



Recursive games: uniform value, Tauberian theorem and the Mertens conjecture “ $\text{Maxmin}=\lim v_n=\lim v_{\{\lambda\}}$ ”

Post-print version of the following publication: | Versione post-print della seguente pubblicazione:

Original Citation/Citazione:

Li, X; Venel, Xavier Mathieu Raymond. (2016). Recursive games: uniform value, Tauberian theorem and the Mertens conjecture “ $\text{Maxmin}=\lim v_n=\lim v_{\{\lambda\}}$ ”. INTERNATIONAL JOURNAL OF GAME THEORY, (ISSN: 0020-7276), 45: 155-189. Doi: 10.1007/s00182-015-0496-4.

Availability/Disponibilità:

This version is available at: [11385/197457](#) since: 2020-10-08T09:44:45Z - Questa versione è disponibile alla pagina: [11385/197457](#) dal: 2020-10-08T09:44:45Z

Publisher/Casa editrice:

Published version/Pubblicato:

DOI: <https://dx.doi.org/10.1007/s00182-015-0496-4>

License/Licenza:

DRM (Digital rights management) non definiti

Availability/Termini d'uso:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license. For all terms of use and more information see the publisher's website. | I termini e le condizioni relativi al riutilizzo della presente versione della pubblicazione sono disciplinati dalla politica editoriale. Le opere messe a disposizione con licenze Creative Commons possono essere utilizzate conformemente ai termini e alle condizioni previste da tali licenze. Per l'insieme delle condizioni di utilizzo e per ulteriori informazioni si rinvia al sito web dell'editore.

This item was downloaded from IRIS Luiss (<https://iris.luiss.it/>). When citing, please refer to the published version. | Questo documento è stato scaricato da IRIS Luiss (<https://iris.luiss.it/>). Per la citazione, fare riferimento alla versione pubblicata sul sito dell'editore.

(Article begins on next page | Il contributo inizia nella pagina successiva)

Recursive games: Uniform value, Tauberian theorem and the Mertens conjecture “ $Maxmin = \lim v_n = \lim v_\lambda$ ”

Xiaoxi LI ^{*}; Xavier VENEL [†]

May 30, 2015

Contents

1	Introduction	2
2	Preliminaries: model and notations	3
2.1	The model	3
2.2	Definition of strategies and evaluations	4
2.3	Stopping time and concatenation of strategies	5
2.4	Uniform value	5
3	Main results	6
3.1	Sufficient condition for the existence of the uniform value	6
3.2	A recursive game without uniform value	7
4	Proofs	7
4.1	The case of positive-valued recursive game	8
4.2	Existence of the uniform value (proof of Theorem 3.1)	9
4.3	Pure optimal strategy (proof of Corollary 3.4)	17
5	Application to recursive games with signals	18
5.1	Model	18
5.2	Evaluation	20
5.3	Results	20
5.4	Proof of Theorem 5.8	21

^{*}CNRS, IMJ-PRG, UMR 7586, Sorbonne Universités, UPMC Univ. Paris 06, Univ. Paris Diderot, Sorbonne Paris Cité, Paris, France. Email: xxleewhu@gmail.com.

[†]CES, Université Paris 1 Panthéon Sorbonne, Paris, France. Email: xavier.venel@univ-paris1.fr

Abstract

We study two-player zero-sum recursive games with a countable state space and finite action spaces at each state. When the family of n -stage values $\{v_n, n \geq 1\}$ is totally bounded for the uniform norm, we prove the existence of the uniform value. Together with a result in Rosenberg and Vieille [12], we obtain a uniform Tauberian theorem for recursive game: (v_n) converges uniformly if and only if (v_λ) converges uniformly.

We apply our main result to finite recursive games with signals (where players observe only signals on the state and on past actions). When the maximizer is more informed than the minimizer, we prove the Mertens conjecture $Maxmin = \lim_{n \rightarrow \infty} v_n = \lim_{\lambda \rightarrow 0} v_\lambda$. Finally, we deduce the existence of the uniform value in finite recursive game with symmetric information.

Keywords: Stochastic games, recursive games, asymptotic value, uniform value, Tauberian theorem, maxmin

1 Introduction

Stochastic games were introduced by Shapley [13] to model a multiplayer dynamic interaction, where players' collective decisions influence the current payoff and also the future state. In this article, we focus on two-player zero-sum recursive games introduced by Everett [2]. The specificity of a recursive game is that the state space is divided into two sets: absorbing states and active states. On absorbing states, the process is absorbed and the payoff is fixed. On active (non-absorbing) states, the payoff is always equal to 0.

There are several ways to evaluate the payoff stream in a zero-sum stochastic game. Given a positive integer n , the n -stage payoff is the expected average payoff during the first n stages. Given $\lambda \in (0, 1]$, the λ -discounted payoff is the Abel mean of the infinite stage payoffs with a weight $\lambda(1 - \lambda)^{t-1}$ for stage t . We will focus on the concept of uniform value. A stochastic game admits a uniform value if both players can approximately guarantee the same payoff level in all sufficiently long n -stage games without knowing *a priori* the length of the game.

Mertens and Neyman [7] proved that a stochastic game with a finite state space and finite set of actions where the players observe the current state and the stage payoffs admits a uniform value. Their proof uses the fact that the function $\lambda \mapsto v_\lambda$ has bounded variation, where v_λ is the λ -discounted value (Bewley and Kohlberg [1]). For stochastic games with an infinite state space, this argument in general does not apply.

Markovian decision processes (henceforth MDP) are stochastic games with only one player. Lehrer and Sorin [5] showed that in a MDP, the uniform convergence of (v_λ) (w.r.t. the initial state) as λ tends to zero is equivalent to the uniform convergence of the n -stage values (v_n) as n tends to infinity. Nevertheless, uniform convergence of (v_n) or (v_λ) is not sufficient for the existence of the uniform value (cf. Monderer and Sorin [9] or Lehrer and Monderer [4]).

For recursive games, the situation seems to be different. There are two results giving sufficient conditions for a recursive game with countable state space to have a uniform value. The first one can be derived from Rosenberg and Vieille [12]: if (v_λ) converges uniformly to some function v , then the recursive game has a uniform value, which is equal to v . The second one is due to Solan and Vieille [14]: if, except on a finite subset, the *limsup value*¹ is above a strictly positive

¹The limsup value is the value of the game in which the global payoff to player 1 is the limsup of the stage payoff stream.

constant on the non-absorbing states, then the recursive game has a uniform value, which is equal to the limsup value.

The main result of this paper is that the uniform convergence of the n -stage values is a sufficient condition for the existence of the uniform value. In fact we prove a stronger result: for any recursive game with countable state space, if the family $\{v_n, n \geq 1\}$ is totally bounded for the uniform norm, then the uniform value exists. Our proof follows the same idea as Solan and Vieille [14] and we will use several of their results.

Our result together with the result of Rosenberg and Vieille [12] provides a uniform Tauberian theorem for recursive games: (v_n) converges uniformly if and only if (v_λ) converges uniformly, and in case of convergence, both limits are the same. For general stochastic games, Ziliotto [20] provided recently a direct proof of this result.

Finally, we apply our main result to finite recursive games with signals. In a recursive game with signals, players do not perfectly observe the state and actions at every stage anymore, rather they receive a private signal. Mertens [6] conjectured that in a general model of zero-sum repeated games, if player 1 (the maximizer) is always more informed than player 2 (the minimizer) during the play (in the sense that player 2's private signal can be deduced from player 1's private signal) then $Maxmin = \lim_{n \rightarrow \infty} v_n = \lim_{\lambda \rightarrow 0} v_\lambda$, *i.e.*, both the uniform maxmin and the asymptotic value exist and are equal.

Ziliotto [19] showed that the result is false in general. Nevertheless, several positive results have been obtained for subclasses of games including Sorin [15] and Sorin [16] for Big match with one-sided incomplete information, Rosenberg *et al.* [11], Renault [10] and Gensbittel *et al.* [3] for a more informed controller, and Rosenberg and Vieille [12] for recursive games with one-sided incomplete information.

We prove the Mertens conjecture in finite recursive games with signals, where player 1 is always more informed than player 2 during the play. The proof uses several results from Gensbittel *et al.* [3], concerning the n -stage value functions in a repeated game where player 1 is more informed than player 2. Our result generalizes Rosenberg and Vieille [12], which deals with the model where player 1 is informed of a private signal on the state at the beginning of the game. Moreover, we deduce the existence of the uniform value in finite recursive games with symmetric information.

The organization of the article is as follows: in Section 2 we introduce the model of recursive games; in Section 3 we present the main result and several corollaries; Section 4 is dedicated to the proofs; finally in Section 5 we apply the result to finite recursive games with signals.

2 Preliminaries: model and notations

Notation Given any metric space S , endowed with the Borelian σ -algebra, we denote by $\Delta(S)$ the set of probabilities on S and we denote by $\Delta_f(S)$ the set of probabilities with finite support.

2.1 The model

A two-player zero sum stochastic game $\Gamma = \langle X, A, B, g, q \rangle$ is given by

- a state space X .

- player 1's action set A , and for any $x \in X$, $A(x)$ is a finite subset of A .
- player 2's action set B , and for any $x \in X$, $B(x)$ is a finite subset of B .
- a payoff function: $g : X \times A \times B \rightarrow [-1, +1]$.
- a transition probability function: $q : X \times A \times B \rightarrow \Delta_f(X)$.

Play of the game The stochastic game with initial state $x_1 \in X$ is denoted by $\Gamma(x_1)$, and is played as follows: at each stage $t \geq 1$, after observing $(x_1, a_1, b_1, \dots, a_{t-1}, b_{t-1}, x_t)$, player 1 and player 2 choose simultaneously actions $a_t \in A(x_t)$ and $b_t \in B(x_t)$. The stage payoff is $g(x_t, a_t, b_t)$ and a new state x_{t+1} is drawn according to the probability distribution $q(x_t, a_t, b_t)$. Both players observe the action pair (a_t, b_t) and the state x_{t+1} . The game then proceeds to stage $t + 1$.

Note that we did not make any measurability assumption on the model. As the transition probability distribution is supposed to be finitely supported, given an initial state, the set of actions and states that might appear in the infinite game are in fact countable. Therefore probability distributions are well defined.

Recursive game Γ is a recursive game if there exist a set of active states denoted by X^0 and a set of absorbing states denoted by X^* with $X^0 \cup X^* = X$ and $X^0 \cap X^* = \emptyset$, such that:

- the stage payoff is 0 on active states: $\forall x \in X^0, g(x, a, b) = 0, \forall (a, b) \in A(x) \times B(x)$;
- states in X^* are absorbing: $\forall x \in X^*, q(x, a, b)(x) = 1, \forall (a, b) \in A(x) \times B(x)$, and $g(x, a, b)$ depends only on x .

2.2 Definition of strategies and evaluations

History At stage t , the space of finite histories is $H_t = (X \times A \times B)^{t-1} \times X$. Set $H_\infty = (X \times A \times B)^\infty$ to be the space of infinite *plays*. We consider the discrete topology on X, A and B . For every $t \geq 1$, we identify any $h_t \in H_t$ with a cylinder set in H_∞ and denote by \mathcal{H}_t the σ -field of H_t induced on H_∞ . The product σ -field on H_∞ is $\mathcal{H}_\infty = \sigma(\mathcal{H}_t, t \geq 1)$.

Strategy A (*behavior*) *strategy* for player 1 is a sequence of functions $\sigma = (\sigma_t)_{t \geq 1}$ with each $t \geq 1, \sigma_t : (H_t, \mathcal{H}_t) \rightarrow \Delta(A)$ such that for every $h_t \in H_t, \sigma_t(h_t)(A(x_t)) = 1$. If for every $t \geq 1$ and $h_t \in H_t$, there exists $a \in A(x_t)$ such that $\sigma_t(h_t)[a] = 1$, then the strategy is *pure*. We define similarly a behavior strategy τ for player 2. Denote by Σ and \mathcal{T} respectively player 1's and player 2's sets of behavior strategies. Denote by $\widehat{\Sigma}$ and $\widehat{\mathcal{T}}$ respectively player 1's and player 2's subsets of strategies that depend on the histories only through the states but not on the actions.

Evaluations Let us describe several ways to evaluate the payoff in Γ . By Kolmogorov's extension theorem, any triple $(x_1, \sigma, \tau) \in X \times \Sigma \times \mathcal{T}$ induces a unique probability distribution over $(H_\infty, \mathcal{H}_\infty)$ denoted by $\mathbb{P}_{x_1, \sigma, \tau}$. Let $\mathbb{E}_{x_1, \sigma, \tau}$ be the corresponding expectation.

n-stage average For each positive $n \geq 1$, the expected average payoff up to stage n , induced by the couple of strategies (σ, τ) and the initial state x_1 is given by

$$\gamma_n(x_1, \sigma, \tau) = \mathbb{E}_{x_1, \sigma, \tau} \left(\frac{1}{n} \sum_{t=1}^n g(x_t, a_t, b_t) \right).$$

The game with expected n -stage average payoff and initial state x_1 is denoted as $\Gamma_n(x_1)$.

λ -discounted average For each $\lambda \in (0, 1]$, the expected λ -discounted average payoff, induced by the couple of strategies (σ, τ) and the initial state x_1 is given by

$$\gamma_\lambda(x_1, \sigma, \tau) = \mathbb{E}_{x_1, \sigma, \tau} \left(\lambda \sum_{t=1}^{\infty} (1-\lambda)^{(t-1)} g(x_t, a_t, b_t) \right).$$

The game with expected λ -discounted average payoff and initial state x_1 is denoted as $\Gamma_\lambda(x_1)$.

In either $\Gamma_n(x_1)$ or $\Gamma_\lambda(x_1)$, player 1 maximizes the expected average payoff and player 2 minimizes it. For a fixed x_1 the game $\Gamma_n(x_1)$ is finite, so there exists a value $v_n(x_1)$ by minmax theorem. The existence of the discounted value $v_\lambda(x_1)$ is also standard, and we refer to Mertens *et al.* [8] (Section VII.1.) for a general presentation.

2.3 Stopping time and concatenation of strategies

A function $\theta : (H_\infty, \mathcal{H}_\infty) \rightarrow \mathbb{N}$ is called a *stopping time* if the set $\{h \in H_\infty | \theta(h) = t\}$ is \mathcal{H}_t -measurable for all $t \geq 1$. Explicitly for any $h, h' \in H_\infty$ and $n \geq 1$: if h and h' coincide until stage n and $\theta(h) = n$ then $\theta(h') = n$. Let θ and θ' be two stopping times, we write $\theta \leq \theta'$ if for every $h \in H_\infty$, $\theta(h) \leq \theta'(h)$.

Given a sequence of strategies $(\sigma^{[\ell]})_{\ell \geq 1}$ and a sequence of increasing stopping time $(\theta_\ell)_{\ell \geq 1}$, we define $\sigma^* := \sigma^{[1]\theta_1} \sigma^{[2]\theta_2} \dots$ as the *concatenation* of $(\sigma^{[\ell]})_{\ell \geq 1}$ along $(\theta_\ell)_{\ell \geq 1}$. Given $n \geq t \geq 1$ and $h \in H_\infty$, let h_n be the projection of h on H_n and h_n^t be the history of h between stage t and n . The strategy σ^* is defined by $\sigma_n^*(h_n) = \sigma_n^{[1]}(h_n)$ if $n < \theta_1(h)$; $\sigma_n^*(h_n) = \sigma_{n-\theta_{m-1}}^{[m]}(h_n^{\theta_{m-1}})$ if $\theta_{m-1} \leq n < \theta_m$. Informally, for every $\ell \geq 1$ at stage θ_ℓ , the player forgets the past and starts to play $\sigma_{\ell+1}$ at the current state.

2.4 Uniform value

Uniformly guarantee Player 1 *uniformly guarantees* w if for every $\varepsilon > 0$, there exists σ_ε in Σ and $N_0 \geq 1$ such that for every $x_1 \in X^0$,

$$\gamma_n(x_1, \sigma_\varepsilon, \tau) \geq w(x_1) - \varepsilon, \quad \forall n \geq N_0, \quad \forall \tau \in \mathcal{T}.$$

We say that the strategy σ_ε uniformly guarantees $w - \varepsilon$. Similarly, player 2 uniformly guarantees w if for every $\varepsilon > 0$, there exists τ_ε in \mathcal{T} and $N_0 \geq 1$ such that for every $x_1 \in X^0$,

$$\gamma_n(x_1, \sigma, \tau_\varepsilon) \leq w(x_1) + \varepsilon, \quad \forall n \geq N_0, \quad \forall \sigma \in \Sigma.$$

Uniform value $v_\infty : X \rightarrow \mathbb{R}$ is the uniform value of the game Γ if both players uniformly guarantee v_∞ . A strategy for player 1 (*resp.* player 2) that uniformly guarantees $v_\infty - \varepsilon$ (*resp.* $v_\infty + \varepsilon$) is called *uniform ε -optimal*. If both players can uniformly guarantee v_∞ with pure strategies, Γ has a *uniform value in pure strategies*.

Remark 2.1 *In defining the uniform value, we ask N_0 to be independent of the initial state x_1 . One direct consequence of the existence of the uniform value v_∞ is the uniform convergence of $(v_n)_{n \geq 1}$ to v_∞ . This is stronger than the definition where the existence of the uniform value is considered state by state (see for example Solan and Vieille [14], Definitions 3-4)*

3 Main results

In this section, we present the main result of the paper, namely Theorem 3.1, as well as several corollaries. We also provide an example that does not satisfy the condition of Theorem 3.1 and does not have a uniform value.

3.1 Sufficient condition for the existence of the uniform value

Denote by $\mathbf{B}(X)$ the set of functions from X to $[-1, 1]$ with the uniform norm $\|\cdot\|_\infty$. Recall that a set of functions F in $(\mathbf{B}(X), \|\cdot\|_\infty)$ is *totally bounded* if for every $\varepsilon > 0$, there exists a finite subset $F_R = \{f_r : 1 \leq r \leq R\} \subseteq F$ such that for any $f \in F$, there is $f_r \in F_R$ with $\|f - f_r\|_\infty \leq \varepsilon$.

Theorem 3.1 *Suppose that the space $\{v_n, n \geq 1\}$ is totally bounded for the uniform norm, then the recursive game Γ has a uniform value v_∞ . Moreover both players can uniformly guarantee v_∞ with strategies that depend only on the history of states and not on past actions.*

We deduce from the previous result a uniform Tauberian theorem in recursive games.

Corollary 3.2 *The sequence of n -stage values $(v_n)_{n \geq 1}$ converges uniformly as n tends to infinity if and only if the sequence of λ -discounted values $(v_\lambda)_{\lambda \in (0,1]}$ converges uniformly as λ tends to zero. In case of convergence, both limits are the same.*

On one hand, if (v_n) converges uniformly, the family is totally bounded, thus the uniform value exists, and this implies the uniform convergence of (v_λ) (Sorin [17], Lemma 3.1). On the other hand, the converse result is established in Rosenberg and Vieille [12] (see Remark 6, Theorem 1 and Theorem 3).

Remark 3.3 *The equivalence of the uniform convergences of $(v_n)_{n \geq 1}$ and $(v_\lambda)_{\lambda \in (0,1]}$ has been proven in MDP by Lehrer and Sorin [5]. Ziliotto [20] recently showed that it is also true for stochastic games whenever the Shapley operator is well defined.*

If, in addition, for every $n \geq 1$ the n -stage value $v_n(x)$ exists in pure strategies, then Γ has a uniform value in pure strategies.

Corollary 3.4 *Suppose that for every $n \geq 1$, both players have pure optimal strategies in the n -stage game, and $\{v_n, n \geq 1\}$ is totally bounded for the uniform norm. Then Γ has a uniform value v_∞ in pure strategies. Moreover, both players can uniformly guarantee v_∞ with strategies that depend only on the history of states and not on past actions.*

Remark 3.5 *The result in Corollary 3.4 extends to games with general action sets $A(x)$ and $B(x)$ provided that for any $n \geq 1$, the n -stage game has a value and both players have pure optimal strategies.*

The proof of Corollary 3.4 is similar to that of Theorem 3.1. The key difference involves a technical lemma (Lemma 4.18) for the existence of a (pure) stopping time which is used in the definition of players' optimal strategies (see the proof of Proposition 4.3). We discuss this point and present the proof in Subsection 4.3.


3.2 A recursive game without uniform value

We present here an example of a recursive game with countable state space where $\{v_n, n \geq 1\}$ is not totally bounded and there is no uniform value (See **Figure 3.2** below for illustration). This is an adaptation to our framework of an example in Lehrer and Sorin [5].

The state space is a subset of $\mathbb{Z} \times \mathbb{Z}$. The set of active states is $X^0 = \{(x, y) \in \mathbb{N} \times \mathbb{N} \mid 0 \leq y \leq x\}$ and the set of absorbing states is $X^* = X_1^* \cup X_{-2}^*$ (two types), where $X_1^* = \mathbb{N} \times \{-1\}$ and $X_{-2}^* = \{(x, x+1) \mid x \geq 0\}$. The payoff is 1 on X_1^* and is -2 on X_{-2}^* . There is only one player (maximizer), whose action set is $\{R(\text{ight}), J(\text{ump})\}$. The transition rule is given by:

- at $(x, 0) \in X^0$: $q((x, 0), R)(x+1, 0) = 1$, and $q((x, 0), J)(x, -1) = q((x, 0), J)(x, 1) = \frac{1}{2}$;
- at $(x, y) \in X^0$ with $0 < y \leq x$: $q((x, y), a)(x, y+1) = 1, \forall a \in \{R, J\}$.

Starting at $(0, 0)$, one optimal strategy for an n -stage game is to go *Right* for half of the game, and then to *Jump*. This gives an expected average payoff around $\frac{1}{4}$, thus $\lim_{n \rightarrow \infty} v_n(0, 0) = \frac{1}{4}$. In a λ -discounted game, the optimal stage to *Jump* is approximately $\frac{\ln(\frac{2-\lambda}{4})}{\ln(1-\lambda)}$. It follows that $v_\lambda(0, 0) \approx \frac{2-\lambda}{16}$ and thus $\lim_{\lambda \rightarrow 0} v_\lambda(0, 0) = \frac{1}{8}$. This implies that there is no uniform value. On the other hand, $\{v_n, n \geq 1\}$ is not totally bounded for the uniform norm. Indeed, the convergence of (v_n) is not uniform: for any $x \geq 1$, $\lim_{n \rightarrow \infty} v_n(x, 1) = -2$ while $v_x(x, 1) = 0$.



The figure on the left illustrates a play (R, \dots, R, J) jumping after n steps: with probability $1/2$ the state is absorbed at $(n, n+1) \in X_{-2}^*$; with probability $1/2$ the state is absorbed at $(n, -1) \in X_1^*$.

\longrightarrow : a deterministic transition;
 $- \rightarrow$: a probabilistic transition.

Figure 3.2

4 Proofs

In the first subsection, we introduce and establish preliminary results for a subclass of recursive game, which will be called *positive-valued recursive games*. In the second subsection, we prove Theorem 3.1 by a reduction of any recursive game to a positive-valued recursive game. The proof for Corollary 3.4 is given in the third subsection.

4.1 The case of positive-valued recursive game

Definition 4.1 A recursive game is positive-valued if there exist $M > 0$ and $n_0 \geq 1$ such that for every non-absorbing state $x \in X^0$, there exists $n(x) \leq n_0$ such that $v_{n(x)}(x) \geq M$.

In order to state the next proposition, we first introduce the notion of uniformly terminating strategy.

Definition 4.2 Denote by ρ the stopping time of absorption in X^* : $\rho = \inf\{n \geq 1, x_n \in X^*\}$. The strategy σ is said to be uniformly terminating if for any $\varepsilon > 0$, there exists $N \geq 1$ such that for every $x_1 \in X^0$ and for every $\tau \in \mathcal{T}$, $\mathbb{P}_{x_1, \sigma, \tau}(\rho \leq N) \geq 1 - \varepsilon$.

Proposition 4.3 Let Γ be a positive-valued recursive game. We fix the numbers $M > 0, n_0 \geq 1$ and the mapping $n(\cdot) : X^0 \rightarrow \{1, \dots, n_0\}$ such that $v_{n(x)}(x) \geq M, \forall x \in X^0$. Then player 1 uniformly guarantees $v_{n(\cdot)}(\cdot)$ with uniformly terminating strategies that depends only on states: for all $\varepsilon > 0$, there exists σ^* in $\widehat{\Sigma}$ and $N_0 \geq 1$ such that for every $x_1 \in X^0$ and every τ in \mathcal{T} ,

$$(i) \mathbb{P}_{x_1, \sigma^*, \tau}(\rho \leq N_0) \geq 1 - \varepsilon \text{ and } (ii) \gamma_n(x_1, \sigma^*, \tau) \geq v_{n(x_1)}(x_1) - \varepsilon, \forall n \geq N_0.$$

Proof. Let $\hat{\sigma}$ be a profile of strategies such that for every $x \in X^0$, $\hat{\sigma}(x)$ is optimal in the $n(x)$ -stage game $\Gamma_{n(x)}(x)$. Let $\tilde{k} := \tilde{k}(x)$ be a random stage uniformly chosen in $\{1, \dots, n(x)\}$. For any $\tau \in \mathcal{T}$ and $x \in X^0$, $(x, \hat{\sigma}, \tau)$ and \tilde{k} induce a probability distribution over $H_\infty \times \{1, \dots, n(x)\}$, which we denote by $\tilde{\mathbb{P}}_{x, \hat{\sigma}, \tau}$. Let $\tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau}$ be the corresponding expectation. We obtain:

$$\tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau}[g(x_{\tilde{k}})] = \mathbb{E}_{x, \hat{\sigma}, \tau}\left[\frac{1}{n(x)} \sum_{l=1}^{n(x)} g(x_l)\right] \geq \inf_{\tau'} \mathbb{E}_{x, \hat{\sigma}, \tau'}\left[\frac{1}{n(x)} \sum_{l=1}^{n(x)} g(x_l)\right] \geq v_{n(x)}(x) \geq M.$$

It follows that

$$\tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau}\left[g(x_{\tilde{k}})\mathbb{1}_{\rho \leq \tilde{k}} + g(x_{\tilde{k}})\mathbb{1}_{\rho > \tilde{k}}\right] \geq M.$$

On the event $\{\rho > \tilde{k}\}$, $g(x_{\tilde{k}}) = 0$, whereas on the event $\{\rho \leq \tilde{k}\}$, we have $g(x_{\tilde{k}}) = g(x_\rho)$. This implies that

$$\tilde{\mathbb{P}}_{x, \hat{\sigma}, \tau}(\rho \leq \tilde{k})\tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau}[g(x_\rho) | \rho \leq \tilde{k}] = \tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau}[g(x_{\tilde{k}})] \geq v_{n(x)}(x) \geq M. \quad (4.1)$$

Using the fact that the payoff function g has maximal norm 1, we deduce from (4.1):

$$\tilde{\mathbb{P}}_{x, \hat{\sigma}, \tau}(\rho \leq \tilde{k}) \geq M. \quad (4.2)$$

Define the strategy ² σ^* as concatenations of $(\hat{\sigma}(x_{u_\ell}))_{l \geq 0}$ at the random stages $(u_\ell)_{\ell \geq 0}$, where u_ℓ is defined inductively along the play by $u_0 = 1$ and $u_{\ell+1} - u_\ell = \tilde{k}(x_{u_\ell})$ follows the uniform distribution over $\{1, \dots, n(x_{u_\ell})\}$. Let $\tilde{\mathbb{P}}_{x, \sigma^*, \tau}$ be the (product) probability distribution over $H_\infty \times \{1, \dots, n_0\}^{\mathbb{N}}$ induced by (x, σ^*, τ) , and $\tilde{\mathbb{E}}_{x, \sigma^*, \tau}$ the corresponding expectation. Let $\varepsilon > 0$.

(i) We show that σ^* is uniformly terminating. By (4.2), the conditional probability of absorbing on each block $\{u_{l-1}, \dots, u_l - 1\}$ is no smaller than M . Thus for any τ and $x_1 \in X^0$,

$$\tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(\rho \geq u_l) \leq (1 - M)^l, \forall l \geq 1.$$

²The strategy σ^* is a generalized mixed strategy, which is equivalent to a behavior strategy by Kuhn's theorem.

The length of each block is uniformly bounded by n_0 , thus if we put $l^* \geq \frac{\ln(\varepsilon)}{\ln(1-M)}$:

$$\tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(\rho \leq n_0 l^*) \geq \tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(\rho \leq u_{l^*}) \geq 1 - (1-M)^{l^*} \geq 1 - \varepsilon. \quad (4.3)$$

(ii) We now argue that σ^* uniformly guarantees $v_{n(x_1)}(x_1) - 3\varepsilon$. Let $N_0 = n_0 l^* / \varepsilon$.

For any $\tau \in \mathcal{T}$, $x_1 \in X^0$ and $n \geq n_0 l^*$, we have

$$\begin{aligned} \mathbb{E}_{x_1, \sigma^*, \tau}[g(x_n)] &= \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau} \left[\sum_{l=0}^{\ell^*-1} g(x_n) \mathbb{1}_{u_l \leq \rho < u_{l+1}} + g(x_n) \mathbb{1}_{u_{\ell^*} \leq \rho} \right] \\ &= \sum_{l=0}^{\ell^*-1} \tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(u_l \leq \rho < u_{l+1}) \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau}[g(x_\rho) | u_l \leq \rho < u_{l+1}] \\ &\quad + \tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(u_{\ell^*} \leq \rho) \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau}[g(x_n) | \rho \geq u_{\ell^*}]. \end{aligned}$$

According to (4.3), $\mathbb{P}_{x_1, \sigma^*, \tau}(\rho \geq u_{l^*}) \leq \varepsilon$, thus we focus on an absorption before u_{ℓ^*} :

$$\mathbb{E}_{x_1, \sigma^*, \tau}[g(x_n)] \geq \sum_{l=0}^{\ell^*-1} \tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(u_l \leq \rho < u_{l+1}) \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau}[g(x_\rho) | u_l \leq \rho < u_{l+1}] - \varepsilon \quad (4.4)$$

For each $l \geq 0$, σ^* is following $\hat{\sigma}(x_{u_l})$ for $u_{l+1} - u_l = \tilde{k}(x_{u_l})$ stages. Thus (4.1) applies, and we obtain: for $\ell \geq 1$,

$$\tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(u_l \leq \rho < u_{l+1}) \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau}[g(x_\rho) | u_l \leq \rho < u_{l+1}] \geq \tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(\rho > u_\ell) M > 0,$$

and for $l = 0$,

$$\tilde{\mathbb{P}}_{x_1, \sigma^*, \tau}(1 \leq \rho < u_1) \tilde{\mathbb{E}}_{x_1, \sigma^*, \tau}[g(x_\rho) | 1 \leq \rho < u_1] \geq v_{n(x_1)}(x_1).$$

By substituting the two previous inequalities into (4.4), we obtain that

$$\forall n \geq n_0 l^*, \forall x_1 \in X^0, \mathbb{E}_{x_1, \sigma^*, \tau}[g(x_n)] \geq v_{n(x_1)}(x_1) - \varepsilon. \quad (4.5)$$

Now for $n \geq N_0$, we deduce that $\gamma_n(x_1, \sigma^*, \tau) \geq v_{n(x_1)}(x_1) - 3\varepsilon$. ■

One can deduce from Proposition 4.3 a first result on recursive games with the condition that the sequence of n -stage values converges uniformly to a function bounded away from 0.

Corollary 4.4 *Assume that in a recursive game Γ , the sequence of n -stage values $(v_n)_{n \geq 1}$ converges uniformly to a function v satisfying for every $x \in X^0$, $v(x) \geq M' > 0$ for some M' . Then Γ is positive-valued and player 1 uniformly guarantees v with uniformly terminating strategies.*

4.2 Existence of the uniform value (proof of Theorem 3.1)

This subsection is devoted to the proof of Theorem 3.1: the total boundedness of $\{v_n, n \geq 1\}$ implies the existence of the uniform value v_∞ . We prove that player 1 guarantees the point-wise limit superior value $x \mapsto v(x) := \limsup_n v_n(x)$. By symmetry, player 2 guarantees $\liminf_n v_n(x)$, and the result follows.

The uniform ε -optimal strategy will use alternatively two different types of strategies. This approach is classical for recursive games and has been used for example in Rosenberg and Vieille

[12] and in Solan and Vieille [14]. Our construction is close to Solan and Vieille [14] in which some similar "positive-valued recursive game" is introduced to make a reduction for the general case.

The proof is decomposed into three parts. In the first one, we introduce a family of auxiliary positive-valued recursive games and define the first type of strategies. In the second part, we define the second type of strategies. Finally, we construct the strategy σ^* and prove that it is uniform ε -optimal.

Before proceeding to the proof, let us first prove a preliminary result, which shows that due to the total boundedness of $\{v_n\}$, the point-wise limit superior of (v_n) can be realized along uniform convergent subsequences. We fix a recursive game Γ for the rest of this section.

Proposition 4.5 *For every $x \in X$, we have*

$$v(x) = \limsup_n v_n(x) = \max_{f \in F} f(x),$$

where F is the set of limit points of the sequence $(v_n)_{n \geq 1}$ in $(\mathbf{B}(X), \|\cdot\|_\infty)$.

Proof. $(\mathbf{B}(X), \|\cdot\|_\infty)$ is a complete metric space and $(\{v_n\}, \|\cdot\|_\infty)$ is totally bounded, therefore F is compact and non-empty. For every $x \in X$, we denote $w(x) := \max_{f \in F} f(x)$. Fix $x \in X$. Since $v(x)$ is the largest limit point of $(v_n(x))_{n \geq 1}$, we have $w(x) \leq v(x)$. By definition of the limit superior, there exists a subsequence $(v_{n_k}(x))_{k \geq 1}$ which converges to $\limsup v_n(x)$. There exists a subsequence of $(v_{n_k})_{k \geq 1}$ that converges in $(\mathbf{B}(X), \|\cdot\|_\infty)$ to some $f^* \in F$, therefore

$$\max_{f \in F} f(x) \geq f^*(x) = v(x).$$

■

4.2.1 Reduction: auxiliary recursive games

Auxiliary recursive games Let $\theta : X \rightarrow \{0, 1\}$. We define the auxiliary recursive game $\Gamma^\theta = \langle A, B, X = X_\theta^0 \cup X_\theta^*, q_\theta, g_\theta \rangle$ where any active state $x \in X^0$ such that $\theta(x) = 1$ is seen as an absorbing state: the active state space of Γ^θ is $X_\theta^0 = \{x \in X^0, \theta(x) = 0\}$ and the absorbing state space is $X_\theta^* = X^* \cup \{x \in X^0, \theta(x) = 1\}$. The transition q_θ is equal to q and the payoff g_θ is equal to g on all states except $\{x \in X^0, \theta(x) = 1\}$, on which the state is absorbing and the absorbing payoff is $g_\theta = v$. For every $n \geq 1$, let v_n^θ be the value of the n -stage auxiliary game Γ_n^θ .

Proposition 4.6 *Let $\eta > 0$ and $\theta : X \rightarrow \{0, 1\}$. There exists $n_0 \geq 1$ such that for every $x_1 \in X_\theta^0$, there exists $n(x_1) \leq n_0$ with $v_{n(x_1)}^\theta(x_1) \geq v(x_1) - 4\eta$.*

Proof. Let $\eta > 0$ be fixed and $F_R = \{f_1, \dots, f_R\} \subseteq F$ be a finite cover of size $\frac{\eta}{2}$ of the set F . As $\{v_n, n \geq 1\}$ is totally bounded, there exists some stage $n(\eta) \in \mathbb{N}$, after which any n -stage value v_n is $\frac{\eta}{2}$ -close to F its set of accumulation points, hence η -close to F_R :

$$\exists n(\eta) \in \mathbb{N}, \forall n \geq n(\eta), \exists f_r \in \{f_1, \dots, f_R\}, \text{ s.t. } \|v_n - f_r\|_\infty \leq \eta; \quad (4.6)$$

Moreover for every $r \in \{1, \dots, R\}$, f_r is an accumulation point of $\{v_n, n \geq 1\}$, therefore there exists some $n_r > \frac{n(\eta)}{\eta}$ such that v_{n_r} is η -close to f_r :

$$\forall f_r \in F_R, \exists n_r > \frac{n(\eta)}{\eta}, \text{ s.t. } \|v_{n_r} - f_r\|_\infty \leq \eta. \quad (4.7)$$

Finally we take $n_0 = \max\{n_r : 1 \leq r \leq R\}$. The integers n_r are chosen such that when absorption in X_θ^* occurs in the game of length n_r the remaining number of stages is either a fraction smaller than η of the total length of the game or greater than $n(\eta)$ and Equations (4.6) applies.

Let $x_1 \in X_\theta^0$ be any non-absorbing state in the auxiliary game Γ^θ . By compactness of F , there exists $f \in F$ such that $f(x_1) = v(x_1)$ and $f_r \in F_R$ with $\|f - f_r\|_\infty \leq \frac{\eta}{2}$. In particular at state x_1 ,

$$f_r(x_1) \geq f(x_1) - \frac{\eta}{2} = v(x_1) - \frac{\eta}{2},$$

which together with (4.7) implies that

$$v_{n_r}(x_1) \geq f_r(x_1) - \eta \geq v(x_1) - \frac{3}{2}\eta. \quad (4.8)$$

We now prove that

$$v_{n_r}^\theta(x_1) \geq v_{n_r}(x_1) - 2\eta. \quad (4.9)$$

Denote by

$$\rho_\theta = \inf_{t \geq 1} \{x_t \in X_\theta^*\} = \inf_{t \geq 1} \{x_t \in X^* \text{ or } \theta(x_t) = 1\}$$

the stopping time associated to absorption in Γ^θ , and set $\rho_\theta^{n_r} = \min(\rho_\theta, n_r)$. An adaptation of standard proof technique of the Shapley equation gives us:

$$v_{n_r}(x_1) = \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \mathbb{E}_{x_1, \sigma, \tau} \left(\frac{1}{n_r} \left(\sum_{t=1}^{\rho_\theta^{n_r}-1} g(x_t) \right) + \frac{n_r - \rho_\theta^{n_r} + 1}{n_r} v_{n_r - \rho_\theta^{n_r} + 1}(x_{\rho_\theta^{n_r}}) \right).$$

We separate the histories into two sets depending on whether $n_r - \rho_\theta^{n_r}(h) + 1 > n(\eta)$ in which cases Equation (4.6) applies, or $n_r - \rho_\theta^{n_r}(h) + 1 \leq n(\eta)$ in which cases $\frac{n_r - \rho_\theta^{n_r}(h) + 1}{n_r} \leq \eta$ (by definition $n_r \geq \frac{n(\eta)}{\eta}$), and deduce that

$$v_{n_r}(x_1) \leq \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \mathbb{E}_{x_1, \sigma, \tau} \left(\frac{1}{n_r} \sum_{t=1}^{\rho_\theta^{n_r}-1} g(x_t) + \frac{n_r - \rho_\theta^{n_r} + 1}{n_r} f'_h(x_{\rho_\theta^{n_r}}) \right) + 2\eta,$$

with $f'_h \in F_R$ depending on the history given by Equation (4.6) applied to $v_{n_r - \rho_\theta^{n_r} + 1}$ when $n_r - \rho_\theta^{n_r} + 1 > n(\eta)$, and any function in F_r otherwise. Therefore, by considering v as the supremum of $f \in F$ at each point $x_{\rho_\theta^{n_r}} \in X$, we have $f'_h(x_{\rho_\theta^{n_r}}) \leq v(x_{\rho_\theta^{n_r}})$, thus

$$\begin{aligned} v_{n_r}(x_1) &\leq \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \mathbb{E}_{x_1, \sigma, \tau} \left(\frac{1}{n_r} \sum_{t=1}^{\rho_\theta^{n_r}-1} g(x_t) + \frac{n_r - \rho_\theta^{n_r} + 1}{n_r} v(x_{\rho_\theta^{n_r}}) \right) + 2\eta \\ &= v_{n_r}^\theta(x_1) + 2\eta. \end{aligned}$$

This proves inequality (4.9). We now use Equation (4.8) and Equation (4.9) to conclude:

$$v_{n_r}^\theta(x_1) \geq v(x_1) - 4\eta.$$

It means that for each $x_1 \in X_\theta^0$, there exists $n(x_1) := n_r \leq n_0 = \max\{n_r : 1 \leq r \leq R\}$, such that $v_{n(x_1)}^\theta(x_1) \geq v(x_1) - 4\eta$. \blacksquare .

Remark 4.7 *Proposition 4.6 is also true if θ is a deterministic stopping time and not only a function on the state. The auxiliary game would be defined on a larger state space: the set of finite histories of the original game. The proof in itself is similar.*

Fix now any $\varepsilon > 0$ and define $\theta_\varepsilon : X \rightarrow \{0, 1\}$ such that $\{x \in X, \theta_\varepsilon(x) = 1\} = \{x \in X, v(x) < \varepsilon\}$. We denote by $\Gamma^\varepsilon = \langle A, B, X = X_\varepsilon^0 \cup X_\varepsilon^*, q_\varepsilon, g_\varepsilon \rangle$ the auxiliary game associated to Γ defined by the stopping time θ_ε .

Corollary 4.8 *In the game Γ^ε , Player 1 uniformly guarantees v with uniformly terminating strategies that depend only on past states.*

Proof. Let $\eta \in (0, \varepsilon/8]$, by Proposition 4.6 there exists $n_0 \geq 1$ such that for every $x_1 \in X_\varepsilon^0$, there exists $n(x_1) \leq n_0$ with

$$v_{n(x_1)}^\varepsilon(x_1) \geq v(x_1) - 4\eta \geq \varepsilon/2, \quad (4.10)$$

where the second inequality comes from the definition of X_ε^0 . Therefore, Γ^ε is a positive-valued recursive game and by Proposition 4.3, player 1 uniformly guarantees $v_{n(\cdot)}^\varepsilon(\cdot)$ with uniformly terminating strategies in $\widehat{\Sigma}$. By Equation (4.10), it follows that for every $\eta > 0$, player 1 uniformly guarantees $v - 4\eta$ with uniformly terminating strategies. \blacksquare

Fix now a strategy σ_ε^* that is uniformly terminating in Γ^ε , depends only on past states and guarantees $v(x_1) - \varepsilon^2$ in $\Gamma^\varepsilon(x_1)$ for every $x_1 \in X_\varepsilon^0$.

4.2.2 One-shot game

One-shot game G^f For each $f : X \rightarrow [-1, +1]$ and $x_1 \in X$, we define the one-shot game G^f as follows: player 1's action set is $A(x_1)$, player 2's action set is $B(x_1)$, and the payoff is for each $(s, t) \in \Delta(A) \times \Delta(B)$,

$$\mathbb{E}_{q(x_1, s, t)}[f(x_2)] = \sum_{a \in A, b \in B} s(a)t(b) \left(\sum_{x_2 \in X} q(x_1, a, b)(x_2) f(x_2) \right).$$

Lemma 4.9 *For any limit point $f \in F$, the one-shot game G^f has a value equal to f .*

Proof. Let $n \geq 1$, it is known that (cf. Vigerat [18] p.40, Lemma 4.2.2)

$$\|v_n - v_{n+1}\|_\infty \leq \frac{2}{n+1},$$

and by Shapley's formula that

$$\begin{aligned} v_{n+1}(x_1) &= \sup_{s \in \Delta(A(x_1))} \inf_{t \in \Delta(B(x_1))} \mathbb{E}_{q(x_1, s, t)} \left[\frac{1}{n+1} g(x_1) + \frac{n}{n+1} v_n(x_2) \right] \\ &= \inf_{t \in \Delta(B(x_1))} \sup_{s \in \Delta(A(x_1))} \mathbb{E}_{q(x_1, s, t)} \left[\frac{1}{n+1} g(x_1) + \frac{n}{n+1} v_n(x_2) \right]. \end{aligned}$$

We obtain the result by taking the limit along a subsequence converging uniformly to $f \in F$. ■

Following Proposition 4.5, one can take for each $x \in X$ some $f^* \in F$ such that $v(x) = f^*(x) \geq f(x), \forall f \in F$. Then the following result is a direct consequence of Lemma 4.9.

Corollary 4.10 *For every $x_1 \in X$, there exists $s^*(x_1) \in \Delta(A(x_1))$ such that*

$$\forall b \in B(x_1), \mathbb{E}_{q(x_1, s^*(x_1), b)} [v(x_2)] \geq v(x_1).$$

Fix now $s^* := (s^*(x_1))_{x_1 \in X}$ a profile of strategies satisfying the conclusion of Corollary 4.10.

4.2.3 Optimal strategy

Roughly speaking, we build $\bar{\sigma}$ a uniform ε -optimal strategy for player 1 to play σ_ε^* in Γ^ε on the states with value v above 2ε , and to play s^* in G^v on the states with value v below ε . And for the states with value v between ε and 2ε , $\bar{\sigma}$ will be either of the two depending on the regime.

Construction of $\bar{\sigma}$ Define a sequence of stopping times $(u_l)_{l \geq 1}$ and the concatenated strategy $\bar{\sigma} := s^* u_1 \sigma_\varepsilon^* u_2 s^* u_3 \sigma_\varepsilon^* u_4 \dots$ in Γ as follows:

- $\bar{\sigma}$ is to play $s^*(x_n)$ at each stage n up to stage (not included)

$$u_1 = \inf\{n \geq 1, v(x_n) > 2\varepsilon\};$$

and then to play $\sigma_\varepsilon^*(x_{u_1})$ up to stage (not included)

$$u_2 = \inf\{n \geq u_1, v(x_n) < \varepsilon\}.$$

- In general: for each $r \geq 1$, $\bar{\sigma}$ is to play $\sigma_\varepsilon^*(x_{u_{2r-1}})$ from stage u_{2r-1} (*the odd phase*) up to stage (not included)

$$u_{2r} = \inf\{n \geq u_{2r-1}, v(x_n) < \varepsilon\}.$$

and then to play $s^*(x_n)$ at each stage $n \geq u_{2r}$ (*the even phase*), up to stage (not included)

$$u_{2r+1} = \inf\{n \geq u_{2r}, v(x_n) > 2\varepsilon\}.$$

Remark 4.11 *The idea of alternating between two types of strategies is common in Rosenberg and Vieille [12], Solan and Vieille [14] and this article. The main difference is the definition of the target function v used to define how to switch from one type of strategies to the other. Rosenberg and Vieille [12] use the limit of discounted values and σ_ε^* is an optimal strategy in some λ -discounted game (for λ close to zero). Solan and Vieille [14] use the limsup value and introduce an auxiliary positive-valued game. We adopt a similar approach to Solan and Vieille [14] but with v the largest limit point of (v_n) .*

By construction, $\bar{\sigma}$ depends on the histories only through the states and not the actions. Let us show that $\bar{\sigma}$ uniformly guarantees $v - 25\varepsilon$ for player 1, which finishes the proof of Theorem 3.1.

Fix from now on any $x_1 \in X$. Recall that ρ denotes the absorption time in the game Γ . The next result shows that the process $(v(x_{\min(\rho, u_l)}))_{l \geq 1}$, which is the value of v at switching times (u_l) , is almost a submartingale up to an error of ε^2 .

Proposition 4.12 *For every $l \geq 1$ and every $\tau \in \mathcal{T}$:*

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [v(x_{\min(\rho, u_{l+1})}) | \mathcal{H}_{\min(\rho, u_l)}] \geq v(x_{\min(\rho, u_l)}) - \varepsilon^2 \mathbb{1}_{\rho > u_l},$$

on the event $\min(\rho, u_l) < +\infty$.

Proof. Take any τ in \mathcal{T} . The result is true if $\rho \leq u_l$. Suppose that l is even and $\rho > u_l$: by construction the strategy $(s^*(x_n))$ is used during the phrase $n \in \{u_l, \dots, u_{l+1} - 1\}$, thus:

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [v(x_{n+1}) | \mathcal{H}_n] \geq v(x_n), \text{ for all } u_l \leq n < \min(\rho, u_{l+1}).$$

Therefore $(v(x_n))$ is a bounded submartingale and by Doob's stopping theorem,

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [v(x_{\min(\rho, u_{l+1})}) | \mathcal{H}_{\min(\rho, u_l)}] \geq v(x_{\min(\rho, u_l)}).$$

Suppose that l is odd and $\rho > u_l$. By construction, player 1 is using $\sigma_\varepsilon^*(x_{u_l})$, which uniformly guarantees $v(x_{u_l}) - \varepsilon^2$ in the auxiliary game $\Gamma^\varepsilon(x_{u_l})$:

$$\exists N_0 \geq 1, \quad \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{t=u_l+1}^{u_l+n} g_\varepsilon(x_t) | \mathcal{H}_{u_l} \right] \geq v(x_{u_l}) - \varepsilon^2 \text{ for all } n \geq N_0. \quad (4.11)$$

Denote by $\rho^\varepsilon = \min\{m \geq u_l + 1 : x_m \in X_\varepsilon^*\}$ the absorption time in $\Gamma^\varepsilon(x_{u_l})$. Since in recursive games the payoff is zero before absorption, we have

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [g_\varepsilon(x_{\rho^\varepsilon}) | \mathcal{H}_{u_l}] = \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=u_l+1}^{u_l+n} g_\varepsilon(x_t) | \mathcal{H}_{u_l} \right]. \quad (4.12)$$

By the dominated convergence theorem,

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=u_l+1}^{u_l+n} g_\varepsilon(x_t) | \mathcal{H}_{u_l} \right] = \lim_{n \rightarrow \infty} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{t=u_l+1}^{u_l+n} g_\varepsilon(x_t) | \mathcal{H}_{u_l} \right]. \quad (4.13)$$

We deduce from (4.11)-(4.13) that

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [g_\varepsilon(x_{\rho^\varepsilon}) | \mathcal{H}_{u_l}] \geq v(x_{u_l}) - \varepsilon^2.$$

Moreover, $g_\varepsilon(x_{\rho^\varepsilon}) = v(x_{\rho^\varepsilon})$ and conditionally on $\rho > u_l$, $\rho^\varepsilon = \min(u_{l+1}, \rho)$. It follows that

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} [v(x_{\min(\rho, u_{l+1})}) | \mathcal{H}_{u_l}] \geq v(x_{u_l}) - \varepsilon^2.$$

■

Due to the possible error term ε^2 , the sequence $(v(x_{\min(\rho, u_i)}))_{i \geq 1}$ is not a submartingale. Nevertheless, one can prove a lemma similar to the usual upcrossing lemma for submartingale. Indeed, the value is a martingale excepts if it crosses upwards the interval $[\varepsilon, 2\varepsilon]$. When this happens, the value may decrease of at most ε^2 . With the submartingale property established in Proposition 4.12, an easy adaptation of the standard result on upcrossing number of submartingale implies the following result, as was shown in Proposition 3 of Rosenberg and Vieille [12]:

Lemma 4.13 *Let $N = \sup\{p \geq 1 : u_{2p-1} < +\infty\}$ be the number of times the process $(v(x_{u_i}))$ crosses upward the interval $[\varepsilon, 2\varepsilon]$. For every $\tau \in \mathcal{T}$,*

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau}[N] \leq \frac{1}{\varepsilon - \varepsilon^2}.$$

By construction, σ_ε^* is *uniformly terminating* within the auxiliary absorbing states X_ε^* . That is to say, any play between stages u_{2p-1} and u_{2p} (on an odd phase) has bounded length with high probability under the strategy $\sigma_\varepsilon^*(x_{u_{2p-1}})$, uniformly over any starting state $x_{u_{2p-1}} \in X_\varepsilon^0$. Since Lemma 4.13 implies that the number of odd phases is bounded in expectations, the total frequency of stages on all odd phases is negligible for n large. Let us formalize this fact.

Recall that ρ^ε denotes the absorption time in the auxiliary game Γ^ε . It follows that there exists $N_1 > 0$ such that

$$\forall x \in X_\varepsilon^0 \text{ and } \tau \in \mathcal{T} : \mathbb{P}_{x, \sigma_\varepsilon^*(x), \tau}(\rho^\varepsilon > N_1) \leq \varepsilon^3. \quad (4.14)$$

For each $n \in \mathbb{N}$, define $A_n = \{u_{2p-1} \leq n < \min(\rho, u_{2p}), u_{2p-1} < \rho, \text{ for some } p\} \subseteq H_\infty$. These are all infinite plays where stage n is in an odd phrase, *i.e.*, the stages between u_{2p-1} and u_{2p} on which $\sigma_\varepsilon^*(x_{u_{2p-1}})$ is used. We fix for the rest of subsection the uniform stage number N_1 satisfying (4.14).

Lemma 4.14 *For every $\tau \in \mathcal{T}$ and every $n \geq \frac{N_1}{\varepsilon^3}$,*

$$\frac{1}{n} \sum_{k=1}^n \mathbb{P}_{x_1, \bar{\sigma}(x_1), \tau}(A_k) \leq 5\varepsilon.$$

The proof for this lemma relies on the upcrossing property established in Lemma 4.13, and takes the same form as Lemma 27 in Solan and Vieille [14]. Solan and Vieille [14] make some *finiteness* assumption (on the set of non-absorbing states on which the target function is not bounded away from zero) in order to obtain the existence of X_ε^1 a subset of X_ε^0 and a uniform bound $N_1 \geq 1$ such that

$$\forall x \in X_\varepsilon^1 \text{ and } \tau \in \mathcal{T} : \mathbb{P}_{x, \sigma_\varepsilon^*(x), \tau}(\rho^\varepsilon > N_1) \leq \varepsilon^3.$$

Under the assumption that $\{v_n, n \geq 1\}$ is totally bounded, we showed in Section 4.2.1 (cf. the condition defined in (4.14)) that we can consider X_ε^1 to be the whole set X_ε^0 .

The following result is a reformulation of the submartingale property in Lemma 4.12.

Lemma 4.15 *For any $m_0 \geq 1$, we have*

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau}[v(x_{m_0})] \geq v(x_1) - \varepsilon^2 \cdot \mathbb{E}_{x_1, \bar{\sigma}, \tau}[N] - 2\mathbb{P}_{x_1, \bar{\sigma}, \tau}(A_{m_0}) - \varepsilon.$$

Proof. For a proof, we refer to Proposition 28 in Solan and Vieille [14], where our lemma is stated as Equation (4) in their proof. \blacksquare

Now we use Lemma 4.13, Lemma 4.14 and Lemma 4.15 to prove the following proposition, which concludes the proof of Theorem 3.1.

Proposition 4.16 *For any $x_1 \in X^0$ and for any τ ,*

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{m=1}^n g(x_m) \right] \geq v(x_1) - 25\varepsilon, \quad \forall n \geq \frac{N_1}{\varepsilon^3}.$$

Proof. Take $x_1 \in X^0$ and fix any τ . In this proof h will denote a pure play. We use the fact that $g(x_m) \geq v(x_m) - 2\varepsilon$ if $h \notin A_m$: indeed, either the play has absorbed so $g(x_m) = v(x_m)$, or we have $v(x_m) < 2\varepsilon$ and $g(x_m) = 0$. Moreover, if $h \in A_m$, we use $g(x_m) \geq -1$. This gives us:

$$\begin{aligned} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{m=1}^n g(x_m) \right] &\geq \frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n \mathbb{1}_{h \notin A_m} (v(x_m) - 2\varepsilon) + \sum_{m=1}^n \mathbb{1}_{h \in A_m} (-1) \right] \\ &\geq \frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n v(x_m) \right] + \frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n \mathbb{1}_{h \in A_m} (-1 - (v(x_m) - 2\varepsilon)) \right] - 2\varepsilon. \end{aligned} \quad (4.15)$$

Lemma 4.15 (taking average sum on $m_0 = 1, \dots, n$) implies that

$$\frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n v(x_m) \right] \geq v(x_1) - \varepsilon^2 \cdot \mathbb{E}_{x_1, \bar{\sigma}, \tau} [N] - \frac{2}{n} \sum_{m=1}^n \mathbb{P}_{x_1, \bar{\sigma}, \tau}(A_m) - \varepsilon. \quad (4.16)$$

Moreover, the bound $v(x_m) \leq 1$ gives

$$\begin{aligned} \frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n \mathbb{1}_{h \in A_m} (-1 - v(x_m) + 2\varepsilon) \right] &\geq \frac{1}{n} \mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\sum_{m=1}^n \mathbb{1}_{h \in A_m} (-2 + 2\varepsilon) \right] \\ &= (-2 + 2\varepsilon) \frac{1}{n} \sum_{m=1}^n \mathbb{P}_{x_1, \bar{\sigma}, \tau}(A_m). \end{aligned} \quad (4.17)$$

We substitute (4.16) and (4.17) back into (4.15) to obtain

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{m=1}^n g(x_m) \right] \geq v(x_1) - \varepsilon^2 \cdot \mathbb{E}_{x_1, \bar{\sigma}, \tau} [N] - 3\varepsilon + (-4 + 2\varepsilon) \cdot \left(\frac{1}{n} \sum_{m=1}^n \mathbb{P}_{x_1, \bar{\sigma}, \tau}(A_m) \right).$$

Finally, we use Lemma 4.13 and Lemma 4.14 in the equality to have that: $\forall n \geq \frac{N_1}{\varepsilon^3}$ and $\forall \varepsilon \leq \frac{1}{2}$,

$$\mathbb{E}_{x_1, \bar{\sigma}, \tau} \left[\frac{1}{n} \sum_{m=1}^n g(x_m) \right] \geq v(x_1) - \frac{\varepsilon^2}{\varepsilon - \varepsilon^2} - 3\varepsilon - 20\varepsilon \geq v(x_1) - 25\varepsilon.$$

note that N_1 does not depend on the particular choice of x_1 in X^0 , so the strategy $\bar{\sigma}$ uniformly guarantees $v - 25\varepsilon$ in the infinite game Γ . \blacksquare

4.3 Pure optimal strategy (proof of Corollary 3.4)

To prove the result, it is sufficient to show that both the strategy s^* and the strategy σ_ε^* defined in the proof of Theorem 3.1 can be chosen pure and depending only on the history of states.

By assumption, the n -stage game $\Gamma_n(x)$ has a value in pure strategies. It follows that Shapley's equation for any v_n is satisfied with pure strategies, and so is Lemma 4.9. We deduce that there exists a pure action s^* that satisfies the conclusion of Corollary 4.10.

The construction of the strategy σ_ε^* appeared in the proof of Proposition 4.2, where it was defined as the concatenation of a sequence of strategies $(\hat{\sigma}(x_{u_\ell}))_{\ell \geq 1}$ at the random stages $(u_\ell)_{\ell \geq 1}$. As each $\hat{\sigma}(x)$ is optimal in the $n(x)$ -stage game $\Gamma_{n(x)}(x)$, $\hat{\sigma}(x_{u_\ell})$ can be taken pure. The definition of the random stages u_ℓ involved a randomized stopping time $\tilde{k} \in \{1, \dots, n(x)\}$ satisfying:

$$\forall \tau \in \mathcal{T}, \tilde{\mathbb{E}}_{x, \hat{\sigma}, \tau} [g(x_{\tilde{k}})] \geq \min_{\tau' \in \mathcal{T}} \mathbb{E}_{x, \hat{\sigma}, \tau'} \left[\frac{1}{n(x)} \sum_{t=1}^{n(x)} g(x_t) \right].$$

To obtain a pure strategy σ_ε^* , we show that the random stopping time \tilde{k} can be replaced by a stopping time (pure one), which depends only on the history of states and not on the actions. In order to build this stopping time, we restrict ourselves to strategies in $\widehat{\Sigma}$, *i.e.*, strategies which depend only on past states. Note that each $\hat{\sigma}(x_{u_\ell})$, as an optimal strategy in $\Gamma_{n(x_{u_\ell})}(x_{u_\ell})$, can be taken in $\widehat{\Sigma}$.

Lemma 4.17 *Fix any $\hat{\sigma} \in \widehat{\Sigma}$ and x_1 . For any $\tau \in \mathcal{T}$, there exists some $\hat{\tau} \in \widehat{\mathcal{T}}$ such that $\mathbb{P}_{x_1, \hat{\sigma}, \hat{\tau}}(x_1, \dots, x_t) = \mathbb{P}_{x_1, \hat{\sigma}, \tau}(x_1, \dots, x_t)$ for any $(x_1, \dots, x_t) \in X^t, t \geq 1$.*

Proof. For all $t \geq 1$, we denote by $s_t := (x_1, \dots, x_t)$ the t first states. For any $\tau \in \mathcal{T}$, define the reduced strategy $\hat{\tau} \in \widehat{\mathcal{T}}$ as:

$$\hat{\tau}_t(s_t) = \sum_{h_t \in H_t(s_t)} \mathbb{P}_{x_1, \hat{\sigma}, \tau}(h_t | s_t) \tau_t(h_t), \quad \forall s_t, \quad \forall t \geq 1.$$

where $H_t(s_t)$ denotes the histories in H_t containing s_t . Then we obtain by definition:

$$\mathbb{P}_{x_1, \hat{\sigma}, \hat{\tau}}(s_{t+1}) = \sum_{h_t \in H_t(s_t)} \mathbb{P}_{x_1, \hat{\sigma}, \tau}(h_t | s_t) \mathbb{P}_{x_1, \hat{\sigma}, \tau}(s_{t+1} | h_t) = \mathbb{P}_{x_1, \hat{\sigma}, \tau}(s_{t+1}).$$

■

Lemma 4.18 *Fix any $x_1 \in X^0$ and $\sigma \in \widehat{\Sigma}$. For any $n \geq 1$, there exists a stopping time $\theta : \cup_{1 \leq t \leq n} X^t \rightarrow \{1, \dots, n\}$ such that for every strategy τ of player 2:*

$$\mathbb{E}_{x_1, \sigma, \tau} [g(x_\theta)] \geq \min_{\tau'} \mathbb{E}_{x_1, \sigma, \tau'} \left[\frac{1}{n} \sum_{t=1}^n g(x_t) \right].$$

Proof. By Lemma 4.17, we can assume that $\tau \in \widehat{\mathcal{T}}$. Let us prove the result by induction. For every $x_1 \in X^0$, the result is true for $n = 1$. Suppose that the claim is true for $n - 1$. Let $x_1 \in X^0$. By applying the inductive assumption to the different states possible at stage 2, we obtain that there is some stopping time $\theta^+ : \cup_{t=1}^{n-1} X^t \rightarrow \{2, \dots, n\}$ such that

$$\mathbb{E}_{x_1, \sigma, \tau} [g(x_{\theta^+}) | x_2] \geq \min_{\tau'} \mathbb{E}_{x_1, \sigma, \tau'} \left[\frac{1}{n-1} \sum_{t=2}^n g(x_t) | x_2 \right] := w_{n-1}(\sigma, x_1, x_2). \quad (4.18)$$

Denote $w_{n-1}(\sigma, x_1) = \inf_{y \in \Delta(J)} \mathbb{E}_{x_1, \sigma, y} [w_{n-1}(\sigma, x_1, x_2)]$. We define the stopping time $\theta : \bigcup_{t=1}^n X^t \rightarrow \{1, \dots, n\}$ by

$$\forall (x_1, \dots, x_t) \in X^t, \quad \theta(x_1, \dots, x_t) = \begin{cases} 1 & \text{if } 0 \geq w_{n-1}(\sigma, x_1), \\ \theta^+(x_2, \dots, x_t) & \text{otherwise.} \end{cases}$$

According to the definition of θ and the inductive assumption (4.18) for θ^+ :

$$\begin{aligned} \mathbb{E}_{x_1, \sigma, \tau} [g(x_\theta)] &= g(x_1) \mathbb{1}_{0 \geq w_{n-1}(\sigma, x_1)} + \mathbb{E}_{x_1, \sigma, \tau} \left[\mathbb{E}_{x_1, \sigma, \tau} [g(x_{\theta^+}) | x_2] \right] \mathbb{1}_{0 < w_{n-1}(\sigma, x_1)} \\ &\geq \max \{0, w_{n-1}(\sigma, x_1)\} \\ &\geq \frac{n-1}{n} w_{n-1}(\sigma, x_1). \end{aligned}$$

Finally $g(x_1) = 0$, therefore

$$\frac{n-1}{n} w_{n-1}(\sigma, x_1) = \left(\frac{n-1}{n} \right) \inf_{\tau'} \mathbb{E}_{x_1, \sigma, \tau'} \left[\frac{1}{n-1} \sum_{t=2}^{n-1} g(x_t) \right] = \min_{\tau'} \mathbb{E}_{x_1, \sigma, \tau'} \left[\frac{1}{n} \sum_{t=1}^n g(x_t) \right].$$

This concludes the inductive proof. ■

Remark 4.19 *Let Γ be a stochastic game where the payoff function depends only on the state but not the actions, the proof for the above result follows the same way.*

5 Application to recursive games with signals

In this last section, we apply our result to the model of finite recursive games with signals where one player is more informed than the other player. Introducing an auxiliary stochastic game similar to the one defined in Gensbittel *et al.* [3], we show that the study of such a recursive game can be reduced to the study of recursive game with a countable state space satisfying the assumption of Corollary 3.4.

5.1 Model

The following model of general repeated game is introduced in Mertens *et al.* [8]. A repeated game $\Gamma = (K, I, J, C, D, g, q)$ is given by

- a finite state space: K .
- two finite action spaces I and J .
- two finite signal spaces C and D .
- a payoff function: $g : K \times I \times J \rightarrow [-1, +1]$.
- a transition probability function (on states and signals): q from $K \times I \times J$ to $\Delta(K \times C \times D)$.

Denote by $\Gamma(\pi)$ the game with an initial probability distribution $\pi \in \Delta(K \times C \times D)$, which is played as follows. Initially, the triple (k_1, c_1, d_1) is drawn according to π . At stage 1: player 1 learns c_1 and player 2 learns d_1 . Then simultaneously player 1 chooses an action $i_1 \in I$ and player 2 chooses an action $j_1 \in J$. The stage payoff is $g(k_1, i_1, j_1)$, and the new triple (k_2, c_2, d_2) is drawn according to $q(k_1, i_1, j_1)$. The game then proceeds to stage 2: player 1 observes c_2 , and player 2 observes d_2 etc...

We assume that each player's signal contains his own action. Formally, there exists $\hat{i}: C \rightarrow I$ and $\hat{j}: D \rightarrow J$ such that

$$\forall k \in K, \sum_{k', c, d} q(k, \hat{i}(c), \hat{j}(d))(k', c, d) = 1.$$

We will focus on repeated games with the following two features: *recursive* and *one player is more informed than the other*.

Definition 5.1 *The repeated game Γ is recursive if there exist K^0 and K^* , a partition of K such that:*

- *the stage payoff is 0 on active states: $\forall (k, i, j) \in K^0 \times I \times J, g(k, i, j) = 0$.*
- *states in K^* are absorbing: $\forall k \in K^*, \sum_{c \in C, d \in D} q(k, i, j)(k, c, d) = 1$ for all $(i, j) \in I \times J$ and $g(k, i, j)$ depends only on k .*

In the rest of the paper, a recursive repeated game will be called a recursive games with signals.

Definition 5.2 *Player 1 is more informed than player 2 in the recursive game Γ if there exists a mapping $\hat{d}: C \rightarrow D$ such that, if E denotes $\{(k, c, d) \in K \times C \times D, \hat{d}(c) = d\}$, then*

$$q(k, i, j)(E) = 1, \quad \forall (k, i, j) \in K \times I \times J.$$

Notation 5.3 *We denote by: $\Delta^1(K \times C \times D) = \{\pi | \pi(E) = 1\}$.*

We define similarly that player 2 is more informed than player 1. Whenever player 1 is more informed than player 2 and player 2 is more informed than player 1, Γ is a repeated game with *symmetric signals*. We denote by $\Delta^*(K \times C \times D)$ the set of symmetric initial distributions.

Remark 5.4 *By assumption, if player 1 is more informed than player 2, he learns especially the action played by player 2 since it is included in the signal of player 2. Player 2 is in general not informed of the action played by player 1.*

In Gensbittel et al. [3], the authors considered a weaker notion of "a more informed player" but they made a different assumption on the transition function, especially that the less informed player has no influence on the evolution of beliefs of both players. It is not clear if our result still holds under this weaker assumption.

5.2 Evaluation

At stage t , the space of past histories of player 1 is $H_t^1 = (C \times I)^{t-1} \times C$ and the space of past histories of player 2 is $H_t^2 = (D \times J)^{t-1} \times D$. Let $H_\infty = (K \times C \times D \times I \times J)^\infty$ be the space of infinite *plays*. For any play $h = (k_s, c_s, d_s, i_s, j_s)_{s \geq 1}$, we denote by h_t its projection on H_t , by h_t^1 its projection on H_t^1 , and by h_t^2 its projection on H_t^2 .

A (*behavior*) *strategy* for player 1 is a sequence $(\sigma_t)_{t \geq 1}$ of functions $\sigma_t : H_t^1 \rightarrow \Delta(I)$. A (*behavior*) *strategy* for player 2 is a sequence $\tau = (\tau_t)_{t \geq 1}$ of functions $\tau_t : H_t^2 \rightarrow \Delta(J)$. We denote by Σ and \mathcal{T} players' respective sets of strategies. An initial distribution $\pi \in \Delta(K \times C \times D)$ and a couple of strategies (σ, τ) define a probability distribution over the set of infinite plays, which we denote by $\mathbb{P}_{\sigma, \tau}^\pi$. Let $\mathbb{E}_{\sigma, \tau}^\pi$ be the expectation w.r.t. to $\mathbb{P}_{\sigma, \tau}^\pi$.

For any given $\pi \in \Delta(K \times C \times D)$, let $\gamma_n(\pi, \sigma, \tau)$ (resp. $\gamma_\lambda(\pi, \sigma, \tau)$) be the expected n -stage payoff (resp. λ -discounted payoff) associated with $(\sigma, \tau) \in \Sigma \times \mathcal{T}$. We denote by $v_n(\pi)$ the n -stage value and by $v_\lambda(\pi)$ the λ -discounted value.

Definition 5.5 *Given an initial distribution $\pi \in \Delta(K \times C \times D)$, the game $\Gamma(\pi)$ has an asymptotic value $v(\pi)$ if:*

$$v(\pi) = \lim_{n \rightarrow \infty} v_n(\pi) = \lim_{\lambda \rightarrow 0} v_\lambda(\pi).$$

Definition 5.6 *Given an initial distribution $\pi \in \Delta(K \times C \times D)$, the game $\Gamma(\pi)$ has a uniform maximin $\underline{v}_\infty(\pi)$ if:*

- *Player 1 can guarantee $\underline{v}_\infty(\pi)$, i.e., for all $\varepsilon > 0$ there exists a strategy $\sigma^* \in \Sigma$ of player 1 and $n_0 \geq 1$ such that*

$$\forall n \geq n_0, \forall \tau \in \mathcal{T}, \gamma_n(\pi, \sigma^*, \tau) \geq \underline{v}_\infty(\pi) - \varepsilon.$$

- *Player 2 can defend $\underline{v}_\infty(\pi)$, i.e., for all $\varepsilon > 0$ and for every strategy $\sigma \in \Sigma$ of player 1, there exists $n_0 \geq 1$ and $\tau^* \in \mathcal{T}$ such that*

$$\forall n \geq n_0, \gamma_n(\pi, \sigma, \tau^*) \leq \underline{v}_\infty(\pi) + \varepsilon.$$

The game $\Gamma(\pi)$ has a *uniform minmax* $\bar{v}_\infty(\pi)$ if similarly player 2 can guarantee $\bar{v}_\infty(\pi)$ and player 1 can defend $\bar{v}_\infty(\pi)$.

Definition 5.7 *Given an initial distribution $\pi \in \Delta(K \times C \times D)$, we say that $\Gamma(\pi)$ has a uniform value if both $\bar{v}_\infty(\pi)$ and $\underline{v}_\infty(\pi)$ exist and are equal. Whenever the uniform value exists, we denote it by $v_\infty(\pi)$.*

5.3 Results

Theorem 5.8 *Let Γ be a recursive game such that player 1 is more informed than player 2. Then for every distribution $\pi \in \Delta^1(K \times C \times D)$, both the asymptotic value and the uniform maximin exist and are equal:*

$$\underline{v}_\infty(\pi) = \lim v_n(\pi) = \lim v_\lambda(\pi)$$

By symmetry, we deduce a similar result by exchanging the roles of player 1 and player 2. When the information is symmetric, both results are true and we obtain the existence of the uniform value.

Corollary 5.9 *Let Γ be a recursive game with symmetric signals. Then for every $\pi \in \Delta^*(K \times C \times D)$, the game $\Gamma(\pi)$ has a uniform value.*

It is known from Ziliotto [19] that stochastic games with symmetric signals may have no uniform value. Therefore recursive games have very particular properties. It is a challenging task to identify the subclass of repeated games with $\underline{v}_\infty(\pi) = \lim v_n(\pi) = \lim v_\lambda(\pi)$.

Remark 5.10 *Note that we have assumed that the stage payoff on absorbing states does not depend on the actions played. Under this assumption, players' strategies have only an influence on non-absorbing plays. Therefore, without loss of generality, we assume in the following that players observe whenever an absorption occurs and in which state it is.*

If we consider that the payoff in absorbing states still depends on the actions played, then our proof does not work. Indeed the auxiliary game introduced in Proposition 5.16 is not recursive anymore. The result $\underline{v}_\infty(\pi) = \lim v_n(\pi) = \lim v_\lambda(\pi)$ is unknown for this general case.

Remark 5.11 *It is not known whether recursive games with any structure of signals have a uniform value. As highlighted in Rosenberg and Vieille [12], the equicontinuity of the λ -discounted value functions is sufficient in order to deduce the existence of the uniform value for recursive games (with perfect observations). For a recursive game with any structure of signals, one can introduce the game associated with a universal belief space but we do not know a metric on this space such that the λ -discounted values or the n -stage values are equicontinuous/totally bounded.*

5.4 Proof of Theorem 5.8

We introduce some notations concerning different belief hierarchies. Denote by $B_1 = \Delta(K)$ the set of beliefs of player 1 on the state variable. Denote by $B_2 = \Delta_f(B_1) = \Delta_f(\Delta(K))$ the set of beliefs of player 2 on the (first-order) beliefs of player 1. Finally, we denote by $\Delta_f(B_2) = \Delta_f(\Delta_f(\Delta(K)))$ the set of probability distributions over the second-order beliefs of player 2.

Overview of the proof

We fix Γ a recursive game with signals such that player 1 is more informed than player 2. The first subsection presents general properties for repeated games with one player more informed than the other. Given any $\pi \in \Delta^1(K \times C \times D)$, we can define the distribution of the beliefs of player 2 on the beliefs of player 1 about the state. This defines a function Φ from $\Delta^1(K \times C \times D)$ to $\Delta_f(B_2)$. Applying results in Gensbittel *et al.* [3], we know that $v_n(\pi)$ depends on π only through $\Phi(\pi)$. This enables us to show that the value function v_n , defined on $\Delta^1(K \times C \times D)$, induces a canonical function \hat{v}_n defined on B_2 such that $v_n(\pi) = \hat{v}_n(\Phi(\pi))$ and the family $\{\hat{v}_n, n \geq 1\}$ is totally bounded.

In the second subsection, we introduce an auxiliary recursive game \mathcal{G} which is defined on B_2 and is played with pure actions. We prove in Proposition 5.17 that the n -stage value of \mathcal{G}_n is equal to \hat{v}_n . Therefore, \mathcal{G} satisfies the conditions of Corollary 3.4 and it has a uniform value w_∞ . It follows that $\Gamma(\pi)$ has an asymptotic value equal to $w_\infty(\Phi(\pi))$.

The third subsection proves that (cf. Proposition 5.22) player 1 can uniformly guarantee $w_\infty(\Phi(\pi))$ in $\Gamma(\pi)$ by mimicking uniform ε -optimal strategies in $\mathcal{G}(\Phi(\pi))$.

The last subsection proves that (cf. Proposition 5.25) player 2 can uniformly defend $w_\infty(\Phi(\pi))$ by introducing a second auxiliary recursive game \mathcal{R} .

5.4.1 Canonical value function \hat{v}_n

We follow in this subsection Gensbittel *et al.* [3] to introduce the canonical function \hat{v}_n . Note that to obtain results in this subsection, the additional assumption that player 1 controls the transition (made later in their paper) is not used in Gensbittel *et al.* [3].

For convenience, we extend the definition of $\Gamma(\pi)$ to a larger family of initial probability distributions. Given any two finite sets C' and D' and $\pi \in \Delta^1(K \times C' \times D')$, $\Gamma(\pi)$ is the game where (k, c', d') is drawn at stage 1 according to π , player 1 observes c' , player 2 observes d' (which is contained in c' π -a.s.) and then from stage 2 on, the game is played as previously described with signals in C and D .

For any random variable ξ defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and \mathcal{F} a sub σ -algebra of \mathcal{A} , let $\mathcal{L}_{\mathbb{P}}(\xi | \mathcal{F})$ denote the conditional distribution of ξ given \mathcal{F} , which is seen as a \mathcal{F} -measurable random variable³ and let $\mathcal{L}_{\mathbb{P}}(\xi)$ denote the distribution of ξ .

Notation 5.12 For every strategy profile $(\sigma, \tau) \in \Sigma \times \mathcal{T}$, we denote the first-order belief of player 1 on K at stage n given h_n^1 by $p_n \in B_1$, the second-order belief of player 2, i.e., his belief about the belief of player 1 on K at stage n given h_n^2 by $x_n \in B_2$, and the distribution of x_n by $\eta_n \in \Delta_f(B_2)$, i.e.,

$$p_n \triangleq \mathcal{L}_{\mathbb{P}_{\sigma\tau}}(k_n | h_n^1), \quad x_n \triangleq \mathcal{L}_{\mathbb{P}_{\sigma\tau}}(p_n | h_n^2), \quad \text{and} \quad \eta_n \triangleq \mathcal{L}_{\mathbb{P}_{\sigma\tau}}(x_n).$$

Notation 5.13 For any $\pi \in \Delta^1(K \times C' \times D')$ where C' and D' are two finite sets, the image of π is given by the following function in $\Delta_f(B_2)$:

$$\begin{aligned} \Phi(\pi) &\triangleq \mathcal{L}_{\pi}(\mathcal{L}_{\pi}(\mathcal{L}_{\pi}(k_1 | c_1) | d_1)), \\ &= \sum_{d \in D'} \pi(d) \delta_{(\sum_{c \in C'} \pi(c|d) \delta_{\pi(\cdot | c, d)})}. \end{aligned}$$

The interpretation of $\Phi(\pi)$ is as follows: with probability $\pi(d)$, player 2 observes the signal d and believes that: player 1 received the signal c with probability $\pi(c|d)$ and therefore player 1's belief over K is $\pi(\cdot | c, d)$.

The assumptions imply that if $\pi \in \Delta^1(K \times C' \times D')$, then π satisfies the following two properties:

P1) $\pi(c)\pi(k, c, d) = \pi(k, c)\pi(c, d), \quad \forall (k, c, d) \in K \times C' \times D'$.

P2) There exists a map $f_1 = f_1^\pi : C' \rightarrow B_2$ such that $x_1 = f_1(c_1)$, π -almost surely.

Under P1) and P2), Proposition 1 of Gensbittel *et al.* [3] applies and we obtain the following result, which states that the value of any n -stage game depends on any initial distribution π only through its image $\Phi(\pi)$.

Proposition 5.14 [Gensbittel *et al.* 2014] Let C' and D' be two finite sets. Let $\pi, \pi' \in \Delta^1(K \times C' \times D')$ and let $n \geq 1$. If $\Phi(\pi) = \Phi(\pi')$, then $v_n(\pi) = v_n(\pi')$.

³All random variables appearing here take only finitely many values so that the definition of conditional laws does not require any additional care about measurability.

Reciprocally, given $\eta \in \Delta_f(B_2)$, let us construct a *canonical distribution* π satisfying $\Phi(\pi) = \eta$.

The canonical game $\hat{\Gamma}(\eta)$. Given $\eta \in \Delta_f(B_2)$. Define two finite sets $D' := \text{supp}(\eta) \subseteq B_2$ and $C' := D' \times (\bigcup_{x \in \text{supp}(\eta)} \text{supp}(x))$, and a probability distribution $\pi(\eta) \in \Delta(K \times C' \times D')$ by

$$\forall (k, p) \in K \times B_1, x, x' \in B_2, \pi(k, (p, x), x') := \begin{cases} \eta(x)x(p)p(k) & \text{if } x = x' \\ 0 & \text{if } x \neq x'. \end{cases}$$

By construction, $\pi(\eta)$ can be seen as an element of $\Delta^1(K \times C' \times D')$, and satisfies $\Phi(\pi(\eta)) = \eta$. The *canonical game* of $\Gamma(\pi)$ is denoted as $\hat{\Gamma}(\eta)$. Its value, denoted by $\hat{v}_n(\eta)$, is equal to $v_n(\pi(\eta))$ the value of $\Gamma_n(\pi(\eta))$. If $\eta = \delta_x$ for some $x \in B_2$, we denote $\hat{v}_n(x)$ for $\hat{v}_n(\delta_x)$.

Informally, the game $\hat{\Gamma}(\eta)$ proceeds as follows: η is common knowledge, player 2 is informed about the realization x of a random variable with law η (player 2 learns his beliefs). Then player 1 is informed about x (his opponent's beliefs) and about the realization p of a random variable with law x (his own beliefs). The state variable is finally chosen according to p , but no player observes it.

By the above construction, one obtains that: $v_n(\pi) = \hat{v}_n(\Phi(\pi))$ for any $\pi \in \Delta^1(K \times C' \times D')$.

The result below follows from Proposition 2 of Gensbittel *et al.* [3]. The Wasserstein metric \mathbf{d} on $B_2 = \Delta(B_1)$ is defined by:

$$\forall x, y \in B_2, \mathbf{d}(x, y) = \sup_{f \in \mathcal{D}} \left| \int_{B_1} f(p)x(dp) - \int_{B_1} f(p)y(dp) \right|,$$

where \mathcal{D} is the set of 1-Lipschitz function from $(B_1, \|\cdot\|_1)$ to $[-1, 1]$.

Proposition 5.15 [Gensbittel *et al.* 2014] *Let $\eta \in \Delta_f(B_2)$, $n \geq 1$ and let $x \in B_2$. Then, $\hat{v}_n(\eta)$ is linear on $\Delta_f(B_2)$ and, as a mapping on B_2 , $\hat{v}_n(x)$ is 1-Lipschitz for the Wasserstein metric \mathbf{d} .*

Since the state space B_2 is totally bounded for the Wasserstein metric, we deduce by Arzela-Ascoli theorem that the set of functions $\{\hat{v}_n, n \geq 1\}$ is totally bounded.

5.4.2 Auxiliary recursive game and asymptotic value

Let $\mathcal{G} = (X, A, B, G, \ell)$ be the stochastic game played in pure strategies, defined by:

- the state space $X = \Delta_f(\Delta(K))$ (endowed with the Wasserstein metric \mathbf{d}).
- the action space $A = \{f : \Delta(K) \rightarrow \Delta(I)\}$ and for all $x \in X$, $A(x) = \{\text{supp}(x) \rightarrow \Delta(I)\}$ for player 1.
- the action space $B = \Delta(J)$ for player 2.
- the payoff function $G : X \rightarrow [-1, 1]$, defined for any $x \in X$ by $G(x) := \sum_{p \in \Delta(X)} g(p)x(p)$.
- the transition function $\ell : X \times A \times B \rightarrow \Delta_f(X)$ defined as $\ell(x, a, b) := \Phi(Q(x, a, b))$. Here, $Q(x, a, b) \in \Delta_f(K \times (\Delta(K) \times C) \times D)$ is the joint distribution of $(k_2, (p, c_2), d_2)$ in the canonical game $\hat{\Gamma}(\delta_x)$ when the players play $(\sigma_1, \tau_1) = (a, b)$ at stage 1. The sets K, C, D and $\text{supp}(x)$ being finite, Q can be seen as an element in $\Delta^1(K \times C' \times D')$ with C' a finite subset of $\Delta(K) \times C$ and $D' = D$.

For any $x \in X$, we denote by $\mathcal{G}(x)$ the game starting at x . We extend the definition to $\mathcal{G}(z)$ for any $z \in \Delta_f(X)$ such that the initial state is chosen randomly along z .

Since players observe when and where absorption occurs, their beliefs (first and second-order) are either supported on K^0 (therefore respectively in $\Delta(K^0)$ and in $\Delta(\Delta(K^0))$) or supported on each single point $k \in K^*$ (to be δ_k and to be δ_{δ_k}).

Proposition 5.16 *Let $X_r = \Delta_f(\Delta(K^0)) \cup \{\delta_{\delta_k} : k \in K^*\}$. The game $\mathcal{G}^r = (X_r, A, B, G, \ell)$ with the state space X_r is well defined and is recursive with the absorbing states $\{\delta_{\delta_k} : k \in K^*\}$.*

In the following, we identify each δ_{δ_k} with k itself for any $k \in K^*$, and write $X_r = \Delta_f(\Delta(K^0)) \cup K^*$. By abuse of notations, we write again X for X_r and \mathcal{G} for \mathcal{G}^r .

Proposition 5.17 *For every $n \geq 1$, the n -stage game \mathcal{G}_n has a value w_n in pure strategies. Moreover, for every $x \in X$, $w_n(x) = \hat{v}_n(x)$.*

Proof. We prove the result by induction on $n \geq 1$. Let $n = 1$. Given $x \in X$, the game $\mathcal{G}_1(x)$ has a value $w_1(x)$ and it is equal to $w_1(x) = G(x) = \sum_p g(p)x(p)$. It is equal to $\hat{v}_1(x)$ by construction. This initializes our induction.

Let $n \geq 1$ such that w_n , the value of \mathcal{G}_n , exists in pure strategies, and for every $x \in X$, $w_n(x) = \hat{v}_n(x)$. Gensbittel *et al.* [3] showed in the proof of their Proposition 5 that the family $\{\hat{v}_n, n \geq 1\}$ satisfies the Shapley equation: for every $x \in X$ and for every $n \geq 1$,

$$\begin{aligned} \hat{v}_{n+1}(x) &= \sup_{a \in A(x)} \inf_{b \in B} \mathbb{E}_{\ell(x,a,b)} \left[\frac{1}{n+1} g(x, a, b) + \frac{n}{n+1} \hat{v}_n(x') \right] \\ &= \inf_{b \in B} \sup_{a \in A(x)} \mathbb{E}_{\ell(x,a,b)} \left[\frac{1}{n+1} g(x, a, b) + \frac{n}{n+1} \hat{v}_n(x') \right], \end{aligned}$$

where the random variable $x' \in X$ is chosen along the law $\ell(x, a, b)(\cdot)$. By the inductive assumption, we can replace \hat{v}_n by w_n on the right hand side of above equation, to obtain that:

$$\begin{aligned} \forall x \in X, \hat{v}_{n+1}(x) &= \sup_{a \in A(x)} \inf_{b \in B} \mathbb{E}_{\ell(x,a,b)} \left[\frac{1}{n+1} g(x) + \frac{n}{n+1} w_n(x') \right] \\ &= \inf_{b \in B} \sup_{a \in A(x)} \mathbb{E}_{\ell(x,a,b)} \left[\frac{1}{n+1} g(x) + \frac{n}{n+1} w_n(x') \right]. \end{aligned}$$

We now use the above equation to show that both players can guarantee $\hat{v}_{n+1}(x)$ in $\mathcal{G}_{n+1}(x)$ in pure strategies. Let $x \in X$ be fixed and a^* be an action of player 1 such that

$$\inf_{b \in B} \mathbb{E}_{\ell(x,a^*,b)} \left[\frac{1}{n+1} g(x) + \frac{n}{n+1} w_n(x') \right] \geq \hat{v}_{n+1}(x). \quad (5.1)$$

Again by inductive assumption, let $\sigma_n^*(x')$ be an optimal pure strategy in $\mathcal{G}_n(x')$, $\forall x' \in X$. We define the strategy $\sigma_{n+1}^*(x)$ to play a^* at the first stage and then $\sigma_n^*(x')$ where x' is the current state at stage 2. $\sigma_{n+1}^*(x)$ is pure and guarantees player 1 the payoff in $\mathcal{G}_{n+1}(x)$ no smaller than the left hand side of Equation (5.1), hence $\hat{v}_{n+1}(x)$. A similar construction for player 2 finishes the inductive proof. \blacksquare

Therefore, by Proposition 5.15 the family of n -stage values $\{w_n\}$ is totally bounded for the uniform norm, and we can apply Corollary 3.4 (with infinite sets of actions) for the game \mathcal{G} .

Proposition 5.18 *For every $z \in \Delta_f(X)$, the game $\mathcal{G}(z)$ has a uniform value denoted by $w_\infty^*(z)$. Moreover both players can uniformly guarantee the value with pure strategies that depend on the history of states but not on the past actions.*

Since $v_n(\pi) = \hat{v}_n(\Phi(\pi)) = w_n(\Phi(\pi))$, and the same construction of the canonical value function \hat{v}_λ implies $v_\lambda(\pi) = \hat{v}_\lambda(\Phi(\pi)) = w_\lambda(\Phi(\pi))$, we deduce the existence of the asymptotic value in $\Gamma(\pi)$ in the game $\Gamma(\pi)$ for every $\pi \in \Delta^1(K \times C \times D)$.

Proposition 5.19 *For every $\pi \in \Delta^1(K \times C \times D)$, we have*

$$\lim_{n \rightarrow \infty} v_n(\pi) = \lim_{\lambda \rightarrow 0} v_\lambda(\pi) = w_\infty^*(\Phi(\pi)).$$

5.4.3 Player 1 uniformly guarantees w_∞^*

We first show that player 1 is able to compute in the original game $(p_t)_{t \geq 1}$ his first-order beliefs and $(x_t)_{t \geq 1}$ the second-order beliefs of player 2 without knowing the strategy of player 2.

Lemma 5.20 *Let (σ, τ) be a pair of strategies in $\Gamma(\pi)$. For every $t \geq 1$, $p_t = \mathcal{L}_{\mathbb{P}_{\sigma, \tau}^\pi}(k_t | h_t^1)$ and $x_t = \mathcal{L}_{\mathbb{P}_{\sigma, \tau}^\pi}(p_t | h_t^2)$ are independent of τ for all h_t^1, h_t^2 .*

Proof. Let (σ, τ) be a pair of strategies and $\pi \in \Delta(K \times C \times D)$, we write $\mathbb{P} := \mathbb{P}_{\sigma, \tau}^\pi$ for short. Let $h = (k_s, c_s, d_s, i_s, j_s)_{s \geq 1} \in H_\infty$. For any $t \geq 1$, we define

$$\beta(h_t) = \pi(k_1, c_1, d_1) \prod_{\ell=1}^{t-1} q(k_\ell, i_\ell, j_\ell)(k_{\ell+1}, c_{\ell+1}, d_{\ell+1})$$

with the convention $\beta(k_1, c_1, d_1) = \pi(k_1, c_1, d_1)$. These notations help to write

$$\mathbb{P}(h_t) = \beta(h_t) \prod_{\ell=1}^{t-1} \sigma_\ell(h_\ell^1)[i_\ell] \tau_\ell(h_\ell^2)[j_\ell].$$

The key point is that under $\mathbb{P}(\cdot)$, i_{t-1}, j_{t-1} and d_t are c_t -measurable whereas j_{t-1} is d_t -measurable. It follows that after observing (c_1, \dots, c_t) , player 1's belief is:

$$\begin{aligned} p_t(k_t) &= \mathbb{P}(k_t | c_1, \dots, c_t) = \mathbb{P}(k_t | c_1, d_1, i_1, j_1, \dots, c_t, d_t) \\ &= \frac{\sum_{k'_1, \dots, k'_{t-1}} \mathbb{P}(k'_1, c_1, d_1, i_1, j_1, \dots, k_t, c_t, d_t)}{\sum_{k'_1, \dots, k'_{t-1}, k'_t} \mathbb{P}(k'_1, c_1, d_1, i_1, \dots, k'_t, c_t, d_t)} = \frac{\sum_{k'_1, \dots, k'_{t-1}} \beta(k'_1, c_1, d_1, i_1, j_1, \dots, k_t, c_t, d_t)}{\sum_{k'_1, \dots, k'_{t-1}, k'_t} \beta(k'_1, c_1, d_1, i_1, j_1, \dots, k'_t, c_t, d_t)}, \end{aligned}$$

which depends on neither σ nor τ . We now consider $x_t = \mathcal{L}_{\mathbb{P}}(p_t | d_1, \dots, d_t)$ for a given observed history $h_t^2 = (d_1, \dots, d_t)$ of player 2, which is decomposed as:

$$x_t = \mathcal{L}_{\mathbb{P}}\left(\mathcal{L}_{\mathbb{P}}(k_{t+1} | c_1, \dots, c_t) | d_1, \dots, d_t\right) = \sum_{c'_1, \dots, c'_t} \mathbb{P}(c'_1, \dots, c'_t | d_1, \dots, d_t) \delta_{\mathcal{L}_{\mathbb{P}}(k_{t+1} | c'_1, \dots, c'_t)}.$$

By the previous result that $\mathcal{L}_{\mathbb{P}}(k_{t+1} | c'_1, \dots, c'_t)$ does not depend on τ , it is sufficient to prove that $\mathbb{P}(c'_1, \dots, c'_t | d_1, \dots, d_t)$ is independent of τ . Let us consider a sequence of signals (c'_1, \dots, c'_t) inducing

(d_1, \dots, d_t) that we complete with (i'_1, \dots, i'_t) the sequence of actions it contains. This gives us

$$\begin{aligned} \mathbb{P}(c'_1, \dots, c'_t | d_1, \dots, d_t) &= \mathbb{P}(c'_1, d_1, i'_1, j_1, \dots, c'_t, d_t | d_1, j_1, \dots, d_t) = \frac{\mathbb{P}(c'_1, d_1, i'_1, j_1, \dots, c'_t, d_t)}{\mathbb{P}(d_1, j_1, \dots, d_t)} \\ &= \frac{\sum_{k'_1, \dots, k'_t} \beta(h'_t) \prod_{\ell=1}^{t-1} \sigma_\ell(h'_\ell) [i'_\ell]}{\sum_{k'_1, \dots, k'_t} \sum_{c'_1, \dots, c'_t} \beta(h'_t) \prod_{\ell=1}^{t-1} \sigma_\ell(h'_\ell) [i'_\ell]}, \end{aligned}$$

where $h'_\ell = (k'_1, c'_1, d_1, i'_1, j_1, \dots, k'_\ell, c'_\ell, d_\ell)$ is the history of stage ℓ and $h'_\ell{}^1 = (c'_1, i'_1, \dots, c'_\ell)$ is the private history of player 1 of stage ℓ . The right hand side of the above equation does not depend on the strategy of player 2 and we obtain the result. \blacksquare

Before building the strategy of player 1, we prove that the transition rule $\ell(\cdot) : X \times A \times B \rightarrow \Delta_f(X)$ of the auxiliary game is linear with respect to $b \in B$ (the action of player 2).

Lemma 5.21 *For any $(x, a) \in X \times A$ and $\bar{b} = \sum_{s \in S} \lambda_s b_s$ a convex combination in B , we have*

$$\ell(x, a, \bar{b}) = \sum_{s \in S} \lambda_s \ell(x, a, b_s).$$

Proof. Let $(x, a) \in X \times A$ and $b \in B$. Recall that $Q := Q(x, a, b)$ denotes a distribution in $\Delta_f(K \times (\Delta(K) \times C) \times D)$, which can be seen as an element in $\Delta^1(K \times C' \times D')$ with $C' = \text{supp}(x) \times C$ a finite subset of $\Delta(K)$ and $D' = D$. We have by definition of the image mapping $\Phi(\cdot)$: $\ell(x, a, b) = \Phi(Q) = \sum_{d' \in D'} Q(d') \delta_{\mathcal{L}_Q(\mathcal{L}_Q(k|c')|d')}$.

Similarly to the previous lemma, for every $(c', d') = ((p, c), d') \in C' \times D'$, $\mathcal{L}_Q(\mathcal{L}_Q(k|(p, c))|d')$ does not depend on b . Indeed, the signal (c', d') contains the action $(i_1, j_1) = (\hat{i}(c'), \hat{j}(d'))$ and c' contains d' a.s. It follows that

$$\mathbb{P}_Q(k|p, c) = \frac{a(p)[\hat{i}(c)]b[\hat{j}(c)]q^{K \times C}(p, \hat{i}(c), \hat{j}(c))(k, c')}{a(p)[\hat{i}(c)]b[\hat{j}(c)]q^C(p, \hat{i}(c), \hat{j}(c))(c')} = \frac{q^{K \times C}(p, \hat{i}(c), \hat{j}(c))(k, c')}{q^C(p, \hat{i}(c), \hat{j}(c))(c')}$$

and

$$\mathbb{P}_Q(p, c|d') = \frac{x(p)q^C(p, a(p), \hat{j}(d'))(c)}{\sum_{p \in \text{supp}(x)} x(p)q^D(p, a(p), \hat{j}(d'))(d')}.$$

Since these quantities do not depend on b , we will not precise b in the following. The application $Q(x, a, b)$ being linear in b , we can easily deduce the announced result:

$$\begin{aligned} \Phi(Q(x, a, \bar{b})) &= \sum_{d' \in D'} Q(x, a, \bar{b})(d') \delta_{\mathcal{L}_{Q(x, a, \cdot)}(\mathcal{L}_{Q(x, a, \cdot)}(k|c')|d')} \\ &= \sum_{d' \in D'} \left(\sum_{s \in S} \lambda_s Q(x, a, b_s)(d') \delta_{\mathcal{L}_{Q(x, a, \cdot)}(\mathcal{L}_{Q(x, a, \cdot)}(k|c')|d')} \right) \\ &= \sum_{s \in S} \lambda_s \Phi(Q(x, a, b_s)). \end{aligned}$$

We now use the two previous lemmas to prove that player 1 uniformly guarantees $w_\infty^*(\Phi(\pi))$ in the game $\Gamma(\pi)$. \blacksquare

Proposition 5.22 *Player 1 uniformly guarantees $w_\infty^*(\Phi(\pi))$ in $\Gamma(\pi)$.*

Proof. Fix any $\varepsilon > 0$. We divide the proof into three steps. First, we define the optimal strategy $\hat{\sigma}$ in $\Gamma(\pi)$. Then we show how to link the distribution over the states in \mathcal{G} to the distribution of second-order beliefs in Γ . Finally, we deduce that the strategy $\hat{\sigma}$ is uniform ε -optimal.

Step I: Defining the strategy.

Consider the auxiliary game $\mathcal{G}(z)$ with $z = \Phi(\pi) \in \Delta_f(X)$. According to Proposition 5.18, player 1 has pure uniform ε -optimal strategies which depend on histories only through the states but not the actions. With a slight abuse of notations, there exists $\hat{\sigma}^* : \bigcup_{t=1}^\infty X^t \rightarrow A = \{a : \Delta(K) \rightarrow \Delta(I)\}$ and $N_0 \geq 1$ such that

$$\hat{\gamma}_n(z, \hat{\sigma}^*, \hat{\tau}) \geq w_\infty^*(z) - \varepsilon \text{ for all } n \geq N_0 \text{ and for all } \hat{\tau} : \bigcup_{t=1}^\infty X^t \rightarrow B = \Delta(J)$$

where $\hat{\gamma}_n(z, \hat{\sigma}^*, \hat{\tau})$ is the expected n -stage average payoff in the auxiliary game $\mathcal{G}(z)$ induced by $(z, \hat{\sigma}^*, \hat{\tau})$.

We define the strategy $\sigma^* \in \Sigma$ in the game $\Gamma(\pi)$ such that for any h_t^1 ,

$$\sigma^*(h_t^1) = \hat{\sigma}^*(x_1, \dots, x_t)[p_t] \text{ with } p_t = \mathcal{L}_{\mathbb{P}_{\sigma^*}^\pi}(k_t|h_t^1) \text{ and } x_t = \mathcal{L}_{\mathbb{P}_{\sigma^*}^\pi}(p_t|h_t^2).$$

By Lemma 5.20, this is a well defined strategy of player 1 since he can compute p_t and x_t at every stage $t \geq 1$. We now check that the strategy σ^* uniformly guarantees $w_\infty^*(z) - \varepsilon$ in $\Gamma(\pi)$.

Step II: Linking the probability law of beliefs

Let $\tau \in \mathcal{T}$ be a strategy in $\Gamma(\pi)$. We define a strategy $\hat{\tau}$ in $\mathcal{G}(\Phi(\pi))$ such that (π, σ^*, τ) and $(\Phi(\pi), \hat{\sigma}, \hat{\tau})$ generate the same probability law for (x_1, \dots, x_t, \dots) . With a slight abuse in notation, we denote by $\hat{\tau}$ the strategy in \mathcal{G} such that for all $(x_1, \dots, x_t) \in X^t$,

$$\hat{\tau}(x_1, \dots, x_t) = \sum_{h_t^2 \in H_t^2(x_1, \dots, x_t)} \mathbb{P}_{\sigma^*, \tau}^\pi(h_t^2|x_1, \dots, x_t)\tau(h_t^2),$$

where $H_t^2(x_1, \dots, x_t) = \{h_t^2 \in H_t^2 | \mathcal{L}_{\mathbb{P}_{\sigma^*, \tau}^\pi}(k_\ell|h_\ell^2) = x_\ell, 1 \leq \ell \leq t\}$ denotes the set of player 2's t -stage histories in Γ that induce the beliefs (x_1, \dots, x_t) .

Lemma 5.23 *Let σ^* and $\hat{\tau}$ be constructed as above given $\hat{\sigma}^*$ and τ , we have:*

$$\forall t \geq 1, \mathcal{L}_{\mathbb{P}_{\sigma^*, \tau}^\pi}(x_1, \dots, x_t) = \mathcal{L}_{\mathbb{P}_{\hat{\sigma}^*, \hat{\tau}}^z}(x_1, \dots, x_t).$$

Proof of Lemma 5.23: We prove the lemma by induction on $t \geq 1$. For $t = 1$, the law of x_1 is independent of the strategy profile. By definition of the image mapping $\Phi(\cdot)$,

$$\mathcal{L}_{\mathbb{P}_{\sigma^*, \tau}^\pi}(x_1) = \mathcal{L}_\pi(\mathcal{L}_\pi(\mathcal{L}_\pi(k_1|c_1)|d_1)) = \Phi(\pi).$$

As $\Phi(\pi) = z$, the probability law to choose the initial state $x_1 \in X$ in $\mathcal{G}(z)$, $\mathcal{L}_{\mathbb{P}_{\hat{\sigma}^*, \hat{\tau}}^z}(x_1) = \Phi(\pi)$.

Suppose now that we have proved that $\mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_1, \dots, x_t) = \mathcal{L}_{\hat{\sigma}^*, \hat{\tau}}^{\mathbb{P}^z}(x_1, \dots, x_t)$ for some $t \geq 1$. It is then sufficient to prove that conditional on any realization⁴ $\tilde{s}_t := (\tilde{x}_1, \dots, \tilde{x}_t) \in (B_2)^t$,

$$\mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_{t+1}|x_1 = \tilde{x}_1, \dots, x_t = \tilde{x}_t) = \mathcal{L}_{\hat{\sigma}^*, \hat{\tau}}^{\mathbb{P}^z}(x_{t+1}|x_1 = \tilde{x}_1, \dots, x_t = \tilde{x}_t).$$

Fix some $\tilde{s}_t = (\tilde{x}_1, \dots, \tilde{x}_t) \in X^t$. By definition of $\hat{\tau}$ and the linearity of ℓ showed in Lemma 5.21, we know that

$$\begin{aligned} \mathcal{L}_{\hat{\sigma}^*, \hat{\tau}}^{\mathbb{P}^z}(x_{t+1}|x_1 = \tilde{x}_1, \dots, x_t = \tilde{x}_t) &= \ell(\tilde{x}_t, \hat{\sigma}^*(\tilde{s}_t), \hat{\tau}(\tilde{s}_t)) \\ &= \sum_{h_t^2 \in H_t^2(\tilde{s}_t)} \mathbb{P}_{\sigma^*, \tau}^\pi(h_t^2|\tilde{s}_t) \ell(\tilde{x}_t, \hat{\sigma}^*(\tilde{s}_t), \tau(h_t^2)). \end{aligned} \quad (5.2)$$

By definition of the conditional expectation, we have in Γ ,

$$\mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_{t+1}|x_1 = \tilde{x}_1, \dots, x_t = \tilde{x}_t) = \sum_{h_t^2 \in H_t^2(\tilde{s}_t)} \mathbb{P}_{\sigma^*, \tau}^\pi(h_t^2|\tilde{s}_t) \mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_{t+1}|h_t^2).$$

Thus, it is sufficient to prove that for every $h_t^2 \in H_t^2(\tilde{s}_t)$, $\ell(\tilde{x}_t, \hat{\sigma}^*(\tilde{s}_t), \tau(h_t^2)) = \mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_{t+1}|h_t^2)$. Let $h_t^2 \in H_t^2(\tilde{x}^t)$ and $Q[h_t^2] := Q(\tilde{x}_t, \hat{\sigma}^*(\tilde{s}_t), \tau(h_t^2)) \in \Delta_f(K \times (\Delta(K) \times C) \times D)$ the joint distribution of $(k_{t+1}, (p_t, c_{t+1}), d_{t+1})$ in the canonical game $\hat{\Gamma}(\delta_{\tilde{x}_t})$ when $(\hat{\sigma}^*(\tilde{s}_t), \tau(h_t^2)) \in A \times B$ is played. By definition of the image mapping $\Phi(\cdot)$ and σ^* , we obtain

$$\mathcal{L}_{\sigma^*, \tau}^{\mathbb{P}^\pi}(x_{t+1}|h_t^2) = \mathcal{L}_{Q[h_t^2]}(\mathcal{L}_{Q[h_t^2]}(k_{t+1}|c_{t+1})|d_{t+1}) = \Phi(Q[h_t^2]) = \ell(\tilde{x}_t, \hat{\sigma}^*(\tilde{s}_t), \tau(h_t^2)).$$

□

Step III: Conclusion of the proof

Finally, let us compare the payoffs in both games. If $k^* \in K^*$, we have $G(k^*) = g(k^*) = \mathbb{E}_{\sigma^*, \tau}^\pi[g(k_t)|x_t = k^*]$. If $x_t \in (\Delta_f(\Delta(K^0)))$, we have $G(x_t) = 0 = \mathbb{E}_{\sigma^*, \tau}^\pi[g(k_t)|x_t]$. It follows that for every $x_t \in X$, we have $G(x_t) = \mathbb{E}_{\sigma^*, \tau}^\pi[g(k_t)|x_t]$. By taking conditional expectation, Lemma 5.23 implies that $\mathbb{E}_{\hat{\sigma}^*, \hat{\tau}}^z[G(x_t)] = \mathbb{E}_{\sigma^*, \tau}^\pi[g(k_t)]$. Since $\hat{\sigma}^*$ is uniform ε -optimal in the auxiliary game $\mathcal{G}(\Phi(\pi))$, we obtain

$$\gamma_n(\pi, \sigma^*, \tau) = \hat{\gamma}_n(\Phi(\pi), \hat{\sigma}^*, \hat{\tau}) \geq w_\infty^*(\Phi(\pi)) - \varepsilon \text{ for all } n \geq N_0.$$

Therefore, the strategy σ^* uniformly guarantees $w_\infty^*(\Phi(\pi)) - \varepsilon$ in $\Gamma(\pi)$. ■

5.4.4 Player 2 uniformly defends w_∞^*

We now prove that player 2 can defend $w_\infty^*(\Phi(\pi)) = \lim v_n(\pi) = \lim v_\lambda(\pi)$. The situation of player 2 is different since he is allowed to know the strategy of player 1. In order to prove this result, we introduce another auxiliary recursive game \mathcal{R} .

For any $n \geq 1$, let $H_n' \subseteq H_n$ be the set of n -stage histories such that player 1 can deduce player 2's private signals, and $H_n^0 \subseteq H_n$ be the set of n -stage histories containing only non-absorbing states. We consider the following game \mathcal{R} where the set of states is almost the set of distribution over all finite histories. It is defined as follows:

⁴For this part of the proof it is convenient to differentiate the random variable describing the second order belief (or the state in \mathcal{G}) that will be denoted by x_t from its realization denoted by \tilde{x}_t .

- the state space is $Z = Z_0 \cup K^*$ where $Z_0 = \bigcup_{n \geq 1} \Delta(H_n^0 \cap H_n')$,
- the action space of player 1 is $A = \bigcup_{n \geq 1} \{f : H_n^1 \rightarrow \Delta(I)\}$ and for any $\pi_n \in \Delta(H_n)$, $A(\pi_n) = \{f : H_n^1 \rightarrow \Delta(I)\}$,
- the action space of player 2 is $B = \bigcup_{n \geq 1} \{f : H_n^2 \rightarrow \Delta(J)\}$ and for any $\pi_n \in \Delta(H_n)$, $B(\pi_n) = \{f : H_n^2 \rightarrow \Delta(J)\}$,
- the transition $Q : Z \times A \times B \rightarrow \Delta_f(Z)$ is given by:

$$\forall (k^*, a, b) \in K^* \times A \times B, Q(k^*, a, b) = \delta_{k^*},$$

and

$$\forall (z, a, b) \in Z_0 \times A \times B, Q(z, a, b) = Q^0(z, a, b)\delta_{\pi^0} + \sum_{k \in K^*} Q(z, a, b)(k^*)\delta_{k^*},$$

where $Q(z, a, b)(k^*)$ is the probability of absorption in state k^* at the next stage given by

$$Q(z, a, b)(k^*) = \sum_{h_n, i, j, c, d} z(h_n) a(h_n^1)[i] b(h_n^2)[j] q(k_n, i, j)(k^*, c, d);$$

$Q^0(z, a, b)$ is the probability of no absorption given by

$$Q^0(z, a, b) = \sum_{h_n, i, j, c, d} \sum_{k \in K^0} z(h_n) a(h_n^1)[i] b(h_n^2)[j] q(k_n, i, j)(k, c, d),$$

and $\pi^0 \in Z_0$ is the conditional probability on not having absorbed, *i.e.*,

$$\forall (h_n, k, i, j, c, d) \in H_n \times K \times I \times J \times C \times D, \pi^0(h_n, i, j, c, d) = \frac{z(h_n) a(h_n^1)[i] b(h_n^2)[j] q(k_n, i, j)(k, c, d)}{Q^0(z, a, b)}.$$

- the stage payoff function $R : Z \times A \times B \rightarrow [-1, +1]$ is given by

$$\forall (\pi_n, a, b) \in Z_0 \times A \times B, R(\pi_n, a, b) = 0,$$

and

$$\forall (k^*, a, b) \in K^* \times A \times B, R(k^*, a, b) = g(k^*).$$

By construction, the game \mathcal{R} is recursive. We denote by $\tilde{\Sigma}$ (resp. $\tilde{\mathcal{T}}$) the set of behavior strategy for player 1 (resp. for player 2) in the game \mathcal{R} .

Proposition 5.24 *For every $\pi \in Z_0$ and every $n \geq 1$, the n -stage game $\mathcal{R}_n(\pi)$ has a value in history independent pure strategies, which is denoted by $\tilde{v}_n(\pi)$ and*

$$\tilde{v}_n(\pi) = v_n(\pi).$$

Moreover, if player 2 can uniformly defend some payoff level v in the game $\mathcal{R}(\pi)$ with pure strategies then he can also uniformly defend v in the game $\Gamma(\pi)$.

Proof. First, a strategy σ of player 1 in Γ is a sequence of applications $(\sigma_n)_{n \geq 1}$ such that σ_n is a mapping from H_1^n to $\Delta(I)$. By definition, this is a sequence of actions in the game \mathcal{R} , *i.e.*, a history independent pure strategy in \mathcal{R} . Similarly, a strategy τ of player 2 in Γ induces a sequence of actions in \mathcal{R} . By definition of Q and R , it follows that for every $\pi \in Z_0$, $\sigma \in \Sigma$ and $\tau \in \mathcal{T}$,

$$\gamma_n(\pi, \sigma, \tau) = \tilde{\gamma}_n(\pi, \sigma, \tau). \quad (5.3)$$

Let σ be an optimal strategy of player 1 in the game $\Gamma_n(\pi)$. Consider now a pure strategy $\tilde{\tau} \in \tilde{\mathcal{T}}$. The triple $(\pi, \sigma, \tilde{\tau})$ generates a probability distribution \mathbb{P} on $(Z \times A \times B)^{\mathbb{N}}$ such that there exists at most one play $(\pi_t, a_t, b_t)_{t \geq 1}$ that is non absorbing \mathbb{P} -*a.s.*, *i.e.*, $(\pi_t, a_t, b_t)_{t \geq 1} \in (Z_0 \times A \times B)^{\mathbb{N}}$. Define the strategy $\tau \in \mathcal{T}$ of player 2 in $\Gamma(\pi)$ by setting $\tau_t = b_t$ for all $t \geq 1$. We obtain

$$\tilde{\gamma}_n(\pi, \sigma, \tilde{\tau}) = \tilde{\gamma}_n(\pi, \sigma, \tau) = \gamma_n(\pi, \sigma, \tau) \geq v_n(\pi).$$

Therefore, player 1 guarantees the payoff $v_n(\pi)$ in $\mathcal{R}(\pi)$ with the history independent pure strategy σ . Similarly, player 2 can guarantee $v_n(\pi)$ with a history independent pure strategy and $\tilde{v}_n(\pi) = v_n(\pi)$.

Finally, let us assume that player 2 can uniformly defend the payoff level v with pure strategies in the game $\mathcal{R}(\pi)$. Let $\varepsilon > 0$ and $\sigma \in \Sigma$. Interpreting σ as an history-independent strategy in \mathcal{R} , there exist $N_0 \geq 1$ and a pure strategy $\tilde{\tau} \in \tilde{\mathcal{T}}$ such that

$$\forall n \geq N_0, \tilde{\gamma}_n(\pi, \sigma, \tilde{\tau}) \leq v + \varepsilon. \quad (5.4)$$

As in the previous paragraph, we can associate to the triple $(\pi, \sigma, \tilde{\tau})$ a unique play $(\pi_t, a_t, b_t)_{t \geq 1}$ in $(Z_0 \times A \times B)^{\mathbb{N}}$ and define the strategy $\tau \in \mathcal{T}$ of player 2 in $\Gamma(\pi)$ by setting $\tau_t = b_t$ for all $t \geq 1$. We obtain

$$\forall n \geq N_0, \gamma_n(\pi, \sigma, \tau) = \tilde{\gamma}_n(\pi, \sigma, \tau) = \tilde{\gamma}_n(\pi, \sigma, \tilde{\tau}) \leq v + \varepsilon.$$

This proves that player 2 can uniformly defend v in $\Gamma(\pi)$. ■

We conclude by showing that the game \mathcal{R} fulfills the conditions of Corollary 3.4.

Proposition 5.25 *Player 2 uniformly defends $w_\infty^*(\Phi(\pi))$ in $\Gamma(\pi)$.*

Proof. We already noticed that the game \mathcal{R} is recursive. Let $\pi \in \Delta(H_n^0 \cap H_n') \subseteq Z_0$ for some $n \geq 1$. Since player 1 is more informed than player 2 (π supported on H_n'), π can be identified as an element in $\Delta^1(K \times C' \times D')$ for some finite C' and D' . By Proposition 5.24, we obtain that for any $\pi \in Z_0$,

$$\tilde{v}_n(\pi) = v_n(\pi) = \hat{v}_n(\Phi(\pi)).$$

According to Corollary 5.15, the family $\{\hat{v}_n, n \geq 1\}$ considered as functions on B_2 is totally bounded, and so is the family of their linear extensions to $\Delta_f(B_2)$.

By Corollary 3.4, $\mathcal{R}(\pi)$ has a uniform value $w_\infty^*(\Phi(\pi))$ in pure strategies for every $\pi \in \Delta^1(K \times C \times D)$. It follows from Proposition 5.24 that player 2 can uniformly defend $w_\infty^*(\Phi(\pi))$ in $\Gamma(\pi)$. ■

Acknowledgements The authors thank Sylvain Sorin for his careful reading of earlier versions of this paper, whose comments have significantly improved its presentation. The authors also thank an associated editor and an anonymous referee for their numerous helpful remarks. The authors gratefully acknowledge the support of the Agence Nationale de la Recherche, under grant ANR JEUDY, ANR-10-BLAN 0112. The second author acknowledges the support of the Israel Science Foundation under Grant #1517/11.

References

- [1] T. Bewley and E. Kohlberg. The asymptotic theory of stochastic games. *Mathematics of Operations Research*, 1:197–208, 1976.
- [2] H. Everett. Recursive games. *Contributions to the Theory of Games III*, 39:47–78, 1957.
- [3] F. Gensbittel, M. Oliu-Barton, and X. Venel. Existence of the uniform value in repeated games with a more informed controller. *Journal of Dynamics and Games*, 1:411–445, 2014.
- [4] E. Lehrer and D. Monderer. Discounting versus averaging in dynamic programming. *Games and Economic Behavior*, 6:97–113, 1994.
- [5] E. Lehrer and S. Sorin. A uniform tauberian theorem in dynamic programming. *Mathematics of Operations Research*, 17:303–307, 1992.
- [6] J.-F. Mertens. Repeated games. In *Proceedings of the International Congress of Mathematicians, (Berkeley, 1986)*, pages 1528–1577. Amer. Math. Soc., 1987.
- [7] J.-F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10:53–66, 1981.
- [8] J.-F. Mertens, S. Sorin, and S. Zamir. *Repeated games*. Cambridge University Press, 2015.
- [9] D. Monderer and S. Sorin. Asymptotic properties in dynamic programming. *International Journal of Game Theory*, 22:1–11, 1993.
- [10] J. Renault. The value of repeated games with an informed controller. *Mathematics of Operations Research*, 37:154–179, 2012.
- [11] D. Rosenberg, E. Solan, and N. Vieille. Stochastic games with a single controller and incomplete information. *SIAM Journal on Control and Optimization*, 43:86–110, 2004.
- [12] D. Rosenberg and N. Vieille. The Maxmin of recursive games with incomplete information on one side. *Mathematics of Operations Research*, 39:23–35, 2000.
- [13] L.S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. U. S. A.*, 39:1095–1100, 1953.
- [14] E. Solan and N. Vieille. Uniform value in recursive games. *The Annals of Applied Probability*, 12:1185–1201, 2002.

- [15] S. Sorin. "Big Match" with lack of information on one side (i). *International Journal of Game Theory*, 13:201–255, 1984.
- [16] S. Sorin. "Big Match" with lack of information on one side (ii). *International Journal of Game Theory*, 14:173–204, 1985.
- [17] S. Sorin. *A First Course on Zero-Sum Repeated Games*. Springer, 2002.
- [18] G Vigerál. *Propriétés asymptotiques des jeux répétés à somme-nulle*. Doctoral thesis of Université Paris 6, 2009.
- [19] B. Ziliotto. Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture $\max\min = \lim v(n)$. *arXiv:1305.4778*, 2013.
- [20] B. Ziliotto. A Tauberian theorem for nonexpansive operators and applications to zero-sum stochastic games. *arXiv:1501.06525*, 2015.