



Optimal strategies in repeated games with incomplete information: the dependent case

Fabien Gensbittel, Miquel Oliu-Barton

► To cite this version:

Fabien Gensbittel, Miquel Oliu-Barton. Optimal strategies in repeated games with incomplete information: the dependent case. 2014. hal-01091666v2

HAL Id: hal-01091666

<https://hal.science/hal-01091666v2>

Preprint submitted on 8 Jun 2021

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

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

Optimal strategies in zero-sum repeated games with incomplete information: the dependent case

FABIEN GENSBITTEL*, MIQUEL OLIU-BARTON†

June 8, 2021

Abstract. Using the duality techniques introduced by De Meyer (1996a, 1996b), Rosenberg (1998) and De Meyer and Marino (2005) provided an explicit construction for optimal strategies in repeated games with incomplete information on both sides, in the independent case. In this note, we extend both the duality techniques and the construction of optimal strategies to the dependent case.

1 Introduction

We consider here a zero-sum repeated game with incomplete information on both sides, in the spirit of Aumann and Maschler [1]. Let K (resp. L) be the finite set of types of Player 1 (resp. 2), and let π be a probability distribution over $K \times L$. To any pair (k, ℓ) corresponds a matrix game $G^{k\ell} : I \times J \rightarrow \mathbb{R}$, where I (resp. J) is the finite set of actions of Player 1 (resp. 2). The game is played as follows. First, a pair $(k, \ell) \in K \times L$ is drawn with the probability distribution π . Player 1 (resp. 2) is informed only of k (resp. ℓ). Then, the game $G^{k\ell}$ is played repeatedly. At each stage $m \geq 1$, the players choose actions $(i_m, j_m) \in I \times J$, which produces a stage-payoff $G^{k\ell}(i_m, j_m)$. Actions are publicly observed after each stage. For any initial distribution π and any sequence of non-negative weights $\theta = (\theta_m)_m$, we consider the game $\mathcal{G}_\theta(\pi)$ in which the overall payoff is the expected θ -weighted sum of the stage-payoffs $\sum_{m \geq 1} \theta_m G^{k\ell}(i_m, j_m)$ and where π stands for the probability distribution of (k, ℓ) . This game has a value, denoted by $v_\theta(\pi)$. The particular case where, for some $n \in \mathbb{N}^*$, one has $\theta_m = \frac{1}{n} \mathbb{1}_{\{m \leq n\}}$ for all $m \geq 1$ corresponds to the classical n -stage repeated games. Similarly, the case where $\theta_m = \lambda(1 - \lambda)^{m-1}$ for all $m \geq 1$ and some $\lambda \in (0, 1]$ corresponds to λ -discounted repeated games. We use then the notation $\mathcal{G}_n(\pi)$ and $v_n(\pi)$, and $\mathcal{G}_\lambda(\pi)$ and $v_\lambda(\pi)$, respectively.

This model was analyzed by Mertens and Zamir in [7]. Their main result was the existence of $v_n(\pi)$ and $v_\lambda(\pi)$ and their convergence (respectively, as n goes to $+\infty$ and as λ vanishes) to

*Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France, email: fabien.gensbittel@tse-fr.eu

†Université Paris-Dauphine, PSL Research University, CNRS, CEREMADE, 75016 Paris, France, email: miquel.oliu.barton@normalesup.org

the unique solution of a system of functional equations. The proof of this result was based on the introduction of the specific notion of I -concavity for the value function $\pi \mapsto v_n(\pi)$, which can be described as follows. Any probability π over the product set $K \times L$ can be decomposed as a pair (p, Q) where p is the marginal probability on K and Q is a matrix of conditional probabilities on L given $k \in K$. This decomposition can be expressed as $\pi = p \otimes Q$ where $\otimes$ denotes the direct product. One may then consider v_n as a function of (p, Q) , and show that $p \mapsto v_n(p, Q) := v_n(p \otimes Q)$ is a concave function for any fixed Q . A dual notion of $\mathbb{I}$ -convexity was also introduced and the notions of I -concave and $\mathbb{I}$ -convex envelopes were the building blocks of the system of functional equations characterizing the limit value. Based on this characterization, a construction of asymptotically optimal strategies (i.e. strategies being almost optimal in $\mathcal{G}_n(\pi)$, with an error term vanishing as the number of stages tends to $+\infty$) was obtained by Heuer [6]. The convergence of the values $v_\theta(\pi)$ for a general evaluation, as $\max_{m \geq 1} \theta_m$ tends to 0, and the construction of asymptotically optimal strategies in this case were obtained by Oliu-Barton [9, 11].

In addition to their main result, Mertens and Zamir [7] also established a recursive formula for $v_n(\pi)$ and $v_\lambda(\pi)$ in terms of the conditional probabilities on $K \times L$ induced by the players' strategies at the first stage, and the extension to a general evaluation θ is straightforward. Though very useful for studying the values, the formula cannot be used by the players for the simple reason that none of them can actually compute these conditional probabilities. There is, however, one important exception: games with incomplete information on one side. Indeed, when Player 2 has no private information, Player 1 controls and observes the conditional probabilities while Player 2 does not. As a consequence, the former, and not the latter, can use the recursive formula satisfied by the values to construct an optimal strategy. This game is denoted by $\mathcal{G}_\theta(p)$ where p is the probability distribution of k . The dual game was introduced by De Meyer in [3, 4] in order to construct an optimal strategy for Player 2. The idea of the dual game is to consider a game with vector payoffs: for each realised pair of actions (i, j) , the uninformed player knows the vector $(G^k(i, j))_{k \in K}$. Like in approachability theory, an optimal strategy for Player 2 is one that ensures that the θ -weighted sum of payoffs lies in an appropriate subset of $\mathbb{R}^K$. This set depends on a well-chosen dual variable $x \in \mathbb{R}^K$ which replaces the unknown type of Player 1 in the following sense: Player 2 can choose his opponent's type to be k at a cost x^k . De Meyer [3] proves that the values of the dual game $w_\theta(x)$ satisfy a recursive formula in terms of the dual variable, and that Player 2 can use this formula to construct an optimal strategy in the dual game. More importantly, this strategy is an optimal strategy in $\mathcal{G}_\theta(p)$ provided that x belongs to the sub-differential of the concave function $p' \mapsto v_\theta(p')$ at p .

The duality techniques were extended by Rosenberg [12], Sorin [14] and De Meyer and Marino [5] for repeated games with incomplete information on both sides, in the special case of independent initial probabilities, i.e. $\pi = p \otimes q$, for some probabilities p over K and q over L . As both players are uninformed about their opponent's type, one needs to consider two dual games, one for each player. The first dual game is related to the Fenchel conjugate of the function $p' \mapsto v_\theta(p', q) := v_\theta(p' \otimes q)$, where q is a fixed probability over L . Rosenberg [12] proved that its value $w_\theta(x, q)$ satisfies a recursive formula in terms of the dual variable x and the conditional probabilities over L induced by the strategy of Player 2. As these two variables are accessible to Player 2, the latter can use this formula to construct an optimal strategy in the dual game, and this strategy is an optimal strategy in the game $\mathcal{G}_\theta(\pi)$ provided that x belongs to the sub-differential of the concave function $p' \mapsto v_\theta(p' \otimes q)$ at p , where p and q are such that $\pi = p \otimes q$. An alternative formula with similar properties was obtained more recently by De

Meyer and Marino [5], who also considered the case of infinite action spaces. The second dual game is constructed in a symmetric manner, and provides an optimal strategy of Player 1. It is worth mentioning that, unlike the asymptotic results from Heuer [6] and Oliu-Barton [9], the constructions of Rosenberg [12] and De Meyer and Marino [5] provide optimal strategies for repeated games with a fixed evaluation (namely, n -stage and λ -discounted games).

In the present paper, we extend the results from Rosenberg [12] to the so-called dependent case. That is, we provide a recursive formula for the values of the dual games, from which we deduce the computation of explicit optimal strategies for the players in the repeated game with incomplete information $\mathcal{G}_\theta(\pi)$, for any evaluation θ and any probability π on $K \times L$. Our construction can be extended, word for word, to stochastic games with incomplete information, as long as the incomplete information concerns the payoff function, but not the transition function. Extending the duality techniques to the dependent case was never done before; albeit not technically difficult, the extension requires some new ideas, such as considering an intermediate step: first, the type of one player is drawn according to the corresponding marginal law, and then the other type is drawn according to the conditional law. These considerations are crucial in the proof of our main result in order to prove the convexity of some auxiliary functions (see Remark 4.2). Besides, contrary to [12], who considered n -stage games and λ -discounted games separately, and games with incomplete information on one and two sides separately too, we present here a unified approach. Let us also point out that the approach proposed by De Meyer and Marino [5], which was designed to handle games with infinite action spaces, does not seem well-suited to analyze the dependent case. Indeed, when applied to this case, their method requires the introduction of an additional dual variable, and thus results in a substantially more complicated dual recursive formula (see [8, Chapter 4]).

1.1 Main results

As both players have symmetric roles, we will only state our results on one side. Namely, we will focus on the optimal strategies of Player 2.

In order to state our main results, we need the following notation:

- For a non-empty finite set X , $\Delta(X)$ denotes the set of probabilities over X , and is identified with the canonical simplex in $\mathbb{R}^X$.
- $\mathbb{N}^*$ denotes the set of positive integers, and $\Delta(\mathbb{N}^*)$ is the set of nonnegative sequences $(a_m)_{m \in \mathbb{N}^*}$ satisfying $\sum_{m \geq 1} a_m = 1$.
- For any $\theta \in \Delta(\mathbb{N}^*)$ satisfying $\theta_1 < 1$ we denote by $\theta^+ \in \Delta(\mathbb{N}^*)$ the sequence $(\frac{\theta_m}{1-\theta_1})_{m \geq 1}$.
- For any $p \in \Delta(K)$ and $Q \in \Delta(L)^K$ we denote by $p \otimes Q$ the probability on $K \times L$ induced by p and Q , i.e. $(p \otimes Q)(k, \ell) = p^k Q(\ell | k)$ for all (k, ℓ) . For any $\tau \in \Delta(J)^L$, we denote by $\mathbb{P}_\tau^{pQ}$ the probability over $K \times L \times J$ induced by (p, Q, τ) . For every $j \in J$, $\mathbb{P}_\tau^{pQ}(j)$ denotes the marginal probability of j and Q_j the matrix of conditional probabilities over L given $k \in K$, conditionally on j .
- For any $(k, i, Q, \tau) \in K \times I \times \Delta(L)^K \times \Delta(J)^L$, we define

$$G_{i\tau}^{kQ} = \sum_{j \in J} \sum_{\ell \in L} G_{ij}^{k\ell} \tau^\ell(j) Q(\ell | k).$$

For any evaluation θ , any dual variable x and any matrix of conditional probabilities Q , we denote by $w_\theta(x, Q)$ the value of the dual game corresponding to the game $\mathcal{G}_\theta(p, Q)$ from the perspective of Player 2. We can now state our main result.

Theorem 1.1 (Dual recursive formula). *For all $(x, Q) \in \mathbb{R}^K \times \Delta(L)^K$ and $\theta \in \Delta(\mathbb{N}^*)$ one has:*

$$w_\theta(x, Q) = \min_{\tau \in \Delta(J)^L} \min_{(x_{ij})_{ij \in (I \times J)}} \max_{(k, i) \in K \times I} \left\{ \theta_1 G_{i\tau}^{kQ} + (1 - \theta_1) \sum_{j \in J} \mathbb{P}_\tau^{kQ}(j) \left(w_{\theta+}(x_{ij}, Q_j) + x_{ij}^k \right) - x^k \right\}.$$

Corollary 1.2. *Player 2 can construct an optimal strategy in $\mathcal{G}_\theta(\pi)$ by using the dual recursive formula, starting from an appropriate pair (x, Q) , namely Q is the matrix of conditionals corresponding to π and x belongs to the sub-differential of $p' \mapsto v_\theta(p' \otimes Q)$ at p such that $\pi = p \otimes Q$.*

1.2 Outline of the paper

Section 2 is devoted to introduce the duality techniques in all its generality. In particular, we show how to deal with the dependent case. In Section 3 we introduce repeated games with incomplete information on both sides. Section 4 is devoted to prove our main results. In Section 5 we provide some comments on the extensions of our results to two classes of dynamics games, stochastic games and differential games.

2 Duality techniques

For any pair of sets S and T and any function $g : S \times T \rightarrow \mathbb{R}$, we denote by (S, T, g) the zero-sum game where S is the set of strategies of Player 1, T is the set of strategies of Player 2 and g is the payoff. The maxmin and minmax of (S, T, g) are given by:

$$v^- := \sup_{s \in S} \inf_{t \in T} g(s, t) \quad \text{and} \quad v^+ := \inf_{t \in T} \sup_{s \in S} g(s, t),$$

and the game is said to have a value if $v^- = v^+$. An ε -optimal strategy for Player 1 is an element $s_\varepsilon \in S$ satisfying $\inf_{t \in T} g(s_\varepsilon, t) \geq v^- - \varepsilon$. Similarly, $t_\varepsilon \in T$ is an ε -optimal strategy for Player 2 if $\sup_{s \in S} g(s, t_\varepsilon) \leq v^+ + \varepsilon$. Note that ε -optimal strategies exist for all $\varepsilon > 0$ but not necessarily for $\varepsilon = 0$.

The aim of this section is to recall some properties of the *dual game*, introduced by De Meyer in [3, 4] to study repeated games with incomplete information on one side. We follow the presentation given in [14, Chapter 2]. Throughout this section, S and T denote two convex sets, K and L are two finite sets and $G^{k\ell} : S \times T \rightarrow \mathbb{R}$ is a payoff function for each $(k, \ell) \in K \times L$ that is bi-linear and bounded, i.e. $\sup_{s, t} |G^{k\ell}(s, t)| < +\infty$.

2.1 Incomplete information on one side

Let us start by considering the case $L = \{\ell\}$, and set $G^k := G^{k\ell}$ for all $k \in K$ to simplify the notation. To the collection of zero-sum games (henceforth, games) $\{(S, T, G^k), k \in K\}$, we

associate two families of games, the so-called primal and dual games, parameterised in terms of a probability $p \in \Delta(K)$ and a vector $x \in \mathbb{R}^K$, respectively.

The primal game $\mathcal{G}(p)$. To every probability distribution $p \in \Delta(K)$ corresponds a game with incomplete information on one side, defined as follows:

- Before the play, $k \in K$ is chosen according to p and told to Player 1.
- Then, the game (S, T, G^k) is played, i.e. Player 1 chooses $s \in S$, Player 2 chooses $t \in T$ (both choices being independent and simultaneous), and the payoff is $G^k(s, t)$.

The set of strategies of Player 1 is S^K , the set of strategies of Player 2 is T , and the payoff function is given by:

$$\gamma(p, \hat{s}, t) := \sum_{k \in K} p^k G^k(\hat{s}^k, t).$$

The game $(S^K, T, \gamma(p, \cdot))$ is denoted by $\mathcal{G}(p)$ and will be referred as the *primal game*. The maxmin and minmax of $\mathcal{G}(p)$ are given, respectively, by:

$$\begin{aligned} v^-(p) &:= \sup_{\hat{s} \in S^K} \inf_{t \in T} \gamma(p, \hat{s}, t), \\ v^+(p) &:= \inf_{t \in T} \sup_{\hat{s} \in S^K} \gamma(p, \hat{s}, t). \end{aligned}$$

Concavity and continuity: The maps $p \mapsto v^\pm(p)$ are concave and continuous on $\Delta(K)$.

The dual game $\mathcal{D}[\mathcal{G}](x)$. To every vector $x \in \mathbb{R}^K$ corresponds the dual game $\mathcal{D}[\mathcal{G}](x)$, a modified version of the primal game $\mathcal{G}(p)$ where Player 1 can *choose* the parameter $k \in K$ at a cost x^k . Formally, the set of strategies of Player 1 in the dual game is $\Delta(K) \times S^K$, the set of strategies of Player 2 is T , and the payoff function is given by:

$$h[x](p, \hat{s}, t) := \sum_{k \in K} p^k G^k(\hat{s}^k, t) - \langle p, x \rangle.$$

Let $w^-(x)$ and $w^+(x)$ denote, respectively, the maxmin and the minmax of $\mathcal{D}[\mathcal{G}](x)$, i.e.

$$\begin{aligned} w^-(x) &:= \sup_{(p, \hat{s}) \in \Delta(K) \times S^K} \inf_{t \in T} h[x](p, \hat{s}, t), \\ w^+(x) &:= \inf_{t \in T} \sup_{(p, \hat{s}) \in \Delta(K) \times S^K} h[x](p, \hat{s}, t). \end{aligned}$$

Convexity and continuity: The maps $x \mapsto w^\pm(x)$ are convex and continuous on $\mathbb{R}^K$.

The values of the primal game and the values of the dual game are essentially Fenchel conjugates from each other. To be more precise about the link between the functions $v^\pm$ and $w^\pm$, one needs to introduce two closely related convex transforms: the lower and the upper conjugates.

Definition 2.1. Let $f : \mathbb{R}^K \rightarrow \overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$. Define its upper and lower conjugates $f^\sharp, f^\flat : \mathbb{R}^K \rightarrow \overline{\mathbb{R}}$ by:

$$\begin{aligned} f^\sharp(x) &:= \sup_{y \in \mathbb{R}^K} f(y) - \langle y, x \rangle, \\ f^\flat(y) &:= \inf_{x \in \mathbb{R}^K} f(x) + \langle x, y \rangle. \end{aligned}$$

Recall that the Fenchel conjugate of f is given by $f^*(x) = \sup_{y \in \mathbb{R}^K} \langle y, x \rangle - f(y)$ so that:

$$f^\sharp(x) = (-f)^*(-x), \quad \text{and} \quad f^\flat(y) = -f^*(-y), \quad \forall x, y \in \mathbb{R}^K.$$

Recall also the usual definition of the superdifferential of f at x :

$$\partial^+ f(x) := \{y \in \mathbb{R}^K \mid f(x) + \langle y, x' - x \rangle \geq f(x'), \forall x' \in \mathbb{R}^K\}.$$

Theorem 2.2. *For all $(x, p) \in \mathbb{R}^K \times \Delta(K)$ one has:*

$$w^\pm(x) = (v^\pm)^\sharp(x) \quad \text{and} \quad v^\pm(p) = (w^\pm)^\flat(p)$$

where the functions $v^\pm$ are extended to $\mathbb{R}^K$ by $-\infty$ in $\mathbb{R}^K \setminus \Delta(K)$.

Corollary 2.3. *Let $(x, p) \in \mathbb{R}^K \times \Delta(K)$ be such that $x \in \partial^+ v^+(p)$. Then, any ε -optimal strategy $t_\varepsilon \in T$ of Player 2 in $\mathcal{D}[\mathcal{G}](x)$ is also an ε -optimal strategy of Player 2 in $\mathcal{G}(p)$.*

2.2 Incomplete information on both sides

Consider now the general case where K and L may contain more than one element. To the collection of games $\{(S, T, G^{k\ell}), (k, \ell) \in K \times L\}$ one can associate a family of games with incomplete information on both sides, as before.

The primal game $\mathcal{G}(\pi)$. For any $\pi \in \Delta(K \times L)$, consider the following primal game, denoted by $\mathcal{G}(\pi)$:

- Before the play, a couple $(k, \ell) \in K \times L$ is chosen according to π , Player 1 is informed of k and Player 2 is informed of ℓ .
- Then, the game $(S, T, G^{k\ell})$ is played, i.e. Player 1 chooses $s \in S$, Player 2 chooses $t \in T$ (both choices being independent and simultaneous), and the payoff is $G^{k\ell}(s, t)$.

In normal-form, $\mathcal{G}(\pi)$ is given by a triplet $(S^K, T^L, \gamma(\pi, \cdot))$ where:

$$\gamma(\pi, \hat{s}, \hat{t}) := \sum_{(k, \ell) \in K \times L} \pi^{k\ell} G^{k\ell}(\hat{s}^k, \hat{t}^\ell).$$

Let $v^\pm(\pi)$ denote the maxmin and minmax of $\mathcal{G}(\pi)$.

In order to apply the duality techniques described above, one needs to reformulate the primal game $\mathcal{G}(\pi)$ in a slightly different way. Let $p \in \Delta(K)$ and $Q \in \Delta(L)^K$ denote, respectively, the marginal of π on K and its matrix of conditional probabilities, so that:

$$\pi = p \otimes Q, \quad \left(\text{i.e. } \pi^{k\ell} = p^k Q(\ell \mid k), \quad \forall (k, \ell) \in K \times L \right)$$

From the perspective of Player 2, the game with incomplete information on both sides $\mathcal{G}(\pi)$, can be seen as a game with incomplete information on one side, where Player 1 is the informed player.

The first primal game $\mathcal{G}_Q(p)$. Let $Q \in \Delta(L)^K$ be fixed. For any $p \in \Delta(K)$, consider the following game:

- Before the play, $k \in K$ is chosen according to p and told to Player 1.
- Then, the game (S, T^L, G_Q^k) is played, i.e. Player 1 chooses $s \in S$, Player 2 chooses $t \in T^L$ (both choices being independent and simultaneous), and the payoff is:

$$G_Q^k(s, \hat{t}) := \sum_{\ell \in L} Q(\ell|k) G^{k\ell}(s, \hat{t}^\ell).$$

The sets of strategies are thus S^K and T^L and the payoff function is given by:

$$\gamma(p, \hat{s}, \hat{t}) := \sum_{k \in K} p^k G_Q^k(\hat{s}^k, \hat{t}).$$

The maxmin and minmax are denoted respectively by $v_Q^-(p)$ and $v_Q^+(p)$, and the maps $v_Q^\pm(p)$ are concave and continuous.

The sets S and T^L are convex and G_Q^k is bi-linear and bounded for all $k \in K$, so that a dual game can be defined, like in Section 2.1.

The first dual game $\mathcal{D}[\mathcal{G}_Q](x)$. As before, to every $x \in \mathbb{R}^K$ corresponds a dual game. The sets of strategies are $\Delta(K) \times S^K$ and T^L and the payoff function is given by:

$$h_Q[x](p, \hat{s}, \hat{t}) := \sum_{k \in K} p^k G_Q^k(\hat{s}^k, \hat{t}) - \langle p, x \rangle.$$

The maxmin and minmax functions $x \mapsto w_Q^\pm(x)$ are convex and continuous.

Theorem 2.2 and Corollary 2.3 can thus be restated accordingly.

Theorem 2.4. *For each $(p, x) \in \Delta(K) \times \mathbb{R}^K$ and $Q \in \Delta(L)^K$ one has:*

$$w_Q^\pm(x) = (v_Q^\pm)^\#(x) \quad \text{and} \quad v_Q^\pm(p) = (w_Q^\pm)^\flat(p).$$

Corollary 2.5. *Let $(p, Q) \in \Delta(K) \times \Delta(L)^K$ and $\pi := p \otimes Q$. Let $x \in \partial^+ v_Q^+(p)$, and let $\hat{t}_\varepsilon \in T^L$ be an ε -optimal strategy of Player 2 in $\mathcal{D}[\mathcal{G}_Q](x)$. Then, $\hat{t}_\varepsilon$ is ε -optimal for Player 2 in $\mathcal{G}(\pi)$.*

The second primal game $\mathcal{G}_P(q)$ and the second dual game $\mathcal{D}[\mathcal{G}_P](y)$. In games with incomplete information on both sides, the two players have similar roles. Thus, by expressing the primal game $\mathcal{G}(\pi)$ from the perspective of Player 1, one similarly defines a primal game $\mathcal{G}_P(q)$ for Player 1 and the corresponding dual game $\mathcal{D}[\mathcal{G}_P](y)$, for all $(q, P) \in \Delta(L) \times \Delta(K)^L$ and $y \in \mathbb{R}^L$. Analogue versions of Theorem 2.4 and Corollary 2.5 can thus be obtained.

3 Repeated games with incomplete information

3.1 Preliminaries

Let K, L, I, J be finite sets. For any $(k, \ell) \in K \times L$, let $G^{k\ell} = (G^{k\ell}(i, j))_{(i, j) \in I \times J}$ be an $I \times J$ matrix. A repeated game with incomplete information is described by the finite collection of matrix games $\{G^{k\ell}, (k, \ell) \in K \times L\}$ and a probability $\pi \in \Delta(K \times L)$. It is played as follows:

- A pair of parameters $(k, \ell) \in K \times L$ is drawn according to $\pi \in \Delta(K \times L)$. Player 1 is informed of k , Player 2 is informed of ℓ .
- Then, the game $G^{k\ell}$ is played repeatedly: at each stage $m \geq 1$, knowing the past actions, the players choose $(i_m, j_m) \in I \times J$ and a stage-payoff $G^{k\ell}(i_m, j_m)$ is produced (though not observed).

The payoff of Player 1 is the expectation of $\sum_{m \geq 1} \theta_m G^{k\ell}(i_m, j_m)$, for some given $\theta \in \Delta(\mathbb{N}^*)$ that is known to both players. The payoff of Player 2 is the opposite amount.

We denote this game by $\mathcal{G}_\theta(\pi)$.

Strategies. For each $m \geq 1$, let $\mathcal{H}_m := (I \times J)^{m-1}$. The information available to the players at stage $m \geq 1$ is given by (k, h_m) and (ℓ, h_m) for some $(k, \ell, h_m) \in K \times L \times \mathcal{H}_m$. Thus, a strategy of Player 1 is a function $\hat{s} : K \times \mathcal{H} \rightarrow \Delta(I)$, where $\mathcal{H} := \bigcup_{m \geq 1} \mathcal{H}_m$, and a strategy of Player 2 is a function $\hat{t} : L \times \mathcal{H} \rightarrow \Delta(J)$. The set of strategies are denoted, respectively, by S^K and T^L .

Payoff and values. A pair of strategies $(\hat{s}, \hat{t}) \in S^K \times T^L$ and an initial probability $\pi \in \Delta(K \times L)$ induce a unique probability over $K \times L \times (I \times J)^{\mathbb{N}^*}$ on the σ -algebra generated by the cylinders, denoted by $\mathbb{P}_{\hat{s}, \hat{t}}^\pi$. One can then write the game $\mathcal{G}_\theta(\pi)$ in normal-form, i.e. $\mathcal{G}_\theta(\pi) = (S^K, T^L, \gamma_\theta(\pi, \cdot))$ where:

$$\gamma_\theta(\pi, \hat{s}, \hat{t}) := \mathbb{E}_{\hat{s}, \hat{t}}^\pi \left[\sum_{m \geq 1} \theta_m G^{k\ell}(i_m, j_m) \right],$$

and where $\mathbb{E}_{\hat{s}, \hat{t}}^\pi$ is the expectation with respect to $\mathbb{P}_{\hat{s}, \hat{t}}^\pi$.

The following result is well-known¹, and we omit its proof.

Lemma 3.1. *For any $\pi \in \Delta(K \times L)$, the game $\mathcal{G}_\theta(\pi)$ has a value, i.e.*

$$v_\theta(\pi) := \sup_{\hat{s} \in S^K} \inf_{\hat{t} \in T^L} \gamma_\theta(\pi, \hat{s}, \hat{t}) = \inf_{\hat{t} \in T^L} \sup_{\hat{s} \in S^K} \gamma_\theta(\pi, \hat{s}, \hat{t}).$$

Moreover, both players have 0-optimal strategies.

¹One may apply Sion's minmax theorem to the game in mixed strategies when pure strategies are endowed with the product topology, and then apply Kuhn's theorem to deduce the result. See e.g. chapter 3 and appendix A in [14] where the same method is applied in the discounted case.

The aim of this paper is to provide an explicit construction for a couple of 0-optimal strategies. Recall that, as the two players have similar roles, we will focus on the construction for Player 2 only. For this reason, from now on the function $v_\theta : \Delta(K \times L) \rightarrow \mathbb{R}$ will be expressed in the following equivalent manner:

$$v_\theta(p, Q) := v_\theta(p \otimes Q), \quad \forall (p, Q) \in \Delta(K) \times \Delta(L)^K,$$

which is more convenient for studying the game from the perspective of Player 2. Let us start by recalling an important result, the so-called *primal recursive formula*, which expresses the values of the repeated game with incomplete information $v_\theta(p, Q)$ in terms of the values of the *continuation game*, that is, the sub-game that the players are facing after the first stage.

3.2 Primal recursive formula

Like in Section 2.2, for each $(p, Q) \in \Delta(K) \times \Delta(L)^K$, let $\mathcal{G}_\theta(p, Q)$ denote the repeated game of incomplete information on both sides $\mathcal{G}_\theta(\pi)$, expressed from the point of view of Player 2.

The aim of this section is to provide a recursive formula satisfied by the values $v_\theta(p, Q)$. The following specific notation will be used to express this result.

- $(\sigma, \tau) \in \Delta(I)^K \times \Delta(J)^L$ denotes a pair of strategies for the first stage of the game.
- For any $(p, Q) \in \Delta(K) \times \Delta(L)^K$, $(\sigma, \tau) \in \Delta(I)^K \times \Delta(J)^L$ and $(i, j) \in I \times J$ define $p_{ij} \in \Delta(K)$ and $Q_j \in \Delta(L)^K$ by setting, for all (k, ℓ) :

$$p_{ij}^k := \mathbb{P}_{\sigma\tau}^{pQ}(k | i, j) \quad \text{and} \quad Q_j(\ell | k) := \mathbb{P}_{\sigma\tau}^{pQ}(\ell | k, j),$$

where $\mathbb{P}_{\sigma\tau}^{pQ}$ is the probability over $K \times L \times I \times J$ induced by (σ, τ, p, Q) .

The following easy result is important, as in particular it shows that Player 2 can compute (and, in fact, controls) the matrix of conditional probabilities $Q_j \in \Delta(L)^K$ for all $j \in J$.

Lemma 3.2. *For all (k, ℓ, i, j) one has:*

$$Q_j(\ell | k) = \frac{Q(\ell | k) \tau^\ell(j)}{\sum_{\ell' \in L} Q(\ell' | k) \tau^{\ell'}(j)} \quad \text{and} \quad \mathbb{P}_{\sigma\tau}^{pQ}(k, \ell | i, j) = p_{ij}^k Q_j(\ell | k).$$

Proof. For any (k, j) such that $\mathbb{P}_{\sigma\tau}^{pQ}(k, j) > 0$, a direct computation gives:

$$\mathbb{P}_{\sigma\tau}^{pQ}(\ell | k, j) = \frac{p^k Q(\ell | k) \tau^\ell(j)}{p^k \sum_{\ell' \in L} Q(\ell' | k) \tau^{\ell'}(j)} = \frac{Q(\ell | k) \tau^\ell(j)}{\sum_{\ell' \in L} Q(\ell' | k) \tau^{\ell'}(j)}.$$

so that the first relation holds. Similarly, for any (k, i, j) such that $\mathbb{P}_{\sigma\tau}^{pQ}(k, i, j) > 0$ one obtains:

$$\mathbb{P}_{\sigma\tau}^{pQ}(\ell | k, i, j) = \mathbb{P}_{\sigma\tau}^{pQ}(\ell | k, j).$$

For any (i, j) such that $\mathbb{P}_{\sigma\tau}^{pQ}(i, j) > 0$, disintegration gives then the second relation:

$$\mathbb{P}_{\sigma\tau}^{pQ}(k, \ell | i, j) = \mathbb{P}_{\sigma\tau}^{pQ}(k | i, j) \mathbb{P}_{\sigma\tau}^{pQ}(\ell | k, i, j) = p_{ij}^k Q_j(\ell | k).$$

□

We are now ready to state the so-called primal recursive formula, due to Mertens and Zamir [7, Section 3]. For convenience, we provide a direct and shorter proof here.

Proposition 3.3 (Primal recursive formula). *For any $(p, Q) \in \Delta(K) \times \Delta(L)^K$ and $\theta \in \Delta(\mathbb{N}^*)$:*

$$\begin{aligned} v_\theta(p, Q) &= \max_{\sigma \in \Delta(I)^K} \min_{\tau \in \Delta(J)^L} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\}, \\ &= \min_{\tau \in \Delta(J)^L} \max_{\sigma \in \Delta(I)^K} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\}. \end{aligned}$$

Proof. Consider the maxmin. Let $\hat{s}$ be a strategy of Player 1. At the first stage, the information available to Player 1 is k , so that $\sigma := (\hat{s}(k))_k \in \Delta(I)^K$ represents the strategy of Player 1 at the first stage. Similarly, the information available to Player 1 at the second stage is (k, i, j) for some couple of actions (i, j) played at the first stage, so that $\hat{s}^+ = (\hat{s}(k, i, j))_{k, i, j}$ represents the strategy of Player 1 at the second stage. One can then write $\hat{s} = (\sigma, \hat{s}^+)$. Similarly, a strategy $\hat{t}$ of Player 2 can be written as $\hat{t} = (\tau, \hat{t}^+)$ where $\tau \in \Delta(J)^L$ and $\hat{t}^+ = (\hat{t}(\ell, i, j))_{\ell, i, j}$.

For each (i, j) , let $\hat{t}_{ij}^+$ be a best reply of Player 2 to the strategy $\hat{s}_{ij}^+ := (\hat{s}(k, i, j))_k$ in the so-called continuation game, i.e. $\mathcal{G}_{\theta+}(p_{ij}, Q_j)$. Then, for all $(\sigma, \tau) \in \Delta(I)^K \times \Delta(J)^L$ one has:

$$\gamma_\theta(p \otimes Q, \hat{s}, \hat{t}) \leq \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j).$$

Player 1 can still maximize over his own first-stage strategy, and Player 2 can again play a best reply. Hence:

$$\begin{aligned} v_\theta(p, Q) &= \max_{(\sigma, \hat{s}^+)} \min_{(\tau, \hat{t}^+)} \gamma_\theta(p \otimes Q, \hat{s}, \hat{t}), \\ &\leq \max_{\sigma \in \Delta(I)^K} \min_{\tau \in \Delta(J)^L} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\}. \end{aligned}$$

Reversing the roles of the players one obtains, symmetrically:

$$v_\theta(p, Q) \geq \min_{\tau \in \Delta(J)^L} \max_{\sigma \in \Delta(I)^K} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\}. \quad (3.1)$$

The result follows then from the inequality “maxmin $\leq$ minmax”. $\square$

Comments. Proposition 3.3 provides a recursive formula satisfied by the values, from the perspective of Player 2. Similarly, one can obtain a recursive formula satisfied by the values from the perspective of Player 1, expressing $v_\theta(q, P) := v_\theta(q \otimes P)$ in terms of $v_{\theta+}(q_{ij}, P_i)$ for any $(q, P) \in \Delta(L) \times \Delta(K)^L$ and suitably defined conditional probabilities $(q_{ij}, P_i) \in \Delta(L) \times \Delta(K)^L$ for all (i, j) . However, neither of these recursive formulae can be used by the players as both probability distributions p_{ij} and q_{ij} depend on the first-stage strategies of the two players. The situation contrasts with the case of repeated games with incomplete information on one side (that is, when L is a singleton), where Player 1 observes and controls the conditional probability $p_i \in \Delta(K)$, so that the primal recursive formula provides an explicit and recursive manner to obtain an optimal strategy for Player 1 (see [14, Section 3]).

Remark 3.4. The sequence of weights θ^+ is not defined when $\theta_1 = 1$, so that neither is the value function v_{θ^+} . This, however, does not matter as the primal formula contains the term $(1 - \theta_1)v_{\theta^+}$, which is 0.

4 Dual recursive formula

The aim of this section is to prove the *dual recursive formula*, stated in Theorem 1.1, and to deduce an explicit construction of an optimal strategy for Player 2 in $\mathcal{G}_\theta(\pi)$.

4.1 The dual game

Consider the game $\mathcal{G}_\theta(\pi)$ described in Section 3 from the point of view of Player 2. For any fixed matrix of conditionals $Q \in \Delta(L)^K$ and any $k \in K$, consider the collection of games $\{\mathcal{G}_\theta(\delta_k \otimes Q), k \in K\}$, where $\delta_k \in \Delta(K)$ is the Dirac mass, i.e. $\delta_k(k) = 1$ and $\delta_k(k') = 0$ for all $k' \neq k$. Let S and T^L denote, respectively, the common sets of strategies of Player 1 and 2 in each of these games. These sets are convex and the payoff functions $(s, \hat{t}) \mapsto \gamma_\theta(\delta_k \otimes Q, s, \hat{t})$ are bi-linear and bounded. Hence, like in Section 2, one can define the corresponding dual game $\mathcal{D}[\mathcal{G}_\theta](x, Q)$ for Player 2, for any $x \in \mathbb{R}^K$.

The dual game $\mathcal{D}[\mathcal{G}_\theta](x, Q)$. By construction, the sets of strategies of this game are $\Delta(K) \times S^K$ and T^L , and the payoff function is given by:

$$h_\theta[x, Q](p, \hat{s}, \hat{t}) = \gamma_\theta(p \otimes Q, \hat{s}, \hat{t}) - \langle p, x \rangle.$$

By Lemma 3.1 and Theorem 2.4, this game has a value $w_\theta(x, Q)$, i.e.

$$\begin{aligned} w_\theta(x, Q) &= \max_{(p, \hat{s}) \in \Delta(K) \times S^K} \min_{\hat{t} \in T^L} h_\theta[x, Q](p, \hat{s}, \hat{t}), \\ &= \min_{\hat{t} \in T^L} \max_{(p, \hat{s}) \in \Delta(K) \times S^K} h_\theta[x, Q](p, \hat{s}, \hat{t}). \end{aligned}$$

and the mappings $x \mapsto w_\theta(x, Q)$ are convex and continuous.

The following notation will be used in the proof of the dual recursive formula:

- $\mathbb{P}_{\mu\tau}^Q$ is the unique probability on $K \times L \times I \times J$ induced by $\mu \in \Delta(K \times I)$, $Q \in \Delta(L)^K$ and $\tau \in \Delta(J)^L$, i.e. $\mathbb{P}_{\mu\tau}^Q(k, \ell, i, j) = \mu(k, i)Q(\ell | k)\tau^\ell(j)$.
- $G_{\mu\tau}^Q$ denotes the expectation of $G^{k\ell}(i, j)$ with respect to the probability $\mathbb{P}_{\mu\tau}^Q$. Note that $G_{i\tau}^{kQ}$ stands for $G_{\mu\tau}^Q$ for the particular case where $\mu = \delta_{(k, i)}$ for some $(k, i) \in K \times I$.
- $\mathbb{P}_\tau^{kQ}(j) := \sum_{\ell \in L} Q(\ell | k)\tau^\ell(j)$ for any $k \in K$, $Q \in \Delta(L)^K$ and $\tau \in \Delta(J)^L$.
- $\|G\| := \max_{k, \ell, i, j} |G^{k\ell}(i, j)|$.
- B denotes the following set $\{x \in \mathbb{R}^K \mid \|x\|_\infty \leq \|G\|\}$.

4.2 Proof of Theorem 1.1

Let us recall the statement of the theorem:

For all $(x, Q) \in \mathbb{R}^K \times \Delta(L)^K$ and $\theta \in \Delta(\mathbb{N}^*)$ one has:

$$w_\theta(x, Q) = \min_{\substack{\tau \in \Delta(J)^L \\ (x_{ij})_{ij \in (\mathbb{R}^K)^{I \times J}}} \max_{(k, i) \in K \times I} \left\{ \theta_1 G_{i\tau}^{kQ} + (1 - \theta_1) \sum_{j \in J} \mathbb{P}_\tau^{kQ}(j) \left(w_{\theta+}(x_{ij}, Q_j) + x_{ij}^k \right) - x^k \right\}.$$

Furthermore, we will prove that the minimum in $(x_{ij})_{ij}$ is reached in the set $B^{I \times J}$.

Remark 4.1. Again, there is no need in defining the value function $w_{\theta+}$ when $\theta_1 = 1$, since in this case the term $(1 - \theta_1)w_{\theta+}$ is equal to 0 by convention.

Proof. On the one hand, by Proposition 3.3 one has the primal recursive formula:

$$v_\theta(p, Q) = \max_{\sigma \in \Delta(I)^K} \min_{\tau \in \Delta(J)^L} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\},$$

where $p_{ij} \in \Delta(K)$ and $Q_j \in \Delta(L)^K$ are defined in (3.2) and where, by Lemma 3.2, Q_j does not depend on (p, σ) . On the other hand, by the duality results presented in Section 2, namely Theorem 2.4, one has:

$$w_\theta(x, Q) = \max_{p \in \Delta(K)} v_\theta(p, Q) - \langle p, x \rangle.$$

Replacing $v_\theta(p, Q)$ by its expression in the primal recursive formula one obtains:

$$\begin{aligned} w_\theta(x, Q) &= \max_{p \in \Delta(K)} \left\{ \max_{\sigma \in \Delta(I)^K} \min_{\tau \in \Delta(J)^L} \left\{ \theta_1 G_{\sigma\tau}^{pQ} + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\sigma\tau}^{pQ}(i, j) v_{\theta+}(p_{ij}, Q_j) \right\} - \langle p, x \rangle \right\}, \\ &= \max_{\mu \in \Delta(K \times I)} \min_{\tau \in \Delta(J)^L} \left\{ \theta_1 G_{\mu\tau}^Q + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) v_{\theta+}(\mu_{ij}, Q_j) - \langle \mu^K, x \rangle \right\}. \end{aligned}$$

where μ^K is the marginal of μ on K and μ_{ij} is the conditional on K given (i, j) , i.e.

$$\mu_{ij}^k = \mathbb{P}_{\mu\tau}^Q(k | i, j), \quad \forall (k, i, j).$$

Consider the one-shot game with action sets $\Delta(K \times I)$ and $\Delta(J)^L$ and payoff function:

$$F[\theta, x, Q](\mu, \tau) := \theta_1 G_{\mu\tau}^Q + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) v_{\theta+}(\mu_{ij}, Q_j) - \langle \mu^K, x \rangle.$$

Clearly, $F[\theta, x, Q]$ is continuous on the compact set $\Delta(K \times I) \times \Delta(J)^L$, its first and last terms are linear in μ and τ and, as we have already shown, one has:

$$w_\theta(x, Q) = \max_{\mu \in \Delta(K \times I)} \min_{\tau \in \Delta(J)^L} F[\theta, x, Q](\mu, \tau).$$

Therefore, one can apply Sion's minmax theorem [13] as soon as we prove that the following function is concave-convex:

$$f(\mu, \tau) := \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) v_{\theta+}(\mu_{ij}, Q_j).$$

First, let us recall that from Theorem 2.4, the following relation holds for all (i, j) :

$$v_{\theta+}(\mu_{ij}, Q_j) = \inf_{x_{ij} \in \mathbb{R}^K} w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle \quad (4.1)$$

Recall that any concave function φ on $\Delta(K)$ which is $\|G\|$ -Lipschitz with respect to the norm $\|\cdot\|_1$ can be extended to a concave Lipschitz function $\tilde{\varphi}$ on the whole space $\mathbb{R}^K$ having the same Lipschitz constant, by defining $\tilde{\varphi}(x) = \sup_{y \in \mathbb{R}^K} \{\varphi(y) - \|G\| \|y - x\|_1\}$. By construction, $\tilde{\varphi}$ admits super-gradients at every point, which belong to the compact convex set $B := \{x \in \mathbb{R}^K \mid \|x\|_\infty \leq \|G\|\}$. Since $\varphi = \tilde{\varphi}$ on $\Delta(K)$, for all $p \in \Delta(K)$, the super-gradients of $\tilde{\varphi}$ at p are super-gradients of φ at p , and therefore φ admits super-gradients in B at every point of $\Delta(K)$.

According to Fenchel's lemma, the set of minimisers of the right-hand side of (4.1) is exactly the set of super-gradients of the concave mapping $p' \mapsto v_{\theta+}(p', Q_j)$ at $p' = \mu_{ij}$. Since this mapping is $\|G\|$ -Lipschitz with respect to the norm $\|\cdot\|_1$ on its domain $\Delta(K)$, it admits a super-gradient in the compact convex set B at μ_{ij} , which is therefore a minimiser of the right-hand side of (4.1). We deduce that

$$v_{\theta+}(\mu_{ij}, Q_j) = \min_{x_{ij} \in B} w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle \quad (4.2)$$

Hence, replacing this expression, one can write:

$$\begin{aligned} f(\mu, \tau) &= \inf_{(x_{ij})_{ij} \in B^{I \times J}} \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) (w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle), \\ &= \inf_{(x_{ij})_{ij} \in B^{I \times J}} \left(\sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(x_{ij}, Q_j) + \sum_{(i,j,k) \in I \times J \times K} \mathbb{P}_{\mu\tau}^Q(i, j) \mu_{ij}^k x_{ij}^k \right), \\ &= \inf_{(x_{ij})_{ij} \in B^{I \times J}} \left(\sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(x_{ij}, Q_j) + \sum_{(i,j,k) \in I \times J \times K} \mathbb{P}_{\mu\tau}^Q(i, j, k) x_{ij}^k \right). \end{aligned}$$

Since $\mu \rightarrow \mathbb{P}_{\mu\tau}^Q$ is affine, we deduce from the above expression that $\mu \rightarrow f(\mu, \tau)$ is concave as an infimum of affine functions.

To prove the convexity of $\tau \mapsto f(\mu, \tau)$ we will consider the primal game $\mathcal{G}_\theta(\pi)$ from the point of view of Player 1. For any $q \in \Delta(L)$ and $P \in \Delta(K)^L$, denote this game by $\mathcal{G}_\theta(q, P)$ and let $v_\theta(q, P) := v_\theta(q \otimes P)$ denote its value. Using this notation, for each (i, j) one can write:

$$v_{\theta+}(\mu_{ij}, Q_j) = v_{\theta+}(q_{ij}, P_i),$$

where $q_{ij} \in \Delta(L)$ and $P_i \in \Delta(K)^L$. In particular, P_i is independent from j and τ (just like Q_j does not depend neither on i nor σ). Explicitly, for all (k, ℓ, i, j) one has:

$$q_{ij}^\ell = \mathbb{P}_{\mu\tau}^Q(\ell \mid i, j) \quad \text{and} \quad P_i(k \mid \ell) = \mathbb{P}_{\mu\tau}^Q(k \mid \ell, i) = \frac{\mu(k, i) Q(\ell \mid k)}{\sum_{k' \in K} \mu(k', i) Q(\ell \mid k')}.$$

Use the duality techniques of Section 2 to define, for each $y \in \mathbb{R}^L$, the dual game $\mathcal{D}[\mathcal{G}_\theta](y, P)$, and denote its value by $w_\theta(y, P) := \sup_{q \in \Delta(L)} v_\theta(q, P) - \langle q, y \rangle$. For each (i, j) one then has:

$$v_{\theta+}(q_{ij} \otimes P_i) = \sup_{y_{ij} \in \mathbb{R}^L} w_{\theta+}(y_{ij}, P_i) - \langle q_{ij}, y_{ij} \rangle,$$

so that:

$$\begin{aligned} f(\mu, \tau) &= \sup_{(y_{ij})_{ij} \in (\mathbb{R}^L)^{I \times J}} \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) (w_{\theta+}(y_{ij}, P_i) - \langle q_{ij}, y_{ij} \rangle), \\ &= \sup_{(y_{ij})_{ij} \in (\mathbb{R}^L)^{I \times J}} \left(\sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(y_{ij}, P_i) - \sum_{(i,j,\ell) \in I \times J \times L} \mathbb{P}_{\mu\tau}^Q(i, j) q_{ij}^\ell y_{ij}^\ell \right), \\ &= \sup_{(y_{ij})_{ij} \in (\mathbb{R}^L)^{I \times J}} \left(\sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(y_{ij}, P_i) - \sum_{(i,j,\ell) \in I \times J \times L} \mathbb{P}_{\mu\tau}^Q(i, j, \ell) y_{ij}^\ell \right). \end{aligned}$$

The mappings $\tau \rightarrow \mathbb{P}_{\mu\tau}^Q(i, j)$ and $\tau \rightarrow \mathbb{P}_{\mu\tau}^Q(i, j, \ell)$ being affine, the previous expression shows that $\tau \rightarrow f(\mu, \tau)$ is convex, as a supremum of affine functions.

Therefore, one can indeed apply Sion's minmax theorem. Exchanging the maximum and the minimum one obtains:

$$w_\theta(x, Q) = \min_{\tau \in \Delta(J)^L} \max_{\mu \in \Delta(K \times I)} F[\theta, x, Q](\mu, \tau).$$

Replacing now $v_{\theta+}(\mu_{ij}, Q_j)$ with $\min_{x_{ij} \in B} w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle$ gives then:

$$\begin{aligned} w_\theta(x, Q) &= \min_{\tau \in \Delta(J)^L} \max_{\mu \in \Delta(K \times I)} \min_{(x_{ij})_{ij} \in B^{I \times J}} \\ &\quad \left\{ \theta_1 G_{\mu\tau}^Q + (1 - \theta_1) \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) (w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle) - \langle \mu^K, x \rangle \right\}. \end{aligned}$$

Again, in order to apply Sion's minmax theorem to exchange the order of the maximum and the infimum one needs to check that the mapping $(\mu, \mathbf{x}) \mapsto g(\mu, \mathbf{x})$ is concave-convex, where $\mathbf{x} := (x_{ij})_{ij} \in B^{I \times J}$ and:

$$\begin{aligned} g(\mu, \mathbf{x}) &:= \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) (w_{\theta+}(x_{ij}, Q_j) + \langle \mu_{ij}, x_{ij} \rangle), \\ &= \sum_{(i,j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(x_{ij}, Q_j) + \sum_{(i,j,k) \in I \times J \times K} \mathbb{P}_{\mu\tau}^Q(i, j, k) x_{ij}^k. \end{aligned}$$

This property follows from the fact that the mappings $\mu \rightarrow \mathbb{P}_{\mu\tau}^Q(i, j)$ and $\mu \rightarrow \mathbb{P}_{\mu\tau}^Q(i, j, k)$ are affine, and that the map $x \rightarrow w_{\theta+}(x, Q_j)$ is convex. We thus obtain:

$$\begin{aligned} w_\theta(x, Q) &= \min_{\tau \in \Delta(J)^L} \min_{(x_{ij})_{ij} \in B^{I \times J}} \max_{\mu \in \Delta(K \times I)} \\ &\quad \left\{ \theta_1 G_{\mu\tau}^Q + (1 - \theta_1) \sum_{i,j} \mathbb{P}_{\mu\tau}^Q(i, j) w_{\theta+}(x_{ij}, Q_j) + (1 - \theta_1) \sum_{i,k} \mu(i, k) \sum_j \mathbb{P}_{\mu\tau}^Q(j|i, k) x_{ij}^k - \langle \mu^K, x \rangle \right\}. \end{aligned}$$

Since the expression above is affine with respect to μ , we can consider without loss of generality the maxima at extreme points:

$$w_\theta(x, Q) = \min_{\tau \in \Delta(J)^L} \min_{(x_{ij})_{ij} \in B^{I \times J}} \max_{(k, i) \in K \times I} \left\{ \theta_1 G_{i\tau}^{kQ} + (1 - \theta_1) \sum_j \mathbb{P}_\tau^{kQ}(j) \left(w_{\theta+}(x_{ij}, Q_j) + x_{ij}^k \right) - x^k \right\}.$$

To conclude, note that the minimum over $B^{I \times J}$ can be replaced by a minimum over $(\mathbb{R}^K)^{I \times J}$ since the above proof is still valid if we replace $v_{\theta+}(\mu_{ij}, Q_j)$ by the expression given by (4.1) instead of the one given by (4.2). $\square$

Remark 4.2. *The proof of Theorem 1.1 follows the main lines of [12], but there is a crucial point where an obstacle arises, namely in proving that the function*

$$f(\mu, \tau) := \sum_{(i, j) \in I \times J} \mathbb{P}_{\mu\tau}^Q(i, j) v_{\theta+}(\mu_{ij}, Q_j)$$

is concave-convex. Unlike the independent case, where the proof relies deeply on the fact that $(p, q) \mapsto v_{\theta+}(p, q)$ is a concave-convex functions of independent probabilities $p \in \Delta(K)$ and $q \in \Delta(L)$, the arguments of $v_{\theta+}(\mu_{ij}, Q_j)$ are not independent from each other, and one thus needs to use the duality techniques for the dependent case, which are more sophisticated. In this point, our proof diverges from the one in [12].

4.3 Construction of an optimal Markovian strategy

In this section, we deduce from Theorem 1.1 and Corollary 2.5 the construction of an optimal strategy for Player 2 in the game $\mathcal{G}_\theta(\pi)$. The strategy is Markovian, in the sense that it depends on the past history only through the (updated) variable (x, Q, θ) .

Let $(p, Q) \in \Delta(K) \times \Delta(L)^K$ be such that $\pi = p \otimes Q$. Let $x \in \partial^+ v_\theta(p, Q)$ be a super-gradient at p of the mapping $v_\theta(\cdot, Q) : \Delta(K) \rightarrow \mathbb{R}$.

For all $(\theta', x', Q') \in \Delta(\mathbb{N}^*) \times \mathbb{R}^K \times \Delta(L)^K$, let us denote by $S_{\theta'}(x', Q')$ the set of minimizers in the dual recursive formula, i.e.

$$S_{\theta'}(x', Q') = \operatorname{argmin}_{(\tau, \mathbf{x}) \in \Delta(J)^L \times \mathbb{R}^{K \times I \times J}} h(\theta', x', Q')[\tau, \mathbf{x}],$$

where $\mathbf{x} = (x_{ij})_{ij}$ and where, using the notation of the previous sections:

$$h(\theta', x', Q')[\tau, \mathbf{x}] = \max_{(k, i) \in K \times I} \left\{ \theta'_1 G_{i\tau}^{kQ'} + (1 - \theta'_1) \sum_j \mathbb{P}_\tau^{kQ'}(j) \left(w_{(\theta')+}(x_{ij}, Q'_j) + x_{ij}^k \right) - (x')^k \right\}.$$

An optimal strategy for Player 2 in $\mathcal{G}_\theta(\pi)$ can be constructed, recursively, as follows.

Case 1. If $\theta_1 = 1$, play $\tau \in \Delta(J)^L$ which is optimal in the formula:

$$w_1(x, Q) = \min_{\tau \in \Delta(J)^L} \max_{(k, i) \in K \times I} \left\{ \theta_1 G_{i\tau}^{kQ} - x^k \right\}.$$

Case 2. If $\theta_1 < 1$, plays as follows:

- Compute $(\tau, \mathbf{x}) \in S_\theta(x, Q)$ optimal in the dual recursive formula at (x, Q, θ) .
- Choose $j \in J$ with probability τ^ℓ , where $\ell \in L$ is Player 2's private type.
- Observe $i \in I$ and update the triplet (x, Q, θ) to (x_{ij}, Q_j, θ^+) .
- Play optimally in $\mathcal{D}[\mathcal{G}_{\theta^+}](x_{ij}, Q_j)$.

This strategy is optimal in the dual game $\mathcal{D}[\mathcal{G}_\theta](x, Q)$, thanks to Theorem 1.1. By Corollary 2.5 and the choice of (x, Q) , the strategy is also optimal in the game $\mathcal{G}_\theta(\pi)$. Furthermore, it is *Markovian*. Indeed, at every stage $m \geq 1$, the mixed action of Player 2 at this stage $\tau_m \in \Delta(J)^L$ depends only on a triplet of variables $(x^{(m-1)}, Q^{(m-1)}, \theta^{(m-1)}) \in \mathbb{R}^K \times \Delta(L)^K \times \Delta(\mathbb{N}^*)$ constructed recursively as follows. For $m = 0$ set:

$$\begin{cases} x^{(0)} := x \\ Q^{(0)} := Q \\ \theta^{(0)} := \theta. \end{cases}$$

For all $m \geq 1$, if $\theta_1^{(m-1)} < 1$, compute $(\tau_m, \mathbf{x}^{(m)}) \in S_{\theta^{(m-1)}}(x^{(m-1)}, Q^{(m-1)})$ and set:

$$\begin{cases} x^{(m)} := (\mathbf{x}^{(m)})_{i_m j_m} \\ Q^{(m)}(\ell|k) = \frac{Q^{(m-1)}(\ell|k) \tau_m^\ell(j_m)}{\sum_{\ell' \in L} Q^{(m-1)}(\ell'|k) \tau_m^{\ell'}(j_m)}, & \forall (k, \ell) \in K \times L, \\ \theta_t^{(m)} := \frac{\theta_{m+t}}{\sum_{t' > m} \theta_{t'}}, & \forall t \geq 1, \end{cases}$$

where (i_m, j_m) is the pair of actions played at stage m .

5 Comments and Extensions

5.1 Stochastic games with incomplete information

Consider a game with incomplete information over a finite set of states S , where each state represents a different state of the world. Formally, let $G^{k\ell} : S \times I \times J \rightarrow \mathbb{R}$ denote a payoff function for each pair of types $(k, \ell) \in K \times L$, depending not only on the players' actions, but also on the state, and let $\rho : S \times I \times J \rightarrow \Delta(S)$ denote a transition kernel. For any $\theta \in \Delta(\mathbb{N}^*)$, any $\pi \in \Delta(K \times L)$ and any initial state $s_1 \in S$, the *stochastic game with incomplete information on both sides*, denoted by $\mathcal{G}_\theta(\pi; s_1)$, is played as follows:

- First, a pair of parameters $(k, \ell) \in K \times L$ is drawn according to $\pi \in \Delta(K \times L)$. Player 1 is informed of k , Player 2 is informed of ℓ .
- At each stage $m \geq 1$, knowing the current state $s_m \in S$ and knowing the past actions, the players choose actions $(i_m, j_m) \in I \times J$. A stage-payoff $G^{k\ell}(s_m, i_m, j_m)$ is produced (though not observed) and a new state $s_{m+1} \in S$ is drawn with the probability distribution $\rho(s_m, i_m, j_m)$.

The payoff of Player 1 is the expectation of $\sum_{m \geq 1} \theta_m G^{k\ell}(s_m, i_m, j_m)$, while the payoff of Player 2 is the opposite amount.

Remark 5.1. *The case where the set of states $S = \{s\}$ is a singleton corresponds to repeated games with incomplete information on both sides. In this sense, stochastic games with incomplete information extend our previous model.*

Let $p \in \Delta(K)$ and $Q \in \Delta(L)^K$ be such that $\pi = p \times Q$ and let $s \in S$ be an initial state. The existence of the value for stochastic games with incomplete information is well-known, and we omit its proof. Let $v_\theta(\pi, s)$ and $w_\theta(x, Q, s)$ denote, respectively, the values of $\mathcal{G}_\theta(\pi; s)$ and of the corresponding first dual game. The dual recursive formula obtained in Theorem 1.1 can be extended, word for word, to stochastic games with incomplete information on both sides in the dependent case.

Notation. In the following result we use the notations introduced earlier. Moreover, for each $(Q, \tau) \in \Delta(L)^K \times \Delta(J)^L$ and $(s, k, i) \in S \times K \times I$, we set:

$$G_{i\tau}^{kQ}(s) = \sum_{(\ell, j) \in L \times J} Q(\ell|k) \tau^\ell(j) G^{k\ell}(s, i, j).$$

Theorem 5.2 (Dual recursive formula). *For all $(x, Q, s) \in \mathbb{R}^K \times \Delta(L)^K \times S$ and $\theta \in \Delta(\mathbb{N}^*)$ one has:*

$$w_\theta(x, Q, s) = \min_{\tau \in \Delta(J)^L} \min_{(x_{ijs})_{ij \in (\mathbb{R}^K)^{I \times J \times S}}} \max_{(k, i) \in K \times I} \left\{ \theta_1 G_{i\tau}^{kQ}(s) + (1 - \theta_1) \sum_{(s', j) \in S \times J} \mathbb{P}_\tau^{kQ}(j) \rho(s'|s, i, j) \left(w_{\theta+}(x_{ijs}, Q_j, s') + x_{ijs}^k \right) - x^k \right\}.$$

Corollary 5.3. *Player 2 can construct an optimal strategy in $\mathcal{G}_\theta(\pi; s_1)$ by using the dual recursive formula, starting from an appropriate pair (x, Q) , namely Q is the matrix of conditionals corresponding to π and x belongs to the sub-differential of $p' \mapsto v_\theta(p' \otimes Q, s)$ at p such that $\pi = p \otimes Q$.*

5.2 Differential games with incomplete information

Differential games with incomplete information were introduced by Cardaliaguet [2]. As in repeated games with incomplete information, before the game starts, a pair of parameters (k, ℓ) is drawn according to some commonly known probability distribution π on $K \times L$. Player 1 is informed of k and Player 2 of ℓ . Then, a differential game is played in which the dynamic and the payoff function depend on both types: each player is thus partially informed about the differential game that is played. The existence and characterisation of the value function was established by Cardaliaguet [2] in the independent case, and extended to the general case by Oliu-Barton [10]. The proof relies on the geometry of the value function (I -concavity and II -convexity) and on a sub-dynamic programming principle satisfied by its Fenchel conjugates (i.e. the values of the first and the second dual games).

Though useful for establishing the existence of the value for differential games with incomplete information, the sub-dynamic programming principles satisfied by the values of the dual games do not yield a construction of optimal strategies for these games, which remains an open problem. Establishing a continuous-time analogue of the dual recursive formula (i.e. Theorem 1.1) would be a natural way to solve it.

6 Acknowledgements

The authors are indebted to Sylvain Sorin and Bernard De Meyer for their insight and suggestions. The authors are also thankful to Mario Bravo for his comments. The first author gratefully acknowledges support from the Artificial and Natural Intelligence Toulouse Institute under Grant ANR-3IA, and funding from the French National Research Agency(ANR), under the Investments for the Future (Investissements d’Avenir) program under grant ANR-17-EURE-0010. The second author gratefully acknowledges funding from the French National Research Agency (ANR), under grant ANR CIGNE (ANR-15-CE38-0007-01).

References

- [1] R.J. Aumann and M. Maschler, *Repeated games with incomplete information*, with the collaboration of R. Stearns, MIT Press, 1995.
- [2] P. Cardaliaguet, *Differential games with asymmetric information*, SIAM journal on Control and Optimization **46** (2007), 816–838.
- [3] B. De Meyer, *Repeated games and partial differential equations*, Mathematics of Operations Research **21** (1996), 209–236.
- [4] ———, *Repeated games, duality and the Central Limit theorem*, Mathematics of Operations Research **21** (1996), 237–251.
- [5] B. De Meyer and A. Marino, *Duality and optimal strategies in the finitely repeated zero-sum games with incomplete information on both sides*, Cahiers de la MSE **27** (2005).
- [6] M. Heuer, *Asymptotically optimal strategies in repeated games with incomplete information*, International Journal of Game Theory **20** (1992), 377–392.
- [7] J.-F. Mertens and S. Zamir, *The value of two-person zero-sum repeated games with lack of information on both sides*, International Journal of Game Theory **1** (1971), 39–64.
- [8] M. Oliu-Barton, *Discrete and continuous time games with incomplete information*, PhD thesis, Pierre and Marie Curie University, Paris 6 (2013).
- [9] ———, *Asymptotically optimal strategies in zero-sum repeated games with incomplete information*, 2015.
- [10] ———, *Differential games with asymmetric and correlated information*, Dynamic Games and Applications **5** (2015), 378–396.
- [11] ———, *The splitting game: Value and optimal strategies*, Dynamic Games and Applications **8** (2018), no. 1, 157–179.
- [12] Dinah Rosenberg, *Duality and markovian strategies*, International Journal of Game Theory **27** (1998), 577–597.
- [13] M. Sion, *On general minimax theorems*, vol. 8, Pacific Journal of Mathematics, 1958.
- [14] S. Sorin, *A First Course on Zero-Sum Repeated Games*, Springer, 2002.