 =$

	$A1$	$B1$
$A1$	0	2
$B1$	1	0

, and $S^2 =$

	$A2$	$B2$
$A2$	0	1
$B2$	1	0

.

In the corresponding one-shot game each player has 4 actions of the form $j_1^\ell j_2^\ell$ (where $j_i^\ell \in I^\ell$) and the corresponding payoff matrix is

	$A2A2$	$A2B2$	$B2A2$	$B2B2$
$A1A1$	$2, 1$	$1, 0.5 - c$	$1, 0.5 - c$	$0, 0$
$A1B1$	$1 - 2c, 0.5$	$1.5 - 2c, 1.5 - c$	$-2c, -c$	$0.5 - 2c, 1$
$B1A1$	$1 - c, 0.5$	$-c, -c$	$1.5 - c, 1.5 - c$	$0.5 - c, 1$
$B1B1$	$0, 0$	$0.5, 1 - c$	$0.5, 1 - c$	$1, 2$

The payoffs and feasible sets are shown in Fig. 1 for $c = 1$.

In this paper, we consider c as a variable and perform different comparative statics on the outcomes of the game for different values of c. For example, we study how the feasible set of payoffs changes as a function of c. One interesting property of the correspondences we are studying is decreasing for the order, formally defined below.

Definition 4. A correspondence $L(c) : \mathbb{R} \to \mathbb{R}^n$ is *decreasing* for the order on $\mathbb{R}^n$ if:

$$\forall c_1 < c_2,\ \forall x \in L(c_2),\ \exists x' \in L(c_1),\ \forall \ell \in N,\ x^\ell \leqslant x'^\ell.$$

Roughly speaking, when this property holds, the correspondence is "decreasing", at least in the sense that for each element x of the mapping under the higher c_2, can be found an element x' of the mapping under the lower c_1, such that x' is larger than x in all coordinates. It is easy to verify that this relation holds between the light gray area ("high c") and the dark gray area ("low c") of Fig. 1.

3. Undiscounted payoffs

We start with the repeated game with undiscounted payoffs. This model is slightly simpler as players are indifferent to things happening in finite times.

Our main result in this section is the Folk Theorem for repeated games with switching costs and undiscounted payoffs, and a full characterization of the long-run feasible set when the switching costs are symmetric.

The two key elements in a Folk Theorem are usually the feasible set and the individually rational set. In the rest of the section, we provide adequate definitions for any parameter c and characterize both sets. We then prove the Folk Theorem for any c.

3.1. The long-run feasible set

For the definition of the feasible set, we follow an approach similar to Dutta (1995) and Marlats (2015). To have some flexibility, a vector of payoffs is feasible if it is an accumulation point of the sequence of mean-average payoffs for a given profile of strategies.

Definition 5. Let $x \in \mathbb{R}^n$. The payoff x is *generated* by the strategy profile σ if there exists a sequence of lengths $(T_k)_{k\geqslant 1}$ such that $(\gamma_{T_k}(\sigma))_{k\geqslant 1}$ converges to x. We can then define the *long-run feasible set*, denoted $\mathcal{F}(c)$ by

$$\mathcal{F}(c) = \{x \in \mathbb{R}^n, \text{ s.t. there exists } \sigma \text{ that generates } x\}.$$

When considering $\mathcal{F}(c)$ as a function of c, we can discuss how the long-run feasible set changes when the relative weight between the stage payoff and the switching costs changes. The key properties of $\mathcal{F}(c)$ are summarized in the following propositions.

Proposition 1. The correspondence $\mathcal{F}(c)$ is Lipschitz,[2] increasing in the sense of inclusion from $\mathcal{F}(0) = F(u)$ and decreasing for the order on $\mathbb{R}^n$.

Proof. See Appendix A.1. □

In particular, since the correspondence $\mathcal{F}(c)$ is increasing in the sense of inclusion, any long-run feasible payoff in the repeated game without switching cost can still be attained in the repeated game with high switching cost. The intuition is that players can change actions very rarely.

Proposition 2. For every $c \geqslant 0$, $\mathcal{F}(c)$ is convex and closed.

Proof. See Appendix A.2. □

Our next result connects the long-run feasible set, $\mathcal{F}(c)$, and the feasible set of the one-shot auxiliary game with utilities g_c, $F(g_c)$. We show that they are equal when all the switching cost matrices are symmetric, and the former is a subset of the latter in the more general case.

Theorem 1. *Fix a repeated game Γ with switching costs and let g_c be its associated auxiliary one-shot game for some c. Let P be the set of distributions over $I \times I$ such that the first and the second marginals are equal:*

$$P = \left\{ \pi \in \Delta(I \times I), \text{ such that for all } i \in I, \sum_{i' \in I} \pi(i,i') = \sum_{i' \in I} \pi(i',i) \right\}.$$

Then the long-run feasible set is the image of the set P by the function g_c:

$$\mathcal{F}(c) = g_c(P) := \{g^c(p),\ p \in P\} \subseteq F(g_c).$$

Moreover, if all of the cost matrices S^ℓ are symmetric then the inclusion is an equality: $\mathcal{F}(c) = F(g_c)$.

Proof. See Appendix A.3. □

The ability to define the long-run feasible set using the convex hull of an auxiliary game which depends solely on the primary data of the game strongly relies on the symmetry of the switching costs. The key difference between the characterization for any cost and the characterization for only symmetric costs is the relaxation on the set of distributions over $I \times I$ that are allowed. In the first case, we restrict to probability distributions such that both marginals are equal whereas in the second one, we allow any probability distribution. The intuition is the following. A history in the original game can always be mapped into a history in the one-shot auxiliary game with the same payoff (up to some error). The converse is in general not true creating a possible gap between the long-run feasible set of the repeated game with switching cost and the feasible set of the one-shot auxiliary game. We prove that in the case of symmetric payoff, this gap does not exist.

Although the assumption of symmetric switching costs seems like a limitation, in fact, it is less restrictive than the common assumption in the literature that all the switching costs are the same (i.e. $s^\ell_{i,j} = 1$ if $i \neq j$ and $s^\ell_{i,j} = 0$ otherwise). Hence, we generalize the literature to a larger family of switching cost functions. The following example shows that this condition is indeed necessary, and without it, the long-run feasible set is only a subset of $F(g_c)$.

---

[2] A correspondence F is Lipschitz if there exists $\theta > 0$ such that for every c_1, c_2, and every $x_1 \in F(c_1)$ there exists $x_2 \in F(c_2)$ such that $|x_2 - x_1| \leqslant \theta |c_2 - c_1|$.

Example 2. *Counterexample with non symmetric switching costs.*

Consider the following two-player game where Player 1 has only one action and Player 2 has two actions

L	R
2,1	0,0

and the switching costs matrix for Player 2 is

$S^2 =$

	L	R
L	0	1
R	0	0

.

In the corresponding one-shot game, Player 2 has 4 actions and the payoff matrix is

	LL	LR	RL	RR
TT	$2,1$	$1,0.5-c$	$1,0.5$	$0,0$

The payoff vector $(1, 0.5 - c)$ is a feasible payoff in the auxiliary game but is not attainable in the long run in the original game. For Player 1 to obtain 1, Player 2 can at most switch half of the time from L to R inducing a mean-average cost of maximum $0.5c$. The sequence where Player 2 repeats LR infinitely often can not be mapped to a sequence in the original game with a similar payoff.

We can also obtain an asymptotic result when the cost goes to infinity. As c increases, more payoff can be dissipated and in the limit $c \to \infty$, any arbitrarily negative payoff can be obtained for players in N^1. Players in N^0 are unaffected by c. This limit set is therefore

$$\overline{F} = \{x \in \mathbb{R}^n | \exists y \in F(u) \text{ such that for all } n \in N^1,\ x \leqslant y \text{ and for all } n \in N^0,\ x = y\},$$

and it includes the feasible set $F(u)$ and all the possible payoffs which are bounded by a payoff inside the feasible set (except for players who never pay switching costs, their payoff is always a convex combination of their one-shot payoffs). In other words, for Players in N^1, the set $\overline{F}$ is the set of all payoffs that are bounded above by the Pareto-Efficient front.

Corollary 1. *The correspondence $\mathcal{F}(c)$ converges to $\overline{F}$ when c goes to ∞.*

The idea is demonstrated using the following example (see also Fig. 2).

Example 3. *"Battle of the Sexes" with large switching costs.*

Consider the game in Example 1 and assume $c \to \infty$. The "payoffs" of the corresponding game are

	$A2A2$	$A2B2$	$B2A2$	$B2B2$
$A1A1$	$2,1$	$1,-\infty$	$1,-\infty$	$0,0$
$A1B1$	$-\infty,0.5$	$-\infty,-\infty$	$-\infty,-\infty$	$-\infty,1$
$B1A1$	$-\infty,0.5$	$-\infty,-\infty$	$-\infty,-\infty$	$-\infty,1$
$B1B1$	$0,0$	$0.5,-\infty$	$0.5,-\infty$	$1,2$

The convex hull of these payoffs is indeed $\overline{F}$, as shown in Fig. 2.

Clearly, the equilibrium payoffs cannot be arbitrarily negative as players can always obtain at least their minimax value in pure strategies, independent of the value of c or the strategy of the others. In the following section, we study how the individually rational levels change with c.

3.2. Individually rational payoffs

To study the individually rational level of a particular Player ℓ, we use the usual approach and assume the rest of the players disregard their own payoffs and only care about minimizing the payoff of Player ℓ. We therefore define n auxiliary zero-sum-like repeated games with switching costs. For every $\ell \in N$, define the game Γ^ℓ as the zero-sum game where Player ℓ has payoff $\widetilde{u}^\ell$ and is facing an imaginary Player $-\ell$. We call the *long-run individually rational level* of Player ℓ the maximum payoff that Player ℓ can defend. It is called the "long-run average minmax" in Dutta (1995).

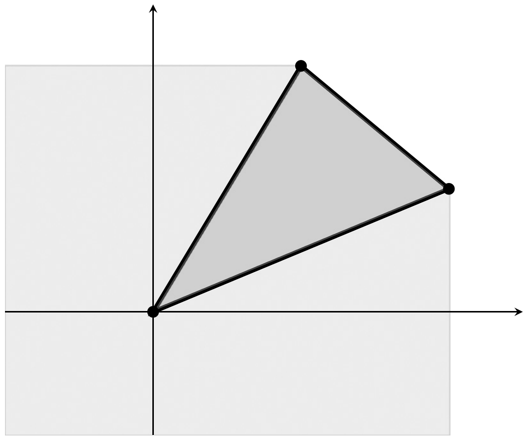

Fig. 2. Feasible set $F(u)$ (dark gray) and $\overline{F}$ (dark and light gray) of Example 3.

Definition 6. The *long-run individually rational level* of Player ℓ is defined by

$$\overline{v}^{\ell}(c) = \inf_{\sigma^{-\ell}\in\Sigma^{-\ell}} \sup_{\sigma^{\ell}} \mathbb{E}_{\sigma^{\ell},\sigma^{-\ell}}\left(\liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \widetilde{u}^{\ell}(i(t-1), i(t))\right).$$

Zero-sum-like repeated games with switching costs can be reformulated as dynamic games where the state is the previous action profile. Here, only the previous action of Player ℓ matters as only he pays switching costs, so this is a special type of stochastic game called a single-controller stochastic game (see, for example, Filar (1981)) where the long-run evaluation does not depend on the state variable.

Our first goal is to characterize the behavior of the long-run individually rational level $\overline{v}^{\ell}(c)$ as a function of c. A similar analysis was done in Tsodikovich et al. (2022) for the two-player case.

Proposition 3. Fix a game Γ with switching costs. For every $\ell \in N$, the long-run individually rational level $\overline{v}^{\ell}(c)$ of Γ^{ℓ} is a continuous, decreasing, and semialgebraic function of c. If there are only 2 players, it is also convex and piece-wise affine.

Proof. See Appendix A.4. □

The following example shows that the assumption on the number of players is needed to obtain the nicer properties of $\overline{v}^{\ell}(c)$. Whenever there are more than two players, convexity and piece-wise linearity are no longer guaranteed.

Example 4. *A game with non-convex and non-affine* $\overline{v}^{\ell}(c)$.

Consider a 3-player game with the following payoff matrices (only the payoff of Player 1 is written):

	L	R
T	−1	0
B	0	0

l

	L	R
T	0	0
B	0	−1

r

Player 1 pays a cost of one if he changes his action. We denote the actions of Player 1 (chooses rows) by $\{T, B\}$, the actions of Player 2 (chooses columns) by $\{L, R\}$ and the actions of Player 3 (chooses matrices) by $\{l, r\}$.

For a switching cost of $c \leqslant 1$, the optimal action for Player 2 and Player 3 to minimize the payoff of Player 1 is to play randomly as follows. After Player 1 played T, Player 2 (resp. Player 3) plays L (resp. l) with probability $\frac{1+c}{2}$ whereas after Player 1 played B, Player 2 (resp. Player 3) plays R (resp. r) with probability $\frac{1+c}{2}$. For a switching cost larger than 1, one can check that the optimal strategies of Player 2 and Player 3 to minimize the payoff of Player 1 are to best reply to the previous action of Player 1. We obtain that the minmax value for Player 1 is equal to $\overline{v}^{1}(c) = \max\left(-1, -\frac{1}{4}(1+c)^2\right)$ and therefore is neither piece-wise linear nor convex. The high switching costs of Player 1 restrict him and allow the other two players to coordinate to punish him.

We now move from the long-run individually rational level of one player to the set of all long-run individually rational payoffs.

Definition 7. The set of *long-run individually rational payoffs* with switching cost c is defined by $\mathcal{IR}(c) = \left\{(x^{\ell})_{\ell\in N} \in \mathbb{R}^n, \text{ for all } \ell \in N,\ x^{\ell} \geqslant \overline{v}^{\ell}(c)\right\}$.

When considering $\mathcal{IR}(c)$ as a function of c, we can discuss how the set of long-run individually rational payoffs changes when the relative weight between the payoff from the strategic interaction and the switching costs changes. These properties are obtained immediately by Proposition 3.

Corollary 2. *The correspondence $\mathcal{IR}(c)$ is semialgebraic, increasing in the sense of inclusion, and decreasing for the order on $\mathbb{R}^n$.*

Moreover, we also know the asymptotic behavior of $\overline{v}^{\ell}(c)$ for large c in two cases: when there are no free switches and when there are no switching costs at all. When there are no free switches, there is a positive switching cost between any pair of actions, and for large enough c, the optimal strategy is to play purely and never switch. This leads to the *pure reservation payoff* defined in Lipman and Wang (2009):

$$\overline{w}^{\ell} = \max_{i^{\ell} \in I^{\ell}} \min_{y^{-\ell} \in \bigtimes_{k \neq \ell} \Delta(I^k)} u^{\ell}(i^{\ell}, y^{-\ell}) = \max_{i^{\ell} \in I^{\ell}} \min_{i^{-\ell} \in \bigtimes_{k \neq \ell} I^k} u^{\ell}(i^{\ell}, i^{-\ell}).$$

Proposition 4. Assume that Player ℓ has no free-switching, then there exists $\overline{c}^{\ell}$ s.t. for every $c \geqslant \overline{c}^{\ell}$, Player ℓ's optimal strategy in the game Γ^{ℓ} is the pure maximin strategy in the game with utilities u^{ℓ} and for all $c \geqslant \overline{c}^{\ell}$, $\overline{v}^{\ell}(c) = \overline{w}^{\ell}$.

Proof. See Appendix A.5. □

This result is closely related to the results of Chakrabarti (1990) in inertia supergames when there is a high enough switching cost such that any one-period gain is offset by the cost of changing.

When there are no switching costs at all, c plays no role, and for every c, the minimax level is the same as in the one-shot game: $\overline{v}^{\ell}(c) = \overline{v}^{\ell}$. In the third option, when there are switching costs but there are also free switches, there is no simple way to calculate the asymptotic long-run individually rational level (which exists, as $\overline{v}(c)$ is bounded by $\overline{w}^{\ell}$ and decreasing).

We obtain that the set of all long-run individually rational payoffs converges to the set $\overline{V}$, and it comprises of all vector payoff above pure reservation payoff for players in N^{nf}, above the minimax value of the one-shot game for players in N^0 and above $\underline{v}^{\ell}$ for the rest:

$$\overline{V} = \left\{ (x^{\ell}) \in \mathbb{R}^N, \text{ for all } \ell \in N, x^{\ell} \geqslant \lim_{c \to \infty} \overline{v}^{\ell}(c) \right\},$$

where $\lim_{c \to \infty} \overline{v}^{\ell}(c)$ may take a simple form as explained above according to the switching costs of Player ℓ.

3.3. The folk theorem for the undiscounted payoffs game

With the long-run feasible and long-run individually rational sets, the statement of our Folk Theorem is identical to the one that appears in Aumann and Shapley (1994) for repeated games without switching costs:

Theorem 2. *The set of SPE is equal to*

$$SPE_{\infty}(c) = \mathcal{IR}(c) \cap \mathcal{F}(c).$$

Moreover, these SPEs are defined without a public randomization device.

Proof. See Appendix A.6. □

The result is a natural one, which combines the definitions of IR and Feasibility already defined in Dutta (1995) while relaxing his assumption of Full Dimensionality, which is necessary to establish the Folk Theorem with discounted payoff.

Let us analyze the two extreme cases. For small c, one obtains that the limiting set coincides with the SPE of the repeated game without switching costs. This is a radically different conclusion than Theorem 6 in Lipman and Wang (2009) for undiscounted repeated games with arbitrarily small switching costs. This is due to the fact that we consider the average of the switching costs instead of their total sum. For large c, Theorem 2 and the previous discussions lead directly to a characterization of the set of equilibrium payoffs.

Corollary 3. *When c goes to infinity, the set of SPE payoffs converges to $\overline{F} \cap \overline{V}$*

One key difference with the literature is that we allow heterogeneity of agents by allowing different switching costs matrices: some pay no switching costs (N^0), some pay switching costs among part of their actions ($N^1 \setminus N^{nf}$), and some pay switching costs between all their actions (N^{nf}). Hence, Corollary 3 is an extension of Chakrabarti (1990) when some players have also off-diagonal zeros in their switching costs matrix. The result is also related to Lipman and Wang (2009). If we assume that all players have no free-switching, i.e. $N^1 = N$, then $\overline{F}$ becomes the negative orthant below the feasible payoffs and $\overline{V}$ becomes

$W = \left\{ (x^\ell) \in \mathbb{R}^N, \text{ for all } \ell \in N, \ x^\ell \geqslant \overline{w}^\ell \right\}$. We obtain that $SPE_\infty(c)$ is the same set as in Theorem 3 in Lipman and Wang (2009) on discounted evaluations for large c.

Let us now investigate the consequence in terms of comparative statics for different costs. From Proposition 1 and Corollary 2, we have the following regularity for the set of subgame perfect equilibrium payoffs.

Corollary 4. *The correspondence $SPE_\infty(c)$ is semialgebraic, increasing in the sense of inclusion and decreasing for the order on $\mathbb{R}^n$.*

Hence, a higher cost can only lead to a decrease in the expected payoff. Moreover, the set of Pareto Efficient allocations is constant in c. Let $c_1 < c_2$ then for any payoff vector $x^*(c_2)$ of Player ℓ in $SPE_\infty(c_2)$, there exists a payoff vector $x^*(c_1)$ in $SPE_\infty(c_1)$ such that the payoff of Player ℓ under $x^*(c_1)$ is higher than under $x^*(c_2)$. Informally, a higher cost can not give Player ℓ any advantage since he can be punished to a lower level, and the higher the switching costs, the lower the possible equilibrium payoffs. It is in particular the case in the extreme case where only Player ℓ has some switching costs.

4. Discounted payoffs

Let us consider the discounted game with the parameter c. By definition, we know that it can be reformulated as a stochastic game where the state space is I (the set of pure action profiles) and deterministic transitions. It is therefore possible to apply the results of Dutta (1995). More formally, we consider a stochastic game where the state space is $K = I \cup \{\emptyset\}$ (to take into account the initial stage of the game with switching costs where no action has been played yet) and the action set is I in every state. The transition q kernel is defined such that for all (k, i') in $K \times I$, $q(i'|k)$ is the Dirac mass at i'. The payoff function $\widetilde{u}^\ell(k, i')$ to be equal to u if $k = \emptyset$ and $u^\ell(i') - cs^\ell_{k^\ell i'^\ell}$ otherwise.

This stochastic game satisfies immediately Assumptions (A1) and (A2) from Dutta (1995):

Assumption 1 *(A1).* The set of long-run feasible payoffs is independent of the state.

Assumption 2 *(A2).* The long-run individually rational level of Player $\ell \in N$ is independent of the original state.

In order to establish a Folk Theorem for discounted evaluation, it is usual to assume in addition that the one-shot game has *Full Dimensionality*. Adapted to the stochastic framework, we obtain the following assumption.

Assumption 3 *(A3).* The dimension of the long-run feasible set is equal to the number of players, i.e., $dim(\mathcal{F}(c)) = n$.

Notice that there is a separation between the case $c = 0$ where only the strategic interactions are taken into account and the case $c > 0$ where both the strategic interactions and the switching costs are playing a role.

Proposition 5. Let $c > 0$. If every player has some switching costs, i.e. $N^1 = N$, then the repeated game with switching cost c has Full Dimensionality.

Proof. See Appendix A.7. □

Clearly, Full Dimensionality can also be obtained if some players have no switching costs but the payoffs of the original game are sufficiently rich.

Theorem 3. *Let $c \geqslant 0$. If the repeated game with switching cost c has Full Dimensionality, then for all $x \in \mathcal{IR}(c) \cap \mathcal{F}(c)$, for all $\varepsilon > 0$, there exists $\underline{\delta} < 1$, s.t. for any $\delta \geqslant \underline{\delta}$, there is a perfect equilibrium strategy whose payoff is within ε of x. Equivalently*

$$\lim_{\delta \to 1} SPE_\delta(c) = \mathcal{IR}(c) \cap \mathcal{F}(c),$$

in the sense that the Hausdorff distance goes to 0.

The proof of Theorem 3 is a direct application of Theorem 9 in Dutta (1995). We need to check the three assumptions of the theorem: independence of the long-run feasible set, independence of the minmax value, and the Full Dimensionality of the feasible set. As mentioned previously, the first two hold by definition for games with switching costs whereas the third one is by assumption.

Let us comment on this result. First, the proof of Theorem 9 in Dutta (1995) relies on the use of public correlation devices to prevent some deviations of the players. It is an open question whether it is possible to prove the result without public randomization in the general framework of stochastic games or in our special framework of games with switching costs. Second, the assumption of Full Dimensionality is necessary. Since we considered a general model of games with switching costs where players may have different switching costs or even have no switching costs, the model contains in particular classical repeated games without switching costs. It was shown by Fudenberg and Maskin (1986) that the Folk Theorem may fail without Full Dimensionality. One can notice that if the original repeated game does not satisfy Full Dimensionality (as a one-shot game) but every player has some switching cost, then Full dimensionality is not satisfied for $c = 0$ but is for every strictly positive cost. Hence creating the possibility for a discontinuity of the set of SPE at $c = 0$.

Let us now investigate the consequence in terms of comparative statics for different costs. Under the assumptions of Theorem 3, we know by Theorem 2 and Theorem 3 that $SPE_\delta(c)$ is arbitrary close to $\mathcal{IR}(c) \cap \mathcal{F}(c)$. It follows by Corollary 4 that a higher cost can only have (asymptotically) a negative impact in terms of equilibrium payoff. More precisely, let $\varepsilon > 0$ and take two switching costs $c_1 < c_2$. There exists $\delta^* < 1$ such that for every $\delta > \delta^*$ and for any payoff vector $x^*_\delta(c_2)$ in $SPE_\delta(c_2)$, there exists a payoff vector $x^*_\delta(c_1)$ in $SPE_\delta(c_1)$ such that the payoff of Player ℓ under $x^*_\delta(c_1)$ is higher than under $x^*_\delta(c_2)$ up to an error ε. Informally, a higher cost can only give Player ℓ a smaller than ε benefit. As shown by Example 5, the monotonicity does not hold for a fixed discount factor. This also happens in the extreme case where only Player ℓ has some switching costs.

This result is obtained asymptotically and under the condition of Full Dimensionality. We now present a counter-example showing that the result fails for a fixed discount factor. It is an open question what is the impact on arbitrarily patient players without the Full Dimensionality assumption. As seen in the proof of Proposition 5, failure of Full Dimensionality is closely related to the absence of switching costs. Hence, assuming that the game does not satisfy the Full Dimensionality condition imposes some restrictions on the structure of the switching costs.

Example 5. *Positive benefit of switching costs for fixed discount factor.*

Consider the two-player game

	L	R
T	1, 1	0, 0
M	0, 1	4, 0
B	0, 1	2, 1

, $S^1 =$

	T	M	B
T	0	1	1
M	1	0	1
B	1	1	0

, and $S^2 =$

	L	R
L	0	0
R	0	0

.

In this example, Player 1 can guarantee a payoff of 2 with a switching cost of $c_1 = 2$ whereas it is impossible with a switching cost of $c_2 = 0$ (without switching cost). The intuition is the following. First along any SPE, Player 2 needs to obtain a payoff of 1. It follows that the only possibility along an SPE to obtain a payoff of 2 for Player 1 is to play almost surely the pair of actions (B, R). It is not possible without switching costs since Player 1 has an incentive to deviate to play Middle. With switching costs, Player 1 can commit to play Bottom. Indeed, the following profile of strategies is an SPE generating the payoffs vector $(2, 1)$: If nothing has been played or Bottom was played at the previous stage, Player 1 plays Bottom whereas Player 2 plays Right. If Middle or Top was played at the previous stage, Player 1 plays Top whereas Player 2 plays Left.

5. Finite games

The formalization of a game with switching costs as a stochastic game allows us to study games with finite horizons too. For finite horizon repeated games without switching cost, the Folk Theorem was established by Benoit and Krishna (1985). They introduced new weak conditions such that any feasible and individually rational payoff vector of the one-shot game can be approximated by the average payoff in a subgame perfect equilibrium of a repeated game with a sufficiently long horizon. The idea behind this condition relies on the fact that in the last stage of the game, the players always play a Nash Equilibrium of the one-shot game. The authors assume the existence of a good Equilibrium and a bad Equilibrium for each player (good and bad in terms of payoff to the player) and choose the equilibrium according to the history.

The SPE in the T-stage game is intuitively the following: play a cooperative profile for the main part of the game, and, approaching the end, finish the game depending on the past. If all players cooperate, then the good equilibrium is played in the last stages whereas if an agent deviates he is punished by everybody playing his bad Equilibrium. The result of Benoit and Krishna (1985) was extended in Marlats (2015) to stochastic games. Marlats (2015) uses the same idea to assume the existence of a good SPE for every player from every state (that can then be used to finish the T-stage game and to reward a player for having cooperated) and the existence of bad end-of-game for each player and each state (that can then be used to finish the T-stage game and to punish one player for having deviated). We adopt this approach to our model.

Definition 8. Let $k \in K = I \cup \{\emptyset\}$. We say that a payoff vector $x(k) \in \mathbb{R}^N$ is *a limiting SPE payoff vector at* k for the repeated game with switching cost c if there exists a sequence of SPE whose payoffs converges to $x(k)$, i.e. there exists $(\sigma_T)_{T\geqslant 1}$ such that:

- σ_T is an SPE in the T-stage game,
- $\gamma_T(\sigma_T, k)$ converges to $x(k)$.

The set of limiting SPE payoffs at k is denoted by $\Pi_c(k)$.

Each limiting SPE payoff is a payoff that can be approximately sustained and used for the last stages of a T-stage game since they are themselves obtained by a subgame perfect equilibrium. The next assumption states that there is a good one and bad ones, one for every player and every state.

Assumption 4 *(A4).* There exist $(n+1)^2$ payoff vectors denoted $x(k), x_{[1]}(k), \cdots, x_{[n]}(k)$ for $k \in K$ such that

- $x(k) \in \Pi_c(k)$ and for every $\ell \in N$ and $k \in K$, $x_{[\ell]} \in \Pi_c(k)$,
- $x^{\ell}_{[\ell]}(k) < x^{\ell}(k')$ for all $\ell \in N$ and $k, k' \in K$.
- $x(k) = x(k')$ for all k, k',

When this assumption holds, we obtain the Folk Theorem for these settings:

Theorem 4. *Let $c \geqslant 0$. If the repeated game with switching costs c has Full Dimensionality and satisfies Assumption A4, then for every $x \in \mathcal{IR}(c) \cap \mathcal{F}(c)$, for all $\varepsilon > 0$, there exists $T^* < \infty$ s.t. for any $T \geqslant T^*$, there is a subgame perfect equilibrium such that the payoff is within ε of x. Equivalently,*

$$\lim_{T \to \infty} SPE_T(c) = \mathcal{IR}(c) \cap \mathcal{F}(c),$$

in the sense that the Hausdorff distance goes to 0.

It is tempting to weaken the richness Assumption A4 to the weaker assumption of Benoit and Krishna (1985) that there exists in the one-shot game a good Nash equilibrium and a family of bad equilibria (one for each player). This is possible when the equilibria are strict and the costs are small.

Definition 9. A profile of strategies $(y^{\ell})_{\ell \in N}$ is a strict Nash equilibrium in the one-shot game if

$$\forall z^{\ell} \neq y^{\ell} \in \Delta(I^{\ell}),\ u^{\ell}(z^{\ell}, y^{-\ell}) < u^{\ell}(y^{\ell}, y^{-\ell}).$$

By definition, a strict Nash equilibrium is necessarily pure. Repeating a strict equilibrium independently of the past yields an SPE of the T stage game for small costs. Assumption A4 becomes:

Assumption 5 *(A5).* There exists $(n+1)$ payoffs vectors denoted $x, x_{[1]}, \cdots, x_{[n]}$ such that

- for every $\ell \in N$, $x_{[\ell]}$ and x are the Nash equilibrium payoffs in the one-shot game obtained by a strict Nash equilibrium,
- $x^{\ell}_{[\ell]} < x^{\ell}$ for all $\ell \in N$.

We obtain the following result for finitely repeated games and small costs.

Proposition 6. Assume that the one-shot game satisfies Assumption A5. Then there exists $c_0 > 0$ such that for all $0 \leqslant c < c_0$, the repeated game with switching cost c satisfies Assumption A4.

Proof. See Appendix A.8. □

One can deduce from Theorem 4 and Proposition 6 the following corollary:

Corollary 5. *If the repeated game with switching costs has Full Dimensionality and satisfies Assumption A5, then for every $x \in IR(c) \cap \mathcal{F}(c)$, for all $\varepsilon > 0$, there exists $T^* < \infty$ and $c_0 > 0$ s.t. for any $T \geqslant T^*$ and for any $c < c_0$, there is a subgame perfect equilibrium such that the payoff is within ε of x. Equivalently,*

$$\lim_{T \to \infty} SPE_T(c) = IR(c) \cap \mathcal{F}(c),$$

in the sense that the Hausdorff distance goes to 0.

This result is quite different from the conclusion in Lipman and Wang (2000) and in particular of Theorem 6. The authors proved there is the possibility to change completely the set of SPE with finite horizon with small costs. The key difference is the ratio between the cost and the weight of each period. They assume that the weight of the strategic interaction is the inverse of the number of periods whereas the switching cost has a constant weight of 1. This is not the case in our formulation.

Proposition 6 fails as shown by the following example if some of the equilibrium payoffs are obtained by a non-strict Nash equilibrium.

Example 6. *Counterexample for games with a non-strict Nash Equilibrium.*
Consider the following example with two players. Player 1 has three actions $\{T, M, B\}$ and pays some asymmetric cost given by the matrix

$S^1 =$

	T	M	B
T	0	1	0
M	1	0	0
B	1	1	0

Informally, switching to B is always costless whereas switching to T (resp. M) from another action has a unitary cost. Player 2 has three actions $\{L, C, R\}$ and pays no costs for switching. The payoff is given by

	L	C	R
T	0, 2	2, 0	0, 0
M	2, 0	0, 2	0, 0
B	1, 0	1, 0	3, 4

This one-shot game admits two Nash equilibria: a mixed equilibrium that yields a payoff of $(1, 1)$ (both players play uniformly respectively on T, M and on L, C) and a pure equilibrium (B, R) that yields a payoff of $(3, 4)$.

In particular, it satisfies the condition of Benoit and Krishna (1985). Nevertheless, as soon as some costs are introduced the mixed equilibrium disappears. Indeed, to keep Player 2 indifferent between his actions, Player 1 has to mix his actions with strictly positive weight on T and M hence ensuring a cost c. This diminishes his payoff to a payoff strictly below 1 and therefore B becomes a profitable deviation. We obtain that any T-stage game only admits (B, R) as an SPE.

Let us now investigate the consequence in terms of comparative statics for different costs. Under the assumptions of Theorem 4, we know by Theorem 2 and Theorem 4 that

$$\lim_{T\to\infty} SPE_T(c) = SPE_\infty(c).$$

It follows by Corollary 4 that a higher cost can only have asymptotically a negative impact in terms of equilibrium payoff like for discounted SPE.

More precisely, let $\varepsilon > 0$ and take two switching costs $c_1 < c_2$. There exists $T^* \in \mathbb{N}$ such that for every $T > T^*$ for any payoff vector $x^*_T(c_2)$ in $SPE_T(c_2)$, there exists a payoff vector $x^*_T(c_1)$ in $SPE_T(c_1)$ such that the payoff of Player ℓ under $x^*_T(c_1)$ is higher than under $x^*_T(c_2)$ up to an error ε.

Informally, a higher cost can only give Player ℓ a benefit up to ε. It is in particular the case where only Player ℓ has some switching costs. This result is obtained asymptotically and under the assumptions of Full-dimensionality and Assumption A4. We now present a counter-example showing that the result fails for a fixed length and even asymptotically without Assumption A4.

Example 7. *Counterexample when Assumption A4 is not fulfilled.*

Consider the two-player game

	L	R
T	3, 3	6, 1
B	2, 1	5, 2

, $S^1 =$

	T	B
T	0	1
B	1	0

, and $S^2 =$

	L	R
L	0	0
R	0	0

.

Here, Player 2 does not pay any switching costs. We will compare two values for c: $c_1 = 2$ and $c_2 = 0$. Notice that the one-shot game admits only one Nash equilibrium which is (T, L) due to strict dominance. For $c_2 = 0$, the unique SPE vector payoff for every length is $(3, 3)$.

For $c_1 = 2$, the unique SPE vector payoff is $\left(5 - \frac{4}{T}, 2 + \frac{1}{T}\right)$ if $T = 3k + 1$, and $(5, 2)$ otherwise. Let us focus on a game where the number of stages is a multiple of 3 and consider the following profile of strategies

- If Player 1 played Top at the previous stage and the remaining number of stages is not a multiple of 3: Player 1 plays Top and Player 2 plays Left.
- In all other cases: Player 1 plays Bottom and Player 2 plays Right.

One can check that this profile of strategies is indeed an SPE. When the length of the game is not a multiple of 3, the reasoning is similar apart from one variation when the game has $3k + 1$ stages. In this case, the players have an incentive to play in the first stage the one-shot Nash equilibrium before using the previous strategies. Hence, Player 1 incurs a one-time loss of 4 (2 due to the switching cost and $2 = 5 - 3$ due to the difference in stage payoffs at the first stage) whereas Player 2 incurs a one-time gain of $1 = 3 - 2$.

6. Final remarks

In this paper, we study finite and infinite horizon repeated games with switching costs. In each of the models, we prove a version of the Folk Theorem and provide a characterization of the set of equilibrium payoffs. Our work departs from previous papers by assuming a general structure for the switching cost and that the stage payoff is the weighted sum of the payoff coming from the strategic interaction and the switching cost. Thus, we are able to study how the equilibrium payoffs change with a possible change in the relative weights between the two types of payoffs. To the best of our knowledge, this is the most extensive study and most general study of repeated games with switching costs. Moreover, we note that our proof technique for some of the results (e.g. Theorem 1 and Proposition 5) relies on playing in "blocks" of pure actions of increasing lengths. Switching costs are paid only between the blocks, so their impact on the payoff vanishes as the block length increases. We deduce that our results can be easily extended to other switching cost models, such as games with an "Evergreen Clause" (Dutta, 2021), where the payoff is not given in stages where switches are made.

Several generalizations could be interesting. First, we assume that Player i's switching cost at stage t only depends on what this player played at $t-1$. One could define a more general model where the cost would depend on the action profile at stage $t-1$. The Folk Theorems and the characterization would extend without problem to this more general framework. The extension of the comparative static is not clear. Indeed, the game behind the long-run individually rational level is no longer single-controlled and we can not use the results of Tsodikovich et al. (2022). Another extension is to restrict the players to have partial observation of the action played. It will require them to use an estimate to decide if a player has deviated. Several authors have studied Folk Theorems under imperfect monitoring. For example, Renault and Tomala (2004) study the existence of communication equilibrium in repeated games with imperfect monitoring. The proof also relies on splitting the stages into blocks, therefore we conjecture that it is possible to extend our result to imperfect monitoring.

As it is customary in the literature, to prove the Folk Theorem for the discounted and finite case, we assumed the existence of a public correlation device (as well as assumption A4 in the finite case). At the same time, when every player has some strictly positive switching cost, we do not need to assume the Full Dimensionality of the stage game.

We conjecture that the existence of a public correlation device and Assumption A4 can both be relaxed by using the structure of the game and possibly weaker assumptions (see for example, Assumption A5 and Proposition 6), and leave this matter for future study.

We provide a characterization of the feasible set for patient players through an auxiliary one-shot game. When the switching costs are symmetric, the feasible payoff of this one-shot game is equal to the feasible set in the undiscounted repeated game with switching costs. Although this limits the applicability of our result, in fact, it generalizes previous works that assumed constant switching costs. We postulate that it is possible to obtain an alternative characterization in the asymmetric approach of the set by considering bigger one-shot games than the one presented in Section 2.2. The idea is that as our auxiliary game considers all possible average payoffs and switching costs of a combination of two stages of the game, we can consider combining three stages, or four stages, and so forth. The more stages we combine, the closer the resulting set to the one of the repeated games with asymmetric switching costs. Such characterization is not practical therefore we did not pursue this path in this paper and left for future research the search for a better approximation of the feasible set in this case.

Declaration of competing interest

The authors declare that they have no conflict of interest.

Data availability

No data was used for the research described in the article.

Appendix A. Proofs

A.1. Proof of Proposition 1

Decreasing for the order on $\mathbb{R}^n$: Let $c_1 < c_2$. Consider a long-term feasible payoff in $x_2 \in \mathcal{F}(c_2)$. By definition, there exists a profile of strategies σ that generates x_2. Consider the same profile of strategies in the game with cost c_1. We have for every $T \geqslant 1$,

$$\mathbb{E}_\sigma\left(\frac{1}{T}\sum_{t=1}^{T} u^\ell(i(t)) - c_2 \cdot s^\ell_{i^\ell(t-1)i^\ell(t)}\right) \leqslant \mathbb{E}_\sigma\left(\frac{1}{T}\sum_{t=1}^{T} u^\ell(i(t)) - c_1 \cdot s^\ell_{i^\ell(t-1)i^\ell(t)}\right).$$

The payoff for every player in this new game is higher since switching costs are smaller whereas the payoff from strategic interactions is equal. It follows by considering the subsequence $(T_k)_{k\geqslant 1}$ that generates x_2 that x_2 is smaller than the liminf on a subsequence on the right hence there is an accumulation point and hence a feasible payoff $x_1 \in \mathcal{F}(c_1)$ such that for all $\ell \in N$, $x_2^\ell \leqslant x_1^\ell$.

Lipschitz: Let $\|S\|_\infty = \max_{\ell\in N} \|S^\ell\|_\infty$. The correspondence is Lipschitz in the sense that for every $c_1, c_2 \in \mathbb{R}$, and for every $x_1 \in \mathcal{F}(c_1)$ there exists $x_2 \in \mathcal{F}(c_2)$ such that $|x_2 - x_1| \leqslant \|S\|_\infty |c_2 - c_1|$. Indeed, by definition, there exists a profile of strategies that generates x_1. Consider the same profile of strategies in the game with cost c_1. The stage payoffs of player ℓ are bounded above by $\|S^\ell\|_\infty |c_2 - c_1|$ and, as a consequence also the T-stage average payoffs. Consider a subsequence of length such that the T_k-average payoffs converge to x_1 for the cost c_1. One can consider a subsequence such that the average payoff for the cost c_2 converges to some x_2. We obtain by construction the correct inequality.

Increasing for inclusion: Let $c_1 < c_2$ and $x_1 \in \mathcal{F}(c_1)$. By construction, there exists a profile of strategies generating x_1 in the repeated game with switching cost c_1. Let us first notice the following: for every p natural number $x_1 \in \mathcal{F}(pc_1)$. It can be obtained by simply considering the actions induced by the profile of strategies generating x_1 and repeating p times each profile of pure actions before switching. When repeating, switching costs appear once every p stage. Thus, we obtained the same switching cost as this strategy profile induces for the coefficient c_1.

One can also alternate blocks where each action is repeated a different number of times. For example, for $1 < p < q$ natural numbers, by alternating between p blocks of size m and $q - p$ blocks of size $m + 1$, we obtain that $x_1 \in \mathcal{F}\left(\left(m + \frac{p}{q}\right) c_1\right)$. One can reach like that any multiple of c_1 by a rational number greater than 1. It is then possible to approach any cost c_2 larger than c_1 since $\mathcal{F}$ is Lipschitz.

A.2. Proof of Proposition 2

Convex: Consider $c \geqslant 0$, and two distinct payoffs $x_1, x_2 \in \mathcal{F}(c)$. It implies that one player has at least two actions. Without loss of generality, let us assume that it is Player 1 and let us denote his actions by 1 and 2. By definition, there exist two strategy profiles σ_1 and σ_2 such that the payoff under σ_1 is x_1 and the payoff under σ_2 is x_2. Let $\lambda \in (0, 1)$ and consider the following profile of strategies where at stage 1, Player 1 plays randomly his two actions with probabilities λ and $1 - \lambda$. Then depending on the action played, the players play for the rest of the game the profile of strategies σ_1 or the profile of strategies σ_2.

Closed: Let $c \geqslant 0$, and $(x_n)_{n\in\mathbb{N}}$ be a sequence of feasible payoffs in $\mathcal{F}(c)$ converging to x. By definition, there exists for every $n \in \mathbb{N}$, a profile of strategies σ_n that generates x_n. One can concatenate these profiles of strategies to obtain a strategy generating x by successively playing each of them with the proper length.

A.3. Proof of Theorem 1

The proof is decomposed into a preliminary lemma and then four parts. In the first lemma, we rewrite the payoff in a T-stage game under σ in a different way up to some small error.

Lemma 1. *For every $T \geqslant 1$ and for every profile of strategies σ, there exists a probability distribution*[3] π_T *in* $\Delta(I \times I)$ *such that*

$$\left|\mathbb{E}_{\pi_T}\left(g^{\ell}\right) - \gamma^{\ell}(\sigma)\right| \leqslant \frac{\|u^{\ell}\|_\infty}{T},\ \forall \ell \in N,$$

and

$$\|\pi_T(.,I) - \pi_T(.,I)\|_1 \leqslant \frac{2}{T},$$

where $\pi_T(.,I)$ (resp. $\pi_T(.,I)$) is the first (resp. second) marginal of π_T.

Proof. A profile of strategies induces a distribution over I^T. For every $t < T$, let π_T^t be the distribution of the couple $(t, t+1)$-coordinates and π_T^T be the distribution of the couple (T, T)-coordinates (the last coordinate repeated two times). We define π_T as the average of the distributions $(\pi_T^t)_{1\leqslant t\leqslant T}$. Let us check that π_T satisfies the first condition.

$$\begin{aligned}
g_c^{\ell}(\pi_T) := \mathbb{E}_{\pi_T}\left(g_c^{\ell}\right) &= \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{\pi_T^t}\left(g_c^{\ell}\right), \\
&= \frac{1}{T}\sum_{t=1}^{T-1}\mathbb{E}_{\sigma}\left(g_c^{\ell}(i_t, i_{t+1})\right) + \frac{1}{T}\mathbb{E}_{\sigma}\left(g_c^{\ell}(i_T, i_T)\right), \\
&= \mathbb{E}_{\sigma}\left(\frac{1}{T}\sum_{t=1}^{T-1}\left(\frac{1}{2}u^{\ell}(i_t) + \frac{1}{2}u^{\ell}(i_{t+1}) - c s^{\ell}_{i_t^{\ell} i_{t+1}^{\ell}}\right) + \frac{1}{T}u^{\ell}(i_T)\right), \\
&= \mathbb{E}_{\sigma}\left(\frac{1}{T}\sum_{t=1}^{T}\widetilde{u}^{\ell}(i_{t-1}, i_t) - \frac{1}{2T}u^{\ell}(i_1) + \frac{1}{2T}u^{\ell}(i_T)\right).
\end{aligned}$$

One obtains the third equality by splitting the sum into two parts and regroup the utilities obtained at stage t. Hence, $g_c^{\ell}(\pi_T)$ is equal to the expected payoff under σ up to $\frac{\|u^{\ell}\|_\infty}{T}$ proving the first point.

Let us now check that π_T satisfies the second condition. For every $t \geqslant 1$, let Q^t be the law of the t-coordinate under σ. We have

$$\forall i \in I,\ \pi_T(\{i\} \times I) - \pi_T(I \times \{i\}) = \frac{1}{T}\sum_{t=1}^{T}\pi_T^t(\{i\} \times I) - \pi_T^t(I \times \{i\}),$$

[3] As the strategy σ remains constant in the proof, we have decided not to explicitly show its dependence to simplify the notation.

$$= \frac{1}{T}\sum_{t=1}^{T-1}\left(Q^t(i) - Q^{t+1}(i)\right) + \frac{1}{T}(Q^T(i) - Q^T(i)),$$
$$= \frac{1}{T}(Q^1(i) - Q^T(i)).$$

Hence,

$$\|\pi_T(.,I) - \pi_T(I,.)\|_1 \leqslant \frac{1}{T}\|Q^1 - Q^T\|_1 \leqslant \frac{2}{T}. \quad \square$$

Writing a feasible payoff as the expectation of g_c. Let x be a feasible payoff. We prove that x can be expressed as the expected value of g_c for a well-chosen probability distribution in P. Since x is feasible, there exists a profile of strategies σ and a sequence of times $(T_k)_{k\geqslant 1}$ such that the sequence of payoffs in the T_k-stage game converges to x. Let us consider the subsequence $(T_k)_{k\geqslant 1}$ and the sequence $(\pi_{T_k})_{k\geqslant 1}$ associated to it by Lemma 1. Up to a second extraction, one can assume that $(\pi_{T_k})_{k\geqslant 1}$ converges to some $\pi \in \Delta(I\times I)$. We will still denote the subsequence as $(T_k)_{k\geqslant 1}$. By the second property in Lemma 1, we know that

$$\|\pi_{T_k}(.,I) - \pi_{T_k}(I,.)\|_1 \leqslant \frac{2}{T_k}.$$

Hence, at the limit $\|\pi(.,I) - \pi(I,.)\|_1 = 0$ and therefore π is in the set P. Moreover, the sequence $(g_c^\ell(\pi_{T_k}))_{k\in\mathbb{N}}$ converges to x^ℓ for all $\ell \in N$ by the first condition. By continuity of the payoff vector function on $\Delta(I\times I)$, it follows that $g_c(\pi) = x$ proving the result.

Proving that given a probability π in P, $g_c(\pi)$ is a feasible payoff. Reciprocally, given a distribution $\pi \in P$. Let us consider the Markov chain induced on I by the initial distribution obtained with the first marginal of π and the transition for all $i,i' \in I\times I$, $q(i'|i) = \pi(i,i')/\left(\sum_{i'}\pi(i,i')\right)$. Clearly, the first marginal of π is an invariant measure of the Markov chain q. If the Markov chain is irreducible and aperiodic, then it is well known that with probability one the time average is equal to the invariant measure. In particular, there exists one such play $(i_1^*,...,i_{N^1},...)$. Define the joint profile of strategies where the players follow this specific play (deterministically). This generates the correct payoff under g_c and therefore in the repeated game with switching payoffs. If the Markov chain is not irreducible, one can decompose the state space in ergodic classes (C_k), apply the previous reasoning on each class and use the convexity of the set of feasible payoffs.

The long-run feasible set is a subset of the feasible set for the auxiliary game with payoff g_c. The second statement is an immediate consequence of the first one since P is a subset of the set of all probability distributions $\Delta(I\times I)$ and $g_c(\Delta(I\times I)) = F(g_c)$.

If the cost matrices are symmetric, there is an equality. We now prove the third statement. Let $x \in F(g_c)$. We want to prove that it is possible to generate this payoff in the original game whenever all the S^ℓ are symmetric. By definition, there exists a probability distribution π over $I\times I$ such that $x = g_c(\pi)$. It follows that

$$x = \sum_{j\in(I\times I)^D}\pi(j)g_c(j).$$

Let us assume first that for every $j \in I\times I$, $\pi(j)$ is a rational number $\frac{p_j}{Q}$ with p_j an even number and Q a natural number. Then, one can generate the payoff vector x in the repeated game with payoff g^ℓ by fixing θ a natural number and playing successively: θp_j times the action profile $j = (i,i')$ for every $j \in I\times I$ and repeating from the top. To be well-defined, we need to choose an order on the profiles in $I\times I$. For the next computation, it is interesting to specify this order. Let D be the cardinal of $I\times I$ and $(j_d)_{d\in\{1,\cdots,D\}}$ an enumeration of the profiles in $I\times I$. With this notation, the strategy becomes: play θp_1 times the action profile $j_1 = (i_1,i_1')$, θp_2 times the action profile $j_2 = (i_2,i_2')$, $\cdots$, θp_D times the action profile $j_D = (i_D,i_D')$, repeat from the top.

By symmetry of s, one has for every $j = (i,i') \in I\times I$, $g_c^\ell((i,i')) = g_c^\ell((i',i))$ for all $\ell \in N$. Hence repeating j or alternating between $j = (i,i')$ and $j' = (i',i)$ where the two coordinates are interverted yield the same payoff. This allows us to turn back to the original game, by considering the following profile of strategies denoted σ' in the switching cost game: start with i_1 and then alternate between i_1 and i_1' for a total of θp_1 periods, switch to i_2 and then alternate between i_2 and i_2' for θp_2 periods, $\cdots$, switch to i_D and then alternate between i_D and i_D' for θp_D periods, repeat from the top.

Let us compute the payoff of this strategy. For every $d \in \{0,\ldots,D\}$, define $q_d = \sum_{d'\leqslant d}\theta p_{d'}$ with the convention that $q_0 = 1$ and let $T = \theta Q$. We have

$$\gamma_T^\ell(\sigma') = \frac{1}{T}\sum_{t=1}^{T} u^\ell(i(t)) - c\cdot s^\ell_{i^\ell(t-1)i^\ell(t)},$$
$$= \frac{1}{T}\sum_{d=1}^{D}\sum_{t=q_{d-1}+1}^{q_d} u^\ell(i(t)) - c\cdot s^\ell_{i^\ell(t-1)i^\ell(t)},$$
$$= \frac{1}{T}\sum_{d=1}^{D}\left(\frac{\theta p_d}{2}u^\ell(i_d) + \frac{\theta p_d}{2}u^\ell(i_d') - \frac{\theta p_d}{2}c\cdot s^\ell_{i_d^\ell,i_d'^\ell} - \left(\frac{\theta p_d}{2}-1\right)c\cdot s^\ell_{i_d'^\ell,i_d^\ell} - c\cdot s^\ell_{i_{d-1}'^\ell,i_d^\ell}\right),$$
$$= \frac{1}{T}\sum_{d=1}^{D}\left(\frac{\theta p_d}{2}\cdot u^\ell(i_d) + \frac{\theta p_d}{2}\cdot u^\ell(i_d') - c(\theta p_d - 1)s^\ell_{i_d^\ell,i_d'^\ell} - c\cdot s^\ell_{i_{d-1}'^\ell,i_d^\ell}\right),$$

$$= \frac{1}{T}\sum_{d=1}^{D}\left(\theta p_d g_c(i_d, i'_d) + c\cdot(s^{\ell}_{i^{\ell}_d, i'^{\ell}_d} - s^{\ell}_{i'^{\ell}_{d-1}, i^{\ell}_d})\right),$$
$$= x + c\cdot\frac{1}{T}\left(\sum_{d=1}^{D} s^{\ell}_{i^{\ell}_d, i'^{\ell}_d} - s^{\ell}_{i'^{\ell}_{d-1}, i^{\ell}_d}\right).$$

The right-term has an upper bound of $\frac{c\cdot D\cdot\|S\|_\infty}{T}$. By considering larger and larger θ, the limits are equal and x is indeed a feasible payoff in the repeated game with switching cost. The case where the weights are not rational numbers is obtained similarly by approximating the weights by rational numbers closer and closer along the game.

A.4. Proof of Proposition 3

The function $\overline{v}^{\ell}(c)$ is continuous and decreasing: The proofs for these results are standard. Both rely on extending properties satisfied on the initial utilities to the min max values. Fix a pair of profiles of actions $i(t-1)$ and $i(t)$. For every $\ell \in N$, and for every $0 \leqslant c_1 \leqslant c_2$:

$$|u^{\ell}(i(t-1), i(t)) - c_1 s_{i^{\ell}(t-1)i^{\ell}(t)} - u^{\ell}(i(t-1), i(t)) + c_2 s_{i^{\ell}(t-1)i^{\ell}(t)}| \leqslant \|S\|_\infty |c_2 - c_1|$$

and

$$u^{\ell}(i(t-1), i(t)) - c_2 s_{i^{\ell}(t-1)i^{\ell}(t)} \leqslant u^{\ell}(i(t-1), i(t)) - c_1 s_{i^{\ell}(t-1)i^{\ell}(t)}$$

These two inequalities can be extended to the expected T stage payoff limit and then to the minmax values with standard arguments.

The function $\overline{v}^{\ell}(c)$ is semialgebraic: The proof will rely on the link between the minmax value with undiscounted payoff and its equivalent for discounted payoff. This is a standard approach since the discounted evaluation is simpler to characterize. By Mertens and Neyman (1981), it is known that the first one is the limit of the second one when the players become patient. Let δ be a fixed discount factor. We consider the minmax discounted value associated with Player ℓ as a function of the previous action (or no action) of Player ℓ. It satisfies the following fixed point equation

$$\overline{v}^{\ell}_{\delta}(i, c) = \min_{y^{-\ell}\in\Sigma^{-\ell}} \max_{i^{\ell}\in I^{\ell}} (1-\delta)\left(u^{\ell}(y^{-\ell}, i^{\ell}) - c * s_{i,i^{\ell}}\right) + \delta\overline{v}^{\ell}_{\delta}(i^{\ell}, c).$$

Let $m = (m^{\ell}) + 1$ be the number of profiles of pure actions of Player ℓ plus one (for the initialization where no player has played yet). Denote by T^{ℓ} the operator from $\mathbb{R}\times\mathbb{R}^m$ to $\mathbb{R}^m$ defined by

$$\forall (c, x) \in \mathbb{R}\times\mathbb{R}^m,\ T^{\ell}(c, x) = \min_{y^{-\ell}\in\Sigma^{-\ell}} \max_{i^{\ell}\in I^{\ell}} \left(u^{\ell}(y^{-\ell}, i^{\ell}) - c * s_{i,i^{\ell}}\right) + x.$$

Then $\overline{v}^{\ell}_{\delta}(\cdot, c) = (1-\delta)\cdot f(\delta, c)$ where $f(\delta, c)$ is the unique solution of $f = T^{\ell}(c, \delta\cdot f)$. The operator T^{ℓ} is a semialgebraic function since its graph is a semialgebraic set. It follows that $f(\delta, c)$ and $\overline{v}^{\ell}_{\delta}(\cdot, c)$ are semialgebraic.

As a semialgebraic and bounded function in δ, we know that the limit when δ goes to 1 exits. Moreover, by Mertens and Neyman (1981), it is equal to the undiscounted minmax.

The function $\overline{v}^{\ell}(c)$ can be expressed as a first-order formula from semialgebraic functions, it is therefore semialgebraic.

A.5. Proof of Proposition 4

Let us consider that Player ℓ is the one punished and we want to show that there exists $\overline{c}^{\ell}$ s.t. for every $c \geqslant \overline{c}^{\ell}$, $\overline{v}^{\ell}(c) = \overline{w}^{\ell}$. First, it is clear that Player ℓ can always guarantee $\overline{w}^{\ell}$ by playing his pure maximin action. Second, let us define the joint strategy of the other players that plays as a function of the last action played by ℓ: if Player ℓ played i^{ℓ}, at the next stage Players $-\ell$ play the best-reply $y_{-\ell}(i^{\ell})$ to the pure action i^{ℓ}. We now obtain a decision problem controlled by Player ℓ. When the cost to switch becomes high, then at state i^{ℓ}, Player ℓ can either repeat i^{ℓ} and obtain a payoff larger than $\overline{w}^{\ell}$ with no cost or change to i'^{ℓ}, obtain some bounded profit (according to u^{ℓ}) and pay a huge cost making the global payoff arbitrary negative.

A.6. Proof of Theorem 2

First, we check that the set of equilibrium payoffs is a subset of $\mathcal{F}(c) \cap IR(c)$. Let σ^* be an SPE, then the T-stage average of stage payoffs converges with probability one. By dominated convergence, we therefore have

$$\mathbb{E}_{\sigma^*}\left(\liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \widetilde{u}^{\ell}(i(t-1), i(t))\right) = \lim_{T\to\infty} \mathbb{E}_{\sigma^*}\left(\frac{1}{T}\sum_{t=1}^{T} \widetilde{u}^{\ell}(i(t-1), i(t))\right). \quad \text{(A.1)}$$

Therefore, the undiscounted payoff under σ^* is indeed in $\mathcal{F}(c)$. Moreover, let $(\sigma')^{\ell}$ a deviation of Player ℓ which plays a best reply to $\sigma_*^{-\ell}$. By definition of an SPE and of the individually rational level, one has $\gamma^{\ell}(\sigma_*) \geqslant \gamma^{\ell}((\sigma')^{\ell}, \sigma_*^{-\ell}) \geqslant \overline{v}^{\ell}(c)$. Hence, the vector payoff is in $\mathcal{F}(c) \cap IR(c)$.

Let us now prove that any vector $x \in \mathcal{F}(c) \cap IR(c)$ can be obtained as an SPE payoff of the infinitely repeated game with switching cost. First, let us prove that x can be generated by a profile of pure strategies. By Theorem 1, we know that there exists a probability distribution π over pairs of actions such that $x = g_c(\pi) = \widetilde{u}(\pi)$. π generates a finite Markov chain on $\Delta(I \times I)$ where the second coordinate becomes the first one and a new state is chosen along the Markov chain q.

Assume for the moment that this Markov chain is irreducible, the time-average evaluation converges almost surely to x, hence, there exists at least one play $(i_t)_{t\geqslant 1}$ such that the payoff converges to x along this play. If the Markov chain is not irreducible, then there exists a partition of $I \times I$ in sets C_k, some real numbers $\mu_k \in [0,1]$ and some payoffs vector $x_k \in \mathbb{R}^N$ that satisfies the following. First, x is the convex combination of $(x_k)_k$ with weight $(\mu_k)_k$. Second, for every k, there exists a play $(i_t)_{t\geqslant 1}$ such that $i_1 \in C_k$ and the payoffs along the play converges to x_k. It is possible to reconstruct a play generating x by combining adequately these individual plays.

We only defined a path of the strategy that we will call the main path and denote by σ_m. Since players are playing pure, any deviation is immediately observed. Moreover, for every $\ell \in N$ and every $\varepsilon > 0$, there exists a profile of strategies of other players than ℓ, $\sigma^{-\ell}_{\ell,\varepsilon}$, such that

$$\sup_{\sigma^\ell} \mathbb{E}_{\sigma^\ell,\sigma^{-\ell}_{\ell,\varepsilon}} \left(\liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \widetilde{u}^\ell(i(t-1), i(t)) \right) \leqslant \overline{v}^\ell(c) + \varepsilon.$$

Contrary to the strategy σ_m that is playing only pure, these strategies may be mixed. Notice that by the special structure of the transition function, the strategy σ_m (resp. $\sigma^{-\ell}_{\ell,\varepsilon}$ for every $\ell \in N$) guarantees the same payoffs in the long-run after any history. In the previous formula, we control the liminf evaluation. In order to construct our strategies, we need to replace the liminf evaluation with an explicit finite-stage evaluation. We show more precisely that for every $\varepsilon > 0$ and for every minimal time T_{min}, there exists a common time $T(\varepsilon, T_{\min})$ such that all the players can be punished in the $T(\varepsilon, T_{\min})$-stage game down to their long run individually rational level (with an approximation of 2ε). The precise lemma (Lemma 2) and its proof are delayed after the end of this proof. Define for every $n \geqslant 1$, $\varepsilon_n = \frac{1}{n}$. We can construct a sequence of times $(T_n)_{n\geqslant 1}$ as follows: $T_1 = 1$, for every $n \geqslant 2$, $T_{n+1} = \min\left(T(\varepsilon_n), \frac{2(M+c)T_n}{\varepsilon_n}\right)$. Let us call the block of stage $T_n + 1, ..., T_{n+1}$ as block B_n. We decompose the strategy by defining it by block as follows:

- Inside the block B_1, every player follow σ_m,
- Inside the block B_n, there are three possibilities depending on B_{n-1}:
  - Players were supposed to play the main path and nobody deviated: Continue the cooperative play σ_m (from where it was stopped).
  - Players were supposed to play the main path and a set of agents $\mathcal{A}$ has deviated: punish the player with a lower index in $\mathcal{A}$ in B_n.
  - Players were supposed to punish someone: Return to the main path (independently of the existence of a deviator or not inside the block B_{n-1}).

By construction of the sizes of the blocks, we know that the average payoff at the end of block B_n is approximately equal to the average payoff inside the last block.

We now show that this profile of strategies is an SPE. First, if nobody deviates then by construction the undiscounted payoff under the profile of strategies is equal to x. Let us now consider an off-path trajectory. After any path, we need only to consider two cases:

- If Player i only deviates a finite number of times, then his payoff in the long run is the one defined by σ_m and he has no gain.
- If Player i deviates an infinite number of times, then there is an infinite number of stages n such that its payoff is smaller than $\overline{v}^\ell(c) + 3\varepsilon_n$. He obtained his minmax payoff in the long run, hence the deviation is not profitable.

Furthermore, a punisher does not have any incentive to deviate on the punishing block since it has no consequences on the long-run payoff. We can now state and prove Lemma 2.

Lemma 2. *Let $\varepsilon > 0$ and $T_{min} \geqslant 1$. There exists $T(\varepsilon, T_{\min})$ such that for all $\ell \in N$,*

- $T(\varepsilon, T_{\min}) \geqslant T_{min}$,
- *for all σ^ℓ strategy of Player $\ell \in N$,*

$$\mathbb{E}_{\sigma^\ell,\sigma^{-\ell}_{\ell,\varepsilon}} \left(\frac{1}{T(\varepsilon, T_{\min})} \sum_{t=1}^{T(\varepsilon, T_{\min})} \widetilde{u}^\ell(i(t-1), i(t)) \right) \leqslant \overline{v}^\ell(c) + 2\varepsilon.$$

Proof. Denote by M a bound on the stage payoff and $\|S\|_\infty$ the maximal switching cost (before multiplication by c). Let us prove the result by contradiction. Assume that there exists $\varepsilon > 0$ and $T_{\min} \geqslant 1$ such that for all T, there exists $\ell(T) \in N$ such that either $T < T_{min}$, or there exists $\sigma^{\ell(T)}(T)$ a strategy of Player $\ell(T)$ such that

$$\mathbb{E}_{\sigma^\ell,\sigma^{-\ell}_{\ell,\varepsilon}}\left(\frac{1}{T}\sum_{t=1}^{T}\widetilde{u}^\ell(i(t-1),i(t))\right) > \overline{v}^\ell(c)+2\varepsilon.$$

The dependence of σ^ℓ in T was ignored in the previous equation for readability. Consider $T \geqslant \frac{2(M+c\|S\|_\infty)}{\varepsilon}$ and $\ell^* = \ell(T)$. Consider the following strategy of Player ℓ^*: repeat on block of T stages, the strategy $\sigma^{\ell(T)}(T)$. We denote this strategy as $\sigma^{\ell^*}_*$.

By definition, the only difference between the payoff on two blocks is the first stage where the players pay a different cost at stage 1 and stage $kT+1$, but since stage payoffs are bounded by $M + c\|S\|_\infty$ and the weight of one stage is $\frac{1}{T}$, we have for every history,

$$\begin{aligned}\mathbb{E}_{\sigma^{\ell^*}_*,\sigma^{-\ell^*}_{\ell^*,\varepsilon}}\left(\frac{1}{T}\sum_{t=kT+1}^{T(k+1)}\widetilde{u}^\ell(i(t-1),i(t))\right) &> \overline{v}^\ell(c)+2\varepsilon-\frac{M+c\|S\|_\infty}{T},\\ &> \overline{v}^\ell(c)+3/2\varepsilon.\end{aligned}$$

For every $t' \in \mathbb{R}$, let $n = \lfloor t'/T \rfloor$, then the mean average between 1 and t' is bounded below by the mean average between 1 and nT minus a possible loss of maximum $T(M + c\|S\|_\infty)$ on the last block. Hence,

$$\begin{aligned}\mathbb{E}_{\sigma^{\ell^*}_*,\sigma^{-\ell^*}_{\ell^*,\varepsilon}}\left(\frac{1}{t'}\sum_{t=1}^{t'}\widetilde{u}^\ell(i(t-1),i(t))\right) &> \overline{v}^\ell(c)+3/2\varepsilon-\frac{T(M+c\|S\|_\infty)}{t'},\\ &> \overline{v}^\ell(c)+\varepsilon.\end{aligned}$$

Hence, we obtained a strategy of Player ℓ^* that guarantees strictly more than his minmax value contradicting the definition of $\overline{v}^\ell(c)$. □

A.7. Proof of Proposition 5

Assume that every player has a switching cost. Formally, for every $\ell \in N$, there exists $i^\ell_1, i^\ell_2 \in I$ such that $s^\ell_{i^\ell_1,i^\ell_2} > 0$. Let us prove that the game has Full Dimensionality. Denote by $\overline{x}$ the vector payoff obtained by playing each profile of pure strategies with equal weight. For every $\ell \in N$, there exists $\varepsilon > 0$ such that the payoff vector $\overline{x} - \varepsilon e_\ell$ is a feasible payoff, where e_ℓ is the unitary vector with coordinate ℓ. Let us recall that it is possible to obtain $\overline{x}$ by playing large blocks of each action profile in a cycling way and switching more and more rarely. Let $\ell \in N$ and a profile of strategy $i^{-\ell}$ for the other players. Consider a cycle such that the profile $(i^\ell_2, i^{-\ell})$ is the successor in the cycle of the profile $(i^\ell_1, i^{-\ell})$ generating $\overline{x}$. We can now define the alternative play where instead of playing these two profiles successively, all players except ℓ follow $i^{-\ell}$ and Player ℓ alternates between i^ℓ_1 and i^ℓ_2. By construction, the payoff of Player ℓ' is $\overline{x}^{\ell'}$ for every $\ell' \neq \ell$ and the payoff of Player ℓ is $\overline{x}^\ell - \frac{1}{\sharp I!} s^\ell_{i^\ell_1,i^\ell_2}$. Since this is true for every Player $\ell \in N$, the game with switching cost has Full Dimensionality.

A.8. Proof of Proposition 6

By Assumption A5, each payoff is attained by a strict Nash equilibrium. Let m be the minimal loss obtained by a player by deviating from one of these Nash equilibria to a pure strategy. Since they are strict Nash-equilibria, m is strictly positive. Let $\|S\|_\infty$ be the maximal switching cost, let $c_0 \leqslant \frac{m}{\|S\|_\infty}$ and $c \leqslant c_0$. Let us pick one of the payoff vectors, x and the strict Nash-equilibrium generating it denoted by $(i^\ell_*)_{\ell\in N}$ (pure since it is a strict Nash equilibrium). Let us show that repeating $(i^\ell_*)_{\ell\in N}$ independently of the past is an SPE of the T-stage game with cost c_0. By the one-shot deviation principle, it is sufficient to check that after any history there is not a profitable deviation. Let $\ell \in N$ and i^ℓ_1 be the previous action played by Player ℓ. We compare the stage payoff today to play i^ℓ_2 instead of i^ℓ_*:

$$\begin{aligned}u^\ell(i^\ell_*, i^{-\ell}_*) - c\cdot s^\ell_{i^\ell_1,i^\ell_*} &\geqslant u^\ell(i^\ell_*, i^{-\ell}_*) - c\cdot\|S\|_\infty \geqslant u^\ell(i^\ell_2, i^{-\ell}_*) + m - c\cdot\|S\|_\infty,\\ &\geqslant u^\ell(i^\ell_2, i^{-\ell}_*) \geqslant u^\ell(i^\ell_2, i^{-\ell}_*) - c s^\ell_{i^\ell_1,i^\ell_2}.\end{aligned}$$

Hence, there is a loss in payoff today. Moreover, there is an additional potential loss tomorrow since Player ℓ has to pay the cost from i^ℓ_2 to i^ℓ_*. There is no one-shot profitable deviation. We have indeed an SPE.

A.9. Proof of Example 7

We will compare two values for c: $c_1 = 2$ and $c_2 = 0$. Notice that the one-shot game admits only one Nash equilibrium which is (T, L) due to strict dominance. It implies that the only Nash-equilibrium payoff for c_2 is equal to $(3, 3)$.

We now study the game with switching cost c_1. There is a unique SPE where the actions played depend on the last action played and on the number of stages remaining modulo 3. Consider the following profile of strategies,

- If Player 1 played Top at the previous stage and the remaining number of stages is not a multiple of 3: Player 1 plays Top and Player 2 plays Left.
- In all other cases: Player 1 plays Bottom and Player 2 plays Right.

Notice that if the remaining number of stages is a multiple of three, then Player 1 plays Bottom and Player 2 plays Right independently of the past. It is therefore sufficient to focus on the last three stages, check that it is the unique SPE, and then do an induction. Let $T = 3$ and let us prove that we described the unique SPE.

Last stage: $t = 3$

- If Player 1 plays Top at the second stage, then the payoff at stage 3 is

$$\frac{1}{3}\begin{pmatrix} 3,3 & 6,1 \\ 0,1 & 3,2 \end{pmatrix}$$

Hence, we have the same domination as in the game without switching costs.
- If Player 1 plays Bottom at the second stage then the payoff at stage 3 is

$$\frac{1}{3}\begin{pmatrix} 1,3 & 4,1 \\ 2,1 & 5,2 \end{pmatrix}$$

The domination is inverted: Bottom now strictly dominates Top. After eliminating Top, Right dominates Left for Player 2.

Second to last stage: $t = 2$

- If Player 1 plays Top at the first stage, then the payoff at stage 2 is obtained based on the previous paragraph to

$$\frac{1}{3}\begin{pmatrix} 3,3 & 6,1 \\ 2,1 & 5,2 \end{pmatrix} - \frac{1}{3}\begin{pmatrix} 0,0 & 0,0 \\ 2,0 & 2,0 \end{pmatrix} + \frac{1}{3}\begin{pmatrix} 3,3 & 3,3 \\ 5,2 & 5,2 \end{pmatrix} = \frac{1}{3}\begin{pmatrix} 6,6 & 9,4 \\ 5,3 & 8,4 \end{pmatrix}$$

Hence, Top is dominating Bottom.
- If Player 1 plays Bottom at the first stage, then the payoff at stage 2 is obtained based on the previous paragraph to

$$\frac{1}{3}\begin{pmatrix} 3,3 & 6,1 \\ 2,1 & 5,2 \end{pmatrix} - \frac{1}{3}\begin{pmatrix} 2,0 & 2,0 \\ 0,0 & 0,0 \end{pmatrix} + \frac{1}{3}\begin{pmatrix} 3,3 & 3,3 \\ 5,2 & 5,2 \end{pmatrix} = \frac{1}{3}\begin{pmatrix} 4,6 & 7,4 \\ 7,3 & 10,4 \end{pmatrix}$$

Bottom is indeed dominating Top. After eliminating Top, Right dominates Left for Player 2.

Initial stage: $t = 1$

Since no action has been played, there is no cost at the first stage, hence based on the previous paragraphs, one obtains the following matrix game:

$$\frac{1}{3}\begin{pmatrix} 3+2\cdot 3, 3+2\cdot 3 & 6+2\cdot 3, 1+2\cdot 3 \\ 2+2\cdot 5, 1+2\cdot 2 & 5+2\cdot 5, 2+2\cdot 2 \end{pmatrix} = \frac{1}{3}\begin{pmatrix} 9,9 & 12,7 \\ 12,5 & 15,6 \end{pmatrix}$$

Hence, the only equilibrium is for Player 1 to play Bottom and for Player 2 to play Right.

We want to show something stronger. Let us now assume that we are not looking at the first stage of a game but at some stage after a history. Let us check that even in the worst case where action Top was played at the previous stage, it is still the unique Nash equilibrium. If Top has been played, one obtains the following matrix game:

$$\frac{1}{3}\begin{pmatrix} 3+2\cdot 3, 3+2\cdot 3 & 6+2\cdot 3, 1+2\cdot 3 \\ 0+2\cdot 5, 1+2\cdot 2 & 3+2\cdot 5, 2+2\cdot 2 \end{pmatrix} = \frac{1}{3}\begin{pmatrix} 9,9 & 12,7 \\ 10,5 & 13,6 \end{pmatrix}$$

Hence, the only equilibrium is still for Player 1 to play Bottom and for Player 2 to play Right.

One can check that the profile of strategies is indeed an SPE equilibrium by extending the arguments to any length.

References

Akerlof, G.A., Yellen, J.L., 1985a. Can small deviations from rationality make significant differences to economic equilibria? Am. Econ. Rev. 75 (4), 708–720.

Akerlof, G.A., Yellen, J.L., 1985b. A near-rational model of the business cycle, with wage and price inertia. Q. J. Econ. 100, 823–838.

Ashkenazi-Golan, Galit, Flesch, János, Predtetchinski, Arkadi, Solan, Eilon, 2022. Existence of equilibria in repeated games with long-run payoffs. Proc. Natl. Acad. Sci. USA 119 (11), e2105867119. https://doi.org/10.1073/pnas.2105867119. https://www.pnas.org/doi/abs/10.1073/pnas.2105867119.

Aumann, Robert J., Shapley, Lloyd S., 1994. Long-term competition—a game-theoretic analysis. In: Essays in Game Theory. Springer, pp. 1–15.

Avenhaus, Rudolf, Von Stengel, Bernhard, Zamir, Shmuel, 2002. Inspection games. In: Handbook of Game Theory with Economic Applications, vol. 3, pp. 1947–1987.

Beggs, A., Klemperer, P., 1992. Multi-period competition with switching costs. Econometrica, 651–666.

Benoit, Jean-Pierre, Krishna, Vijay, 1985. Finitely repeated games. Econometrica 53 (4), 905–922. ISSN 00129682, 14680262.

Chakrabarti, S.K., 1990. Characterizations of the equilibrium payoffs of inertia supergames. J. Econ. Theory 51 (1), 171–183.

Darlington, Matthew, Glazebrook, Kevin D., Leslie, David S., Shone, Rob, Szechtman, Roberto, 2022. A stochastic game framework for patrolling a border. arXiv preprint. arXiv:2205.10017.
Dutta, Prajit K., 1995. A folk theorem for stochastic games. J. Econ. Theory 66 (1), 1–32.
Dutta, Prajit K., 2021. Compromise is key in infinitely repeated bargaining with an evergreen clause. J. Econ. Theory 193, 105170.
Filar, J.A., 1981. Ordered field property for stochastic games when the player who controls transitions changes from state to state. J. Optim. Theory Appl. 34 (4), 503–515.
Filar, J.A., 1980. Algorithms for solving some undiscounted stochastic games. Ph.D. Thesis.
Filar, J.A., Raghavan, T.E.S., 1984. A matrix game solution of the single-controller stochastic game. Math. Oper. Res. 9 (3), 356–362.
Filar, J.A., Schultz, Todd A., 1986. The traveling inspector model. OR Spektrum 8 (1), 33–36.
Filar, Jerzy, 1985. Player aggregation in the traveling inspector model. IEEE Trans. Autom. Control 30 (8), 723–729.
Friedman, James W., 1971. A non-cooperative equilibrium for supergames. Rev. Econ. Stud. 38 (1), 1–12.
Friedman, James W., 1974. Non-cooperative equilibria in time-dependent supergames. Econometrica, 221–237.
Friedman, James W., 1990. A modification of the Folk theorem to apply to time-dependent supergame. Oxf. Econ. Pap. 42 (2), 317–335. https://doi.org/10.1093/oxfordjournals.oep.a041949. ISSN 0030-7653.
Fudenberg, Drew, Maskin, Eric, 1986. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 54 (3), 533–554.
Fudenberg, Drew, Yamamoto, Yuichi, 2011. The folk theorem for irreducible stochastic games with imperfect public monitoring. J. Econ. Theory 146 (4), 1664–1683.
Guney, Begum, Richter, Michael, 2018. Costly switching from a status quo. J. Econ. Behav. Organ. 156, 55–70.
Guney, Begum, Richter, Michael, 2022. Games with switching costs and endogenous references. Theor. Econ. 17 (2), 617–650.
Hörner, Johannes, Sugaya, Takuo, Takahashi, Satoru, Vieille, Nicolas, 2011. Recursive methods in discounted stochastic games: an algorithm for $\delta \to 1$ and a folk theorem. Econometrica 79 (4), 1277–1318.
Jørgensen, Steffen, Martín-Herrán, Guiomar, Zaccour, Georges, 2010. Dynamic games in the economics and management of pollution. Environ. Model. Assess. 15 (6), 433–467.
Klemperer, Paul, 1987. Markets with consumer switching costs. Q. J. Econ. 102 (2), 375–394.
Klemperer, Paul, 1995. Competition when consumers have switching costs: an overview with applications to industrial organization, macroeconomics, and international trade. Rev. Econ. Stud. 62 (4), 515–539.
Lipman, B.L., Wang, R., 2009. Switching costs in infinitely repeated games. Games Econ. Behav. 66 (1), 292–314.
Lipman, B.L., Wang, R., 2000. Switching costs in frequently repeated games. J. Econ. Theory 93 (2), 149–190.
Marlats, Chantal, 2015. A folk theorem for stochastic games with finite horizon. Econ. Theory 58 (3), 485–507. ISSN 09382259, 14320479.
Maschler, Michael, Solan, Eilon, Zamir, Shmuel, 2013. Game Theory. Cambridge University Press, Cambridge.
Mertens, J-F., Neyman, Abraham, 1981. Stochastic games. Int. J. Game Theory 10 (2), 53–66.
O'Neill, Barry, 1994. Game theory models of peace and war. In: Handbook of Game Theory with Economic Applications, vol. 2, pp. 995–1053.
Rass, S., Rainer, B., 2014. Numerical computation of multi-goal security strategies. In: International Conference on Decision and Game Theory for Security. Springer, pp. 118–133.
Rass, S., Konig, S., Schauer, S., 2017. Defending against advanced persistent threats using game-theory. PLoS ONE 12 (1), e0168675.
Renault, Jérôme, Tomala, Tristan, 2004. Communication equilibrium payoffs in repeated games with imperfect monitoring. Games Econ. Behav. 49 (2), 313–344.
Schoenmakers, G., Flesch, J., Thuijsman, F., Vrieze, O.J., 2008. Repeated games with bonuses. J. Optim. Theory Appl. 136 (3), 459–473.
Tsodikovich, Yevgeny, Venel, Xavier, Zseleva, Anna, 2022. The regularity of the value function of repeated games with switching costs. Math. Oper. Res. 48 (4), 1899–1905.
Tsodikovich, Yevgeny, Venel, Xavier, Zseleva, Anna, 2024. The price of history-independent strategies in games with inter-temporal externalities. Dyn. Games Appl. https://doi.org/10.1007/s13235-024-00555-w. ISSN 2153-0793.
Yavuz, Mesut, Jeffcoat, David, 2007. An analysis and solution of the sensor scheduling problem. In: Advances in Cooperative Control and Optimization. Springer, pp. 167–177.