



HAL
open science

Folk Theorems in Repeated Games with Switching Costs

Yevgeny Tsodikovich, Xavier Venel, Anna Zseleva

► **To cite this version:**

Yevgeny Tsodikovich, Xavier Venel, Anna Zseleva. Folk Theorems in Repeated Games with Switching Costs. 2022. hal-03888188

HAL Id: hal-03888188

<https://hal.science/hal-03888188>

Preprint submitted on 7 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Folk Theorems in Repeated Games with Switching Costs*

Yevgeny Tsodikovich[†], Xavier Venel[‡], Anna Zseleva[§]

December 7, 2022

Abstract

We study how switching costs affect the sub-game 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.

Journal classification numbers: C72, C73

Keywords: Switching Costs, Repeated Games, Folk Theorem, Stochastic Games, Non-Zero-sum games.

*The authors wish to thank Stephane Gaubert, Chantal Marlats, Guillaume Vigerat for their highly valuable comments. YT acknowledges the support of the Israel Science Foundation grants #2566/20, #1626/18 & #448/22. XV acknowledges the support of the Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni.

[†]Department of Economics, Bar Ilan University, Israel. e-mail: yevgets@gmail.com

[‡]Dipartimento di Economia e Finanza, LUISS, Viale Romania 32, 00197 Rome, Italy, xvenel@luiss.it, <https://orcid.org/0000-0003-1150-9139>

[§]Maastricht University, School of Business and Economics, Dept. of Quantitative Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands, e-mail: anna.zseleva@maastrichtuniversity.nl

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.

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 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 to change 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 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 some 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 switching costs (more precisely, 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. The symmetry of the switching costs is essential for this result to hold. 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.

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]. In fact, under the same conditions of Marlats [2015], we 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)* are 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 the analysis of 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, upto an error of ε)*. On the contrary, in cases that 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 to 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 games.

In Section 3 we discuss the scenario where the payoffs are undiscounted and present our main results for this case. 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 counter-examples 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 to 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 in the study of their model [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. They 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, when considering the infinitely repeated game with time-averaging, infinite switching leads to an infinitely negative (time-averaged) payoff. As a consequence, we obtain substantially different results. We discuss these differences along our results.

In Lipman and Wang [2000], they showed that the introduction of switching cost has two effects and that both can lead the result: a switching cost may serve as 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 the previous 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.

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

We note that under some algebraic manipulation, this model (and games with switching costs in general) can describe other situations too. 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.

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 payoff that is feasible (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 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 Sub-game Perfect Equilibria (SPE) by using the two following ideas: punish only for a finite time instead of forever, and punish for longer and longer times 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 Dutta [1995], who established the main Folk Theorem for stochastic games. 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 (see Assumption A4) on the limit set of finite SPE payoffs.

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, it changes 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, it plays a role in the second part of the story, i.e. the minimax payoffs. The presence of cost 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 afterwards. 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 of 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 multiplicity of equilibria in a setting where without switching costs there is a

unique equilibrium, 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 non-zero-sum 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 $I = \prod_{\ell \in N} I^\ell$ to \mathbb{R} and $S^\ell = (s_{ij}^\ell)$ 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 $c \geq 0$.

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 $\tilde{u}^\ell(i(t-1), i(t)) := u^\ell(i(t)) - c \cdot s_{i^\ell(t-1) i^\ell(t)}^\ell$, so Player ℓ is penalized for switching the previous action $i^\ell(t-1)$ to the action $i^\ell(t)$ by $c \cdot s_{i^\ell(t-1) i^\ell(t)}^\ell$. 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 $\tilde{u}^\ell(i(0), i(1)) = u^\ell(i(1))$.

We assume that for all i, j and for all $\ell \in N$, $s_{ij}^\ell \geq 0$ and $s_{ii}^\ell = 0$. 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 $\tilde{u}^1 \neq -\tilde{u}^2$. 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 Sub-game 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 Sub-game Perfect Equilibrium (SPE). Informally, in these equilibria, after each history the continuation strategies are best-responses to each other and are equilibrium strategies in the sub-game that starts at this point. These all are formally defined in the rest of this subsection.

2.1.1 Finite Game of Length T

The game is played for $T \geq 1$ 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

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

For $T = 1$ this is simply the one-shot game without switching costs ($\tilde{u}^\ell = u^\ell$), 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 *individual 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 $F(u) = \text{Conv}(\{u(i), i \in I\})$. 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:

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

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

$$IR = \{(y^\ell)_{\ell \in N} \in \mathbb{R}^n, \text{ for all } \ell \in N, y^\ell \geq \bar{v}^\ell\}.$$

If $n = 2$, then \bar{v}^ℓ 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 sub-game starting after each history of length $t > 1$ includes an initial state and possible payment of switching costs, while the original game is defined 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 sub-game 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 Sub-game 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 ,

$$\gamma_T^\ell((\sigma^\ell, (\sigma_*)^{-\ell}), h_t) \leq \gamma_T^\ell(\sigma_*, h_t),$$

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

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

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 Sub-game 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 ,*

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

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

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

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 sub-game 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 Sub-game 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. For every $\ell \in N$, the following limit exists with probability 1 under σ_* :*

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

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 ,

$$\gamma^\ell((\sigma^\ell, (\sigma_*(h))^{-\ell})) \leq \gamma^\ell(\sigma_*(h)). \quad (4)$$

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.

When $c = 0$, the above definitions reduce to the regular finite and infinite games without switching costs, with the vector of utilities $(u^\ell)_{\ell \in N}$. The aim of this article is 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 $J = \prod_{\ell \in N} J^\ell$. We define naturally the projections of each $j \in J$ on the first and second coordinate 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

$$g_c^\ell(j) := \frac{1}{2} (u^\ell(j_1) + u^\ell(j_2)) - c \cdot s_{j_1 j_2}^\ell. \quad (5)$$

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 payoff 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

$$\begin{pmatrix} 2, 1 & 0, 0 \\ 0, 0 & 1, 2 \end{pmatrix}, \quad S^1 = \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad S^2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

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

| | LL | LR | RL | RR |
|------|-----------------|-------------------------|-----------------------|-----------------|
| TT | 2, 1 | 1, 0.5 - c | 1, 0.5 - c | 0, 0 |
| TB | 1 - 2 c , 0.5 | 1.5 - 2 c , 1.5 - c | -2 c , - c | 0.5 - 2 c , 1 |
| BT | 1 - c , 0.5 | - c , - c | 1.5 - c , 1.5 - c | 0.5 - c , 1 |
| BB | 0, 0 | 0.5, 1 - c | 0.5, 1 - c | 1, 2 |

The payoffs and feasible sets are shown in Figure 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 c s. For example, we study how the feasible set of payoffs change 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 $L(c) : \mathbb{R} \rightarrow \mathbb{R}^n$ is decreasing for the order on \mathbb{R}^n if:*

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

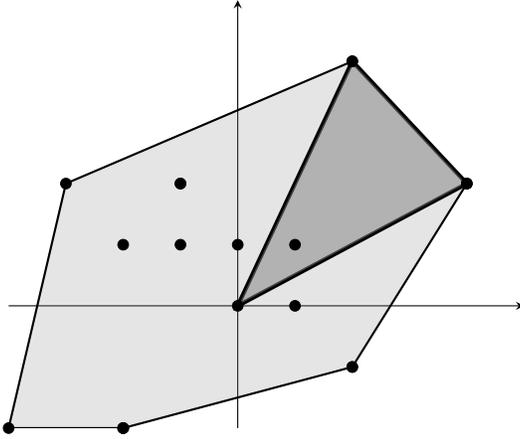


Figure 1: The feasible set of the standard one-shot “Battle of the Sexes” game (dark grey) and the feasible set of the auxiliary one-shot game for $c = 1$ (dark and light grey), discussed in Example 1. Dots correspond to payoffs obtained by pure actions in this one-shot game.

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 grey area (“high c ”) and the dark grey area (“low c ”) of Figure 1.

3 Undiscounted Payoffs

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

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. Moreover, we show that every vector in the feasible set can be reached by an SPE without public randomization.

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]. In order 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 n -tuple of strategy σ if there exists a sequence of lengths $(T_k)_{k \geq 1}$ such that $(\gamma_{T_k}(\sigma))_{k \geq 1}$ 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 Lipschitz², 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 $c \geq 0$, $\mathcal{F}(c)$ 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 costs 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*

²A 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 $|x_2 - x_1| \leq \theta |c_2 - c_1|$.

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 $\mathcal{F}(c) = g_c(P) \subseteq F(g_c)$. 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 distribution such that both marginals are equal whereas in the second one we allow any probability distribution. When considering the feasible set of the auxiliary game, we allow as a consequence histories without any constraint on the sequence of pair of actions (i, i') . Some sequences can be “translated” into the original game (e.g., playing forever a pair of the same actions (i, i)) but most of the sequences can not be “translated” (playing forever a fixed pair of actions (i, i')). This leads to a gap between the long-run feasible set of the repeated game with switching cost and the feasible set of the one-shot auxiliary game. In the symmetric case, repeating the same profile (i, i') is equal in the auxiliary game to alternating between (i, i') and (i', i) . This latter sequence is a sequence that corresponds to a valid history in the original game, hence closing the gap and yielding the equality.

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. $s_{i,j}^\ell = 1$ if $i \neq j$ and $s_{i,j}^\ell = 0$ otherwise). Hence, we generalize the literature to a larger family of switching costs functions. The following example shows that this condition is indeed necessary, and without it the long-run feasible set is only a subset of $F(g_c)$.

Example 2. *Counterexample with not symmetric switching costs.*

Consider the following two-player game where Player 1 has only one action: $((2,1), (0,0))$, and the switching costs matrix for Player 2 is $S^2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. 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. In order 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 $0.5c$. \triangle

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 $c \rightarrow \infty$, 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

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

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 \bar{F} is the set of all payoffs which are bounded above by the Pareto-Efficient front.

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

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

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

Consider the game in Example 1 and assume $c \rightarrow \infty$. The “payoffs” of the corresponding game are

| | LL | LR | RL | RR |
|------|----------------|--------------------|--------------------|--------------|
| TT | 2, 1 | 1, $-\infty$ | 1, $-\infty$ | 0, 0 |
| TB | $-\infty, 0.5$ | $-\infty, -\infty$ | $-\infty, -\infty$ | $-\infty, 1$ |
| BT | $-\infty, 0.5$ | $-\infty, -\infty$ | $-\infty, -\infty$ | $-\infty, 1$ |
| BB | 0, 0 | 0.5, $-\infty$ | 0.5, $-\infty$ | 1, 2 |

The convex hull of these payoffs is indeed \bar{F} , as shown in Figure 2. \triangle

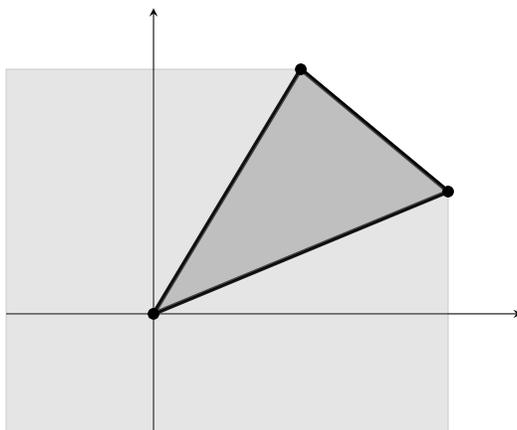


Figure 2: Feasible set $F(u)$ (dark grey) and \bar{F} (dark and light grey) of Example 3.

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 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 \tilde{u}^ℓ and is facing an imaginary Player $-\ell$. We call the *the long-run individually rational level* of Player ℓ the maximum payoff that Player ℓ can defend.

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

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

Note that the strategies of the players different from ℓ may depend on the past hence the best-reply of Player ℓ may not be constant and therefore $\bar{v}^\ell(c)$ does depend on c in general (see Example 4).

Zero-sum-like repeated games with switching cost 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 $\bar{v}^\ell(c)$ 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 $\bar{v}^\ell(c)$ 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 $\bar{v}^\ell(c)$. 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 $\bar{v}^\ell(c)$.*

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

$$l : \begin{pmatrix} -1 & 0 \\ 0 & 0 \end{pmatrix} \quad r : \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix}$$

And the switching costs matrix $S^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. We denote the actions of Player 2 (chooses columns) by $\{L, R\}$ and the actions of Player 3 (chooses matrices) by $\{l, r\}$. Player 1 chooses among the two rows, top or bottom.

For a switching cost of $c \leq 1$, 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. For a switching cost smaller than 1, they have to randomized. We obtain that the value is equal to $\bar{v}^1(c) = \max(-1, -\frac{1}{4}(1+c)^2)$ 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 $\mathcal{IR}(c) = \{(x^\ell)_{\ell \in N} \in \mathbb{R}^n, \text{ for all } \ell \in N, x^\ell \geq \bar{v}^\ell(c)\}$.*

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 $\bar{v}^\ell(c)$ 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]:

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

Proposition 4. *Assume that Player ℓ has no free-switching, then there exists \bar{c}^ℓ s.t. for every $c \geq \bar{c}^\ell$, 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 $c \geq \bar{c}^\ell$, $\bar{v}^\ell(c) = \bar{w}^\ell$.*

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: $\bar{v}^\ell(c) = \bar{v}^\ell$. 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 $\bar{v}(c)$ is bounded by \bar{w}^ℓ and decreasing). We obtain that the set of all long-run individually rational payoffs converges to the set \bar{V} , and it comprises of all vector payoff above pure reservation payoff for player in N^{nf} , above the minimax

value of the one-shot game for players in N^0 and above \underline{u}^ℓ for the rest:

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

where $\lim_{c \rightarrow \infty} \bar{v}^\ell(c)$ is one of $\bar{w}^\ell, \bar{u}^\ell, \bar{v}^\ell$, according to the switching costs of Player ℓ , as explained above.

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 the $c = 0$ case:

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, that 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 cost. This is a radically different conclusion than Theorem 6 in Lipman and Wang [2009] for undiscounted repeated games with arbitrarily small switching cost. This is due to the fact that we consider the average of the switching costs instead of their total sum.

As before, we can directly characterize the set of equilibrium payoffs when the switching costs are large. This is a natural corollary of Theorem 2 and the previous discussions.

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

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 \bar{F} becomes the negative orthant below the feasible payoffs and \bar{V} becomes $W = \{(x^\ell) \in \mathbb{R}^N, \text{ for all } \ell \in N, x^\ell \geq \bar{w}^\ell\}$. 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, the following regularity for the set of sub-game 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 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 a fixed cost 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 $\tilde{u}^\ell(k, i')$ to be equal to u if $k = \emptyset$ and $u^\ell(i') - cs_{k^\ell i' \ell}^\ell$ 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 $c \geq 0$. 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 $\delta \geq \underline{\delta}$, there is a perfect equilibrium strategy whose payoff is within ε of x . Equivalently*

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

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 two firsts are by definition of games with switching cost 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 in order 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 cost. 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 cost. 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.

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 arbitrarily close from $\mathcal{IR}(c) \cap \mathcal{F}(c)$. It follows by Corollary 4 that a higher cost can only have (asymptotically) 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 $x_\delta^*(c_2)$ is higher than under $x_\delta^*(c_1)$ up to an error ε . Informally, a higher cost can only give Player ℓ a smaller than ε benefit. This 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

$$(u^1, u^2) = \begin{pmatrix} 1, 1 & 0, 0 \\ 0, 1 & 4, 0 \\ 0, 1 & 2, 1 \end{pmatrix}, \quad S^1 = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \text{ and } S^2 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

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 cost 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 a 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. \triangle

5 Finite Games

The formalization of a game with switching costs as a stochastic game allows us to study games with finite horizon 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 sub-game 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 of 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 cooperated, then the good equilibrium is played in the last stages whereas if an agent deviated 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 $(\sigma_T)_{T \geq 1}$ 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 payoff 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 games since they are themselves obtained by a Sub-game Perfect Equilibrium. The next assumption states that there is a good one and bad ones for every player and every state.

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

- $x(k) \in \Pi_c(k)$ and for every $\ell \in N$ and $k \in K$, $x_{[\ell]} \in \Pi_c(k)$,
- $x_{[\ell]}^k < x_{[\ell]}^{k'}$ 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 $c \geq 0$. 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 $T \geq T^*$, there is a sub-game perfect equilibrium such that the payoff is within ε of x . Equivalently,

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

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 $x, x_{[1]}, \dots, x_{[n]}$ 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 $0 \leq c < c_0$, 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 $T \geq T^*$ and for any $c < c_0$, there is a sub-game perfect equilibrium such that the payoff is within ε of x . Equivalently,

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

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 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

$$S^1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 0 \end{pmatrix}$$

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

$$\begin{pmatrix} 0, 2 & 2, 0 & 0, 0 \\ 2, 0 & 0, 2 & 0, 0 \\ 1, 0 & 1, 0 & 3, 4 \end{pmatrix}$$

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, in order 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, T) as an SPE. \triangle

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

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

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 $x_T^*(c_2)$ is higher than under $x_T^*(c_1)$ 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

$$(u^1, u^2) = \begin{pmatrix} 3, 3 & 6, 1 \\ 2, 1 & 5, 2 \end{pmatrix}, \quad S^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \text{ and } \quad S^2 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

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 $(5 - \frac{4}{T}, 2 + \frac{1}{T})$ if $T = 3k + 1$, and $(5, 2)$ otherwise. Let us describe the subgame perfect equilibrium in a game where the number of stages is a multiple of 3:

- 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 a 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 of 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.

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 a 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 cost. Although this limits the applicability of our result, in fact it generalizes previous works which assumed constant switching cost. We postulate that it is possible to obtain an alternative characterization in the asymmetric approach 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 game 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.

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 $T \geq 1$,

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

The payoff of every player in this new game is higher since switching cost are smaller whereas the payoff from strategic interactions are equal. It follows by considering the subsequence $(T_k)_{k \geq 1}$ 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 $\ell \in N$, $x_2^\ell \leq x_1^\ell$.

Lipschitz: 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 $|x_2 - x_1| \leq \|S\|_\infty |c_2 - c_1|$. 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 are bounded above by $\|S\|_\infty |c_2 - c_1|$ and, as a consequence also the T -stage average payoffs. Consider a subsequence of length such that the T_k -average payoffs converges 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 repeat p times each profile of pure actions before switching. When repeating, switching costs appear once every p stages. 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 $c \geq 0$, 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 action by 1 and 2. By definition, there exists 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 $c \geq 0$, 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 profile of strategies to obtain a strategy generating x by successively playing each of them with proper length.

A.3 Proof of Theorem 1

The proof is decomposed into four parts. We first prove the two inclusions and then we proceed with the two other results.

Let x be a feasible payoff. By definition, there exists a profile of strategies σ and a sequence of times $(T_k)_{k \geq 1}$ such that the sequence of payoffs in the T_k -stage game converges to x . Let us prove that x can be expressed as the expected value of g_c for a well-chosen probability distribution in P . For every $T \geq 1$ and for every profile of strategies σ , one can define a probability distribution π_T in $\Delta(I \times I)$ such that the expected payoff under σ until stage T is approximately equal to the expectation of g_c under π_T . Formally, for every $T \geq 1$ and every play (i_1, \dots, i_T) , we define the occupation measure along the play as

$$f_T(i_1, \dots, i_T) = \frac{1}{T} \sum_{t=1}^{T-1} \Delta_{i_t, i_{t+1}} + \frac{1}{T} \Delta_{i_T, i_T},$$

where $\Delta_{i, i'}$ is the Dirac mass at the pair (i, i') . Given a profile of strategies of the

players σ , define z_T to be the image of the probability distribution over histories generated under σ by the mapping f_T . By construction z_T is a probability distribution over probability distributions ($z_T \in \Delta(\Delta(I \times I))$). Define π_T as the barycenter of the measure z_T on $\Delta(I \times I)$ formally defined as $\pi_T((i, i')) = \int_{\pi \in \Delta(I \times I)} \pi((i, i')) dz_T(\pi)$. By construction of π_T , one obtains that

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

By definition of x and considering the subsequence $(T_k)_{k \geq 1}$, one obtains that the sequence $(g_c^\ell(\pi_{T_k}))_{k \in \mathbb{N}}$ converges to x^ℓ for all $\ell \in N$. Moreover, for every play $p = ((i_1, i_2), \dots, (i_{T-1}, i_T), (i_T, i_T))$, denote by $f_T^1(p)$ (resp. $f_T^2(p)$) the first (resp. the second) marginal of $f(p)$. By definition

$$z^1(p) - z^2(p) = \frac{1}{T} \sum_{t=1}^T \Delta_{i_t} - \frac{1}{T} \left(\sum_{t=1}^{T-1} \Delta_{i_{t+1}} + \Delta_{i_T} \right) = \frac{1}{T} \Delta_{i_1} - \frac{1}{T} \Delta_{i_T}.$$

Hence, the norm between the two marginals is smaller than $2/T$. Denote by π_T^1 (resp. π_T^2), the first (resp. second) marginal of π_T . It is also the barycenter of the image of the probability distribution over histories by f_T^1 (resp. f_T^2). It follows that norm between the two probability distributions is also smaller than $2/T$. Consider a converging subsequence of $(\pi_{T_k})_{k \geq 1}$ to a probability distribution π_* . We obtain that $(\pi_{T_k}^1)_{k \geq 1}$ and $(\pi_{T_k}^2)_{k \geq 1}$ both converge (along the subsequence) to the same limit hence $\pi_*^1 = \pi_*^2$ and π_* is in P . By continuity of the payoff vector function on $\Delta(I \times I)$, the payoff generated by σ is equal to the expected payoff of g_c under a probability distribution in P .

Reciprocally, given a distribution $\pi \in P$. Let us consider the Markov chain induced on I by the initial distribution π^1 and the transition for all $i, i' \in I$, $q(i'|i) = \pi(i, i') / (\sum_{i'} \pi(i, i'))$. Clearly the distribution π^1 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 $(i_1^*, \dots, i_N^1, \dots)$. Define the joint profile of strategies where the players follows 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 second statement is an immediate consequence of the first one since P is a subset of the set of all probability distribution $\Delta(I \times I)$ and $g_c(\Delta(I \times I)) = F(g_c)$.

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 vector of pair of actions $(j_d)_{d \in \{1, \dots, D\}} \in (I \times I)^D$ and a vector of weights $(\alpha_d)_{d \in \{1, \dots, D\}} \in [0, 1]^D$ such that

$$\sum_d \alpha_d = 1 \text{ and for all } \ell \in N, x^\ell = \sum_d \alpha_d g^\ell(j_d).$$

Let us assume first that for every $d \in \{1, \dots, D\}$, α_d is a rational number $\frac{p_d}{Q}$ with p_d and Q two natural numbers. 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_1 times the action profile $j_1 = (i_1, i'_1)$, θp_2 times the action profile $j_2 = (i_2, i'_2)$, \dots , 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_1 or alternating between $j_1 = (i_1, i'_1)$ and $j'_1 = (i'_1, i_1)$ 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 in the switching cost game: start with i_1 and then alternate between i_1 and i'_1 for θp_1 periods, switch to i_2 and then alternate between i_2 and i'_2 for θp_2 periods, \dots , switch to i_D and then alternate between i_D and i'_D for θp_D periods, repeat from the top. The payoff generated by this profile of strategies in the switching cost game is equal to x up to an error $\frac{\|S\| + \|u\|}{\theta}$ (a mistake each time there is a change from an alternating pattern to the next one). 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 $\bar{v}^\ell(c)$ is continuous and decreasing: The proofs for this results are standard. Both relies on extending properties satisfied on the initial utilities to the min max values. Fix a pair of profile of actions $i(t-1)$ and $i(t)$. For every $\ell \in N$, and for every $c_1 \leq c_2 \geq 0$:

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

and

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

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

The function $\bar{v}^\ell(c)$ 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 more simple 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 to Player ℓ as a function of the previous action (or no action) of Player ℓ . It satisfies the following fixed point equation

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

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

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

Then $\bar{v}_\delta^\ell(\cdot, c) = (1 - \delta) \cdot f(\delta, c)$ 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 $\bar{v}_\delta^\ell(\cdot, c)$ are semialgebraic. As a semiagebraic and bounded function in δ , we know that the limit when δ goes to 1 exists. Moreover, by Mertens and Neyman [1981], it is equal to the undiscounted minmax. The function $\bar{v}^\ell(c)$ 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 \bar{c}^ℓ s.t. for every $c \geq \bar{c}^\ell$, $\bar{v}^\ell(c) = \bar{w}^\ell$. First, it is clear that Player ℓ can always guarantee \bar{w}^ℓ 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 \bar{w}^ℓ 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

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

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 $\gamma^\ell(\sigma_*) \geq \gamma^\ell((\sigma')^\ell, \sigma_*^{-\ell}) \geq \bar{v}^\ell(c)$. 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 $x = g_c(\pi) = \tilde{u}(\pi)$. 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 $(i_t)_{t \geq 1}$ 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 \mathcal{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 $(i_t)_{t \geq 1}$ such that $i_1 \in \mathcal{C}_k$ 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 , $\sigma_{\ell, \varepsilon}^{-\ell}$, such that

$$\sup_{\sigma^\ell} \mathbb{E}_{\sigma^\ell, \sigma_{\ell, \varepsilon}^{-\ell}} \left(\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \tilde{u}^\ell(i(t-1), i(t)) \right) \leq \bar{v}^\ell(c) + \varepsilon.$$

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. $\sigma_{\ell, \varepsilon}^{-\ell}$ 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 by 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 longrun individually rational level (with an approximation of 2ε). The precise lemma (Lemma 1) and its proof are delayed after the end of this proof. Define for every $n \geq 1$, $\varepsilon_n = \frac{1}{n}$. We can construct a sequence of times $(T_n)_{n \geq 1}$ has follows: $T_1 = 1$, for every $n \geq 2$, $T_{n+1} = \min \left(T(\varepsilon_n), \frac{2(M+c)T_n}{\varepsilon_n} \right)$. Let us call the block of stage $T_n + 1, \dots, T_{n+1}$ 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 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 $\bar{v}^\ell(c) + 3\varepsilon_n$. He obtained his minmax payoff in the long run, hence the deviation is not profitable.

We can now state and prove Lemma 1.

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

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

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

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 $T_{\min} \geq 1$ 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

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

The dependence of σ^ℓ in T was ignored in the previous equation for readability. Consider $T \geq \frac{2(M+c\|S\|_\infty)}{\varepsilon}$ 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 at 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,

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

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,

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

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

A.7 Proof of Proposition 5

Assume that every player has a switching cost. Formally, for every $\ell \in N$, there exists $i_1^\ell, i_2^\ell \in I$ such that $s_{i_1^\ell, i_2^\ell}^\ell > 0$. Let us prove that the game has Full Dimensionality. Denote by \bar{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 $\bar{x} - \varepsilon e_\ell$ is a feasible payoff, where e_ℓ is the unitary vector with coordinate ℓ . Let us recall that it is possible to obtain \bar{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 $(i_2^\ell, i^{-\ell})$ is the successor in the cycle of the profile $(i_1^\ell, i^{-\ell})$ generating \bar{x} . We can now define the alternative play where instead of playing these two profiles successively, all players except ℓ follows $i^{-\ell}$ and Player ℓ alternates between i_1^ℓ and i_2^ℓ . By construction, the payoff of Player ℓ' is $\bar{x}^{\ell'}$ for every $\ell' \neq \ell$ and the payoff of Player ℓ is $\bar{x}^\ell - \frac{1}{\#I} s_{i_1^\ell, i_2^\ell}^\ell$. 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 $c_0 \leq \frac{m}{\|S\|_\infty}$ and $c \leq c_0$. 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_1^ℓ be the previous action played by Player ℓ . We compare the stage payoff today to play i_2^ℓ instead of i_*^ℓ :

$$\begin{aligned} u^\ell(i_*^\ell, i_*^{-\ell}) - c \cdot s_{i_1^\ell, i_*^\ell}^\ell &\geq u^\ell(i_*^\ell, i_*^{-\ell}) - c \cdot \|S\|_\infty \geq u^\ell(i_2^\ell, i_*^{-\ell}) + m - c \cdot \|S\|_\infty, \\ &\geq u^\ell(i_2^\ell, i_*^{-\ell}) \geq u^\ell(i_2^\ell, i_*^{-\ell}) - c s_{i_1^\ell, i_2^\ell}^\ell. \end{aligned}$$

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_2^ℓ to i_*^ℓ . There is no one-shot profitable deviation. We have indeed an SPE.

References

- G.A. Akerlof and J.L. Yellen. Can small deviations from rationality make significant differences to economic equilibria? *The American Economic Review*, 75(4):708–720, 1985a.
- G.A. Akerlof and J.L. Yellen. A near-rational model of the business cycle, with wage and price inertia. *The Quarterly Journal of Economics*, 100:823–838, 1985b.
- Robert J Aumann and Lloyd S Shapley. Long-term competition—a game-theoretic analysis. In *Essays in game theory*, pages 1–15. Springer, 1994.
- Rudolf Avenhaus, Bernhard Von Stengel, and Shmuel Zamir. Inspection games. *Handbook of game theory with economic applications*, 3:1947–1987, 2002.

- A. Beggs and P. Klemperer. Multi-period competition with switching costs. *Econometrica: Journal of the Econometric Society*, pages 651–666, 1992.
- Jean-Pierre Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53(4):905–922, 1985. ISSN 00129682, 14680262.
- S.K. Chakrabarti. Characterizations of the equilibrium payoffs of inertia supergames. *Journal of Economic Theory*, 51(1):171–183, 1990.
- Matthew Darlington, Kevin D Glazebrook, David S Leslie, Rob Shone, and Roberto Szechtman. A stochastic game framework for patrolling a border. *arXiv preprint arXiv:2205.10017*, 2022.
- Prajit K Dutta. A folk theorem for stochastic games. *Journal of Economic Theory*, 66(1):1–32, 1995.
- J. A. Filar. Ordered field property for stochastic games when the player who controls transitions changes from state to state. *Journal of Optimization Theory and Applications*, 34(4):503–515, 1981.
- J.A. Filar. Algorithms for solving some undiscounted stochastic games[ph. d. thesis]. 1980.
- J.A. Filar and T.E.S. Raghavan. A matrix game solution of the single-controller stochastic game. *Mathematics of Operations Research*, 9(3):356–362, 1984.
- JA Filar and Todd A Schultz. The traveling inspector model. *Operations-Research-Spektrum*, 8(1):33–36, 1986.
- Jerzy Filar. Player aggregation in the traveling inspector model. *IEEE Transactions on Automatic Control*, 30(8):723–729, 1985.
- James W Friedman. A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12, 1971.
- James W Friedman. Non-cooperative equilibria in time-dependent supergames. *Econometrica: Journal of the Econometric Society*, pages 221–237, 1974.

- James W Friedman. A Modification of the Folk Theorem to Apply to Time-Dependent Supergame. *Oxford Economic Papers*, 42(2):317–335, 04 1990. ISSN 0030-7653. doi: 10.1093/oxfordjournals.oep.a041949.
- Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.
- Drew Fudenberg and Yuichi Yamamoto. The folk theorem for irreducible stochastic games with imperfect public monitoring. *Journal of Economic Theory*, 146(4):1664–1683, 2011.
- Begum Guney and Michael Richter. Costly switching from a status quo. *Journal of Economic Behavior & Organization*, 156:55–70, 2018.
- Begum Guney and Michael Richter. Games with switching costs and endogenous references. *Theoretical Economics*, 17(2):617–650, 2022.
- Johannes Hörner, Takuo Sugaya, Satoru Takahashi, and Nicolas Vieille. Recursive methods in discounted stochastic games: An algorithm for $\delta \rightarrow 1$ and a folk theorem. *Econometrica*, 79(4):1277–1318, 2011.
- Steffen Jørgensen, Guiomar Martín-Herrán, and Georges Zaccour. Dynamic games in the economics and management of pollution. *Environmental Modeling & Assessment*, 15(6):433–467, 2010.
- Paul Klemperer. Markets with consumer switching costs. *The quarterly journal of economics*, 102(2):375–394, 1987.
- Paul Klemperer. Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The review of economic studies*, 62(4):515–539, 1995.
- B. L. Lipman and R. Wang. Switching costs in infinitely repeated games. *Games and Economic Behavior*, 66(1):292–314, 2009.
- B.L Lipman and R. Wang. Switching costs in frequently repeated games. *Journal of Economic Theory*, 93(2):149–190, 2000.
- Chantal Marlats. A folk theorem for stochastic games with finite horizon. *Economic Theory*, 58(3):485–507, 2015. ISSN 09382259, 14320479.

- Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, Cambridge, 2013.
- J-F Mertens and Abraham Neyman. Stochastic games. *International Journal of Game Theory*, 10(2):53–66, 1981.
- Barry O’Neill. Game theory models of peace and war. *Handbook of game theory with economic applications*, 2:995–1053, 1994.
- S. Rass and B. Rainer. Numerical computation of multi-goal security strategies. In *International Conference on Decision and Game Theory for Security*, pages 118–133. Springer, 2014.
- S. Rass, S. Konig, and S. Schauer. Defending against advanced persistent threats using game-theory. *PloS one*, 12(1):e0168675, 2017.
- G. Schoenmakers, J. Flesch, F. Thuijsman, and O.J. Vrieze. Repeated games with bonuses. *Journal of optimization theory and applications*, 136(3):459–473, 2008.
- Yevgeny Tsodikovich, Xavier Venel, and Anna Zseleva. The regularity of the value function of repeated games with switching costs. *Mathematics of Operations Research*, 2022.
- Mesut Yavuz and David Jeffcoat. An analysis and solution of the sensor scheduling problem. In *Advances in Cooperative Control and Optimization*, pages 167–177. Springer, 2007.