



HAL
open science

HSVI Can Solve Zero-Sum Partially Observable Stochastic Games

Aurélien Delage, Olivier Buffet, Jilles Dibangoye, Abdallah Saffidine

► **To cite this version:**

Aurélien Delage, Olivier Buffet, Jilles Dibangoye, Abdallah Saffidine. HSVI Can Solve Zero-Sum Partially Observable Stochastic Games. *Dynamic Games and Applications*, 2023, pp.58. 10.1007/s13235-023-00519-6 . hal-04382756v2

HAL Id: hal-04382756

<https://inria.hal.science/hal-04382756v2>

Submitted on 16 Jan 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

HSVI can solve zero-sum Partially Observable Stochastic Games

Aurélien Delage^{1*}, Olivier Buffet², Jilles S. Dibangoye⁴
and Abdallah Saffidine³

¹Universite de Lyon, INSA Lyon and Inria, CITI, Villeurbanne,
F-69621, France.

² Universite de Lorraine, INRIA, CNRS, LORIA, F-54000,
Villers-lès-Nancy, F-54600, France.

³University of New South Wales, Sidney, 2052, Australia.

⁴ Bernoulli Institute for Mathematics, Computer Science and
Artificial Intelligence, University of Groningen, Groningen, 9712,
CP, The Netherlands .

*Corresponding author(s). E-mail(s): aurelien.delage@insa-lyon.fr;
Contributing authors: olivier.buffet@inria.fr;
j.s.dibangoye@rug.nl; abdallahs@cse.unsw.edu.au;

Abstract

State-of-the-art methods for solving 2-player zero-sum imperfect information games rely on linear programming or regret minimization, though not on dynamic programming (DP) or heuristic search (HS), while the latter are often at the core of state-of-the-art solvers for other sequential decision-making problems. In partially observable or collaborative settings (*e.g.*, POMDPs and Dec-POMDPs), DP and HS require introducing an appropriate statistic that induces a fully observable problem as well as bounding (convex) approximators of the optimal value function. This approach has succeeded in some subclasses of 2-player zero-sum partially observable stochastic games (zs-POSGs) as well, but how to apply it in the general case still remains an open question. We answer it by (i) rigorously defining an equivalent game to work with, (ii) proving mathematical properties of the optimal value function that allow deriving bounds that come with solution strategies, (iii) proposing for the first time an HSVI-like solver that provably converges to

This work has been conducted while Jilles Dibangoye was at INSA Lyon.

an ϵ -optimal solution in finite time, and (iv) empirically analyzing it. This opens the door to a novel family of promising approaches complementing those relying on linear programming or iterative methods.

Keywords: Game Theory, zs-POSGs, Multi-Agent Systems, Heuristic Search, Dynamic Programming

1 Introduction

Solving imperfect information sequential games is a challenging field with many applications from playing Poker [30] to security games [2]. We focus on finite-horizon 2-player zero-sum partially observable stochastic games (zs-POSGs), an important class of games that is equivalent to that of zero-sum extensive-form games (zs-EFGs) [34]¹. From the viewpoint of (maximizing) player 1, we aim at finding a strategy with a worst-case expected return (*i.e.*, whatever player 2’s strategy) within ϵ of the Nash equilibrium value (NEV).

A first approach to solving a zs-POSG is to turn it into a zs-EFG addressed as a *sequence form* linear program (SFLP) [6, 24, 44], giving rise to an exact algorithm. A second approach is to use an iterative game solver, *i.e.*, either a counterfactual-regret-based method (CFR) [7, 49], or a first-order method [18, 29], both coming with asymptotic convergence properties. CFR-based approaches now incorporate deep reinforcement learning and search, some of them winning against top human players at heads-up no limit hold’em poker [7, 9, 33]. A third approach, proposed by Wiggers [45], is to use two parallel searches in strategy space, one per player, so that the gap between both strategies’ security levels (*i.e.*, the values of their opponent’s best responses) bounds the distance to the NEV.

In contrast, dynamic programming and heuristic search have not been applied to general zs-POSGs, while often at the core of state-of-the-art solvers in other problem classes that involve Markovian dynamics, partial observability and multiple agents (POMDP [1, 38], Dec-POMDP [14, 41], or subclasses of zs-POSGs with simplifying observability assumptions [3, 12, 13, 15, 21, 22]). They all rely on some statistic that induces a fully observable problem whose value function (V^*) exhibits continuity properties that allow deriving bounding approximations. Wiggers et al. [46, 47] progress in this direction for zs-POSGs by demonstrating an important continuity property of the optimal value function, and turning a zs-POSG into a game with public strategies.² We work in a similar direction, 1. using a game with different observability hypotheses, 2. proving theoretical results they implicitly rely on, and 3. building on some of their results to derive an HSVI-like algorithm solving the zs-POSG.

¹Note: POSGs are equivalent to the large class of “well-behaved” EFGs as defined by Kovařík et al. [28].

²As discussed in the following, due to global-consistency issues, this game is not equivalent to the original zs-POSG.

Section 2 presents some necessary background, including the concept of *occupancy state* [14, 47] (*i.e.*, the probability distribution over the players' past action-observation histories), and properties that rely on it. Then, Section 3 describes theoretical contributions. First, we rigorously reformulate the problem as a non-observable game, and demonstrate that the Nash equilibrium value can be expressed with a recursive formula, which is a required tool for DP and HS (Sec. 3.1). Second, we exhibit novel continuity properties of optimal value functions and derive bounding approximators, a second tool made necessary due to the continuous state space of the new game, before showing that these approximators come with valid solution strategies for the zs-POSG (Sec. 3.2). Third, we adapt Smith and Simmons' [39] HSVI's algorithmic scheme to ϵ -optimally solve the problem in finitely many iterations (Sec. 3.3). Section 4 presents an empirical analysis of the approach. Section 5 discusses similarities and differences of our approach with most closely related work.

2 Background

Here, we first give basic definitions about zs-POSGs, including the solution concept at hand. Then we introduce an equivalent game where a state corresponds to a statistic summarizing past behaviors, which leads to some important properties of the game's optimal value.

2.1 zs-POSGs

Definition 1 (zs-POSGs) As illustrated through a dynamic influence diagram in Figure 1, a (2-player) zero-sum partially observable stochastic game (zs-POSG) is defined by a tuple $\langle \mathcal{S}, \mathcal{A}^1, \mathcal{A}^2, \mathcal{Z}^1, \mathcal{Z}^2, P, r, H, \gamma, b_0 \rangle$, where

- \mathcal{S} is a finite set of states;
- \mathcal{A}^i is (player) i 's finite set of actions;
- \mathcal{Z}^i is i 's finite set of observations;
- $P_{a^1, a^2}^{z^1, z^2}(s'|s)$ is the probability to transition to state s' and receive observations z^1 and z^2 when actions a^1 and a^2 are performed while in state s ;
- $r(s, a^1, a^2)$ is a (scalar) reward function;
- $H \in \mathbb{N}$ is a (finite) temporal horizon;
- $\gamma \in [0, 1]$ is a discount factor; and
- b_0 is the initial belief state, *i.e.*, a probability distribution over states at $t = 0$.

From the Dec-POMDP, POSG and EFG literature, we use the following concepts and definitions:

$\theta_\tau^i = (a_0^i, z_1^i, \dots, a_{\tau-1}^i, z_\tau^i)$ is a length- τ *action-observation history* (AOH) for i . We note Θ_τ^i the set of all AOHs for player i at horizon τ such that any AOH θ_τ^i is in $\cup_{t=0}^{H-1} \Theta_t^i$.

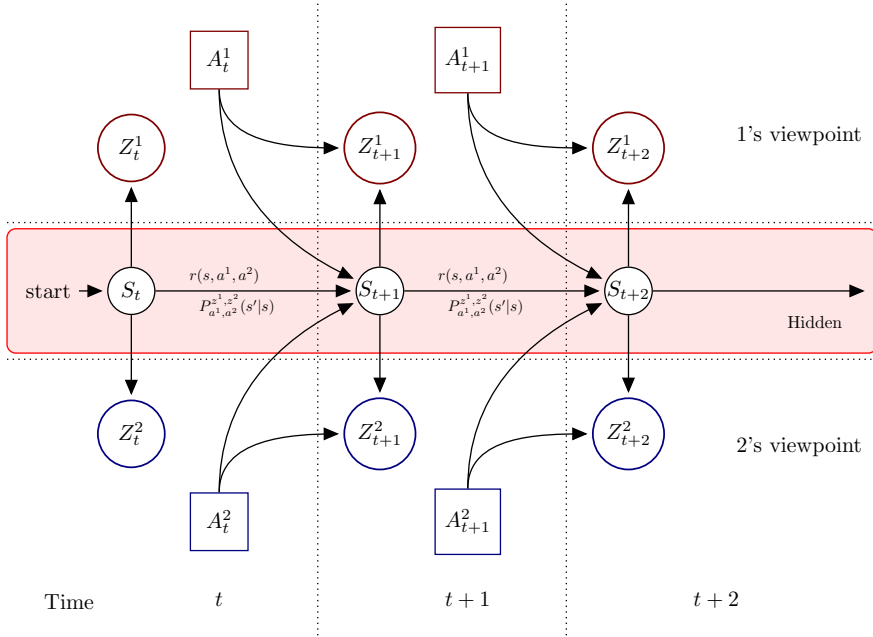
4 *HSVI can solve zero-sum Partially Observable Stochastic Games*

Fig. 1 Dynamic influence diagram representing the evolution of a zs-POSG

β_τ^i is a (*behavioral*) *decision rule* (DR) at τ for i , *i.e.*, a mapping from private AOHs in Θ_τ^i to *distributions* over private actions. $\beta_\tau^i(\theta_\tau^i, a^i)$ is the probability to pick a^i when facing θ_τ^i .

$\beta_{\tau:\tau'}^i = (\beta_\tau^i, \dots, \beta_{\tau'}^i)$ is a *behavioral strategy* for i from time step τ to τ' (included).

When considering both players, the last 3 concepts become:

$\theta_\tau = (\theta_\tau^1, \theta_\tau^2) \in \Theta = \bigcup_{t=0}^{H-1} \Theta_t$, a *joint AOH* at τ ,

$\beta_\tau = \langle \beta_\tau^1, \beta_\tau^2 \rangle \in \mathcal{B} = \bigcup_{t=0}^{H-1} \mathcal{B}_t$, a *decision rule profile*, and

$\beta_{\tau:\tau'} = \langle \beta_{\tau:\tau'}^1, \beta_{\tau:\tau'}^2 \rangle$, a *behavioral strategy profile*.

Nash Equilibria

Here, player 1 (respectively 2) wants to maximize (resp. minimize) the expected return, or *value*, of strategy profile $\beta_{0:H-1}$, defined as the discounted sum of future rewards, *i.e.*,

$$V_0(\beta_{0:H-1}) = E \left[\sum_{t=0}^{H-1} \gamma^t R_t \mid \beta_{0:H-1} \right],$$

where R_t is the random variable associated to the instant reward at t . This leads to the solution concept of Nash equilibrium strategy (NES).

Definition 2 (Nash Equilibrium) The strategy profile $\beta_{0:H-1}^* = \langle \beta_{0:H-1}^{1*}, \beta_{0:H-1}^{2*} \rangle$ is a NES if no player has an incentive to deviate, which can be written:

$$\begin{aligned} \forall \beta_{0:H-1}^1, V_0(\beta_{0:H-1}^{1*}, \beta_{0:H-1}^{2*}) &\geq V_0(\beta_{0:H-1}^1, \beta_{0:H-1}^{2*}) \text{ and} \\ \forall \beta_{0:H-1}^2, V_0(\beta_{0:H-1}^{1*}, \beta_{0:H-1}^{2*}) &\leq V_0(\beta_{0:H-1}^{1*}, \beta_{0:H-1}^2). \end{aligned}$$

In such a game, all NESs have the same Nash-equilibrium value (NEV), $V_0^* \stackrel{\text{def}}{=} V_0(\beta_{0:H-1}^{1*}, \beta_{0:H-1}^{2*})$. Our specific objective is to find an ϵ -NES, *i.e.*, a behavioral strategy profile such that any player would gain at most ϵ by deviating.

Why writing a Bellman Optimality Equation is Hard

Our approach requires writing Bellman optimality equations. The main obstacle to achieve this is to find an appropriate characterization of a *subproblem* that allows

1. *predicting* both the immediate reward and the next possible subproblems given an immediate decision;
2. *connecting* a subproblem's solution with solutions of its own (lower-level) subproblems; and
3. *prescribing* a solution strategy for the subproblem built on solutions of lower-level subproblems.

In our setting, a player's AOH does not characterize a subproblem since her opponent's strategy is also required to predict the expected reward and the next AOHs. For their part, joint AOHs allow predicting next joint AOHs given both player's immediate decision rules, but would not be appropriate either, since player i cannot decide how to act when facing some individual AOH θ_τ^i without considering all possible AOHs of his opponent $\neg i$.

Partial behavioral strategy profiles (sequences of behavioral decision rule profiles from $t = 0$ to some τ) contain enough information to completely describe the situation at τ , and are thus necessarily *predictive*. We still need to demonstrate that they are *connected*, despite decision rules not being public, and *prescriptive*, despite the need to address global-consistency issues illustrated in the following example.

Example 1 Matching pennies is a well-known 2-player zero-sum game in which each player has a penny and secretly chooses one side (head or tail). Then, both penny's sides are revealed, and player 1 wins (payoff of +1) if both chosen sides match and loses (payoff of -1) if not.

We here formalize this game as a zs-POSG (as illustrated in Figure 2) where player 1 actually picks his action at $t = 0$, and player 2 at $t = 1$. Hence the tuple $\langle \mathcal{S}, \mathcal{A}^1, \mathcal{A}^2, \mathcal{Z}^1, \mathcal{Z}^2, P, r, H, \gamma, b_0 \rangle$ where:

- $\mathcal{S} = \{s_i, s_h, s_t\}$, where s_i is the initial state, and s_h and s_t represent a memory of 1's last move: respectively "head" or "tail";
- $\mathcal{A}^1 = \mathcal{A}^2 = \{a_h, a_t\}$ for playing "head" (a_h) or "tail" (a_t);
- $\mathcal{Z}^1 = \mathcal{Z}^2 = \{z_n\}$ a "none" trivial observation;

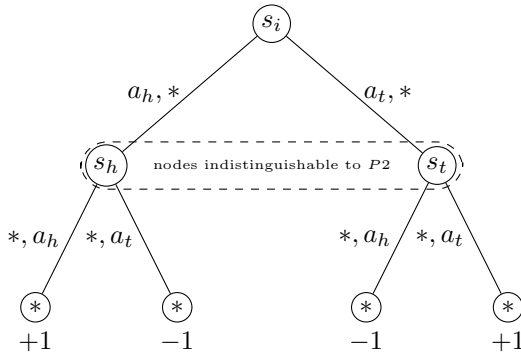


Fig. 2 Simplified tree representation of the sequentialized Matching Pennies game. Irrelevant actions, noted *, allow merging edges with the same action for (i) player 2 at $t = 0$, and (ii) player 1 at $t = 1$. Notes: (a) Due to irrelevant actions, this game can be seen as an Extensive Form Game, despite players acting simultaneously. (b) Players only know about their past action history (in this observation-free game).

- $P_{\mathbf{a}}^{\mathbf{z}}(s'|s) = T(s, \mathbf{a}, s') \cdot \mathcal{O}(\mathbf{a}, s', \mathbf{z})$, using the next two definitions;
- T is deterministic and such that (\cdot is used to denote "for all")
 - $T(\cdot, \cdot, a_h) = s_h$,
 - $T(\cdot, \cdot, a_t) = s_t$;
- \mathcal{O} is deterministic and always returns " z_n ";
- r is such that
 - $r(s_i, \cdot, \cdot) = 0$,
 - $r(s_t, \cdot, a_t) = r(s_h, \cdot, a_h) = +1$,
 - $r(s_t, \cdot, a_h) = r(s_h, \cdot, a_t) = -1$;
- $H = 2$;
- $\gamma = 1$;
- b_0 is such that the initial state is s_i with probability 1.

Let us then assume that both players' DRs at $t = 0$ are fixed, with β_0^1 randomly picking a_t or a_h (*i.e.*, it induces a NES whatever his DR at $t = 1$). Then, we face a "subgame" at $t = 1$ where any strategy profile $\langle \beta_{1:1}^1, \beta_{1:1}^2 \rangle$ is a NES profile with Nash equilibrium value 0. In particular, 2 can pick a deterministic strategy $\beta_{1:1}^2$, which will be said to be *locally consistent*. Yet, for 2, such a NES in the subgame at $\tau = 1$ is not necessarily *globally consistent*, *i.e.*, it may not be part of a NES for the original game (*i.e.*, starting from $\tau = 0$). Intuitively, in such global-consistency issues [27, 36] (also called *safety issues* [10]), the choices made at latter time steps do not account for possible deviations from the opponent at earlier time steps.

As detailed in the next section, we will characterize a subproblem not with the raw data of partial strategy profiles, but with a sufficient statistic, and this characterization will be used as the state of a new dynamic game equivalent to the zs-POSG.

2.2 Occupancy State and Occupancy Markov Game

We now introduce an equivalent game, in which trajectories correspond to behavioral strategy profiles, and which we will be able to decompose temporally (and recursively), a first key tool for DP and HS.

To cope with the necessarily continuous nature of its state space, we will set this game in occupancy space, *i.e.*, a statistic that sums up past DR profiles. This will let us derive continuity properties on which to build point-based approximators.

As Wiggers et al. [47], let us formally define an *occupancy state* (OS) $\sigma_{\beta_{0:\tau-1}}$ as the probability distribution over joint AOHs θ_τ given partial strategy profile $\beta_{0:\tau-1}$. This statistic exhibits the usual Markov and sufficiency properties:

Proposition 1 (Adapted from Dibangoye et al. [14, Thm. 1] – Proof in App. B.1) $\sigma_{\beta_{0:\tau-1}}$, together with β_τ , is a sufficient statistic to compute (i) the next OS, $T(\sigma_{\beta_{0:\tau-1}}, \beta_\tau) \stackrel{\text{def}}{=} \sigma_{\beta_{0:\tau}}$, and (ii) the expected reward at τ : $r(\sigma_{\beta_{0:\tau-1}}, \beta_\tau) \stackrel{\text{def}}{=} \mathbb{E}[R_\tau | \beta_{0:\tau-1} \oplus \beta_\tau]$, where \oplus denotes a concatenation.

Writing from now on σ_τ , as short for $\sigma_{\beta_{0:\tau-1}}$, the OS associated with some prefix strategy profile $\beta_{0:\tau-1}$, the proof essentially relies on deriving the following formulas: $\forall \theta_\tau^1, a^1, z^1, \theta_\tau^2, a^2, z^2$,

$$\begin{aligned} T(\sigma_\tau, \beta_\tau)((\theta_\tau^1, a^1, z^1), (\theta_\tau^2, a^2, z^2)) & \quad (1) \\ & \stackrel{\text{def}}{=} Pr((\theta_\tau^1, a^1, z^1), (\theta_\tau^2, a^2, z^2) | \sigma_\tau, \beta_\tau) \\ & = \beta_\tau^1(\theta_\tau^1, a^1) \beta_\tau^2(\theta_\tau^2, a^2) \sigma_\tau(\theta_\tau) \sum_{s, s'} P_{\mathbf{a}}^z(s' | s) b(s | \theta_\tau), \end{aligned}$$

where $b(s | \theta_\tau)$ is a *belief state* obtained by Hidden Markov Model filtering; and

$$\begin{aligned} r(\sigma_\tau, \beta_\tau) & \stackrel{\text{def}}{=} E[r(S, A^1, A^2) | \sigma_\tau, \beta_\tau] & (2) \\ & = \sum_{s, \theta_\tau, \mathbf{a}} \sigma_\tau(\theta_\tau) b(s | \theta_\tau) \beta_\tau^1(a^1 | \theta_\tau^1) \beta_\tau^2(a^2 | \theta_\tau^2) r(s, \mathbf{a}). \end{aligned}$$

We can then derive, from a zs-POSG, a non-observable zero-sum game in which players' strategies are private.

Definition 3 (zero-sum occupancy Markov Game (zs-oMG)) A *zero-sum occupancy Markov game* (zs-oMG)³ is defined by the tuple $\langle \mathcal{O}^\sigma, \mathcal{B}, T, r, H, \gamma \rangle$, where:

- $\mathcal{O}^\sigma (= \cup_{t=0}^{H-1} \mathcal{O}_t^\sigma)$ is the set of OSS induced by the zs-POSG;
- \mathcal{B} is the set of DR profiles of the zs-POSG;
- T is the deterministic transition function in Eq. (1);

³We use (i) “Markov game” instead of “stochastic game” because the dynamics are not stochastic, and (ii) “partially observable stochastic game” to stick with the literature.

8 *HSVI can solve zero-sum Partially Observable Stochastic Games*

- r is the reward function in Eq. (2); and
- H and γ are as in the zs-POSG

(b_0 is not in the tuple but serves to define T and r).

In this game, as in the zs-POSG, a player's solution is a behavioral strategy. Besides, the value of a strategy profile $\beta_{0:H-1}$ is the same for both zs-oMG and zs-POSG, so that they share the same ϵ -NEV and ϵ -NESs. We can thus work with zs-oMGs as a means to solve zs-POSGs.

The following aims at deriving a recursive expression of V_0^* , as well as continuity properties.

Bellman Optimality Equation

Despite the OS at $\tau > 0$ not being accessible to any player, let us define a *subgame* at σ_τ as the restriction starting from time step τ under this particular occupancy state, meaning that we are seeking strategies $\beta_{\tau:H-1}^1$ and $\beta_{\tau:H-1}^2$. σ_τ tells us which AOHs each player could be facing with non-zero probability, and are thus relevant for planning. We can then define the value function in any OS σ_τ for any strategy profile $\beta_{\tau:H-1}$ as follows:

$$V_\tau(\sigma_\tau, \beta_{\tau:H-1}) \stackrel{\text{def}}{=} E\left[\sum_{t=\tau}^{H-1} \gamma^{t-\tau} r(S_t, A_t) \mid \sigma_\tau, \beta_{\tau:H-1}\right]. \quad (3)$$

We will see in Section 3.1 that (i) the subgame rooted at σ_τ has a NEV, and thus a unique optimal value $V^*(\sigma_\tau)$, and (ii) it corresponds to the value of Wiggers et al.'s subgames. With this observation, their following key continuity result also applies to our zs-oMG.

Concavity and Convexity Results

As a preliminary step, Wiggers et al. decompose an occupancy state σ_τ into a *marginal term* $\sigma_\tau^{m,1}$ and a *conditional term* $\sigma_\tau^{c,1}$, where

- $\sigma_\tau^{m,1}(\theta_\tau^1) = \sum_{\theta_\tau^2} \sigma_\tau(\theta_\tau^1, \theta_\tau^2)$ is the probability of 1 facing θ_τ^1 under σ_τ , and
- $\sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1) = \frac{\sigma_\tau(\theta_\tau^1, \theta_\tau^2)}{\sigma_\tau^{m,1}(\theta_\tau^1)}$ is the probability of 2 facing θ_τ^2 under σ_τ given that 1 faces θ_τ^1 ,

so that $\sigma_\tau(\theta_\tau^1, \theta_\tau^2) = \sigma_\tau^{m,1}(\theta_\tau^1) \cdot \sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1)$. (Symmetric definitions apply by swapping players 1 and 2.) In addition, let us denote $T_m^1(\sigma_\tau, \beta_\tau)$ and $T_c^1(\sigma_\tau, \beta_\tau)$ the marginal and conditional terms associated to $T(\sigma_\tau, \beta_\tau)$.

Now, if 1 faces AOH θ_τ^1 , knows 2's future strategy $\beta_{\tau:H-1}^2$, and has access to $\sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1)$ for any θ_τ^2 , then she faces a POMDP whose optimal value we denote $\nu_{[\sigma_\tau^{c,1}, \beta_{\tau:H-1}^2]}^2(\theta_\tau^1)$. This leads to defining the best-response value vector $\nu_{[\sigma_\tau^{c,1}, \beta_{\tau:H-1}^2]}^2$, which contains one component per AOH θ_τ^1 , and writing the value of 1's best response against $\beta_{\tau:H-1}^2$ under σ_τ as $\sigma_\tau^{m,1} \cdot \nu_{[\sigma_\tau^{c,1}, \beta_{\tau:H-1}^2]}^2$. But then,

because 2 also knows σ_τ , she can in fact pick $\beta_{\tau:H-1}^2$ to minimize this value, so that we get the following theorem.

Theorem 2 ([47, Thm. 2]) *For any $\tau \in \{0 \dots H-1\}$, V_τ^* is (i) concave w.r.t. $\sigma_\tau^{m,1}$ for a fixed $\sigma_\tau^{c,1}$, and (ii) convex w.r.t. $\sigma_\tau^{m,2}$ for a fixed $\sigma_\tau^{c,2}$. More precisely,*

$$V_\tau^*(\sigma_\tau) = \min_{\beta_{\tau:H-1}^2} \left[\sigma_\tau^{m,1} \cdot \nu_{[\sigma_\tau^{c,1}, \beta_{\tau:H-1}^2]}^2 \right] = \max_{\beta_{\tau:H-1}^1} \left[\sigma_\tau^{m,2} \cdot \nu_{[\sigma_\tau^{c,2}, \beta_{\tau:H-1}^1]}^1 \right], \text{ where}$$

$$\nu_{[\sigma_\tau^{c,1}, \beta_{\tau:H-1}^2]}^2(\theta_\tau^1) \stackrel{\text{def}}{=} \max_{\beta_{\tau:H-1}^1} \mathbb{E}_{\theta_\tau^2 \sim \sigma_\tau^{c,1}(\theta_\tau^1)} \left\{ \sum_{t=\tau}^{H-1} \gamma^{t-\tau} r(S_t, A_t^1, A_t^2) \mid \beta_{\tau:H-1}^1, \beta_{\tau:H-1}^2 \right\}.$$

An important observation that ensues from this theorem is that V_τ^* is concave in $\sigma_\tau^{m,1}$ and convex in $\sigma_\tau^{m,2}$. In practice, however, such continuity properties alone only allow upper-bounding V_τ^* for finitely many conditional terms $\sigma_\tau^{c,i}$, thus *not* for the whole occupancy space, as required to enable DP and HS in our game.

In the following, we contribute novel results : properties of V^* in subgames, plus continuity properties that help designing bounding approximators, which will lead us to an HSVI-like solver.

Note:

To help the reader, Appendix A provides three synthetic tables: Table A1 (p. 30) to sum up various theoretical properties that are stated in this paper (assuming a finite temporal horizon), Table A2 (p. 30) and Table A3 (p. 32) to respectively sum up the notations and abbreviations used in this paper, including some that are used only in the appendix.

Also, for convenience, we may replace in the following: (i) subscript “ τ : $H-1$ ” with “ τ :”, (ii) any function $f(\mathbf{x})$ linear in vector \mathbf{x} with either $f(\cdot) \cdot \mathbf{x}$ or $\mathbf{x}^\top \cdot f(\cdot)$, (iii) a full tuple with its few elements of interest, and (iv) an element (a “field”) x of a specific tuple t by $x[t]$.

3 Theoretical Contributions

In this section, we demonstrate how to implement dynamic programming and heuristic search by (1) rigorously showing that Bellman optimality equation (Sec. 3.1) holds, (2) deriving bounding approximators of two novel optimal value functions, which come with solution strategies (Sec. 3.2), and (3) proposing a variant of HSVI that computes (in finite time) a player’s strategy whose value is within ϵ of the zs-POSG’s NEV (Sec. 3.3).

3.1 The Optimal Value Function V^* and its Recursive Expression

Let us first recall that, contrary to Wiggers et al. [47, Section 5, Lemma 4], we do not make the strong assumption that past decision rules can be considered as public (and, thus, we do not assume that any player knows σ_τ).

Indeed, while it is valid in Dec-POMDPs because the players are willing to coordinate their behaviors, it is *a priori* not valid in zs-POSGs, since players are, in the contrary, willing to deceive one another. Safety issues as presented in Example 1 illustrate the possible flaws of such an assumption.

We now discuss the existence of an optimal value function V_τ^* and its properties. These important results are implicitly used by Wiggers et al., who rely on the problematic assumption that player's strategies are public. We state and establish these results while considering private strategies. A first step is to demonstrate that von Neumann's minimax theorem [43] applies when in σ_τ , thus justifying the definition of the optimal (Nash equilibrium) value of a subgame.

Theorem 3 (Minimax theorem – Proof in App. C.1.3) *The subgame defined in Eq. (3) admits a unique NEV*

$$V_\tau^*(\sigma_\tau) \stackrel{\text{def}}{=} \max_{\beta_{\tau:H-1}^1} \min_{\beta_{\tau:H-1}^2} V_\tau(\sigma_\tau, \beta_{\tau:H-1}^1, \beta_{\tau:H-1}^2). \quad (4)$$

$V(\sigma_\tau, \cdot, \cdot)$ not being bilinear in the space of behavioral strategies (Appendix C.1.1), the proof requires reasoning with mixed strategies (for which the bilinearity holds), *i.e.*, distributions over pure (deterministic) strategies defined over *all* time steps. Yet, when in a subgame, we have to reason only on mixed strategies *compatible* with the associated occupancy state σ_τ (*i.e.*, which ensure that the OS at τ is σ_τ), one step being to extend Kuhn's equivalence results between behavioral and mixed strategies [30] to the subgames.

Then, defining the optimal action-value function:

$$Q_\tau^*(\sigma_\tau, \beta_\tau) \stackrel{\text{def}}{=} r(\sigma_\tau, \beta_\tau) + \gamma V_{\tau+1}^*(T(\sigma_\tau, \beta_\tau)), \quad (5)$$

we can now prove that a Bellman optimality equation exists, which justifies reasoning on subgames despite the non-observability.

Theorem 4 (Bellman optimality equation – Proof in App. C.1.3) $V_\tau^*(\sigma_\tau)$ *satisfies the following functional equation:*

$$V_\tau^*(\sigma_\tau) = \max_{\beta_\tau^1} \min_{\beta_\tau^2} r(\sigma_\tau, \beta_\tau) + \gamma V_{\tau+1}^*(T(\sigma_\tau, \beta_\tau)) = \max_{\beta_\tau^1} \min_{\beta_\tau^2} Q_\tau^*(\sigma_\tau, \beta_\tau).$$

The proof requires decomposing min and max operators over different time steps before swapping them appropriately to end up recognizing the optimal value function at the next time step ($V_{\tau+1}^*$).

Theorems 3 and 4 together show that Theorem 2 holds even without player's strategies being public so that we can now build on the convex-concave property to solve zs-oMGs.

3.2 Towards Solving zs-OMGs

This section aims at providing the second tool for DP and HS with continuous state spaces, *i.e.*, bounding approximators of optimal value functions which will allow generalization across occupancy space. Their update and selection operators are written as linear programs, and they turn out to come with solution strategies.

3.2.1 Bounding value functions

So far, several issues prevented to apply the HSVI scheme to zs-POSGs, starting with the continuous spaces of 1. occupancy states (zs-OMG states) and 2. decision rules (zs-OMG actions). One can address (1) by introducing the bounding functions $\bar{V}_\tau(\sigma_\tau)$ and $\underline{V}_\tau(\sigma_\tau)$ of $V_\tau^*(\sigma_\tau)$ (*cf.* App. D.2), for instance:

$$\bar{V}_\tau(\sigma_\tau) = \min_{\langle \tilde{\sigma}_\tau^{c,1}, \langle \bar{\nu}_\tau^2, \beta_\tau^2 \rangle \rangle \in \bar{\mathcal{J}}_\tau} [\sigma_\tau^{m,1} \cdot \bar{\nu}_\tau^2 + \lambda_\tau \|\sigma_\tau - \sigma_\tau^{m,1} \tilde{\sigma}_\tau^{c,1}\|_1],$$

where $\bar{\nu}_\tau^2$ component-wise upper-bounds $\nu_{[\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2]}^2$ for some β_τ^2 . They allow generalizing knowledge from the subgame rooted at σ_τ to any other one rooted at $\tilde{\sigma}_\tau$. To do so, we use V^* 's Lipschitz-Continuity proven below.

Theorem 5 (Lipschitz-Continuity of V^* - proof in App. D.1.3) *Let $h_\tau \stackrel{\text{def}}{=} \frac{1-\gamma^{H-\tau}}{1-\gamma}$ (or $h_\tau \stackrel{\text{def}}{=} H - \tau$ if $\gamma = 1$). Then $V_\tau^*(\sigma_\tau)$ is λ_τ -Lipschitz continuous in σ_τ at any depth $\tau \in \{0 \dots H - 1\}$, where $\lambda_\tau = \frac{1}{2}h_\tau (r_{\max} - r_{\min})$.*

Yet, this yields (generally non-convex) Lipschitz-continuous functions whose max-min optimization would be intractable, so that (2) remains an issue. Also, we do not know how to retrieve valid solution strategies. In particular, and as illustrated in Example 1, simply concatenating decision rules backwards from $\tau = H - 1$ to 0 would not guarantee globally-consistent solutions, and could result in exploitable strategies.

But then, combining Theorems 2 and 4 leads to introducing a novel value function (denoted $W_\tau^{1,*}$) through writing, for any OS σ_τ :

$$V_\tau^*(\sigma_\tau) = \max_{\beta_\tau^1} \min_{\beta_{\tau:H-1}^2 \in \mathcal{B}_\tau^2} \underbrace{\left[r(\sigma_\tau, \beta_\tau) + \gamma \sigma_{\tau+1}^{m,1} \cdot \nu_{[\sigma_{\tau+1}^{c,1}, \beta_{\tau+1:H-1}^2]}^2 \right]}_{\stackrel{\text{def}}{=} W_\tau^{1,*}(\sigma_\tau, \beta_\tau^1)}.$$

Assuming that player 2 can only respond with one of finitely many stored strategies, the concavity and λ_τ -Lipschitz-continuity of $W_\tau^{1,*}$ allow upper-bounding it with finitely many tuples $w = \langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{\nu}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle$ stored in sets $\bar{\mathcal{L}}_\tau$, and where $\bar{\nu}_{\tau+1}^2$ upper-bounds $\nu_{[\tilde{\sigma}_{\tau+1}^{c,1}, \beta_{\tau+1}^2]}^2$.

Proposition 6 (proof in App. D.2.2) *Let $\bar{\mathcal{I}}_\tau$ be a set of tuples $w = \langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{v}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle$. Then,*

$$\begin{aligned} \bar{W}_\tau(\sigma_\tau, \beta_\tau^1) \stackrel{\text{def}}{=} & \min_{\langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{v}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle \in \bar{\mathcal{I}}_\tau} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) \cdot \bar{v}_{\tau+1}^2 \right. \\ & \left. + \lambda_{\tau+1} \|T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) - T_m^1(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) T_C^1(\tilde{\sigma}_\tau^c, \beta_\tau^2)\|_1 \right] \end{aligned} \quad (6)$$

upper-bounds $W_\tau^{1,*}$ over the whole space $\mathcal{O}_\tau^\sigma \times \mathcal{B}_\tau^1$.

Symmetrically, we define \underline{W}_τ as the lower bound of the symmetrically defined $W_\tau^{2,*}$. As explained in the next two sections, \bar{W}_τ will be easier to deal with compared to \bar{V}_τ , allowing 1 to seek for decision rules optimistically, and providing valid solution strategies for 2 for the subgame at τ , *i.e.*, ignoring consistency with higher-level subgames.

3.2.2 Action Selection and Backup Operators

We now detail the decision rule selection for 1 using \bar{W}_τ to optimistically guide a trajectory in occupancy space, and how to update \bar{W}_τ by providing backup operations.

To that end, first note that linearities in β_τ^1 within Eq. (6) allow writing $\bar{W}_\tau(\sigma_\tau, \beta_\tau^1) = \min_{w \in \bar{\mathcal{I}}_\tau} \beta_\tau^{1\top} \cdot M_{(\cdot, w)}^{\sigma_\tau}$, where β_τ^1 and $M_{(\cdot, w)}^{\sigma_\tau}$ (for each w) are column vectors of dimension $|\Theta^1 \times \mathcal{A}^1|$. M^{σ_τ} (see developed formula in App. D.3.1) is thus a $|\Theta_\tau^1 \times \mathcal{A}^1| \times |\bar{\mathcal{I}}_\tau|$ matrix. Then, \bar{W}_τ being a lower envelope of hyperplanes leads to a convenient way of computing $\max_{\beta_\tau^1} \bar{W}_\tau(\sigma_\tau, \beta_\tau^1)$.

Proposition 7 (Proof in App. D.3.1) *For any given σ_τ and any set $\bar{\mathcal{I}}_\tau$ of tuples $w = \langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{v}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle$, $\max_{\beta_\tau^1} \bar{W}_\tau(\sigma_\tau, \beta_\tau^1)$ is equivalent to the LP and dual LP:*

$$\begin{aligned} \text{LP } \bar{W}_\tau(\sigma_\tau) : \quad & \max_{\beta_\tau^1, v} v \text{ s.t. } (i) \quad \forall w \in \bar{\mathcal{I}}_\tau, \quad v \leq \beta_\tau^{1\top} \cdot M_{(\cdot, w)}^{\sigma_\tau} \quad \text{and} \\ & (ii) \quad \forall \theta_\tau^1 \in \Theta_\tau^1, \quad \sum_{a^1} \beta_\tau^1(a^1 | \theta_\tau^1) = 1, \\ \text{DLP } \bar{W}_\tau(\sigma_\tau) : \quad & \min_{\psi_\tau^2, v} v \text{ s.t. } (i) \quad \forall (\theta_\tau^1, a^1), \quad v \geq M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2 \text{ and} \\ & (ii) \quad \sum_{w \in \bar{\mathcal{I}}_\tau} \psi_\tau^2(w) = 1. \end{aligned} \quad (7)$$

Remark 1 (Outcomes of this game) Since \bar{W}_τ upper-bounds $W_\tau^{1,*}$, solving this LP provides 1 with an *optimistically* selected immediate decision rule β_τ^1 . For 2, ψ_τ^2 is a probability distribution over tuples containing strategies $\beta_\tau^2 \oplus \beta_{\tau+1:H-1}^2$, thus recursively induces a strategy, as illustrated by Fig. 3, which can be turned into a behavioral strategy $\beta_{\tau:H-1}^2$ (more details in App. D.3.3) whose value is *at worst* (from 2's viewpoint) the LP's value, *i.e.*, against 1's best response to it.

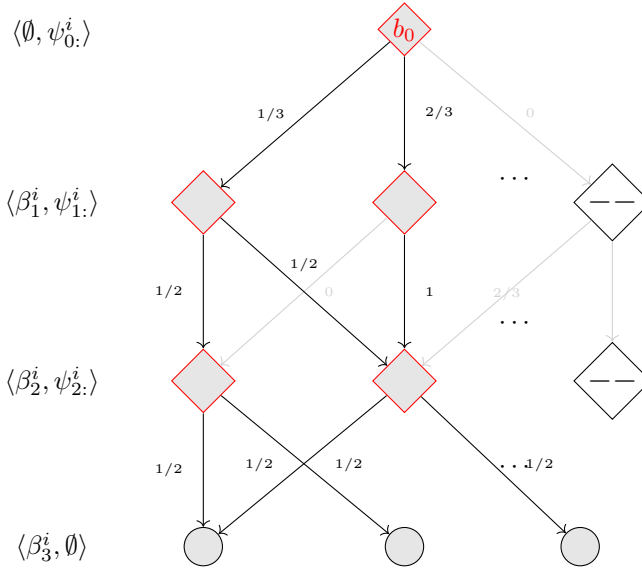


Fig. 3 Representation of the strategy recursively induced by some ψ_0^1 . At each time step τ , one must (i) sample a next tuple/node w_τ^1 from current distribution ψ_τ^1 , (ii) apply DR $\beta_\tau^1[w_\tau^1]$, and (iii) make $\psi_{\tau+1}^1[w_\tau^1]$ the new current distribution (unless reaching a leaf).

Then, the following properties allow performing backups, *i.e.*, filling up the set $\bar{\mathcal{L}}_{\tau-1}$ with new tuples w containing, in particular, vectors \bar{v}_τ^2 .

Lemma 1 (Proof in App. D.3.2). *For any $\psi_\tau^2 = \text{DLP}\bar{W}_\tau(\sigma_\tau)$, the vector $v_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2$ is component-wise upper-bounded by*

$$\bar{v}_\tau^2 \stackrel{\text{def}}{=} \frac{1}{\sigma_\tau^{m,1}} M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2.$$

Proposition 8 (update) *Let us assume that*

- a transition $\sigma_{\tau-1} \rightarrow \sigma_\tau$ has been performed through playing $\langle \beta_{\tau-1}^1, \beta_{\tau-1}^2 \rangle$, and
- solving $\text{DLP}\bar{W}_\tau(\sigma_\tau)$ provides both
 - a tree strategy ψ_τ^2 (as the main solution of the DLP), and
 - a vector $\bar{v}_\tau^2 = \frac{1}{\sigma_\tau^{m,1}} M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2$ (as a by-product).

Then,

1. $\bar{\mathcal{L}}_{\tau-1} \leftarrow \bar{\mathcal{L}}_{\tau-1} \cup \{ \langle \sigma_{\tau-1}^{c,1}, \beta_{\tau-1}^2, \langle \bar{v}_\tau^2, \psi_\tau^2 \rangle \rangle \}$ is a valid update operator in the sense that it preserves \bar{W}_τ 's upper-bounding property, and
2. similarly, $\bar{\mathcal{J}}_\tau \leftarrow \bar{\mathcal{J}}_\tau \cup \{ \langle \sigma_\tau^{c,1}, \langle \bar{v}_\tau^2, \psi_\tau^2 \rangle \rangle \}$ is a valid update operator for \bar{V}_τ .

3.2.3 Initialization

To initialize the bounds \overline{W}_τ and \overline{V}_τ for any time step, we begin by generating a trajectory in a forward phase. At each time step, a uniform decision rule is picked for both players to derive a sequence of occupancy states $\sigma_0, \dots, \sigma_{H-1}$. Then, during a backward phase, for each time step $\tau = H-1, \dots, 1$, we create a tuple $w_{\tau-1,init} = \langle \sigma_{\tau-1}^{c,1}, \beta_{\tau-1}^2, \langle \overline{v}_\tau^2, \psi_\tau^2 \rangle \rangle$, where

- $\sigma_{\tau-1}^{c,1}$ is the conditional term associated to $\sigma_{\tau-1}$;
- $\beta_{\tau-1}^2$ is a uniform decision rule;
- ψ_τ^2 is
 - a degenerate distribution over the only next tuple $w_{\tau+1}$ if $\tau < H-1$ (which induces a concatenation of uniform decision rules for all future time steps);
 - undefined if $\tau = H-1$;

and

- $\overline{v}_\tau^2(\theta_\tau^1) = r_{max} \cdot (H - \tau)$ for any history θ_τ^1 that player 1 could face.

Tuples $w_{\tau-1,init}$ are added to sets $\overline{\mathcal{I}}_{\tau-1}$. For any time step $\tau \geq 0$, we similarly create tuples $\langle \sigma_\tau^{c,1}, \langle \overline{v}_\tau^2, \psi_\tau^2 \rangle \rangle$ and add them to sets $\overline{\mathcal{J}}_\tau$. The lower bounds are initialized symmetrically.

We now show that occupancy states can also be prescriptive, allowing one to retrieve an ϵ -NES for the subgame at occupancy state σ_τ once the bounds are within ϵ from each other, in particular at $\tau = 0$.

3.2.4 Extracting a NES

Vectors \overline{v}_0^2 upper bounding the value of their associated strategies, the following result tells when and how to extract an ϵ -optimal solution strategy for this player.

Theorem 9 *If sets $\overline{\mathcal{J}}_0$ and $\underline{\mathcal{J}}_0$ are such that $\overline{V}_0(\sigma_0) - \underline{V}_0(\sigma_0) \leq \epsilon$, then $\arg \max_{\underline{w} \in \underline{\mathcal{J}}_0} \underline{\nu}_0^2$ and $\arg \min_{\overline{w} \in \overline{\mathcal{J}}_0} \overline{\nu}_0^2$ respectively provide strategies ψ_0^1 and ψ_0^2 that form an ϵ -NES of the zs-POSG.*

Proof First, let us notice that, at $\tau = 0$, the occupancy-state space is reduced to a singleton, $\{\sigma_0 = \langle 1 \rangle\}$, because of the single (empty) joint AOH. The value vectors ν are thus one-dimensional, and here considered as scalar numbers.

Let us assume that sets $\overline{\mathcal{J}}_0$ and $\underline{\mathcal{J}}_0$ are such that

$$\overline{V}_0(\sigma_0) - \underline{V}_0(\sigma_0) \leq \epsilon,$$

and let $\underline{w}^* = \langle \sigma_0^{c,1}, \langle \underline{\nu}_0^*, \psi_0^{1,*} \rangle \rangle$ and $\overline{w}^* = \langle \sigma_0^{c,1}, \langle \overline{\nu}_0^*, \psi_0^{2,*} \rangle \rangle$ be the tuples returned by $\arg \max_{\underline{w} \in \underline{\mathcal{J}}_0} \underline{\nu}_0^2$ and $\arg \min_{\overline{w} \in \overline{\mathcal{J}}_0} \overline{\nu}_0^1$. Then, noting that $\sigma_0 = \langle 1 \rangle$,

$$\begin{aligned} \nu_{[\sigma_0^{c,2}, \psi_0^{1,*}]}^1 - \nu_{[\sigma_0^{c,1}, \psi_0^{2,*}]}^2 &\leq \overline{\nu}_0^* - \underline{\nu}_0^* \\ &= \max_{\underline{w} \in \underline{\mathcal{J}}_0} \underline{\nu}_0^2 - \min_{\overline{w} \in \overline{\mathcal{J}}_0} \overline{\nu}_0^1 \end{aligned}$$

$$\begin{aligned}
&= \bar{V}_0(\sigma_0) - \underline{V}_0(\sigma_0) \\
&\leq \epsilon.
\end{aligned}$$

Thus, ψ_0^1 and ψ_0^2 are two strategies whose security levels (values against best-responding opponents) are ϵ -close, and thus form an ϵ -NES of the zs-POSG. \square

Note: This result can be generalized to any σ_τ at later time steps, but this generalization is not used in practice.

Distributions ψ_0^2 are stored and could be executed as is. Appendix D.3.3 still presents a conversion process to retrieve a behavioral strategy $\beta_{0:H-1}^2$ from a distribution ψ_0^2 over tuples $w \in \bar{\mathcal{I}}_0$. Next, we see how to design a practical HSVI-based algorithm that provably returns sets $\bar{\mathcal{J}}_0$ and $\underline{\mathcal{J}}_0$ satisfying Theorem 9 after finitely many iterations.

3.3 HSVI for zs-POSGs

This section details our adaptation of the general HSVI scheme for ϵ -optimally solving zs-POSGs, and presents a theoretical finite-time convergence property.

3.3.1 Algorithm

HSVI for zs-POSGs is described in Algorithm 1. As vanilla HSVI, it relies on (i) generating trajectories while acting optimistically (lines 10+11), *i.e.*, player 1 (resp. 2) acting “greedily” w.r.t. \bar{W}_τ (resp. \underline{W}_τ), and (ii) locally updating the upper and lower bounds (lines 17+18). Both phases rely on solving the same games described by LP (7). At $\tau = H - 1$, line 14 selects DRs by solving an exact game, and line 20 returns a distribution reduced to the single element added in line 15.

A key difference with Smith and Simmons’s HSVI algorithm [39] lies in the criterion for stopping trajectories. The branching factor for zs-oMGs being infinite, we make use of V^* ’s Lipschitz-continuity to implement the same adaptations as [22] used for zs-OS-POSGs. The Lipschitz-continuity allows controlling the variations of the value function within small balls of radius ρ around a previously visited occupancy state. A finite number of such balls is sufficient to cover the whole space. Then, Theorem 10 (below) ensures ϵ -optimality in finite time if stopping trajectories when $\bar{V}_\tau(\sigma_\tau) - \underline{V}_\tau(\sigma_\tau) \leq thr(\tau)$, with the threshold function $thr(\tau) \stackrel{\text{def}}{=} \gamma^{-\tau}\epsilon - \sum_{i=1}^{\tau} 2\rho\lambda_{\tau-i}\gamma^{-i}$.

Setting ρ

As can be observed, this threshold function should always return positive values, which requires a small enough (but > 0) ρ . For a given problem (*cf.* Thm. 16, App. E.1.1), the maximum possible value ρ_{\max} depends on the Lipschitz constants at each time step, which themselves depend on the initial upper and lower bounds of the optimal value function. Setting $\rho \in (0, \rho_{\max})$ means making a trade-off between generating many trajectories (small ρ) and long ones (large ρ).

Algorithm 1: zs-oMG-HSVI($b_0, [\epsilon, \rho]$)

 [here returning a tuple w_0 containing a solution strategy $\underline{\psi}_0^1$ for player 1]

```

1 Fct zs-oMG-HSVI( $b_0 \simeq \sigma_0$ )
2   foreach  $\tau \in 0 \dots H - 1$  do
3     Initialize  $\overline{V}_\tau, \underline{V}_\tau, \overline{W}_\tau,$  &  $\underline{W}_\tau$ 
4     while  $[\overline{V}_0(\sigma_0) - \underline{V}_0(\sigma_0) > thr(0)]$  do
5       Explore( $\sigma_0, 0, -, -$ )
6     return  $\arg \max_{w_0 \in \mathcal{J}_0} \underline{V}_0^1$ 

7 Fct Explore( $\sigma_\tau, \tau, \sigma_{\tau-1}, \beta_{\tau-1}$ )
8   if  $[\overline{V}_\tau(\sigma_\tau) - \underline{V}_\tau(\sigma_\tau) > thr(\tau)]$  then
9     if  $\tau < H - 1$  then
10       $\overline{\beta}_\tau^1 \leftarrow \text{LP}\overline{W}_\tau(\sigma)$ 
11       $\underline{\beta}_\tau^2 \leftarrow \text{LP}\underline{W}_\tau(\sigma)$ 
12      Explore( $T(\sigma_\tau, \overline{\beta}_\tau^1, \underline{\beta}_\tau^2), \tau + 1, \sigma_\tau, \langle \overline{\beta}_\tau^1, \underline{\beta}_\tau^2 \rangle$ )
13    else ( $\tau = H - 1$ )
14       $(\overline{\beta}_\tau^1, \underline{\beta}_\tau^2) \leftarrow \text{NES}(r(\sigma, \beta_\tau^1, \beta_\tau^2))$ 
15       $\overline{\mathcal{I}}_\tau^1 \leftarrow \overline{\mathcal{I}}_\tau^1 \cup \{\langle \sigma_\tau^{c,1}, \underline{\beta}_\tau^2, - \rangle\}$ 
16       $\underline{\mathcal{I}}_\tau^2 \leftarrow \underline{\mathcal{I}}_\tau^2 \cup \{\langle \sigma_\tau^{c,2}, \overline{\beta}_\tau^1, - \rangle\}$ 
17       $\overline{\text{Update}}(\overline{W}_{\tau-1}, \langle \sigma_\tau, \sigma_{\tau-1}^{c,1}, \underline{\beta}_{\tau-1}^2 \rangle)$ 
18       $\underline{\text{Update}}(\underline{W}_{\tau-1}, \langle \sigma_\tau, \sigma_{\tau-1}^{c,2}, \overline{\beta}_{\tau-1}^1 \rangle)$ 

19 Fct  $\overline{\text{Update}}(\overline{W}_{\tau-1}, \langle \sigma_\tau, \sigma_{\tau-1}^{c,1}, \underline{\beta}_{\tau-1}^2 \rangle)$ 
20    $\langle \overline{v}_\tau^2, \psi_\tau^2 \rangle \leftarrow \text{DLP}\overline{W}_\tau(\sigma_\tau, )$ 
21    $\overline{\mathcal{I}}_{\tau-1} \leftarrow \overline{\mathcal{I}}_{\tau-1} \cup \{\langle \sigma_{\tau-1}^{c,1}, \underline{\beta}_{\tau-1}^2, \langle \overline{v}_\tau^2, \psi_\tau^2 \rangle \rangle\}$ 
22    $\overline{\mathcal{J}}_\tau \leftarrow \overline{\mathcal{J}}_\tau \cup \{\langle \sigma_\tau^{c,1}, \langle \overline{v}_\tau^2, \psi_\tau^2 \rangle \rangle\}$ 

```

3.3.2 Finite-Time Convergence

Theorem 10 (Proof in App. E.2.1) *zs-oMG-HSVI (Alg. 1) terminates in finite time with an ϵ -approximation of $V_0^*(\sigma_0)$ that satisfies Theorem 9.*

The finite time complexity suffers from the same combinatorial explosion as for Dec-POMDPs, and is even worse as we have to handle “infinitely branching” trees of possible futures. More precisely, the bound on the number of iterations depends on the number of balls of radius ρ required to cover occupancy simplexes at each depth.

Also, the following proposition allows solving infinite horizon problems as well (when $\gamma < 1$) by bounding the length of HSVI’s trajectories using the boundedness of $\overline{V} - \underline{V}$ and the exponential growth of $thr(\tau)$.

Proposition 11 (Proof in App. E.2.2) *When $\gamma < 1$, the length of trajectories is upper bounded by $T_{\max} \stackrel{\text{def}}{=} \left\lceil \log_{\gamma} \frac{\epsilon - \frac{2\rho\lambda^{\infty}}{1-\gamma}}{W - \frac{2\rho\lambda^{\infty}}{1-\gamma}} \right\rceil$, where λ^{∞} is a depth-independent Lipschitz constant and $W \stackrel{\text{def}}{=} \|\bar{V}^{(0)} - \underline{V}^{(0)}\|_{\infty}$ is the maximum width between initializations.*

4 Experiments

Experiments presented in this section aim at validating the proposed approach and comparing its behavior to the behavior of some reference algorithms.

4.1 Setup

Benchmark Problems

Five benchmark problems were used. Adversarial Tiger and Competitive Tiger were introduced by Wiggers [45]. Mabc and Recycling Robot are well-known Dec-POMDP benchmark problems (cf. <http://masplan.org>) and were adapted to our competitive setting by making player 2 minimize (rather than maximize) the objective function. The fifth benchmark is the adaptation of the well-known Matching Pennies game detailed in Example 1, with a small difference in that $r(s_h, \cdot, a_h) = +2$ instead of $+1$; this change breaks the symmetry in the optimal strategy, so that HSVI can not find the NES by "chance" by trying uniform strategies. We only consider finite horizons H and $\gamma = 1$. Table 1 gives the cardinal of the state, action and observation sets for each of these problems.

Table 1 Number of states/actions/observations for each benchmark problem

| | S | A^1 | A^2 | O^1 | O^2 |
|-------------------|-----|-------|-------|-------|-------|
| Competitive Tiger | 2 | 4 | 4 | 3 | 3 |
| Adversarial Tiger | 2 | 3 | 2 | 2 | 2 |
| Recycling Robot | 4 | 3 | 3 | 2 | 2 |
| Mabc | 4 | 2 | 2 | 2 | 2 |
| Matching Pennies | 3 | 2 | 2 | 1 | 1 |

Algorithms

For conciseness, Algorithm 1 is here denoted HSVI, and compared against *Random* search and *Informed* search [45] (both using Wiggers's implementation (unlicensed and unreleased)), SFLP [24], and CFR+ [42] (both using open_spiegel [32] (Apache license)).

All algorithms (but SFLP, which is exact) used a target error $\epsilon = 1\%$ of the initial gap $H \cdot (r_{\max} - r_{\min})$. HSVI ran with $\lambda_{\tau} = (H - \tau) \cdot (r_{\max} - r_{\min})$, and ρ the middle of its feasible interval. We also use FB-HSVI's LPE lossless compression of probabilistically equivalent action-observation histories

in occupancy states, so as to reduce their dimensionality [14]. Experiments ran on an Ubuntu machine with i7-10810U 1.10 GHz Intel processor and 16 GB available RAM, and the code is available under MIT license at <https://gitlab.com/aureliendelage1/hsviforzspogs>.

Random and Informed, only ran once, providing fairly representative results.

4.2 Results

Performance Measures

A common performance measure in 2-player zero-sum games is the *exploitability* of a strategy $\beta_{0\cdot}^i$, *i.e.*, the difference between the strategy's *security level* (the value of $-i$'s best response to $\beta_{0\cdot}^i$) and the Nash equilibrium value $V_0^*(\sigma_0)$:

$$\begin{aligned} \text{exploitability}(\beta_{0\cdot}^i) &= |V^*(\sigma_0) - \sigma_0^{m,1} \cdot \nu_{[\sigma_0^{c,-i}, \beta_{0\cdot}^i]}^i| \\ &= |V^*(\sigma_0) - \nu_{[\sigma_0^{c,-i}, \beta_{0\cdot}^i]}^i|, \end{aligned}$$

noting that σ_0 is a degenerate distribution over a single element, the pair of empty action-observation histories. In our setting, it will be convenient to look at the (average) *exploitability of a strategy profile* $\langle \beta_{0\cdot}^1, \beta_{0\cdot}^2 \rangle$:

$$\begin{aligned} \text{exploitability}(\beta_{0\cdot}^1, \beta_{0\cdot}^2) &= \frac{(V^*(\sigma_0) - \nu_{[\sigma_0^{c,2}, \beta_{0\cdot}^1]}^1) + (\nu_{[\sigma_0^{c,1}, \beta_{0\cdot}^2]}^2 - V^*(\sigma_0))}{2} \\ &= \frac{\nu_{[\sigma_0^{c,1}, \beta_{0\cdot}^2]}^2 - \nu_{[\sigma_0^{c,2}, \beta_{0\cdot}^1]}^1}{2}. \end{aligned}$$

This quantity is a more concise statistic than both individual exploitabilities, and can be obtained by solving two POMDPs (fixing one player's strategy or the other) without requiring to know the actual NEV.

This exploitability can also be defined as half of the *gap between security levels* (SL-gap). To analyze the convergence of algorithms with respect to the initial gap, we will look at the *SL-gap percentage*, *i.e.*,

$$\begin{aligned} \text{SL-gap percentage}(\beta_{0\cdot}^1, \beta_{0\cdot}^2) &= \frac{\nu_{[\sigma_0^{c,1}, \beta_{0\cdot}^2]}^2 - \nu_{[\sigma_0^{c,2}, \beta_{0\cdot}^1]}^1}{H \cdot (R_{\max} - R_{\min})} \\ &= \frac{2 \cdot \text{exploitability}(\beta_{0\cdot}^1, \beta_{0\cdot}^2)}{H \cdot (R_{\max} - R_{\min})}. \end{aligned}$$

4.2.1 Comparison with the state of the art

Table 2 gives the convergence time of Wiggers's two heuristic algorithms, CFR, CFR+, SFLP, and HSVI on the benchmark problems with various horizons, or the SL-gap percentage when reaching a 1 h time limit. Executions not returning any result (*i.e.*, for Random, Informed, CFR and CFR+, not performing a single iteration) are noted out-of-time [OOT].

This table first shows that HSVI always outperforms the heuristic baseline provided by Wiggers’s algorithms, thus proving the interest of an HSVI scheme. However, HSVI is outperformed by SFLP, CFR and CFR+, unless they run out of time. As can be noted, HSVI is able to keep improving even when the horizon grows thanks to the LPE compression, taking advantage of underlying structure in some games (*e.g.*, Recycling Robot, a problem with transition+observation independence (TOI), when scaling to larger horizons).

Table 2 Comparison of different solvers on various benchmark problems. Reported values are the running times until the algorithm’s error gap (based on bounds for HSVI) is lower than 1%, or, if the timeout limit is reached, the security-level gap percentages (100% if $\text{gap} = H \cdot (R_{\max} - R_{\min})$). Notes: (1) Horizons with a star exponent (H^*) are those for which the security-level computations ran out of time so that, for HSVI, we give the gap between the pessimistic bounds. (2) Even though Random and Informed contain randomness, we ran them only once, getting fairly representative results.

| Domain | H | Wiggers | | HSVI | SFLP | CFR | CFR+ |
|---------------|----|---------|-------|--------|-------------|-------|-------|
| | | Rand. | Inf. | | | | |
| Comp Tiger | 2 | 2.6% | 8.3% | 6 s | 1 s | 18 s | 2 s |
| | 3 | 7.0% | 6.1% | 3.8 % | 48 s | 57 m | 14 m |
| | 4 | 12.1% | 7.7% | 4.8 % | 14 m | [OOT] | [OOT] |
| | 5* | [OOT] | [OOT] | 53.3 % | [OOT] | [OOT] | [OOT] |
| Rec. Robot | 2 | 3.4% | 5.1% | 5 s | 1 s | 10 s | 1 s |
| | 3 | 9.2% | 15.2% | 4 m | 1 s | 10 m | 1 m |
| | 4 | 14.1% | 19.6% | 4.9 % | 13 s | 2.5 % | 24 m |
| | 5 | [OOT] | [OOT] | 10.7 % | [OOT] | [OOT] | [OOT] |
| | 6* | [OOT] | [OOT] | 45.5 % | [OOT] | [OOT] | [OOT] |
| Adv Tiger | 2 | 1 s | 3.7% | 1 s | 1 s | 1 s | 1 s |
| | 3 | 1.5% | 4.4% | 2 m | 1 s | 13 s | 8 s |
| | 4 | 2.9% | 5.6% | 2.6 % | 8 s | 15 m | 4 m |
| MABC | 2 | 45 s | 18.8% | 8 s | 3 s | 2 m | 1 s |
| | 3 | 4.2% | 9.2% | 27 s | 1 s | 15 m | 10 s |
| | 4 | 18.1% | 36.3% | 4.4 % | 3 s | 1.9 % | 4 m |
| MP | 4 | 2 m | 46.7% | 5 s | 1 s | 1 m | 1 s |
| | 5 | 9 m | 45.8% | 1 m | 1 s | 7 m | 5 s |
| | 6 | 2.2% | 44.6% | 8 m | 2 s | 35 m | 17 s |

We now study the dynamic behavior of the algorithms at hand by providing and analyzing the bounds and exploitability graphs for the same benchmarks.

4.2.2 Bounding Graphs

Left-side graphs in Figures 4 to 9 show how the computed upper- and lower-bounding values $\bar{V}_0(\sigma_0)$ and $\underline{V}_0(\sigma_0)$ (respectively the *dotted* dark and light

green curves) evolve as a function of computation time (always given in seconds). The *solid* dark and light green curves show the security levels $\nu_{[\sigma_0, \psi_0^2]}^1$ and $\nu_{[\sigma_0, \psi_0^1]}^2$ of the current returned strategies ψ_0^2 and ψ_0^1 . Note that, when best-response computations to obtain security levels are expensive (*e.g.*, for the competitive tiger problem, with $H = 4$), they are performed either periodically (*e.g.* every 10 iterations) or only once, at the end. In the captions, we indicate the (arbitrary) frequency of the POMDP evaluations. For example, (1, 1, *once*) means that, for the first two horizons, the POMDP evaluations were done after each iteration, and, for the last one, only once (at the end).

Overall, we observe consistent curves with (i) security levels in-between bounds and around the NEV, and (ii) bounds converging monotonically. Note that HSVI stops when the gap between bounds is small enough, while the gap between SLs (used by Informed, Random and CFR, and whose computation can be time-consuming) can be much smaller. As a matter of fact, one can notice that strategies ψ_0^i returned at each iteration by HSVI are often better (in terms of security level) than their pessimistic lower- or upper-bounding guarantees $\underline{\nu}_{[\sigma_0, \psi_0^i]}^1$ and $\bar{\nu}_{[\sigma_0, \psi_0^i]}^2$.

4.2.3 Exploitability Graphs

Right-side graphs in Figures 4 to 9 show the exploitability of the returned strategy profile as a function of computation time for HSVI, Random, Informed, and CFR+ for the different benchmarks considered. A limit precision of 10^{-7} (chosen empirically, according to the LP solver's precision) was applied to HSVI's exploitability.

As can be observed, Random and Informed tend to produce reasonable strategies quickly, but struggle to improve them so as to converge towards an ϵ -NES with $\epsilon \simeq 0$. In contrast, our algorithm keeps improving as computation time increases. The exploitation graphs support the observed behavior in Table 2 that HSVI converges in reasonable time compared to [Wiggers's](#) algorithms. However, the graphs also show that CFR and CFR+ essentially outperforms HSVI when the problems are difficult enough (*i.e.*, when the temporal horizon grows) but the traversal of the whole tree still remains tractable (thus allowing CFR and CFR+ to perform iterations). An interesting observation is that, on small enough problems, HSVI achieves very low exploitabilities earlier than CFR and CFR+. Finally, HSVI's exploitability graph shares strong similarities with those of [Bořanský et al.](#)'s double-oracle algorithms [6, Fig. 8 and 11]. This can be understood as HSVI iteratively building two sets of strategies, one per player, until they are sufficient to support NES profiles, so that the average exploitability is almost zero. But note that [Bořanský et al.](#) construct LPs using pure strategies (deterministic best responses), while HSVI's strategies are stochastic.

Having empirically studied the behavior of HSVI compared to other basic offline solvers, we now provide insight about the connections between HSVI and continual (thus online) resolving methods.

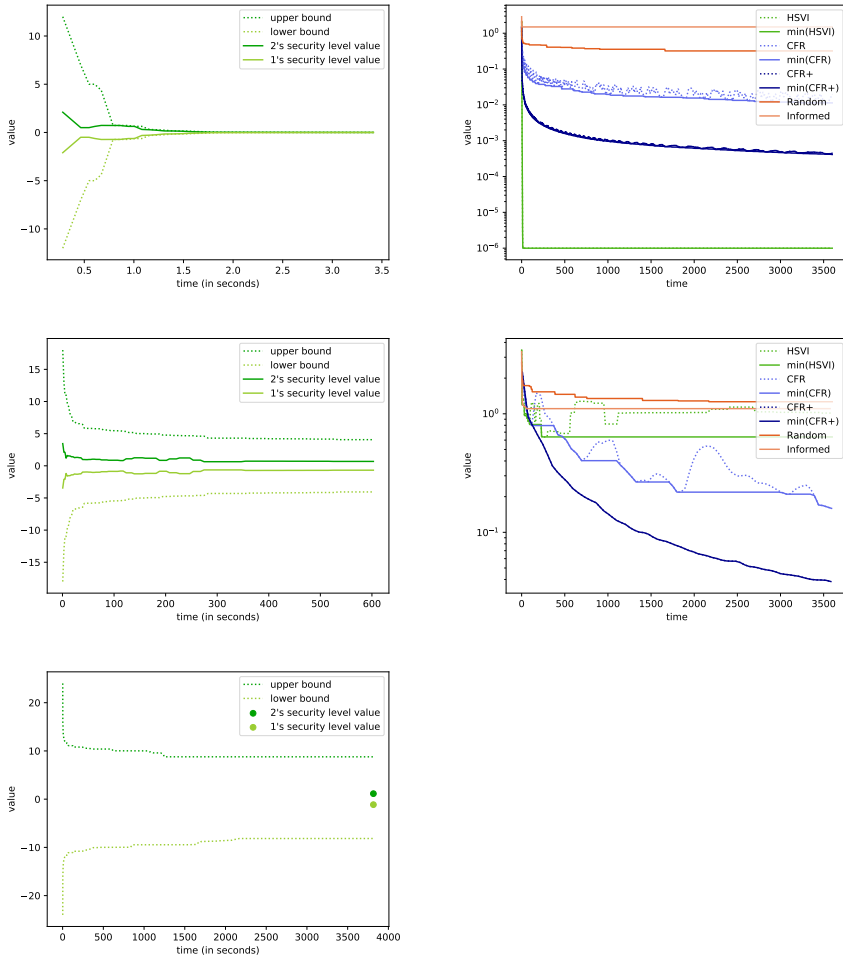


Fig. 4 Competitive Tiger ($H = 2, 3$) (1,1): (left) Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms). (right) Exploitability ($= \frac{SL\text{-gap}}{2}$) as a function of time (s) for Random, Informed, CFR+, and HSVI.

5 Related work

As already pointed out in the introduction, one can solve a zs-POSG by turning it into a zero-sum EFG and then using a solution technique such as SFLP [24, 44] or CFR [49]. In the following, we present and discuss in more details related work closest to our contribution.

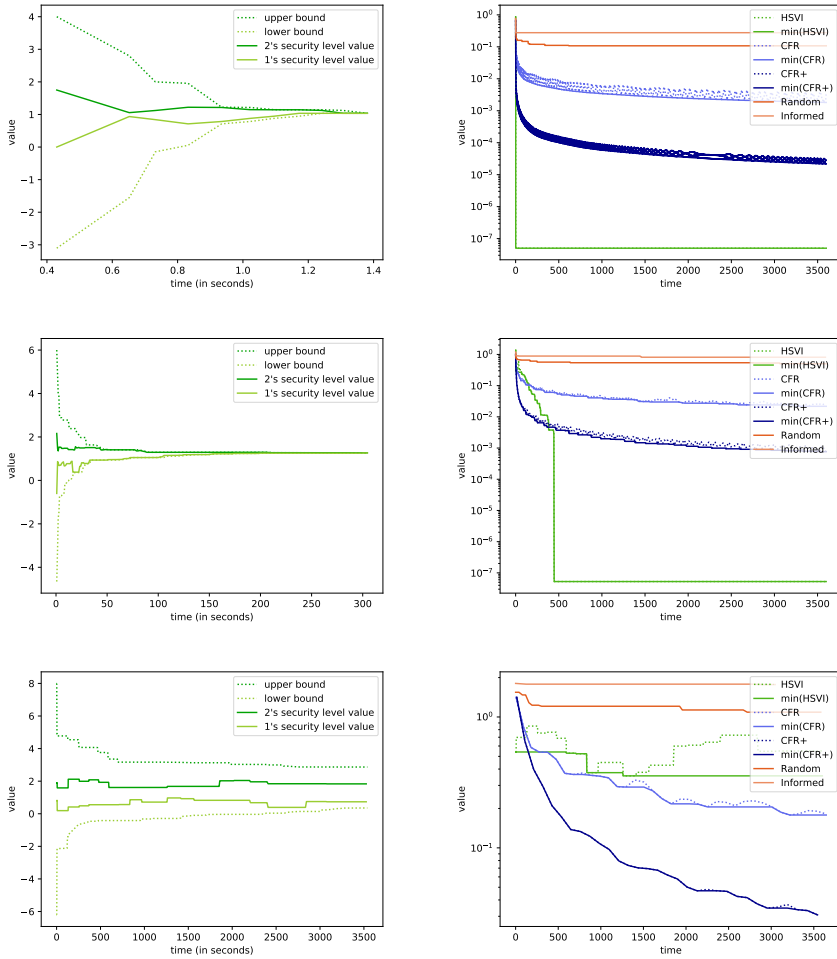


Fig. 5 Recycling Robot ($H = 2, 3, 4$) (1,1,10): (left) Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms). (right) Exploitability ($= \frac{SL\text{-gap}}{2}$) as a function of time (s) for Random, Informed, CFR+, and HSVI.

5.1 Structure in the Value Function of Two-Player Zero-Sum Games of Incomplete Information

Motivated by the transformation of a Dec-POMDP into a non-observable MDP, which is a special case of POMDP [35], Wiggers et al. [47] define a *plan-time non-observable stochastic game* (NOSG) as a tuple $\langle I, \hat{S}, \mathcal{A}^1, \mathcal{A}^2, \hat{O}, \hat{T}, \hat{O}, \hat{R}, \hat{b}_0 \rangle$, where:

- $I = \{1, 2\}$ is the set of players;

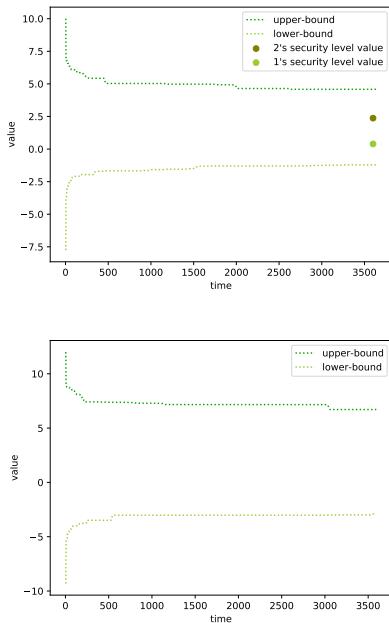


Fig. 6 Recycling Robot ($H = 5, 6$) (once,none): (left) Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms).

- \dot{S} is the set of augmented states \dot{s}_t , each corresponding to a joint AOH θ_t ;
- \mathcal{A}^1 and \mathcal{A}^2 are the decision rules sets as in zs-oMGs;
- \dot{O} is the set of public observations (set to $\{NULL\}$);
- \dot{T} is the transition function between states \dot{s}_t and \dot{s}_{t+1} ;
- \dot{O} is the observation function that specifies that observation $NULL$ is received with probability 1;
- \dot{R} the reward function : $\dot{R} : \dot{S} \times \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}$; and
- \dot{b}_0 is the initial belief ($\dot{b}_0 \in \Delta(\dot{S})$).

zs-oMGs having no observations is equivalent to plan-time NOSGs only working with a single trivial $NULL$ observation. The main difference between plan-time NOSGs and zs-oMGs lies in the state space. The “augmented states” of plan-time NOSGs, *i.e.*, tuples of hidden state and joint histories (s, θ_τ) are very different from occupancy states.⁴

This game is then viewed as a zero-sum *shared observation stochastic game* (zs-SOSG) with public actions by problematically arguing that, assuming that players are rational, their decision rules can be considered as public [47, Lemma 4]. In a very pedagogical discussion regarding public actions, described as the “non-correspondance problem”, Sokota et al. show the flaws of such a

⁴Note that the same distinction between two possible transformations exists for Dec-POMDPs that are translated either in Occupancy MDPs or in Non-Observable MDPs [14].

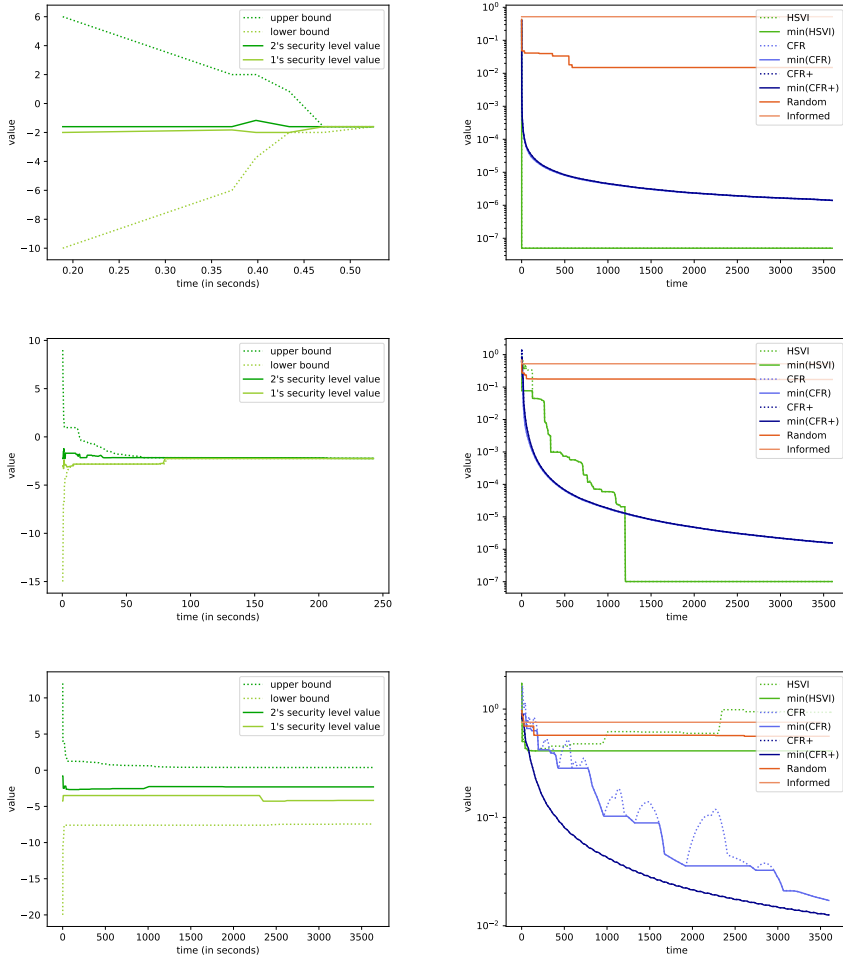


Fig. 7 Adversarial Tiger ($H = 2, 3, 4$) (1,1,10): (left) Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms). (right) Exploitability ($= \frac{SL\text{-gap}}{2}$) as a function of time (s) for Random, Informed, CFR+, and HSVI.

reasoning [40].⁵ In the contrary, we assume actions to be private, which requires proving several important results that are immediate with public actions.

5.2 Solving zero-sum One-Sided Partially Observable Stochastic Game

Inspired by techniques developed to solve POMDPs [39], and later Dec-POMDPs [14], Horák et al. [22] succeeded in designing an algorithm that

⁵In the Dec-POMDP setting, assuming that decision rules are public is valid because agents collaborate.

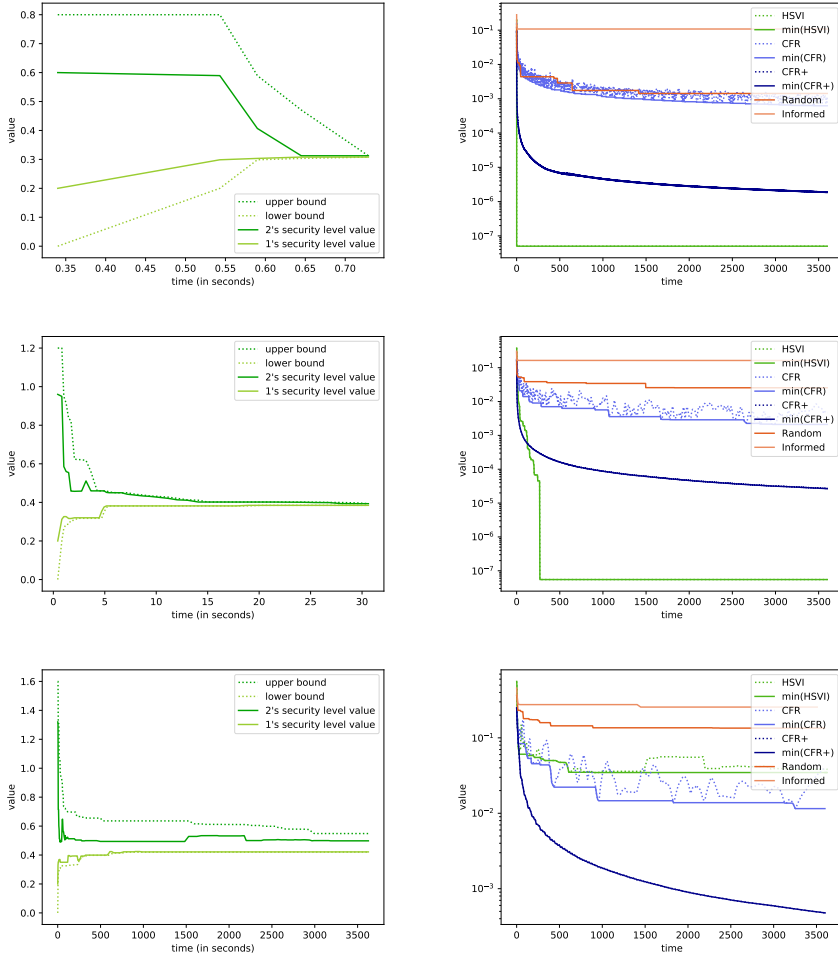


Fig. 8 **Mabc** ($H = 2, 3, 4$) **(1,1,10)**: **(left)** Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms). **(right)** Exploitability ($= \frac{SL\text{-gap}}{2}$) as a function of time (s) for Random, Informed, CFR+, and HSVI.

provably converges to an ϵ -NES of any zs-OS-POSG in finite time. It relies on a statistic (a belief-state) that induces a fully observable problem whose value function (V^*) exhibits continuity properties that allow deriving bounding approximations. Then, greedy selection and update operators are defined and included in an HSVI-like scheme.

We work in a similar direction for general zs-POSGs. As the state space, the strategy space and the observability assumptions (*e.g.* one player having perfect information about the current state of the game, including its opponent's past actions and observations) are different from their counterparts

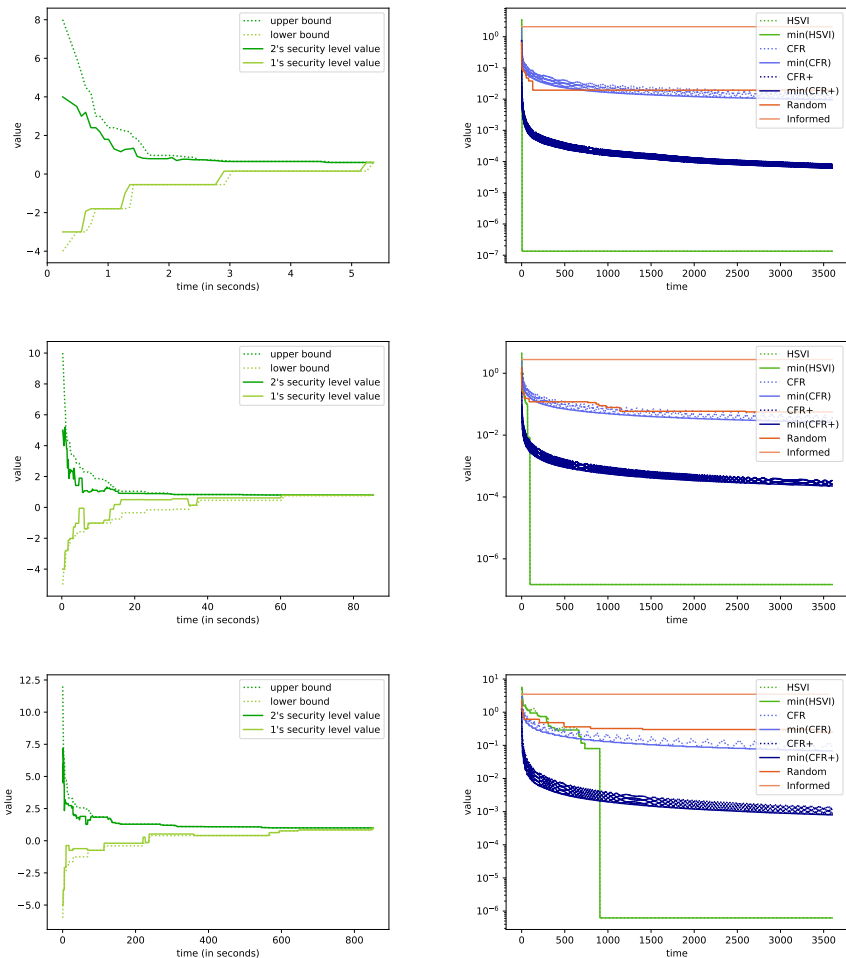


Fig. 9 **Matching Pennies ($H = 4, 5, 6$) (1,1,1):** (left) Evolution of (in dotted lines) the upper- and lower-bound values, and (in solid lines) the security levels of the returned strategies for HSVI as a function of time (ms). (right) Exploitability ($= \frac{SL\text{-gap}}{2}$) as a function of time (s) for Random, Informed, CFR+, and HSVI.

in zs-OS-POSGs, building bounding approximators for zs-POSGs relies on several technical novelties. A parallel can be made with dec-POMDPs with one-sided information sharing [48], whose observability hypotheses for one agent resulted in significant differences in the HSVI-like solving scheme w.r.t. HSVI for general dec-POMDPs.

Furthermore, an important distinction to make with Horák et al.’s work [19] is about extracting a strategy to execute. While they [19] rely on a continual resolving method [33] (see also Section 5.3) that is run online after ϵ -convergence of HSVI bounds, in our work, guaranteed strategies are intrinsically built when

updating bounding approximations, and can then be extracted at the end of the planning phase.

5.3 Comparison with Continual Resolving

Continual Resolving techniques share some similarities with our approach, but also important differences. The purpose of this section is to clarify these points. It starts with a quick description of CFR, on which Continual Resolving is built.

Counterfactual Regret Minimization (CFR) [49] belongs to the self-play family of algorithms, which gave rise to several CFR-based approaches [8, 11, 31, 42]. It iteratively traverses the whole game tree and applies, in each private history, a regret-matching update rule based on a specific type of regret called counterfactual regrets. Iteratively updating an initially uniform strategy asymptotically converges towards a NES. However, the tree traversal becomes intractable when the tree size is large.

Built on top of the CFR framework, approaches based on *limited-lookahead continual resolving* (LLCR) (inspired by Burch et al.'s decomposition [10]) such as DeepStack [33], Libratus [7], ReBeL [9] and Player of Games [37], perform well by exploiting a temporal decomposition in subgames, which are specified through knowledge about both players' past strategies. Our approach thus shares similarities with these works.

Yet, as we will see in the next sections, a closer look at LLCR [36] demonstrates how fundamentally different they are, starting with the fact that LLCR is an *online* search algorithm, *i.e.*, is meant to make good decisions at each time step, based on the current knowledge about the state of the game, while HSVI, as SFLP or CFR (and its variants), is an *offline* algorithm returning a complete solution strategy.

5.3.1 Continual Resolving

(Continual) Resolving techniques have been the first ingredient to adapt CFR to online settings. They assume that, at some time step τ of an actual trajectory, a Nash equilibrium strategy profile is being followed by both players (which induces the subgame at hand), and the objective is to incrementally complete it while preserving global consistency (aka safety) of the whole strategy, *i.e.*, not making choices that could encourage the opponent to deviate *in the past*, before τ . This is achieved by introducing, in a preliminary stage of the subgame, constraints, called *gadgets*, that represent possible deviations and their values, but increase the size of the game tree so that it is practically intractable [33]. For its part, HSVI solves similar subgames (offline), but at prefix strategy profiles that are not necessarily part of a Nash equilibrium, and ensuring only local consistency, *i.e.*, only considering the subgame. Global consistency comes from the way HSVI combines "lower-level" solutions in its backtracking process, without adding any gadget modifying the game.

As a comparison with classical (mono-agent, deterministic) search, Continual Resolving can be seen as calling Anytime A* [16] during an online resolution process, but without constraints to avoid safety issues, while HSVI can be seen as using LRTA* [25], which also works by optimistically generating trajectories until convergence.

Note that neither Resolving nor HSVI require knowing the opponent's actual strategy (which is not public). For Resolving, any opponent prefix strategy of a NES is appropriate to verify that the opponent has no incentive to deviate from a Nash equilibrium strategy. For HSVI, we have shown that reasoning on subgames assuming that the opponent's strategy is known still allows solving a zs-POSG. In practice, both algorithms reason not on actual prefix strategy profiles, but on different statistics. In Resolving, these statistics need to represent *complete* prefix strategy profiles (given the current public information), so that decisions are anticipated for player i even in AOHs (infosates) not reachable given player $\neg i$'s strategy. This leads to using *ranges* [26], a vector that gives, for each player, her contribution to the probability of any history she could face at time step τ . In contrast, HSVI's occupancy state, which is not necessarily related to a Nash equilibrium strategy in any manner, leads to ignoring unreachable AOHs, which helps to reduce the size of the decision-making (sub)problem.

5.3.2 Limited Lookahead

Continual Resolving alone solves complete subgames, thus larger problems at early stages of the game than at the end, which is not appropriate in an online setting. To address this issue through limiting the lookahead of subgames, one needs to estimate the value of the leaves of any truncated subgame. This is achieved through learning offline, for each player i , deep networks that, given the current public belief state, map each AOH to its value under some Nash equilibrium strategy profile. Note that the target function is not unique [26, Proposition A.1], since each NES profile maps to different value vectors. Still, according to [26], this does not seem to cause problems in practice.

In contrast, the individual value functions HSVI considers (the " ν " functions) are uniquely defined since they correspond to the best responses to given (not necessarily Nash equilibrium) strategies of the opponent.

5.3.3 Limited Lookahead Continual Resolving as a General Scheme?

The previous subsections highlight to what extent HSVI and LLCR are fundamentally different, in particular because they are not on the same algorithmic level. LLCR should be seen as a general scheme in which the subgame solver used, namely CFR, could be replaced by other "basic" offline algorithms such as HSVI or SFLP. But we leave further investigation on this topic for future work.

6 Discussion

This paper addresses the problem of ϵ -optimally solving zs-POSGs. In contrast to SFLP or CFR+, we provide the necessary foundational building blocks to apply dynamic programming (in tandem with heuristic search) to solve zs-POSGs. We introduce Bellman optimality equations and uniform-continuity properties of the optimal value function. Next, we exhibit rules for updating value functions while preserving uniform continuity and the ability to extract globally-consistent solutions. Finally, we describe the first effective DP algorithm for zs-POSGs, zs-oMG-HSVI, with finite-time convergence to an ϵ -optimal solution. Experiments support our theoretical findings.

We believe our approach complements existing ones, e.g., SFLP and CFR+, in two dimensions. First, it breaks the original zs-POSG into subgames. Second, it generalizes values from visited subgames to unvisited ones. Our performances are as good as or better than those from SFLP and CFR+ for small-dimensional subgames (e.g., with TOI structure). Unfortunately, the advantage of breaking the original problem into subgames and exploiting uniform continuity properties often fails to fully manifest in the overall computational time.

Despite some similarities, our (offline) approach is fundamentally different from (online) continual resolving approaches. The latter could even possibly be adapted to use other offline methods than CFR-based ones, including HSVI.

We hope that this approach will lay the foundation for further work in the area of both exact and approximate DP solutions for zs-POSGs. In the short term, we shall investigate pruning techniques, better Lipschitz constants, and improved initial bounding approximators using solutions from relaxations of zs-POSGs, e.g., zs-OS-POSGs. In the long term, we shall investigate (deep) RL for zs-POSGs, similarly to a recent approach for Dec-POMDPs [5]. The latter shall investigate the trade-off between the update-rule accuracy and the computational efficiency when facing high-dimensional subgames, hence providing competitive solvers.

Supplementary information. We provide the implementation of the HSVI algorithm as supplementary material.

Funding. This research was funded, in whole or in part, by ANR project PLASMA (ANR-19-CE23-0018). A CC-BY public copyright license has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript arising from this submission, in accordance with the grant's open access conditions. <https://creativecommons.org/licenses/by/4.0/>

Conflict of Interest. The authors have no conflicts of interest to declare that are relevant to the content of this article.

Appendix A Synthetic Tables

For convenience, we provide two synthetic tables: Table A1 to sum up various theoretical properties that are stated in this paper (assuming a finite temporal horizon), Table A2 and Table A3 to respectively sum up the notations and abbreviations used in this paper, adding some notations that appear only in the appendix.

More precisely, Table A1 indicates, for various functions f and variables x , properties that f is known to exhibit with respect to x . We denote by

- a function with no known (or used) property (see also comment below);

N/A a non-applicable case;

Lin a linear function;

LC a Lipschitz-continuous function;

Cv (resp. *Cc*) a convex (resp. concave) function;

PWLCv (resp. *PWLCc*) a piecewise linear and convex (resp. concave) function;

$\perp\!\!\!\perp$ the function being independent of the variable;

$\neg P$ the negation of some property P (*i.e.*, P is known not to hold).

Note also that, as $\sigma_\tau = \sigma_\tau^{c,1} \sigma_\tau^{m,1}$, the linearity or Lipschitz-continuity properties of any function w.r.t. σ_τ extends to both $\sigma_\tau^{c,1}$ and $\sigma_\tau^{m,1}$. Reciprocally, related negative results extend from $\sigma_\tau^{c,1}$ or $\sigma_\tau^{m,1}$ to σ_τ . In these three columns, we just indicate results that cannot be derived from one of the two other columns.

Table A1 Known properties of various functions appearing in this work

| | σ_τ | $\sigma_\tau^{m,1}$ | $\sigma_\tau^{c,1}$ | β_τ^i | β_τ^{-i} |
|---|--|------------------------------------|---------------------------|---|----------------------------|
| $T(\sigma_\tau, \beta_\tau)$ | <i>Lin</i> (thm. 1, p. 7) | - | - | <i>Lin</i> (thm. 1, p. 7) | <i>Lin</i> (thm. 1, p. 7) |
| $T_m^i(\sigma_\tau, \beta_\tau)$ | <i>Lin</i> (lem. 3, p. 39) | - | - | <i>Lin</i> (lem. 3, p. 39) | <i>Lin</i> (lem. 3, p. 39) |
| $T_c^i(\sigma_\tau, \beta_\tau)$ | - | $\perp\!\!\!\perp$ (lem. 4, p. 39) | - | $\perp\!\!\!\perp$ (lem. 4, p. 39) | - |
| $V_\tau^*(\sigma_\tau)$ | <i>LC</i> (app. D.1.3, p. 41) | <i>PWLCv</i> (thm. 2, p. 9) | - | N/A | N/A |
| $W_\tau^{i,*}(\sigma_\tau, \beta_\tau^i)$ | <i>LC</i> (from $Q_{\tau+1}^*$ <i>LC</i>) | - | - | \neg <i>Lin</i> (from Q^* \neg <i>Lin</i>) | N/A |
| $\nu_{[\sigma_\tau^{c,1}, \beta_\tau^i]}$ | N/A | N/A | <i>LC</i> (lem. 7, p. 41) | N/A | - |

Table A2: Various notations used in this work

$\neg i \stackrel{\text{def}}{=} i$'s opponent. Thus: $\neg 1 = 2$, and $\neg 2 = 1$.

Histories and occupancy states

$\theta_\tau^i \stackrel{\text{def}}{=} (a_0^i, z_1^i, \dots, a_{\tau-1}^i, z_\tau^i) (\in \Theta^i = \cup_{t=0}^{H-1} \Theta_t^i)$ is a length- τ *action-observation history* (AOH) for i .

$\theta_\tau \stackrel{\text{def}}{=} (\theta_\tau^1, \theta_\tau^2) (\in \Theta = \cup_{t=0}^{H-1} \Theta_t)$ is a *joint* AOH at τ .

- $\sigma_\tau(\boldsymbol{\theta}_\tau) \stackrel{\text{def}}{=} \text{Occupancy state (OS) } \sigma_\tau (\in \mathcal{O}^\sigma = \cup_{t=0}^{H-1} \mathcal{O}_t^\sigma, \text{ where } \mathcal{O}_\tau^\sigma \stackrel{\text{def}}{=} \Delta(\boldsymbol{\Theta}_\tau)), \text{ i.e., probability distribution over joint AOHs } \boldsymbol{\theta}_\tau \text{ (typically for some applied } \boldsymbol{\beta}_{0:\tau-1}).$
- $\sigma_\tau^{m,i}(\theta_\tau^i) \stackrel{\text{def}}{=} \text{Marginal term of } \sigma_\tau \text{ from player } i\text{'s point of view } (\sigma_\tau^{m,i} \in \Delta(\Theta_\tau^i)).$
- $\sigma_\tau^{c,i}(\theta_\tau^{-i}|\theta_\tau^i) \stackrel{\text{def}}{=} \text{Conditional term of } \sigma_\tau \text{ from } i\text{'s point of view } (\sigma_\tau^{c,i} : \Theta_\tau^i \mapsto \Delta(\Theta_\tau^{-i})).$
- $b(s|\boldsymbol{\theta}_\tau) \stackrel{\text{def}}{=} \text{Belief state, i.e., probability distribution over states given a joint AOH } (b(s|\boldsymbol{\theta}_\tau) : \mathcal{S} \times \boldsymbol{\Theta}_\tau \mapsto \mathbb{R}). \text{ Can be computed by an HMM filtering process.}$
- $o_\tau \stackrel{\text{def}}{=} \text{Full occupancy state } o_\tau (\in \Delta(\mathcal{S} \times \boldsymbol{\Theta}_\tau)), \text{ i.e., } Pr(s, \boldsymbol{\theta}_\tau) \text{ for the current } \boldsymbol{\beta}_{0:\tau-1}, \text{ and thus verifies } \sigma_\tau(\boldsymbol{\theta}_\tau) = \sum_{s \in \mathcal{S}} o_\tau(s, \boldsymbol{\theta}_\tau). \text{ Is used in the implementation to simplify computations (e.g., of } r_t \text{ and } \sigma_{\tau+1} \text{ through } b).$

Decision rules and strategies

- $\pi_{0:\tau}^i \stackrel{\text{def}}{=} \text{A pure strategy for } i \text{ is a mapping } \pi_{0:\tau}^i \text{ from private histories in } \Theta_t^i (\forall t \in \{0.. \tau\}) \text{ to single private actions in } \mathcal{A}^i. \text{ By default, } \pi^i \stackrel{\text{def}}{=} \pi_{0:H-1}^i.$
- $\mu_{0:\tau}^i \stackrel{\text{def}}{=} \text{A mixed strategy } \mu_{0:\tau}^i \text{ for } i \text{ is a probability distribution over pure strategies. It is used by first sampling one of the pure strategies (at } t = 0), \text{ and then executing it until } t = \tau.$
- $\mu_{0:\tau'}^i|\sigma_\tau \stackrel{\text{def}}{=} (\tau \leq \tau') \text{ is a mixed strategy compatible with some OS } \sigma_\tau, \text{ i.e., that could induce this OS at } \tau \text{ (assuming an appropriate complementary } \mu_{0:\tau'}^i|\sigma_\tau).$
- $\beta_\tau^i \stackrel{\text{def}}{=} \text{A (behavioral) decision rule (DR) at time } \tau \text{ for } i \text{ is a mapping } \beta_\tau^i \text{ from private AOHs in } \Theta_\tau^i \text{ to distributions over private actions. We note } \beta_\tau^i(\theta_\tau^i, a^i) \text{ the probability to pick } a^i \text{ when facing } \theta_\tau^i.$
- $\beta_{\tau:\tau'}^i \stackrel{\text{def}}{=} (\beta_\tau^i, \dots, \beta_{\tau'}^i) \text{ is a behavioral strategy for } i \text{ from time step } \tau \text{ to } \tau' \text{ (included).}$
- $rw^i(\theta_\tau^i, a_\tau^i) \stackrel{\text{def}}{=} \prod_{t=0}^{\tau} \beta_0^i(a_t^i|a_0^i, z_1^i, a_1^i, \dots, z_t^i) \text{ is the realization weight (RW) of sequence } a_0^i, z_1^i, a_1^i, \dots, a_\tau^i (= \theta_\tau^i, a_\tau^i) \text{ under strategy } \beta_{0:\tau}^i.$
- $rw^i(\phi_{\tau:\tau'}^i|\theta_\tau^i) \stackrel{\text{def}}{=} \prod_{t=\tau}^{\tau'} \beta_0^i(a_t^i|\theta_\tau^i, a_\tau^i, \dots, z_t^i) \text{ is the RW of a suffix sequence } \phi_{\tau:\tau'}^i = a_\tau^i, \dots, a_{\tau'}^i, \text{ "conditioned" on a prefix sequence/AOH } \theta_\tau^i.$
- $\boldsymbol{\pi}_{0:\tau} \stackrel{\text{def}}{=} \text{is a pure strategy profile.}$
- $\boldsymbol{\mu}_{0:\tau} \stackrel{\text{def}}{=} \text{is a mixed strategy profile.}$
- $\boldsymbol{\mu}_{0:\tau'}|\sigma_\tau \stackrel{\text{def}}{=} (\tau \leq \tau') \text{ is a mixed strategy profile compatible with some OS } \sigma_\tau, \text{ i.e., that could induce this OS at } \tau.$
- $\boldsymbol{\beta}_\tau \stackrel{\text{def}}{=} \langle \beta_\tau^1, \beta_\tau^2 \rangle (\in \mathcal{B} = \cup_{t=0}^{H-1} \mathcal{B}_t) \text{ is a decision rule profile.}$
- $\boldsymbol{\beta}_{\tau:\tau'} \stackrel{\text{def}}{=} \langle \beta_{\tau:\tau'}^1, \beta_{\tau:\tau'}^2 \rangle \text{ is a behavioral strategy profile.}$

Rewards and value functions

| | | |
|--|--|---|
| r_{\max} | $\stackrel{\text{def}}{=} \max_{s, \mathbf{a}} r(s, \mathbf{a})$ | Maximum possible reward. |
| r_{\min} | $\stackrel{\text{def}}{=} \min_{s, \mathbf{a}} r(s, \mathbf{a})$ | Minimum possible reward. |
| $V_{\tau}(\sigma_{\tau}, \beta_{\tau:})$ | $\stackrel{\text{def}}{=} E[\sum_{t=\tau}^{H-1} \gamma^t R_t \mid \sigma_{\tau}, \beta_{\tau:}]$, | Value of $\beta_{\tau:H-1}$ in OS σ_{τ} . where R_t is the random var. for the reward at t . |
| $V_{\tau}^*(\sigma_{\tau})$ | $\stackrel{\text{def}}{=} \max_{\beta_{\tau:}^1} \min_{\beta_{\tau:}^2} V_{\tau}(\sigma_{\tau}, \beta_{\tau:})$ | Optimal value function |
| $Q_{\tau}^*(\sigma_{\tau}, \beta_{\tau:})$ | $\stackrel{\text{def}}{=} r(\sigma_{\tau}, \beta_{\tau:}) + \gamma V_{\tau+1}^*(T(\sigma_{\tau}, \beta_{\tau:}))$ | Opt. (joint) action-value fct. |
| $W_{\tau}^{i,*}(\sigma_{\tau}, \beta_{\tau:}^i)$ | $\stackrel{\text{def}}{=} \text{opt}_{\beta_{\tau:}^i} Q_{\tau}^*(\sigma_{\tau}, \beta_{\tau:})$, | Opt. (individual) action-value fct. where $\text{opt} = \max$ if $i = 1$, \min otherwise. |
| $\nu_{[\sigma_{\tau}^{c,1}, \beta_{\tau:}^2]}^2$ | $\stackrel{\text{def}}{=} \text{Vector of values (one component per AOH } \theta_{\tau}^1 \text{ for 1's best response to } \beta_{\tau:}^2 \text{; assuming } \sigma_{\tau}^{c,1} \text{). This solution of a POMDP allows computing } V_{\tau}^* \text{ (see Theorem 2).}$ | |

Approximations

| | |
|--|---|
| $\bar{V}_{\tau}(\sigma_{\tau})$ | $\stackrel{\text{def}}{=} \text{Upper bound approximation of } V_{\tau}^*(\sigma_{\tau})$; relies on data set $\bar{\mathcal{I}}_{\tau-1}$. |
| $\underline{V}_{\tau}(\sigma_{\tau})$ | $\stackrel{\text{def}}{=} \text{Lower bound approximation of } V_{\tau}^*(\sigma_{\tau})$; relies on data set $\underline{\mathcal{I}}_{\tau-1}$. |
| $\bar{W}_{\tau}(\sigma_{\tau}, \beta_{\tau:}^1)$ | $\stackrel{\text{def}}{=} \text{Upper bound approximation of } W_{\tau}^{*,1}(\sigma_{\tau}, \beta_{\tau:}^1)$; relies on data set $\bar{\mathcal{I}}_{\tau}$. |
| $\underline{W}_{\tau}(\sigma_{\tau}, \beta_{\tau:}^2)$ | $\stackrel{\text{def}}{=} \text{Lower bound approximation of } W_{\tau}^{*,2}(\sigma_{\tau}, \beta_{\tau:}^2)$; relies on data set $\underline{\mathcal{I}}_{\tau}$. |
| $\bar{\nu}_{\tau}^2$ | $\stackrel{\text{def}}{=} \text{Vector (with one component per AOH } \theta_{\tau}^1 \text{) used in } \bar{V}_{\tau} \text{ and } \bar{W}_{\tau-1} \text{ (if } \tau \geq 1 \text{).}$ |

Miscellaneous

| | |
|------------------|--|
| w_{τ} | $\stackrel{\text{def}}{=} \text{Denotes a triplet } \langle \sigma_{\tau-1}^{c,1}, \beta_{\tau-1}^1, \bar{\nu}_{\tau}^2 \rangle \in \bar{\mathcal{I}}_{\tau} \text{ (or a triplet in } \underline{\mathcal{I}}_{\tau} \text{).}$ |
| ψ_{τ}^2 | $\stackrel{\text{def}}{=} \text{Distribution over triplets } w_{\tau+1} \in \bar{\mathcal{I}}_{\tau+1} \text{ (inducing a recursively defined strategy from } \tau \text{ to } H-1 \text{). Often denotes the strategy it induces.}$ |
| x^{\top} | $\stackrel{\text{def}}{=} \text{The transpose of a (usually column) vector } x \text{ of } \mathbb{R}^n \text{.}$ |
| $c[y]$ | $\stackrel{\text{def}}{=} \text{Denotes field } c \text{ of object/tuple } y \text{.}$ |
| $\text{Supp}(d)$ | $\stackrel{\text{def}}{=} \text{Support of distribution } d \text{, i.e., set of its non-zero probability elements.}$ |

Table A3: Various abbreviations used in this work

| | |
|--------------|---|
| zs-POSG [17] | $\stackrel{\text{def}}{=} \text{zero-sum Partially Observable Stochastic Game}$ |
| zs-oMG | $\stackrel{\text{def}}{=} \text{zero-sum Occupancy Markov Game}$ |
| DP | $\stackrel{\text{def}}{=} \text{Dynamic Programming}$ |
| HS | $\stackrel{\text{def}}{=} \text{Heuristic Search}$ |
| HSVI [39] | $\stackrel{\text{def}}{=} \text{Heuristic Search Value Iteration}$ |

| | |
|---------------|---|
| Dec-POMDP [4] | $\stackrel{\text{def}}{=} \text{Decentralized POMDP}$ |
| EFG [30] | $\stackrel{\text{def}}{=} \text{Extensive Form Game}$ |
| SFLP [24] | $\stackrel{\text{def}}{=} \text{Sequence Form Linear Program}$ |
| CFR [49] | $\stackrel{\text{def}}{=} \text{Counterfactual Regret Minimization}$ |
| NEV | $\stackrel{\text{def}}{=} \text{Nash Equilibrium Value}$ |
| NES | $\stackrel{\text{def}}{=} \text{Nash Equilibrium Strategy}$ |
| AOH | $\stackrel{\text{def}}{=} \text{Action-Observation History (for player } i: \theta_\tau^i)$ |
| DR | $\stackrel{\text{def}}{=} \text{Decision Rule (for player } i: \beta_\tau^i)$ |

Appendix B Background

B.1 Occupancy States

The following result shows that the occupancy state is (i) Markovian, *i.e.*, its value at τ only depends on its previous value $\sigma_{\tau-1}$, the system dynamics $P_{a^1, a^2}^{z^1, z^2}(s'|s)$, and the last behavioral decision rules $\beta_{\tau-1}^1$ and $\beta_{\tau-1}^2$, and (ii) sufficient to estimate the expected reward. Note that it holds for general-sum POSGs with any number of agents, and as many reward functions; similar results have already been established, *e.g.*, for Dec-POMDPs [14, Theorem 1].

Proposition 12 (originally stated on page 7) $\sigma_{\beta_{0:\tau-1}}$, together with β_τ , is a sufficient statistic to compute (i) the next OS, $T(\sigma_{\beta_{0:\tau-1}}, \beta_\tau) \stackrel{\text{def}}{=} \sigma_{\beta_{0:\tau}}$, and (ii) the expected reward at τ : $r(\sigma_{\beta_{0:\tau-1}}, \beta_\tau) \stackrel{\text{def}}{=} \mathbb{E}[R_\tau | \beta_{0:\tau-1} \oplus \beta_\tau]$, where \oplus denotes a concatenation.

Proof Let us first derive a recursive way of computing $\sigma_{\beta_{0:\tau}}(\theta_\tau, \mathbf{a}_\tau, \mathbf{z}_{\tau+1})$:

$$\begin{aligned}
 \sigma_{\beta_{0:\tau}}(\theta_\tau, \mathbf{a}_\tau, \mathbf{z}_{\tau+1}) &\stackrel{\text{def}}{=} Pr(\theta_\tau, \mathbf{a}_\tau, \mathbf{z}_{\tau+1} | \beta_{0:\tau}) \\
 &= \sum_{s_\tau, s_{\tau+1}} Pr(\theta_\tau, \mathbf{a}_\tau, \mathbf{z}_{\tau+1}, s_\tau, s_{\tau+1} | \beta_{0:\tau}) \\
 &= \sum_{s_\tau, s_{\tau+1}} Pr(\mathbf{z}_{\tau+1}, s_{\tau+1} | \theta_\tau, \mathbf{a}_\tau, s_\tau, \beta_{0:\tau}) Pr(\mathbf{a}_\tau | \theta_\tau, s_\tau, \beta_{0:\tau}) \\
 &\quad Pr(s_\tau | \theta_\tau, \beta_{0:\tau}) Pr(\theta_\tau | \beta_{0:\tau}) \\
 &= \sum_{s_\tau, s_{\tau+1}} \underbrace{Pr(\mathbf{z}_{\tau+1}, s_{\tau+1} | \mathbf{a}_\tau, s_\tau)}_{=P_{\mathbf{a}_\tau}^{\mathbf{z}_{\tau+1}}(s_{\tau+1}|s_\tau)} \underbrace{Pr(\mathbf{a}_\tau | \theta_\tau, \beta_\tau)}_{=\beta(\theta_\tau, \mathbf{a}_\tau)} \underbrace{Pr(s_\tau | \theta_\tau, \beta_{0:\tau})}_{=b(s_\tau | \theta_\tau)} \underbrace{Pr(\theta_\tau | \beta_{0:\tau-1})}_{=\sigma_{\beta_{0:\tau-1}}(\theta_\tau)},
 \end{aligned}$$

(where $b(s | \theta_\tau)$ is the belief over states obtained by a usual HMM filtering process)

$$= \sum_{s_\tau, s_{\tau+1}} P_{\mathbf{a}_\tau}^{\mathbf{z}_{\tau+1}}(s_{\tau+1}|s_\tau) \beta(\theta_\tau, \mathbf{a}_\tau) b(s_\tau | \theta_\tau) \sigma_{\beta_{0:\tau-1}}(\theta_\tau).$$

$\sigma_{\beta_{0:\tau}}$ can thus be computed from $\sigma_{\beta_{0:\tau-1}}$ and β_τ without explicitly using $\beta_{0:\tau-1}$ or earlier occupancy states.

Then, let us compute the expected reward at τ given $\beta_{0:\tau}$:

$$\begin{aligned}
& E[r(S_\tau, A_\tau^1, A_\tau^2) \mid \beta_{0:\tau}] \\
&= \sum_{s_\tau, \mathbf{a}_\tau} r(s_\tau, \mathbf{a}_\tau) Pr(s_\tau, \mathbf{a}_\tau \mid \beta_{0:\tau}) \\
&= \sum_{s_\tau, \mathbf{a}_\tau} \sum_{\theta_\tau} r(s_\tau, \mathbf{a}_\tau) Pr(s_\tau, \mathbf{a}_\tau, \theta_\tau \mid \beta_{0:\tau}) \\
&= \sum_{s_\tau, \mathbf{a}_\tau} \sum_{\theta_\tau} r(s_\tau, \mathbf{a}_\tau) Pr(s_\tau, \mathbf{a}_\tau \mid \theta_\tau, \beta_{0:\tau}) Pr(\theta_\tau \mid \beta_{0:\tau}) \\
&= \sum_{s_\tau, \mathbf{a}_\tau} \sum_{\theta_\tau} r(s_\tau, \mathbf{a}_\tau) \underbrace{Pr(\mathbf{a}_\tau \mid \theta_\tau, \beta_{0:\tau})}_{\beta_\tau(\theta_\tau, \mathbf{a}_\tau)} \underbrace{Pr(s_\tau \mid \theta_\tau, \beta_{0:\tau})}_{b(s_\tau \mid \theta_\tau)} \underbrace{Pr(\theta_\tau \mid \beta_{0:\tau})}_{\sigma_{\beta_{0:\tau-1}}(\theta_\tau)} \\
&= \sum_{s_\tau, \mathbf{a}_\tau} \sum_{\theta_\tau} r(s_\tau, \mathbf{a}_\tau) \beta_\tau(\theta_\tau, \mathbf{a}_\tau) b(s_\tau \mid \theta_\tau) \sigma_{\beta_{0:\tau-1}}(\theta_\tau).
\end{aligned}$$

The expected reward at τ can thus be computed from $\sigma_{\beta_{0:\tau-1}}$ and β_τ without explicitly using $\beta_{0:\tau-1}$ or earlier occupancy states. \square

Appendix C Occupancy Markov Games: Definition and Preliminary Properties

C.1 Properties of V^*

Before proving the postulates implicitly used by [47], we need to show that we can reason with mixed strategies in subgames as is usually done on full games.

C.1.1 V^* is not linear in β (behavioral strategies)

Let us consider the following (finite-horizon, deterministic) Non-Observable MDP:

$$\begin{aligned}
\mathcal{S} &\stackrel{\text{def}}{=} \{-2, -1, 0, +1, +2\}, & b_0(0) &= 1, & & \text{(always start in } s = 0\text{)} \\
\mathcal{A} &\stackrel{\text{def}}{=} \{-1, +1\}, & & & & \text{(moves = add or subtract 1)} \\
T(s, a) &\stackrel{\text{def}}{=} \min\{+2, \max\{-2, s + a\}\}, & & & & \text{(} +1 \text{ or } -1 \text{ move in } \mathcal{S}\text{)} \\
\mathcal{Z} &\stackrel{\text{def}}{=} \{\text{none}\}, & O(\text{none}) &\stackrel{\text{def}}{=} 1, & & \text{(no observation)} \\
r(s) &\stackrel{\text{def}}{=} \begin{cases} +1 & \text{if } s \in \{-2, +2\} \\ 0 & \text{otherwise,} \end{cases} & & & & \text{(} |s| = 2 \text{ : success!)} \\
\gamma &\stackrel{\text{def}}{=} 1, & H &\stackrel{\text{def}}{=} 2.
\end{aligned}$$

Let us then consider two particular behavioral strategies:

$$\begin{aligned}
\forall \theta, \beta^+(A = +1 \mid \theta) &= 1 & & \text{(always } +1\text{), and} \\
\forall \theta, \beta^-(A = -1 \mid \theta) &= 1 & & \text{(always } -1\text{).}
\end{aligned}$$

These two strategies are optimal, with an expected return of $+1$, because, at $t = H = 2$, β^+ reaches $+2$ w.p. 1, and β^- reaches -2 w.p. 1:

$$V(\beta^+) = V(\beta^-) = +1.$$

Let us now consider their linear combination $\beta^\pm \stackrel{\text{def}}{=} \frac{1}{2}\beta^+ + \frac{1}{2}\beta^-$:

$$\forall \theta, \beta^\pm(A = -1|\theta) = 0.5,$$

$$\forall \theta, \beta^\pm(A = +1|\theta) = 0.5.$$

Here, the probability to reach $s = -2$ or $s = +2$ at the last time step is much lower, and gives the value of that strategy:

$$\begin{aligned} V(\beta^\pm) &= Pr(s_2 = +2|\beta^\pm) + Pr(s_2 = -2|\beta^\pm) \\ &= Pr(a_0 = +1|\beta^\pm) \cdot Pr(a_1 = +1|\beta^\pm) \\ &\quad + Pr(a_0 = -1|\beta^\pm) \cdot Pr(a_1 = -1|\beta^\pm) \\ &= \underbrace{0.5 \cdot 0.5}_{0.25} + \underbrace{0.5 \cdot 0.5}_{0.25} = 0.5. \end{aligned}$$

This confirms that, in a POMDP, the value is not linear in the space of behavioral strategies. As a consequence, in a zs-POSG, the value is not (bi)linear in the spaces of behavioral strategies of both players.

C.1.2 Back to Mixed Strategies

We now generalize mixed strategies as a mathematical tool to handle subgames of a zs-OMG as normal-form games, and give some preliminary results.

First, for a given σ_τ and $\tau \leq \tau'$, let $\mu_{0:\tau'-1|\sigma_\tau}$ denote a mixed strategy profile that is defined over $0 : \tau' - 1$, and induces σ_τ at time τ . We will also write that it is *compatible* with σ_τ . Then, to complete a given mixed *prefix* strategy $\mu_{0:\tau'-1|\sigma_\tau}$ (here $\tau = \tau'$), the solver should provide each player with a different *suffix* strategy to execute for each θ_τ^i it could be facing. We now detail how to build an equivalent set of mixed *full* strategies for i . Each of the pure *prefix* strategies $\pi_{0:\tau-1}^i$ used in $\mu_{0:\tau-1|\sigma_\tau}^i$ (belonging to a set denoted $\Pi_{0:\tau-1|\sigma_\tau}^i$) can be extended by appending a different pure *suffix* strategy $\pi_{\tau:H-1}^i$ at each of its leaf nodes, which leads to a large set of pure strategies $\Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)$. Then, let $M_{0:H-1|\sigma_\tau}^i$ be the set of mixed *full* strategies $\mu_{0:H-1|\sigma_\tau}^i$ obtained by considering the distributions over $\bigcup_{\pi_{0:\tau-1}^i \in \Pi_{0:\tau-1|\sigma_\tau}^i} \Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)$ that verify, $\forall \pi_{0:\tau-1}^i$,

$$\sum_{\substack{\pi_{0:H-1}^i \in \\ \Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)}} \mu_{0:H-1|\sigma_\tau}^i(\pi_{0:H-1}^i) = \mu_{0:\tau-1|\sigma_\tau}^i(\pi_{0:\tau-1}^i). \quad (\text{C1})$$

This is the set of mixed strategies compatible with σ_τ .

Lemma 2. $M_{0:H-1|\sigma_\tau}^i$ is convex and equivalent to the set of behavioral strategies $\beta_{0:H-1|\sigma_\tau}^i$, thus sufficient to search for a Nash equilibrium in σ_τ .

Proof Let $\mu_{0:H-1|\sigma_\tau}^i$ and $\nu_{0:H-1|\sigma_\tau}^i$ be two mixed strategies in $M_{0:H-1|\sigma_\tau}^i$, i.e., which are both full and compatible with occupancy state σ_τ at time step τ , and $\alpha \in [0, 1]$. Then, for any $\pi_{0:\tau-1}^i$,

$$\begin{aligned} & \sum_{\pi_{0:H-1}^i \in \Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)} \left[\alpha \cdot \mu_{0:H-1|\sigma_\tau}^i(\pi_{0:H-1}^i) + (1 - \alpha) \cdot \nu_{0:H-1|\sigma_\tau}^i(\pi_{0:H-1}^i) \right] \\ &= \alpha \left[\sum_{\pi_{0:H-1}^i \in \Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)} \mu_{0:H-1|\sigma_\tau}^i(\pi_{0:H-1}^i) \right] \\ & \quad + (1 - \alpha) \left[\sum_{\pi_{0:H-1}^i \in \Pi_{0:H-1}^i(\pi_{0:\tau-1}^i)} \nu_{0:H-1|\sigma_\tau}^i(\pi_{0:H-1}^i) \right] \end{aligned}$$

(because both mixed strategies are compatible with σ_τ (eq. C1, p. 35):)

$$\begin{aligned} &= \alpha \cdot \mu_{0:\tau-1|\sigma_\tau}^i(\pi_{0:\tau-1}^i) + (1 - \alpha) \cdot \mu_{0:\tau-1|\sigma_\tau}^i(\pi_{0:\tau-1}^i) \\ &= \mu_{0:\tau-1|\sigma_\tau}^i(\pi_{0:\tau-1}^i). \end{aligned}$$

Eq. C1 thus also applies to $\alpha \cdot \mu_{0:H-1|\sigma_\tau}^i + (1 - \alpha) \cdot \nu_{0:H-1|\sigma_\tau}^i$, proving that it belongs to $M_{0:H-1|\sigma_\tau}^i$ and, as a consequence, that this set is convex.

The equivalence with the set of behavioral strategies simply relies on the fact that all mixed strategies over $\tau : H - 1$ can be independently generated at each action-observation history $\theta_{0:\tau-1}^i$. \square

While only future rewards are relevant when making a decision at τ , reasoning with mixed strategies defined from $t = 0$ will be convenient because $V_\tau(\sigma_\tau, \cdot, \cdot)$ is linear in $\mu_{0:H-1|\sigma_\tau}^i$, which allows coming back to a standard normal-form game and applying known results.

In the remaining, we simply note μ^i (without index) the mixed strategies in $M_{0:H-1|\sigma_\tau}^i$, set which we now note $M_{|\sigma_\tau}^i$. Also, since we shall work with local game $Q_\tau^*(\sigma_\tau, \beta_\tau)$, let us define: $M_{|\sigma_\tau, \beta_\tau^j}^i$ the set of i 's mixed strategies compatible with occupancy states reachable given σ_τ and β_τ^j (with either $j = i$ or $j = \neg i$). Then, $M_{|T(\sigma_\tau, \beta_\tau)}^i \subseteq M_{|\sigma_\tau, \beta_\tau^j}^i \subseteq M_{|\sigma_\tau}^i$ (inclusion due to the latter sets being less constrained in their definition). As a consequence, if maximizing some function f over i 's mixed strategies compatible with a given σ_τ :

$$\max_{\mu^i \in M_{|\sigma_\tau}^i} f(\sigma_\tau, \mu^i, \dots) \geq \max_{\mu^i \in M_{|\sigma_\tau, \beta_\tau^i}^i} f(\sigma_\tau, \mu^i, \dots) \geq \max_{\mu^i \in M_{|\sigma_\tau, \beta_\tau}^i} f(\sigma_\tau, \mu^i, \dots).$$

C.1.3 Von Neumann's Minimax Theorem for Subgames and a Bellman Optimality Equation

Using the previous results, one can show that [von Neumann's](#) minimax theorem applies in any subgame, allowing to swap max and min operators.

Theorem 3 (originally stated on page 10) *The subgame defined in Eq. (3) admits a unique NEV*

$$V_{\tau}^*(\sigma_{\tau}) \stackrel{\text{def}}{=} \max_{\beta_{\tau:H-1}^1} \min_{\beta_{\tau:H-1}^2} V_{\tau}(\sigma_{\tau}, \beta_{\tau:H-1}^1, \beta_{\tau:H-1}^2). \quad (4)$$

Proof For any occupancy state σ_{τ} ,

$$\max_{\beta_{\tau:}^1} \min_{\beta_{\tau:}^2} V(\sigma_{\tau}, \beta_{\tau:}) = \max_{\mu_{\tau:}^2} \min_{\mu_{\tau:}^1} V(\sigma_{\tau}, \mu_{\tau:}) \quad (\text{Kuhn's theorem (generalized)}) \quad (C2)$$

$$= \min_{\mu_{\tau:}^2} \max_{\mu_{\tau:}^1} V(\sigma_{\tau}, \mu_{\tau:}) \quad (\text{von Neumann's theorem}) \quad (C3)$$

$$= \min_{\beta_{\tau:}^2} \max_{\beta_{\tau:}^1} V(\sigma_{\tau}, \beta_{\tau:}) \quad (\text{again Kuhn's theorem (generalized)}) \quad (C4)$$

□

One can also show that a Bellman optimality equation allows relating optimal values in subgames at τ and $\tau + 1$, leading to a recursive expression of V_0^* .

Theorem 4 (originally stated on page 10) *$V_{\tau}^*(\sigma_{\tau})$ satisfies the following functional equation:*

$$V_{\tau}^*(\sigma_{\tau}) = \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} r(\sigma_{\tau}, \beta_{\tau}) + \gamma V_{\tau+1}^*(T(\sigma_{\tau}, \beta_{\tau})) = \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} Q_{\tau}^*(\sigma_{\tau}, \beta_{\tau}).$$

Proof Focusing, without loss of generality, on player 1, we have (complementary explanations follow for numbered lines in particular):

$$\max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} Q_{\tau}^*(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2) = \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} \left[r(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2) + \gamma V_{\tau+1}^*(T(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2)) \right]$$

($V_{\tau+1}^*(T(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2))$) being the Nash equilibrium value of normal-form game $V_{\tau+1}(T(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2), \mu^1, \mu^2)$):

$$= \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} \left[r(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2) + \gamma \max_{\mu^1 \in M_{|\sigma_{\tau}, \beta_{\tau}}^1} \min_{\mu^2 \in M_{|\sigma_{\tau}, \beta_{\tau}}^2} V_{\tau+1}(T(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2), \mu^1, \mu^2) \right]$$

$$= \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} \max_{\mu^1 \in M_{|\sigma_{\tau}, \beta_{\tau}}^1} \min_{\mu^2 \in M_{|\sigma_{\tau}, \beta_{\tau}}^2} \left[r(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2) + \gamma V_{\tau+1}(T(\sigma_{\tau}, \beta_{\tau}^1, \beta_{\tau}^2), \mu^1, \mu^2) \right]$$

(using the equivalence between maximin and minimax values for the (constrained normal-form) game at $\tau + 1$, the last two max and min operators can be swapped:)

$$= \max_{\beta_\tau^1} \min_{\beta_\tau^2} \min_{\mu^2 \in M_{|\sigma_\tau, \beta_\tau^1}^2} \max_{\mu^1 \in M_{|\sigma_\tau, \beta_\tau^1}^1} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2), \mu^1, \mu^2) \right]$$

(merging both mins (and with explanations thereafter):)

$$= \max_{\beta_\tau^1} \min_{\mu^2 \in M_{|\sigma_\tau, \beta_\tau^1}^2} \max_{\mu^1 \in M_{|\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)}^1} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)), \mu^1, \mu^2) \right] \quad (\text{C5})$$

(since ignoring the opponent's decision rule does not influence the expected return:)

$$= \max_{\beta_\tau^1} \min_{\mu^2 \in M_{|\sigma_\tau}^2} \max_{\mu^1 \in M_{|\sigma_\tau, \beta_\tau^1}^1} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)), \mu^1, \mu^2) \right]$$

(using again the minimax theorem's equivalence between maximin and minimax on an appropriate game:)

$$= \max_{\beta_\tau^1} \max_{\mu^1 \in M_{|\sigma_\tau, \beta_\tau^1}^1} \min_{\mu^2 \in M_{|\sigma_\tau}^2} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2(\mu^2)), \mu^1, \mu^2) \right] \quad (\text{C6})$$

(merging both maxs (and with explanations thereafter):)

$$= \max_{\mu^1 \in M_{|\sigma_\tau}^1} \min_{\mu^2 \in M_{|\sigma_\tau}^2} \left[r(\sigma_\tau, \beta_\tau^1(\mu^1), \beta_\tau^2(\mu^2)) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau^1(\mu^1), \beta_\tau^2(\mu^2)), \mu^1, \mu^2) \right] \quad (\text{C7})$$

(again with the equivalence property discussed before the lemma:)

$$\begin{aligned} &= \max_{\mu^1 \in M_{|\sigma_\tau}^1} \min_{\mu^2 \in M_{|\sigma_\tau}^2} V_\tau(\sigma_\tau, \mu^1, \mu^2) \\ &= \max_{\beta_{\tau:H-1}^1} \min_{\beta_{\tau:H-1}^2} V_\tau(\sigma_\tau, \beta_{\tau:H-1}^1, \beta_{\tau:H-1}^2) \\ &\stackrel{\text{def}}{=} V_\tau^*(\sigma_\tau). \end{aligned}$$

More precisely, line C5 (and, similarly, line C7) is obtained by observing that

- minimizing over both (i) β_τ^2 and (ii) μ^2 constrained by σ_τ and β_τ^1 is equivalent to minimizing over μ^2 constrained by σ_τ and β_τ^1 ; and
- in the reminder of the formula, decision rule β_τ^2 at time τ can be retrieved as a function of μ^2 (noted $\beta_\tau^2(\mu^2)$).

Also, line C6 results from the observation that, while $M_{|\sigma_\tau, \beta_\tau^1}^1$ and $M_{|\sigma_\tau}^2$ allow to actually make decisions over different time intervals, we are here minimizing over μ^2 while maximizing over μ^1 a function that is linear in both input spaces. This amounts to solving some 2-player zero-sum normal-form game, hence the applicability of von Neumann's minimax theorem.

The above derivation tells us that the maximin value (the best outcome player 1 can guarantee whatever player 2's strategy) in the one-time-step game is thus the Nash equilibrium value (NEV) for the complete subgame from τ onwards. \square

Appendix D Solving zs-OMGs

D.1 Preliminary Properties

D.1.1 Properties of T_m^1 and T_c^1

The first two lemmas below present properties of T_m^1 and T_c^1 that will be useful afterwards.

Lemma 3. $T_m^1(\sigma_\tau, \beta_\tau)$ is linear in σ_τ , β_τ^1 , and β_τ^2 .

Proof

$$\begin{aligned}
 & T_m^1(\sigma_\tau, \beta_\tau)(\theta_\tau^1, a^1, z^1) & (D8) \\
 &= \sum_{\theta_\tau^2, a^2, z^2} T(\sigma_\tau, \beta_\tau)((\theta_\tau^1, a^1, z^1), (\theta_\tau^2, a^2, z^2)) & \text{(from Eq. (1))} \\
 &= \sum_{s', \theta_\tau^2, a^2, z^2} \beta_\tau^1(\theta_\tau^1, a^1) \beta_\tau^2(\theta_\tau^2, a^2) \sum_s P_{a^1, a^2}^{z^1, z^2}(s'|s) b(s|\theta_\tau^1, \theta_\tau^2) \sigma_\tau(\theta_\tau^1, \theta_\tau^2) \\
 &= \beta_\tau^1(\theta_\tau^1, a^1) \sum_{\theta_\tau^2, a^2} \beta_\tau^2(\theta_\tau^2, a^2) \sum_{s, s', z^2} P_{a^1, a^2}^{z^1, z^2}(s'|s) b(s|\theta_\tau^1, \theta_\tau^2) \sigma_\tau(\theta_\tau^1, \theta_\tau^2). & (D9)
 \end{aligned}$$

□

Lemma 4. $T_c^1(\sigma_\tau, \beta_\tau)$ is independent of β_τ^1 and $\sigma_\tau^{m,1}$.

Proof See [45], Lemma 4.2.3. □

D.1.2 Linearity and Lipschitz-continuity of $T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2)$

Lemma 5. At depth τ , $T(\sigma_\tau, \beta_\tau)$ is linear in β_τ^1 , β_τ^2 , and σ_τ , where $\beta_\tau = \langle \beta_\tau^1, \beta_\tau^2 \rangle$. It is more precisely 1-Lipschitz-continuous (1-LC) in σ_τ (in 1-norm), i.e., for any $\sigma_\tau, \sigma'_\tau$:

$$\|T(\sigma'_\tau, \beta_\tau) - T(\sigma_\tau, \beta_\tau)\|_1 \leq 1 \cdot \|\sigma'_\tau - \sigma_\tau\|_1.$$

Proof Let σ be an occupancy state at time τ and β_τ be a decision rule. Then, as seen in the proof of Theorem 1, the next occupancy state $\sigma' = T(\sigma, \beta_\tau)$ satisfies, for any s' and $(\theta, \mathbf{a}, \mathbf{z})$:

$$\begin{aligned}
 \sigma'(\theta, \mathbf{a}, \mathbf{z}) &\stackrel{\text{def}}{=} Pr(\theta, \mathbf{a}, \mathbf{z} | \sigma, \beta_\tau^1, \beta_\tau^2) \\
 &= \beta_\tau^1(\theta^1, a^1) \beta_\tau^2(\theta^2, a^2) \left[\sum_{s', s \in \mathcal{S}} P_{\mathbf{a}}^{\mathbf{z}}(s'|s) b(s|\theta) \right] \sigma(\theta).
 \end{aligned}$$

$b(s|\theta)$ depending only on the model (transition function and initial belief), the next occupancy state σ' thus evolves linearly w.r.t. (i) *private* decision rules β_τ^1 and β_τ^2 , and (ii) the occupancy state σ .

The 1-Lipschitz-continuity holds because each component of vector σ_τ is distributed over multiple components of σ' . Indeed, let us view two occupancy states as vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and their corresponding next states under β_τ as $M\mathbf{x}$ and $M\mathbf{y}$, where $M \in \mathbb{R}^{m \times n}$ is the corresponding transition matrix (*i.e.*, which turns σ into $\sigma' \stackrel{\text{def}}{=} T(\sigma_\tau, \beta_\tau)$). Then,

$$\begin{aligned}
\|M\mathbf{x} - M\mathbf{y}\|_1 &\stackrel{\text{def}}{=} \sum_{j=1}^m \left| \sum_{i=1}^n M_{i,j}(x_i - y_i) \right| \\
&\leq \sum_{j=1}^m \sum_{i=1}^n |M_{i,j}(x_i - y_i)| && \text{(convexity of } |\cdot| \text{)} \\
&= \sum_{j=1}^m \sum_{i=1}^n M_{i,j} |x_i - y_i| && (\forall i, j, M_{i,j} \geq 0) \\
&= \sum_{i=1}^n \underbrace{\sum_{j=1}^m M_{i,j}}_{=1} |x_i - y_i| && (M \text{ is a transition matrix)} \\
&\stackrel{\text{def}}{=} \|\mathbf{x} - \mathbf{y}\|_1.
\end{aligned}$$

□

D.1.3 Lipschitz-Continuity of V^*

The next two results demonstrate that, in the finite horizon setting, V^* is Lipschitz-continuous (LC) in occupancy space, which allows defining LC upper- and lower-bound approximations.

Lemma 6. *At depth τ , $V_\tau(\sigma_\tau, \beta_\tau)$ is linear w.r.t. σ_τ and β_τ .*

Note: This result in fact applies to any reward function of a general-sum POSG with any number of agents (here N), *e.g.*, to a Dec-POMDP. The following proof handles the general case (with $\beta_\tau \stackrel{\text{def}}{=} \langle \beta_\tau^1, \dots, \beta_\tau^N \rangle$, and $\beta_\tau(\mathbf{a}|\theta) = \prod_{i=1}^N \beta_\tau^i(a^i, \theta^1)$).

Proof This property trivially holds for $\tau = H - 1$ because

$$\begin{aligned}
V_{H-1}(\sigma_{H-1}, \beta_{H-1}) &= r(\sigma_{H-1}, \beta_{H-1}) \\
&= \sum_{s, \mathbf{a}} \left(\sum_{\theta} Pr(s, \mathbf{a}|\theta) \sigma_{H-1}(\theta) \right) r(s, \mathbf{a}) \\
&= \sum_{s, \mathbf{a}} \left(\sum_{\theta} b(s|\theta) \beta_\tau(\mathbf{a}|\theta) \sigma_{H-1}(\theta) \right) r(s, \mathbf{a}) \\
&= \sum_{s, \theta} b(s|\theta) \sigma_{H-1}(\theta) \left(\sum_{\mathbf{a}} \beta_\tau(\mathbf{a}|\theta) r(s, \mathbf{a}) \right).
\end{aligned}$$

Now, let us assume that the property holds for $\tau + 1 \in \{1..H - 1\}$. Then,

$$V_\tau(\sigma_\tau, \beta_\tau) = \sum_{s, \mathbf{a}} \left(\sum_{\theta} b(s|\theta) \beta_\tau(\mathbf{a}|\theta) \sigma_\tau(\theta) \right) r(s, \mathbf{a}) + \gamma V_{\tau+1}(T(\sigma_\tau, \beta_\tau), \beta_{\tau+1})$$

$$= \sum_{s, \theta} b(s|\theta) \sigma_{\tau}(\theta) \left(\sum_{\mathbf{a}} \beta_{\tau}(\mathbf{a}|\theta) r(s, \mathbf{a}) \right) + \gamma V_{\tau+1}(T(\sigma_{\tau}, \beta_{\tau}), \beta_{\tau+1}).$$

As

- $T(\sigma_{\tau}, \beta_{\tau})$ is linear in σ_{τ} (Lemma 5) and
- $V_{\tau+1}(\sigma_{\tau+1}, \beta_{\tau+1})$ is linear in $\sigma_{\tau+1}$ (induction hypothesis),

their composition, $V_{\tau+1}(T(\sigma_{\tau}, \beta_{\tau}), \beta_{\tau+1})$, is also linear in σ_{τ} , and so is $V_{\tau}(\sigma_{\tau}, \beta_{\tau})$. Similarly, $V_{\tau}(\sigma_{\tau}, \beta_{\tau})$ is linear in β_{τ} for any σ_{τ} . \square

Theorem 5 (originally stated on page 10) *Let $h_{\tau} \stackrel{\text{def}}{=} \frac{1-\gamma^{H-\tau}}{1-\gamma}$ (or $h_{\tau} \stackrel{\text{def}}{=} H - \tau$ if $\gamma = 1$). Then $V_{\tau}^*(\sigma_{\tau})$ is λ_{τ} -Lipschitz continuous in σ_{τ} at any depth $\tau \in \{0 \dots H-1\}$, where $\lambda_{\tau} = \frac{1}{2} h_{\tau} (r_{\max} - r_{\min})$.*

Proof At depth τ , the value of any behavioral strategy β_{τ} : is bounded, independently of σ_{τ} , by

$$V_{\tau}^{\max} \stackrel{\text{def}}{=} h_{\tau} r_{\max}, \quad \text{where } r_{\max} \stackrel{\text{def}}{=} \max_{s, \mathbf{a}} r(s, \mathbf{a}), \text{ and}$$

$$V_{\tau}^{\min} \stackrel{\text{def}}{=} h_{\tau} r_{\min}, \quad \text{where } r_{\min} \stackrel{\text{def}}{=} \min_{s, \mathbf{a}} r(s, \mathbf{a}).$$

Thus, $V_{\beta_{\tau}}$ being a linear function defined over a probability simplex ($\mathcal{O}_{\tau}^{\sigma}$) (cf. Appendix D.1.3) and bounded by $[V_{\tau}^{\min}, V_{\tau}^{\max}]$, we can apply Horák's PhD thesis' Lemma 3.5 (p. 33) [20] to establish that it is also λ_{τ} -LC, i.e.,

$$|V_{\beta_{\tau}}(\sigma) - V_{\beta_{\tau}}(\sigma')| \leq \lambda_{\tau} \|\sigma - \sigma'\|_1 \quad (\forall \sigma, \sigma'),$$

$$\text{with } \lambda_{\tau} = \frac{V_{\tau}^{\max} - V_{\tau}^{\min}}{2}.$$

Considering now optimal solutions, this means that, at depth τ and for any $(\sigma, \sigma') \in \mathcal{O}_{\tau}^{\sigma}$:

$$\begin{aligned} & V_{\tau}^*(\sigma) - V_{\tau}^*(\sigma') \\ &= \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} V_{\tau}(\sigma, \beta_{\tau}^1, \beta_{\tau}^2) - \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} V_{\tau}(\sigma', \beta_{\tau}^1, \beta_{\tau}^2) \\ &\leq \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} \left[V_{\tau}(\sigma', \beta_{\tau}^1, \beta_{\tau}^2) + \lambda_{\tau} \|\sigma - \sigma'\|_1 \right] - \max_{\beta_{\tau}^1} \min_{\beta_{\tau}^2} V_{\tau}(\sigma', \beta_{\tau}^1, \beta_{\tau}^2) \\ &= \lambda_{\tau} \|\sigma - \sigma'\|_1. \end{aligned}$$

Symmetrically, $V_{\tau}^*(\sigma) - V_{\tau}^*(\sigma') \geq -\lambda_{\tau} \|\sigma - \sigma'\|_1$, hence the expected result:

$$|V_{\tau}^*(\sigma) - V_{\tau}^*(\sigma')| \leq \lambda_{\tau} \|\sigma - \sigma'\|_1.$$

\square

As it will be used later, let us also present the following lemma.

Lemma 7. *Let us consider $\tau \in \{0 \dots H-1\}$, θ_{τ}^1 , and ψ_{τ}^2 . Then $\nu_{[\sigma_{\tau}^{c,1}, \psi_{\tau}^2]}(\theta_{\tau}^1)$ is λ_{τ} -LC in $\sigma_{\tau}^{c,1}(\cdot|\theta_{\tau}^1)$.*

Equivalently, we will also write that $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2$ is λ_τ -LC in $\sigma_\tau^{c,1}$ in vector-wise 1-norm, i.e.:

$$\overrightarrow{|\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2 - \nu_{[\tilde{\sigma}_\tau^{c,1}, \psi_\tau^2]}^2|}_1 \leq \lambda_\tau \overrightarrow{\|\sigma_\tau^{c,1} - \tilde{\sigma}_\tau^{c,1}\|}_1,$$

where (i) the absolute value of a vector is obtained by taking the absolute value of each component; and (ii) the vector-wise 1-norm of a matrix is a vector made of the 1-norm of each of its component vectors.

Proof For any θ_τ^1 , $\sigma_\tau^{c,1}$ and ψ_τ^2 induce a POMDP for Player 1 from τ on, where (i) the state at any $t \in \{\tau \dots H - 1\}$ corresponds to a pair $\langle s, \theta_t^2 \rangle$, and (ii) the initial belief is derived from $\sigma_\tau^{c,1}(\cdot|\theta_\tau^1)$. The belief state at t thus gives:

$$b_{\theta_t^1}(s, \theta_t^2) \stackrel{\text{def}}{=} Pr(s, \theta_t^2 | \theta_t^1) = \underbrace{Pr(s | \theta_t^2, \theta_t^1)}_{b_{\theta_t^2, \theta_t^1}^{\text{HMM}}(s)} \cdot \underbrace{Pr(\theta_t^2 | \theta_t^1)}_{\sigma_\tau^{c,1}(\theta_t^2 | \theta_t^1)}.$$

So,

- the value function of any behavioral strategy β_τ^1 : is linear at t in $b_{\theta_t^1}$, thus (in particular) in $\sigma_\tau^{c,1}(\cdot|\theta_t^1)$; and
- the optimal value function is LC at t also in $b_{\theta_t^1}$ (with the same depth-dependent upper-bounding Lipschitz constant λ_t as in the proof of Theorem 5),⁶ thus (in particular) in $\sigma_\tau^{c,1}(\cdot|\theta_t^1)$.

Using $t = \tau$, the optimal value function is $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2(\theta_\tau^1)$, which is thus λ_τ -LC in $\sigma_\tau^{c,1}(\cdot|\theta_\tau^1)$. \square

D.2 Bounding Approximations of V^* , $W^{1,*}$ and $W^{2,*}$

D.2.1 \bar{V}_τ and \underline{V}_τ

To find a form that could be appropriate for an upper bound approximation of V_τ^* , let us consider an OS σ_τ and a single tuple $\langle \tilde{\sigma}_\tau, \nu_{[\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2]}^2 \rangle$, and define $\zeta_\tau \stackrel{\text{def}}{=} \sigma_\tau^{m,1} \tilde{\sigma}_\tau^{c,1}$. Then,

$$\begin{aligned} V^*(\sigma_\tau) &\leq V^*(\zeta_\tau) + \lambda_\tau \|\sigma_\tau - \zeta_\tau\|_1 && \text{(LC, cf. Theorem 5)} \\ &= V^*(\sigma_\tau^{m,1} \tilde{\sigma}_\tau^{c,1}) + \lambda_\tau \|\sigma_\tau - \zeta_\tau\|_1 \\ &\leq \sigma_\tau^{m,1} \cdot \nu_{[\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2]}^2 + \lambda_\tau \|\sigma_\tau - \sigma_\tau^{m,1} \tilde{\sigma}_\tau^{c,1}\|_1. && \text{(Cvx, cf. Theorem 2)} \end{aligned}$$

Notes:

- $\tilde{\sigma}_\tau^{m,1}$ does not appear in the resulting upper bound, thus will not need to be specified.

⁶The proof process is similar. The only difference lies in the space at hand, but without any impact on the resulting formulas.

- For $\tau = H - 1$, $\nu_{[\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2]}$ is a simple function of r , $\tilde{\sigma}_\tau^{c,1}$, β_τ^2 , and the dynamics of the system, as described in Eq. (9) of Wiggers et al. [47].

From this, we can deduce the following appropriate forms of upper and (symmetrically) lower bound function approximations for V_τ^* :

$$\bar{V}_\tau(\sigma_\tau) = \min_{\langle \tilde{\sigma}_\tau^{c,1}, \langle \bar{\nu}_\tau^2, \beta_\tau^2 \rangle \rangle \in \bar{\mathcal{J}}_\tau} [\sigma_\tau^{m,1} \cdot \bar{\nu}_\tau^2 + \lambda_\tau \|\sigma_\tau - \sigma_\tau^{m,1} \tilde{\sigma}_\tau^{c,1}\|_1], \text{ and} \quad (\text{D10})$$

$$\underline{V}_\tau(\sigma_\tau) = \max_{\langle \tilde{\sigma}_\tau^{c,2}, \langle \underline{\nu}_\tau^1, \beta_\tau^1 \rangle \rangle \in \underline{\mathcal{J}}_\tau} [\sigma_\tau^{m,2} \cdot \underline{\nu}_\tau^1 - \lambda_\tau \|\sigma_\tau - \sigma_\tau^{m,2} \tilde{\sigma}_\tau^{c,2}\|_1], \quad (\text{D11})$$

which are respectively concave in $\sigma_\tau^{m,1}$ and convex in $\sigma_\tau^{m,2}$, and which both exploit the Lipschitz continuity.

D.2.2 \bar{W}_τ and \underline{W}_τ

Note: We discuss all depths from 0 to $H - 1$, even though we do not need these approximations at $\tau = H - 1$.

Let us first see how concavity-convexity properties affect $W_\tau^{*,1}$.

Lemma 8. *Considering that vectors $\nu_{[\sigma_H^{c,1}, \beta_H^2]}$ are null vectors, we have, for all $\tau \in \{0 \dots H - 1\}$:*

$$\begin{aligned} W_\tau^{1,*}(\sigma_\tau, \beta_\tau^1) &= \min_{\beta_\tau^2, \langle \beta_{\tau+1}^2, \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]} \rangle} \beta_\tau^1 \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) \right. \\ &\quad \left. + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]}^2 \right]. \end{aligned}$$

Proof Considering that vectors $\nu_{[\sigma_H^{c,1}, \beta_H^2]}$ are null vectors, we have, for all $\tau \in \{0 \dots H - 1\}$:

$$W_\tau^{1,*}(\sigma_\tau, \beta_\tau^1) = \min_{\beta_\tau^2} Q_\tau^*(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) = \min_{\beta_\tau^2} [r(\sigma_\tau, \beta_\tau) + \gamma V_{\tau+1}^*(T(\sigma_\tau, \beta_\tau))]$$

(Line below exploits Theorem 2 (p. 9) and T_c^1 's independence from β_τ^1 (Lemma 4).)

$$\begin{aligned} &= \min_{\beta_\tau^2} \left[r(\sigma_\tau, \beta_\tau) + \gamma \min_{\langle \beta_{\tau+1}^2, \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]} \rangle} \left[T_m^1(\sigma_\tau, \beta_\tau) \cdot \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]}^2 \right] \right] \\ &= \min_{\beta_\tau^2, \langle \beta_{\tau+1}^2, \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]} \rangle} \left[r(\sigma_\tau, \beta_\tau) + \gamma T_m^1(\sigma_\tau, \beta_\tau) \cdot \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]}^2 \right] \end{aligned}$$

(Line below exploits r and T_m^1 's linearity in β_τ^1 (Lemma 3).)

$$\begin{aligned} &= \min_{\beta_\tau^2, \langle \beta_{\tau+1}^2, \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]} \rangle} \beta_\tau^1 \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) \right. \\ &\quad \left. + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \nu_{[T_c^1(\sigma_\tau, \beta_\tau^2), \beta_{\tau+1}^2]}^2 \right]. \end{aligned}$$

□

Note that, since $V_H^* = 0$, $\tau = H - 1$ is a particular case which can be simply re-written:

$$W_\tau^{1,*}(\sigma_\tau, \beta_\tau^1) = \min_{\beta_\tau^2} \beta_\tau^{1\top} \cdot r(\sigma_\tau, \cdot, \beta_\tau^2).$$

To find a form that could be appropriate for an upper bound approximation of $W_\tau^{*,1}$, let us now consider an OS σ_τ and a single tuple $\langle \tilde{\sigma}_\tau, \tilde{\beta}_\tau^2, \nu_{[T_c^1(\tilde{\sigma}_\tau, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2 \rangle$. Then,

$$\begin{aligned} W_\tau^{1,*}(\sigma_\tau, \beta_\tau^1) &= \min_{\beta_\tau^2} [r(\sigma_\tau, \beta_\tau) + \gamma V_{\tau+1}^*(T(\sigma_\tau, \beta_\tau))] \\ &\leq r(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) + \gamma V_{\tau+1}^{BR,1}(T(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) | \tilde{\beta}_{\tau+1}^2) \end{aligned}$$

(Use $\tilde{\beta}_\tau^2$ & $\tilde{\beta}_{\tau+1}^2$: instead of mins)

(where $V_{\tau+1}^{BR,1}(T(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) | \tilde{\beta}_{\tau+1}^2)$ is the value of 1's best response to $\tilde{\beta}_{\tau+1}^2$: if in $T(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2)$)

$$= r(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) + \gamma T_m^1(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) \cdot \underbrace{\nu_{[T_c^1(\sigma_\tau^{c,1}, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2}$$

(Lem. 3 of Wiggers et al. [47])

$$\begin{aligned} &\leq r(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) + \gamma T_m^1(\sigma_\tau, \beta_\tau^1, \tilde{\beta}_\tau^2) \cdot \left(\nu_{[T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2 \right. \\ &\quad \left. + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \tilde{\beta}_\tau^2) - T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2)\|_1} \right) \end{aligned} \quad (\text{D12})$$

(Lem. 7: $\lambda_{\tau+1}$ -LC of $\nu_{[T_c^1(\sigma_\tau^{c,1}, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2$)

$$\begin{aligned} &= \beta_\tau^{1\top} \cdot \left[r(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) + \gamma T_m^1(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) \cdot \left(\nu_{[T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2 \right. \right. \\ &\quad \left. \left. + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \tilde{\beta}_\tau^2) - T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2)\|_1} \right) \right] \end{aligned} \quad (\text{D13})$$

(Linearity in β_τ^1)

$$\begin{aligned} &= \beta_\tau^{1\top} \cdot \left[r(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) + \gamma T_m^1(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) \cdot \nu_{[T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2), \tilde{\beta}_{\tau+1}^2]}^2 \right. \\ &\quad \left. + \gamma \lambda_{\tau+1} \cdot \|T(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) - T_m^1(\sigma_\tau, \cdot, \tilde{\beta}_\tau^2) T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2)\|_1 \right] \end{aligned}$$

(Alternative writing)

From this, we can deduce the following appropriate forms of (i) upper bounding approximation for $W_\tau^{1,*}$ and (ii) (symmetrically) of lower bound approximation for $W_\tau^{2,*}$:

$$\begin{aligned}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) &= \min_{\langle \tilde{\sigma}_\tau^{c,1}, \beta_\tau^2, \bar{\nu}_{\tau+1}^2 \rangle \in \overline{\mathcal{I}}_\tau} \beta_\tau^{1\top} \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \bar{\nu}_{\tau+1}^2 \right. \\ &\quad \left. + \gamma \lambda_{\tau+1} \cdot \|T(\sigma_\tau, \cdot, \beta_\tau^2) - T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) T_c^1(\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2)\|_1 \right], \text{ and} \\ \underline{W}_\tau(\sigma_\tau, \beta_\tau^2) &= \max_{\langle \tilde{\sigma}_\tau^{c,2}, \beta_\tau^1, \underline{\nu}_{\tau+1}^1 \rangle \in \underline{\mathcal{I}}_\tau} \beta_\tau^{2\top} \cdot \left[r(\sigma_\tau, \beta_\tau^1, \cdot) + \gamma T_m^2(\sigma_\tau, \beta_\tau^1, \cdot) \cdot \underline{\nu}_{\tau+1}^1 \right. \\ &\quad \left. - \gamma \lambda_{\tau+1} \cdot \|T(\sigma_\tau, \beta_\tau^1, \cdot) - T_m^2(\sigma_\tau, \beta_\tau^1, \cdot) T_c^2(\tilde{\sigma}_\tau^{c,2}, \beta_\tau^1)\|_1 \right],\end{aligned}$$

where $\bar{\nu}_{\tau+1}^2$ and $\underline{\nu}_{\tau+1}^1$ respectively upper and lower bound the actual vectors associated to the players' future strategies (resp. of 2 and 1).

Again, $\tau = H - 1$ is a particular case where only the reward term is preserved.

This constitutes the proof to the following proposition

Proposition 13 (originally stated on page 11) *Let $\overline{\mathcal{I}}_\tau$ be a set of tuples $w = \langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{\nu}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle$. Then,*

$$\begin{aligned}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) &\stackrel{\text{def}}{=} \min_{\langle \tilde{\sigma}_\tau, \beta_\tau^2, \langle \bar{\nu}_{\tau+1}^2, \beta_{\tau+1}^2 \rangle \rangle \in \overline{\mathcal{I}}_\tau} \left[r(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) \cdot \bar{\nu}_{\tau+1}^2 \right. \\ &\quad \left. + \lambda_{\tau+1} \|T(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) - T_m^1(\sigma_\tau, \beta_\tau^1, \beta_\tau^2) T_c^1(\tilde{\sigma}_\tau^{c,1}, \beta_\tau^2)\|_1 \right] \quad (6)\end{aligned}$$

upper-bounds $W_\tau^{1,*}$ over the whole space $\mathcal{O}_\tau^\sigma \times \mathcal{B}_\tau^1$.

D.3 Related Operators

D.3.1 Selection Operator: Solving for β_τ^1 as an LP

Proposition 14 *Using now a distribution ψ_τ^2 over tuples $w = \langle \tilde{\sigma}_\tau^{c,1}, \beta_\tau^2, \bar{\nu}_{\tau+1}^2 \rangle \in \overline{\mathcal{I}}_\tau$, the corresponding upper-bounding value for “profile” $\langle \beta_\tau^1, \psi_\tau^2 \rangle$ when in σ_τ can be written as an expectancy:*

$$\beta_\tau^{1\top} \cdot M^{\sigma_\tau} \cdot \psi_\tau^2,$$

where M^{σ_τ} is an $|\Theta_\tau^1| \times |\mathcal{A}^1| \times |\overline{\mathcal{I}}_\tau^1|$ matrix.

Proof From the right-hand side term in (Equation (D13)), the upper-bounding value associated to σ_τ , β_τ^1 and a tuple $\langle \tilde{\sigma}_\tau^{c,1}, \beta_\tau^2, \bar{v}_{\tau+1}^2 \rangle \in \bar{\mathcal{I}}_\tau^1$ can be written:

$$\beta_\tau^{1\top} \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \left(\bar{v}_{\tau+1}^2 + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2) - T_c^1(\tilde{\sigma}_\tau^{c,1}, \tilde{\beta}_\tau^2)\|_1} \right) \right].$$

Using now a distribution ψ_τ^2 over tuples $w = \langle \tilde{\sigma}_\tau^{c,1}, \beta_\tau^2, \bar{v}_{\tau+1}^2 \rangle \in \bar{\mathcal{I}}_\tau^1$, the corresponding upper-bounding value for “profile” $\langle \beta_\tau^1, \psi_\tau^2 \rangle$ when in σ_τ can be written as an expectancy:

$$\begin{aligned} & \sum_{w \in \bar{W}_\tau} \beta_\tau^{1\top} \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2[w]) + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2[w]) \cdot \left(\bar{v}_{\tau+1}^2[w] \right. \right. \\ & \quad \left. \left. + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2[w]) - T_c^1(\tilde{\sigma}_\tau^{c,1}[w], \beta_\tau^2[w])\|_1} \right) \right] \cdot \psi_\tau^2(w) \end{aligned}$$

(where $x[w]$ denotes the field x of tuple w)

$$= \beta_\tau^{1\top} \cdot M^{\sigma_\tau} \cdot \psi_\tau^2,$$

where M^{σ_τ} is an $|\Theta_\tau^1| \times \mathcal{A}^1 \times |\bar{\mathcal{I}}_\tau^1|$ matrix. \square

For implementation purposes, using Eqs. (2) and (D9) (to develop respectively $r(\cdot, \cdot, \cdot)$ and $T_m^1(\cdot, \cdot, \cdot)$), we can derive the expression of a component, *i.e.*, the upper-bounding value if a^1 is applied in θ_τ^1 while w is chosen:

$$\begin{aligned} M_{(\langle \theta_\tau^1, a^1 \rangle, w)}^{\sigma_\tau} & \stackrel{\text{def}}{=} r(\sigma_\tau, \cdot, \beta_\tau^2[w]) + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2[w]) \cdot \\ & \quad \left(\bar{v}_{\tau+1}^2[w] + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2[w]) - T_c^1(\tilde{\sigma}_\tau^{c,1}[w], \beta_\tau^2[w])\|_1} \right) \\ & = \sum_{s, \theta_\tau^2, a^2} \sigma_\tau(\theta_\tau) b(s|\theta_\tau) \beta_\tau^2[w](a^2|\theta_\tau^2) r(s, \mathbf{a}) \\ & \quad + \gamma \sum_{z^1} \left[\sum_{\theta_\tau^2, a^2} \beta_\tau^2[w](a^2|\theta_\tau^2) \sum_{s, s', z^2} P_{\mathbf{a}}^z(s'|s) b(s|\theta_\tau) \sigma_\tau(\theta_\tau) \right] \cdot \\ & \quad \left(\bar{v}_{\tau+1}^2[w](\theta_\tau^1, a^1, z^1) \right. \\ & \quad \left. + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2[w]) - T_c^1(\tilde{\sigma}_\tau^{c,1}[w], \beta_\tau^2[w])\|_1}(\theta_\tau^1, a^1, z^1) \right) \\ & = \sum_{\theta_\tau^2} \sigma_\tau(\theta_\tau) \sum_{a^2} \beta_\tau^2[w](a^2|\theta_\tau^2) \\ & \quad \cdot \left(\sum_s b(s|\theta_\tau) r(s, \mathbf{a}) \right. \\ & \quad \left. + \gamma \sum_{z^1} \left[\sum_{s, s', z^2} P_{\mathbf{a}}^z(s'|s) b(s|\theta_\tau) \right] \cdot \left(\bar{v}_{\tau+1}^2[w](\theta_\tau^1, a^1, z^1) \right. \right. \\ & \quad \left. \left. + \lambda_{\tau+1} \cdot \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2[w]) - T_c^1(\tilde{\sigma}_\tau^{c,1}[w], \beta_\tau^2[w])\|_1}(\theta_\tau^1, a^1, z^1) \right) \right). \end{aligned}$$

Then, solving $\max_{\beta_\tau^1} \bar{W}_\tau(\sigma_\tau, \beta_\tau^1)$ can be rewritten as solving a zero-sum game where pure strategies are:

- for Player 1, the choice of not 1, but $|\Theta_\tau^1|$ actions (among $|\mathcal{A}^1|$) and,
- for Player 2, the choice of 1 element of $\bar{\mathcal{I}}_\tau^1$.

One can view it as a Bayesian game with one type per history θ_τ^1 for 1, and a single type for 2.

With our upper bound, $\max_{\beta_\tau^1} \bar{W}_\tau(\sigma_\tau, \beta_\tau^1)$ can thus be solved as the LP:

$$\begin{aligned} \max_{\beta_\tau^1, v} v \quad \text{s.t.} \quad & \text{(i)} \quad \forall w \in \bar{\mathcal{I}}_\tau^1, \quad v \leq \beta_\tau^1{}^\top \cdot M_{(\cdot, w)}^{\sigma_\tau} \\ & \text{(ii)} \quad \forall \theta_\tau^1 \in \Theta_\tau^1, \quad \sum_{a^1} \beta_\tau^1(a^1 | \theta_\tau^1) = 1, \end{aligned}$$

whose dual LP is given by

$$\begin{aligned} \min_{\psi_\tau^2, v} v \quad \text{s.t.} \quad & \text{(i)} \quad \forall (\theta_\tau^1, a^1) \in \Theta_\tau^1 \times \mathcal{A}^1, \quad v \geq M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2 \\ & \text{(ii)} \quad \sum_{w \in \bar{\mathcal{I}}_\tau^1} \psi_\tau^2(w) = 1. \end{aligned}$$

As can be noted, M^{σ_τ} 's columns corresponding to 0-probability histories θ_τ^1 in $\sigma_\tau^{m,1}$ are empty (full of zeros), so that the corresponding decision rules (for these histories) are not relevant and can be set arbitrarily. The actual implementation thus ignores these histories, whose corresponding decision rules also do not need to be stored.

Remark 2 (Interpretation of M^{σ_τ}) The content of this matrix can be interpreted by noting that, a given w containing a behavioral strategy $\beta_{\tau:H-1}^2$ and an OS $\tilde{\sigma}_\tau$, a pair $\langle w, \theta_\tau^1 \rangle$ induces a POMDP for player 1 whose state space is made of pairs $\langle s, \theta_t^2 \rangle$, and whose initial belief b_τ depends on $\tilde{\sigma}_\tau$ and θ_τ^1 . Solving this POMDP amounts to finding a best response of player 1 to $\beta_{\tau:H-1}^2$. In this setting, an element $M_{((\theta_\tau^1, a^1), w)}^{\sigma_\tau}$ is an upper-bound of the optimal (POMDP) Q -value when player 1 performs a^1 while facing b_τ ($Q_{\text{POMDP}}^*(b_\tau, a^1)$).

D.3.2 Upper Bounding $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2$

Adding a new complete tuple to $\bar{\mathcal{I}}_\tau^1$ requires a new vector \bar{v}_τ^2 that upper bounds the vector $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2$ associated to the strategy induced by ψ_τ^2 . We can obtain one in a recursive manner (not solving the induced POMDP).

Proposition 15 *For each ψ_τ^2 obtained as the solution of the aforementioned (dual) LP in σ_τ , and each θ_τ^1 , $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2(\theta_\tau^1)$ is upper bounded by a value $\bar{v}_\tau^2(\theta_\tau^1)$ that depends on vectors $\bar{v}_{\tau+1}^2$ in the support of ψ_τ^2 . In particular, if $\theta_\tau^1 \in \text{Supp}(\sigma_\tau^{m,1})$, we have:*

$$\bar{v}_\tau^2(\theta_\tau^1) \stackrel{\text{def}}{=} \frac{1}{\sigma_{\tau,m}^1(\theta_\tau^1)} \max_{a^1 \in \mathcal{A}^1} M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2.$$

Proof For a newly derived ψ_τ^2 , as $\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2(\theta_\tau^1)$ is the value of 1's best action ($\in \mathcal{A}^1$) if 1 (i) observes θ_τ^1 while in $\sigma_\tau^{c,1}$ and (ii) 2 plays ψ_τ^2 , we have:

$$\begin{aligned}
\nu_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^2(\theta_\tau^1) &\stackrel{\text{def}}{=} V_{[\sigma_\tau^{c,1}, \psi_\tau^2]}^*(\theta_\tau^1) \quad (\text{optimal POMDP value function}) \\
&= \max_{\beta_\tau^1} \mathbb{E} \left[\sum_{t=\tau}^H \gamma^{t-\tau} R_t \mid \beta_\tau^1, \theta_\tau^1, \sigma_\tau^{c,1}, \psi_\tau^2 \right] \\
&= \max_{a^1} \mathbb{E} \left[R_\tau + \gamma \max_{\beta_{\tau+1}^1} \mathbb{E} \left[\sum_{t=\tau+1}^H \gamma^{t-(\tau+1)} R_t \mid \beta_{\tau+1}^1, \langle \theta_\tau^1, a^1, Z^1 \rangle, \sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2 \right] \right. \\
&\quad \left. \mid a^1, \theta_\tau^1, \sigma_\tau^{c,1}, \psi_\tau^2 \right] \\
&= \max_{a^1} \mathbb{E} \left[R_\tau + \gamma V_{[\sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2]}^*(\theta_\tau^1, a^1, Z^1) \mid a^1, \theta_\tau^1, \sigma_\tau^{c,1}, \psi_\tau^2 \right] \\
&= \max_{a^1} \sum_{w, \theta_\tau^2, a^2, z^1} \underbrace{Pr(w, \theta_\tau^2, z^1, a^2 \mid a^1, \theta_\tau^1, \sigma_\tau^{c,1}, \psi_\tau^2)} \\
&\quad \cdot \left(r(\theta_\tau, \mathbf{a}_\tau) + \gamma \nu_{[\sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2[w]]}^2(\theta_\tau^1, a^1, z^1) \right)
\end{aligned}$$

(where $\sigma_{\tau+1}^{c,1} = T_c^1(\sigma_\tau^{c,1}, \beta_\tau^1[w])$ (Lemma 4, p. 39))

$$\begin{aligned}
&= \max_{a^1} \sum_{w, \theta_\tau^2, a^2, z^1} \underbrace{Pr(w \mid \psi_\tau^2)} \cdot \underbrace{Pr(\theta_\tau^2 \mid \theta_\tau^1, \sigma_\tau^{c,1})} \cdot \underbrace{Pr(a^2 \mid \beta_\tau^2[w], \theta_\tau^2)} \cdot \underbrace{Pr(z^1 \mid \theta_\tau, \mathbf{a}_\tau)} \\
&\quad \cdot \left(r(\theta_\tau, \mathbf{a}) + \gamma \nu_{[\sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2[w]]}^2(\theta_\tau^1, a^1, z^1) \right) \\
&= \max_{a^1} \sum_w \psi_\tau^2(w) \sum_{\theta_\tau^2} \sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1) \sum_{a^2} \beta_\tau^2[w](a^2 \mid \theta_\tau^2) \\
&\quad \cdot \left(r(\theta_\tau, \mathbf{a}) + \gamma \sum_{z^1} Pr(z^1 \mid \theta_\tau, \mathbf{a}) \underbrace{\nu_{[\sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2[w]]}^2(\theta_\tau^1, a^1, z^1)} \right)
\end{aligned}$$

then, as $\nu_{[\sigma_{\tau+1}^{c,1}, \psi_{\tau+1}^2[w]]}^2$ is $\lambda_{\tau+1}$ -LC in (any) $\sigma_{\tau+1}^{c,1}$ (Lemma 7),

$$\begin{aligned}
&\leq \max_{a^1} \sum_w \psi_\tau^2(w) \sum_{\theta_\tau^2} \sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1) \sum_{a^2} \beta_\tau^2[w](a^2 \mid \theta_\tau^2) \cdot \left(r(\theta_\tau, \mathbf{a}) + \gamma \sum_{z^1} Pr(z^1 \mid \theta_\tau, \mathbf{a}) \right. \\
&\quad \left. \cdot \underbrace{\left[\nu_{[\tilde{\sigma}_{\tau+1}^{c,1}[w], \psi_{\tau+1}^2[w]]}^2(\theta_\tau^1, a^1, z^1) + \lambda_{\tau+1} \|\sigma_{\tau+1}^{c,1} - \tilde{\sigma}_{\tau+1}^{c,1}[w]\|_1}_{\rightarrow}(\theta_\tau^1, a^1, z^1) \right]} \right) \\
&\leq \max_{a^1} \sum_w \psi_\tau^2(w) \sum_{\theta_\tau^2} \sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1) \sum_{a^2} \beta_\tau^2[w](a^2 \mid \theta_\tau^2) \cdot \left(\underbrace{r(\theta_\tau, \mathbf{a}) + \gamma \sum_{z^1} Pr(z^1 \mid \theta_\tau, \mathbf{a})}_{\rightarrow} \right. \\
&\quad \left. \cdot \underbrace{\left[\bar{\nu}_{\tau+1}^2[w](\theta_\tau^1, a^1, z^1) + \lambda_{\tau+1} \|\sigma_{\tau+1}^{c,1} - \tilde{\sigma}_{\tau+1}^{c,1}[w]\|_1}_{\rightarrow}(\theta_\tau^1, a^1, z^1) \right]} \right) \\
&= \max_{a^1} \sum_w \psi_\tau^2(w) \sum_{\theta_\tau^2} \sigma_\tau^{c,1}(\theta_\tau^2 \mid \theta_\tau^1) \sum_{a^2} \beta_\tau^2[w](a^2 \mid \theta_\tau^2) \cdot \left(\overbrace{\sum_s b(s \mid \theta_\tau) r(s, \mathbf{a})} \right)
\end{aligned}$$

$$\begin{aligned}
& + \gamma \sum_{z^1} \left(\overbrace{\sum_s b(s|\theta_\tau) Pr(z^1|s, \mathbf{a})} \right) \cdot \left[\bar{v}_{\tau+1}^2[w](\theta_\tau^1, a^1, z^1) \right. \\
& \left. + \lambda_{\tau+1} \overbrace{\|\sigma_{\tau+1}^{c,1} - \bar{\sigma}_{\tau+1}^{c,1}[w]\|_1}(\theta_\tau^1, a^1, z^1) \right] \\
& = \max_{a^1} \sum_w \psi_\tau^2(w) \sum_{\theta_\tau^2} \sigma_\tau^{c,1}(\theta_\tau^2|\theta_\tau^1) \sum_{a^2} \beta_\tau^2[w](a^2|\theta_\tau^2) \\
& \quad \cdot \left(\sum_s b(s|\theta_\tau) r(s, \mathbf{a}) + \gamma \sum_{z^1} \left(\overbrace{\sum_{s, s', z^2} b(s|\theta_\tau) P_{\mathbf{a}}^z(s'|s)} \right) \cdot \left[\bar{v}_{\tau+1}^2[w](\theta_\tau^1, a^1, z^1) \right. \right. \\
& \quad \left. \left. + \lambda_{\tau+1} \overbrace{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2[w]) - T_c^1(\bar{\sigma}_\tau^{c,1}[w], \beta_\tau^2[w])\|_1}(\theta_\tau^1, a^1, z^1) \right] \right) \\
& = \frac{1}{\sigma_{\tau, m}^1(\theta_\tau^1)} \max_{a^1 \in \mathcal{A}^1} M_{((\theta_\tau^1, a^1), \cdot)}^{\sigma_\tau} \cdot \psi_\tau^2.
\end{aligned}$$

□

D.3.3 Strategy Conversion

Firstly, we give details regarding solutions of Dual LPs (Equation (7)) inducing behavioral strategies. As suggested in Section 3.2.2, one can show by induction that for any timestep τ , each ψ_τ^2 is actually equivalent to an element of $\Delta(\mathcal{B}_\tau^2)$. The following lemma shows that for any timestep τ , each element of $\Delta(\mathcal{B}_\tau^2)$ induces an element of \mathcal{B}_τ^2 .

Lemma 9. *Each $\psi_\tau^i \in \Delta(\mathcal{B}_\tau^i)$ induces a behavioral strategy. More precisely, we prove that (i) there is a natural injection from the set \mathcal{B}_τ^i to the set of distributions $\Delta(\mathcal{B}_\tau^i)$ and (ii) there is a surjection from the set $\Delta(\mathcal{B}_\tau^i)$ to \mathcal{B}_τ^i .*

Proof By induction on $\tau \in \{0, \dots, H-1\}$, we prove that $\cup_{t=\tau}^H \mathcal{B}_t^i \subset \cup_{t=\tau}^H \Delta(\mathcal{B}_t^i)$. Firstly, for $\tau = H-1$, for all $\beta_{H-1} \in \mathcal{B}_{H-1}$, one can pick the degenerate distribution $\psi_{H-1} = \beta_{H-1}$ which is in $\Delta(\mathcal{B}_{H-1}^i)$. Next, assume that $\cup_{t=\tau+1}^{H-1} \mathcal{B}_t^i \subset \cup_{t=\tau+1}^{H-1} \Delta(\mathcal{B}_t^i)$ for some $\tau \in \{0, \dots, H-2\}$, then for all $\beta_{\tau:H-1} \in \mathcal{B}_{\tau:H-1}$, $\beta_{\tau:H-1} = \beta_\tau \oplus \beta_{\tau+1:H-1}$. By the induction hypothesis, there is $\psi_{\tau+1:H-1} \in \Delta(\mathcal{B}_{\tau+1:H-1}^i)$ equal to $\beta_{\tau+1:H-1}$. Thus, we define $\psi_{\tau:H-1} = \beta_\tau \oplus \psi_{\tau+1:H-1} = \beta_\tau \oplus \beta_{\tau+1:H-1} = \beta_{\tau:H-1}$ which is in $\mathcal{B}_{\tau:H-1}$. From this follows a natural injection from the behavioral strategies' set to the set of distributions over behavioral strategies.

The surjection from $\Delta(\mathcal{B}_\tau^i)$ to \mathcal{B}_τ^i is given by the realization weight computation algorithm detailed in Algorithm 2. □

As discussed in Appendix D.3.3, no effort is required to extract a solution strategy for a player from the lower bound (for 1) or the upper bound (for 2), but that strategy is in an unusual recursive form. We will here see (in the finite horizon setting) how to derive a (unique) equivalent behavioral strategy β_0^i using realization weights [23] in intermediate steps. To that end, we first define

these realization weights in the case of a behavioral strategy (rather than for a mixed strategy as done by [Koller et al.](#)) and present some useful properties.

About Realization Weights

Let us denote $rw^i(a_0^i, z_1^i, a_1^i, \dots, a_\tau^i)$ the *realization weight* (RW) of sequence $a_0^i, z_1^i, a_1^i, \dots, a_\tau^i$ under strategy β_0^i , defined as

$$rw^i(a_0^i, z_1^i, a_1^i, \dots, a_\tau^i) \stackrel{\text{def}}{=} \prod_{t=0}^{\tau} \beta_0^i(a_t^i | a_0^i, z_1^i, a_1^i, \dots, z_t^i) \quad (\text{D14})$$

$$= rw^i(a_0^i, z_1^i, a_1^i, \dots, a_{\tau-1}^i) \cdot \beta_0^i(a_\tau^i | \underbrace{a_0^i, z_1^i, a_1^i, \dots, z_\tau^i}_{\theta_\tau^i}). \quad (\text{D15})$$

This definition already leads to useful results such as:

$$\beta_0^i(a_\tau^i | \theta_\tau^i) = \frac{rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i, z_\tau^i, a_\tau^i)}{rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i)}, \quad (\text{D16})$$

and

$$\forall z_\tau^i, \quad rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i) = rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i) \cdot \underbrace{\sum_{a_\tau^i} \beta(a_\tau^i | \theta_{\tau-1}^i, a_{\tau-1}^i, z_\tau^i)}_{=1} \quad (\text{D17})$$

$$= \sum_{a_\tau^i} rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i) \cdot \beta(a_\tau^i | \theta_{\tau-1}^i, a_{\tau-1}^i, z_\tau^i) \quad (\text{D18})$$

$$= \sum_{a_\tau^i} rw^i(\theta_{\tau-1}^i, a_{\tau-1}^i, z_\tau^i, a_\tau^i). \quad (\text{D19})$$

We now extend [Koller et al.](#)'s definition by introducing *conditional realization weights*, where the realization weight of a *suffix sequence* is “conditioned” on a *prefix sequence*:

$$rw^i(\underbrace{a_\tau^i, \dots, a_{\tau'}^i}_{\text{suffix seq.}} | \underbrace{a_0^i, \dots, z_\tau^i}_{\text{prefix seq.}}) \stackrel{\text{def}}{=} \prod_{t=\tau}^{\tau'} \beta_0^i(a_t^i | a_0^i, \dots, z_\tau^i, a_\tau^i, \dots, z_t^i) \quad (\text{D20})$$

$$= \beta_0^i(a_\tau^i | a_0^i, \dots, z_\tau^i) \cdot rw^i(a_{\tau+1}^i, \dots, a_{\tau'}^i | a_0^i, \dots, z_{\tau+1}^i). \quad (\text{D21})$$

As can be noted, this definition only requires the knowledge of a partial strategy β_τ^i , rather than a complete strategy β_0^i .

Mixing Realization Weights

Let $\tau' \geq \tau + 1$, and $rw^i[w]$ denote the realization weights of some element w at $\tau + 1$. Then, for some ψ_τ^i , we have

$$rw[\psi_\tau^i](a_{\tau+1}^i, \dots, a_{\tau'}^i | a_0^i, \dots, z_{\tau+1}^i) \quad (\text{D22})$$

$$= \sum_w \psi_\tau^i(w) \cdot rw[w](a_{\tau+1}^i, \dots, a_{\tau'}^i | a_0^i, \dots, z_{\tau+1}^i). \quad (\text{D23})$$

From w_0^i to β_0^i

Using the above results, function **Extract** in Algorithm 2 derives a behavioral strategy β_0^i : equivalent to the recursive strategy induced by some tuple w_0^i in 3 steps as follows:

1. **From w_0^i to $rw(\theta_{0:H-1}^i, a_{H-1}^i)$** ($\forall(\theta_{0:H-1}^i, a_{H-1}^i)$) — These (classical) realization weights are obtained by recursively going through the directed acyclic graph describing the recursive strategy, computing *full length* (conditional) realization weights $rw(\theta_{t:H-1}^i, a_{H-1}^i | \theta_{0:t}^i)$ (for $t = H - 1$ down to 0).

When in a leaf node, at depth $H - 1$, the initialization is given by Equation (D20) when $\tau = \tau' = H - 1$:

$$\begin{aligned} rw^i(a_{H-1}^i | a_0^i, \dots, z_{H-1}^i) &\stackrel{\text{def}}{=} \prod_{t=H-1}^{H-1} \beta^i(a_t^i | a_0^i, \dots, z_t^i) \\ &= \beta^i(a_{H-1}^i | a_0^i, \dots, z_{H-1}^i). \end{aligned}$$

Then, in the backward phase, we can compute full length realization weights $rw(\theta_{t+1:H-1}^i, a_{H-1}^i | \theta_{0:t}^i)$ with increasingly longer suffixes (thus shorter prefixes) using (i) Equation (D23) (in function **RecGetRWMix**, line 25) to “mix” several strategies using the distribution ψ_t^i attached to the current w , and (ii) Equation (D21), with $\tau' = H - 1$, (in function **RecGetRWCat**, line 16) to concatenate the behavioral decision rule β_t^i attached to the current w in front of the strategy induced by the distribution ψ_t^i also attached to w . Note: Memoization can here be used to avoid repeating the same computations.

2. **Retrieving (classical) realization weights $rw(\theta_{0:t}^i, a_t^i | -)$** ($\forall t$) — We can now compute realization weights $rw(\theta_{0:t}^i, a_t^i | -)$ for all t 's using Equation (D19) (line 6).
3. **Retrieving behavioral decision rules β_t^i** — Applying Equation (D16) (line 9) then provides the expected behavioral decision rules.

In practice, lossless compressions are used to reduce the dimensionality of the occupancy state (*cf.* Section 4.1), which are currently lost in the current implementation of the conversion. Ideally, one would like to preserve compressions whenever possible or at least retrieve them afterwards, and possibly identify further compressions in the solution strategy.

Algorithm 2: Extracting β_0^i from w_0^i

```

1 Fct Extract( $w_0^i$ )
   /* Step 1., keeping only  $rw(\theta_{0:H-1}^i)$  for all  $\theta_{0:H-1}^i$  */
2    $(rw(\theta_{0:H-1}^i))_{\theta_{0:H-1}^i} \leftarrow \mathbf{RecGetRWMix}(0, w_0^i)$ 
   /* Step 2. */
3   for  $t = H - 2, \dots, 0$  do
4     forall  $\theta_{0:t}^i, a_t^i$  do
5        $z_{t+1}^i \leftarrow z_t^i$  s.t.  $\beta_t(\cdot | \theta_{0:t}^i, a_t^i, z_t^i)$  is defined
6        $rw(\theta_{0:t}^i, a_t^i) \leftarrow \sum_{a_{t+1}^i} rw(\theta_{0:t}^i, a_t^i, z_{t+1}^i, a_{t+1}^i | -)$ 
   /* Step 3. */
7   for  $t = H - 1, \dots, 0$  do
8     forall  $\theta_{0:t}^i, a_t^i$  do
9        $\beta_t^i(a_t^i | \theta_{0:t}^i) \leftarrow \frac{rw^i(\theta_{0:t-1}^i, a_{t-1}^i, z_t^i, a_t^i)}{rw^i(\theta_{0:t-1}^i, a_{t-1}^i)}$ 
10  return  $\beta_0^i$ 
11 Fct RecGetRWMix( $t, w = \langle \beta_t^i, \psi_t^i \rangle$ )
12  for  $w'$  s.t.  $\psi_t^i(w') > 0$  do
13     $rwCat[w'] \leftarrow \mathbf{RecGetRWCat}(t, w')$ 
14  forall  $(a_0^i, \dots, a_{H-1}^i)$  do
15     $rwMix[w](a_t^i, \dots, a_{H-1}^i | a_0^i, \dots, z_t^i)$ 
16     $\leftarrow \sum_{w'} \psi_t^i(w') \cdot rwCat[w'](a_{t+1}^i, \dots, a_{H-1}^i | a_0^i, \dots, z_{t+1}^i)$ 
17  return  $rwMix[w]$ 
18 Fct RecGetRWCat( $t, w = \langle \beta_t^i, \psi_t^i \rangle$ )
19  if  $t = H - 1$  then
20    forall  $(a_0^i, \dots, a_{H-1}^i)$  do
21       $rwCat[w](a_{H-1}^i | a_0^i, \dots, z_{H-1}^i) \leftarrow \beta_t^i(a_{H-1}^i | a_0^i, \dots, z_{H-1}^i)$ 
22  else
23     $rwMix[w] \leftarrow \mathbf{RecGetRWMix}(t, w)$ 
24    forall  $(a_0^i, \dots, a_{H-1}^i)$  do
25       $rwCat[w](a_t^i, \dots, a_{H-1}^i | a_0^i, \dots, z_t^i) \leftarrow$ 
26       $\beta_t^i(a_t^i | a_0^i, \dots, z_t^i) \cdot rwMix[w](a_{t+1}^i, \dots, a_{H-1}^i | a_0^i, \dots, z_{t+1}^i)$ 
  return  $rwCat[w]$ 

```

Appendix E HSVI for zs-POSGs

This section presents results that help (i) tune zs-OMG-HSVI's radius parameter ρ , ensuring that trajectories will always stop, and (ii) then demonstrate the finite time convergence of this algorithm.

E.1 Algorithm

E.1.1 Setting ρ

Proposition 16 *Bounding λ_τ by $\lambda^\infty = \frac{1}{2} \frac{1}{1-\gamma} [r_{\max} - r_{\min}]$ when $\gamma < 1$, and noting that*

$$\begin{aligned} thr(\tau) &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \frac{\gamma^{-\tau} - 1}{1 - \gamma} \quad \text{if } \gamma < 1 \\ &(\text{=} \epsilon - \rho(r_{\max} - r_{\min})(2H + 1 - \tau)\tau \quad \text{if } \gamma = 1), \end{aligned} \quad (\text{E24})$$

one can ensure positivity of the threshold at any $\tau \in 1..H - 1$ by enforcing $0 < \rho < \frac{1-\gamma}{2\lambda^\infty} \epsilon$ (or $0 < \rho < \frac{\epsilon}{(r_{\max} - r_{\min})(H+1)H}$ if $\gamma = 1$).

Proof Let us first consider the case $\gamma < 1$.

We have (for $\tau \in \{1..H - 1\}$):

$$\begin{aligned} thr(\tau) &= \gamma^{-\tau} \epsilon - \sum_{i=1}^{\tau} 2\rho\lambda^\infty \gamma^{-i} \\ &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \sum_{i=1}^{\tau} \gamma^{-i} \\ &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty (\gamma^{-1} + \gamma^{-2} + \dots + \gamma^{-\tau}) \\ &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \gamma^{-1} (\gamma^0 + \gamma^{-1} + \dots + \gamma^{-(\tau-1)}) \\ &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \gamma^{-1} \frac{\gamma^{-\tau} - 1}{\gamma^{-1} - 1} \\ &= \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \frac{\gamma^{-\tau} - 1}{1 - \gamma}. \end{aligned}$$

Then, let us derive the following equivalent inequalities:

$$\begin{aligned} 0 &< thr(\tau) \\ 2\rho\lambda^\infty \frac{\gamma^{-\tau} - 1}{1 - \gamma} &< \gamma^{-\tau} \epsilon \\ \rho &< \frac{1}{2\lambda^\infty} \frac{1 - \gamma}{\gamma^{-\tau} - 1} \gamma^{-\tau} \epsilon \\ \rho &< \frac{1}{2\lambda^\infty} \frac{1 - \gamma}{1 - \gamma^\tau} \epsilon. \end{aligned}$$

To ensure positivity of the threshold for any $\tau \geq 1$, one thus just needs to set ρ as a positive value smaller than $\frac{1-\gamma}{2\lambda^\infty} \epsilon$.

Let us now consider the case $\gamma = 1$.

We have (for $\tau \in \{1, \dots, H-1\}$):

$$\begin{aligned}
thr(\tau) &\stackrel{\text{def}}{=} \epsilon - \sum_{i=1}^{\tau} 2\rho\lambda_{\tau-i} \\
&= \epsilon - \sum_{i=1}^{\tau} 2\rho(H - (\tau - i)) \cdot (r_{\max} - r_{\min}) \\
&= \epsilon - 2\rho(r_{\max} - r_{\min}) \left[\tau(H - \tau) + \sum_{i=1}^{\tau} i \right] \\
&= \epsilon - 2\rho(r_{\max} - r_{\min}) \left[\tau H - \tau^2 + \frac{1}{2}\tau(\tau + 1) \right] \\
&= \epsilon - 2\rho(r_{\max} - r_{\min}) \left[\left(H + \frac{1}{2}\right)\tau - \frac{1}{2}\tau^2 \right] \\
&= \epsilon - \rho(r_{\max} - r_{\min}) \left[(2H + 1)\tau - \tau^2 \right] \\
&= \epsilon - \rho(r_{\max} - r_{\min}) [(2H + 1 - \tau)\tau].
\end{aligned}$$

Then, let us derive the following equivalent inequalities:

$$\begin{aligned}
0 &< thr(\tau) \\
\rho(r_{\max} - r_{\min})(2H + 1 - \tau)\tau &< \epsilon && \text{(holds when } \tau = 0 \text{ and } \tau = H + 1) \\
\rho &< \frac{\epsilon}{(r_{\max} - r_{\min})(2H + 1 - \tau)\tau} && \text{(when } \tau \in \{0 \dots H + 1\}).
\end{aligned}$$

The function $f : \tau \mapsto \frac{\epsilon}{(r_{\max} - r_{\min})(2H + 1 - \tau)\tau}$ reaches its minimum (for $\tau \in (0, H + 1)$) when $\tau = H + \frac{1}{2}$. To ensure positivity of the threshold for any $\tau \in \{1 \dots H - 1\}$, one thus just needs to set ρ as a positive value smaller than $\frac{\epsilon}{(r_{\max} - r_{\min})(H + 1)H}$. \square

E.2 Finite-Time Convergence

E.2.1 Convergence Proof

Proving the finite-time convergence of zs-oMG-HSVI to an error-bounded solution requires some preliminary lemmas.

Lemma 10. *Let $(\sigma_0, \dots, \sigma_{\tau+1})$ be a full trajectory generated by zs-oMG-HSVI and β_{τ} the behavioral DR profile that induced the last transition, i.e., $\sigma_{\tau+1} = T(\sigma_{\tau}, \beta_{\tau})$. Then, after updating \overline{W}_{τ} and \underline{W}_{τ} , we have that $\overline{W}_{\tau}(\sigma_{\tau}, \beta_{\tau}^1) - \underline{W}_{\tau}(\sigma_{\tau}, \beta_{\tau}^2) \leq \gamma thr(\tau + 1)$.*

Proof By definition,

$$\begin{aligned}
\overline{W}_{\tau}(\sigma_{\tau}, \beta_{\tau}^1) &= \min_{(\tilde{\sigma}_{\tau}^{c,1}, \tilde{\beta}_{\tau}^2, \langle \tilde{\nu}_{\tau+1}^2, \psi_{\tau+1}^2 \rangle)} \beta_{\tau}^1 \cdot \left(r(\sigma_{\tau}, \cdot, \tilde{\beta}_{\tau}^2) + \gamma T_m^1(\sigma_{\tau}, \cdot, \tilde{\beta}_{\tau}^2) \right. \\
&\quad \left. \cdot \left[\overline{\nu}_{\tau+1}^2 + \lambda_{\tau+1} \overrightarrow{\|T_c^1(\sigma_{\tau}^{c,1}, \beta_{\tau}^2) - T_c^1(\tilde{\sigma}_{\tau}^{c,1}, \beta_{\tau}^2)\|_1} \right] \right).
\end{aligned}$$

Therefore, after the update (β_τ^2 and β_τ^1 being added to their respective bags ($\overline{\mathcal{I}}_\tau^1$ and $\underline{\mathcal{I}}_\tau^2$) along with vectors $\overline{\nu}_{\tau+1}^2$ and $\underline{\nu}_{\tau+1}^1$),

$$\begin{aligned}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) &\leq \beta_\tau^1 \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \overline{\nu}_{\tau+1}^2 \right], \text{ and} \\ \underline{W}_\tau(\sigma_\tau, \beta_\tau^2) &\geq \beta_\tau^2 \cdot \left[r(\sigma_\tau, \beta_\tau^1, \cdot) + \gamma T_m^2(\sigma_\tau, \beta_\tau^1, \cdot) \cdot \underline{\nu}_{\tau+1}^1 \right].\end{aligned}$$

Then,

$$\begin{aligned}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) - \underline{W}_\tau(\sigma_\tau, \beta_\tau^2) &\leq \left[r(\overline{\sigma}_\tau, \beta_\tau^1, \beta_\tau^2) + \gamma T_m^1(\sigma_\tau, \beta_\tau) \cdot \overline{\nu}_{\tau+1}^2 \right] \\ &\quad - \left[r(\underline{\sigma}_\tau, \beta_\tau^1, \beta_\tau^2) + \gamma T_m^2(\sigma_\tau, \beta_\tau) \cdot \underline{\nu}_{\tau+1}^1 \right] \\ &= \gamma \left[\overline{V}(T(\sigma_\tau, \beta_\tau)) - \underline{V}(T(\sigma_\tau, \beta_\tau)) \right] \\ &\leq \gamma thr(\tau + 1) \quad (\text{Holds at the end of any trajectory.})\end{aligned}$$

□

Lemma 11 (Monotonic evolution of \overline{W}_τ and \underline{W}_τ). *Let $K\overline{W}_\tau$ and $K\underline{W}_\tau$ be the approximations after an update at σ_τ with behavioral DR $\langle \overline{\beta}_\tau^1, \underline{\beta}_\tau^2 \rangle$ (respectively associated to vectors $\overline{\nu}_{\tau+1}^2$ and $\underline{\nu}_{\tau+1}^1$). Let also $K^{(n+1)}\overline{W}_\tau$ and $K^{(n+1)}\underline{W}_\tau$ be the same approximations after n other updates (in various OSS). Then,*

$$\begin{aligned}\max_{\beta_\tau^1} K^{(n+1)}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) &\leq \max_{\beta_\tau^1} K\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) \leq \overline{W}_\tau(\sigma_\tau, \overline{\beta}_\tau^1) \quad \text{and} \\ \min_{\beta_\tau^2} K^{(n+1)}\underline{W}_\tau(\sigma_\tau, \beta_\tau^2) &\geq \min_{\beta_\tau^2} K\underline{W}_\tau(\sigma_\tau, \beta_\tau^2) \geq \underline{W}_\tau(\sigma_\tau, \underline{\beta}_\tau^2).\end{aligned}$$

Proof Starting from the definition,

$$\begin{aligned}&\max_{\beta_\tau^1} K\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) \\ &= \max_{\beta_\tau^1} \min_{\langle \overline{\sigma}_\tau^{c,1}, \beta_\tau^2, \langle \overline{\nu}_{\tau+1}^2, \psi_{\tau+1}^2 \rangle \rangle \in \overline{\mathcal{I}}_\tau^1 \cup \{ \langle \sigma_\tau^{c,1}, \underline{\beta}_\tau^2, \langle \overline{\nu}_{\tau+1}^2, \psi_{\tau+1}^2 \rangle \rangle \}} \beta_\tau^1 \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) \right. \\ &\quad \left. + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \left(\overline{\nu}_{\tau+1}^2 + \lambda_{\tau+1} \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2) - T_c^1(\overline{\sigma}_\tau^{c,1}, \beta_\tau^2)\|_1} \right) \right] \\ &\leq \max_{\beta_\tau^1} \min_{\langle \overline{\sigma}_\tau^{c,1}, \beta_\tau^2, \langle \overline{\nu}_{\tau+1}^2, \psi_{\tau+1}^2 \rangle \rangle \in \overline{\mathcal{I}}_\tau^1} \beta_\tau^1 \cdot \left[r(\sigma_\tau, \cdot, \beta_\tau^2) \right. \\ &\quad \left. + \gamma T_m^1(\sigma_\tau, \cdot, \beta_\tau^2) \cdot \left(\overline{\nu}_{\tau+1}^2 + \lambda_{\tau+1} \overrightarrow{\|T_c^1(\sigma_\tau^{c,1}, \beta_\tau^2) - T_c^1(\overline{\sigma}_\tau^{c,1}, \beta_\tau^2)\|_1} \right) \right] \\ &= \max_{\beta_\tau^1} \overline{W}_\tau(\sigma_\tau, \beta_\tau^1) \\ &= \overline{W}_\tau(\sigma_\tau, \overline{\beta}_\tau^1).\end{aligned}$$

Then, this upper bound approximation can only be refined, so that, for any $n \in \mathbf{N}$,

$$\forall \beta_\tau^1, \quad K^{(n+1)}\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) \leq K\overline{W}_\tau(\sigma_\tau, \beta_\tau^1),$$

$$\text{thus, } \min_{\beta_\tau^1} K^{(n+1)} \overline{W}_\tau(\sigma_\tau, \beta_\tau^1) \leq \min_{\beta_\tau^1} K \overline{W}_\tau(\sigma_\tau, \beta_\tau^1).$$

The expected result thus holds for \overline{W}_τ , and symmetrically for \underline{W}_τ . □

Lemma 12. *After updating, in order, \overline{W}_τ and \overline{V}_τ , we have*

$$K \overline{V}_\tau(\sigma_\tau) \leq \max_{\beta_\tau^1} K \overline{W}_\tau(\sigma_\tau, \beta_\tau^1).$$

After updating, in order, \underline{W}_τ and \underline{V}_τ , we have

$$K \underline{V}_\tau(\sigma_\tau) \geq \min_{\beta_\tau^2} K \underline{W}_\tau(\sigma_\tau, \beta_\tau^2).$$

Proof After updating $\overline{\mathcal{L}}_\tau^1$, the algorithm computes (Algorithm 1, line 20) a new solution $\overline{\psi}_\tau^2$ of the dual LP (at σ_τ^1) and the associated vector \overline{v}_τ^2 , so that

$$\max_{\beta_\tau^1} K \overline{W}_\tau(\sigma_\tau^1, \beta_\tau^1) = \sigma_\tau^{m,1} \cdot \overline{v}_\tau^2.$$

This vector will feed $\overline{\text{bag}}\overline{V}_\tau$ along with σ_τ^1 , so that

$$K \overline{V}_\tau(\sigma_\tau) \leq \sigma_\tau^{m,1} \cdot \overline{v}_\tau^2.$$

As a consequence,

$$K \overline{V}_\tau(\sigma_\tau) \leq \max_{\beta_\tau^1} K \overline{W}_\tau(\sigma_\tau, \beta_\tau^1).$$

The symmetric property holds for $K \underline{V}_\tau$ and $K \underline{W}_\tau$, which concludes the proof. □

Theorem 10 (originally stated on page 16) *zs-oMG-HSVI (Alg. 1) terminates in finite time with an ϵ -approximation of $V_0^*(\sigma_0)$ that satisfies Theorem 9.*

Proof We will prove by induction from $\tau = H$ to 0, that the algorithm stops expanding OSS at depth τ after finitely many iterations (/trajectories).

First, by definition of horizon H , no OS σ_H is ever expanded. The property thus holds at $\tau = H$.

Let us now assume that the property holds at depth $\tau+1$ after $N_{\tau+1}$ iterations. By contradiction, let us assume that the algorithm generates infinitely many trajectories of length $\tau+1$. Then, because $\mathcal{O}_\tau^\sigma \times \mathcal{B}_\tau$ is compact, after some time the algorithm will have visited $\langle \sigma_\tau, \beta_\tau \rangle$, then, some iterations later, $\langle \sigma'_\tau, \beta'_\tau \rangle$, such that $\|\sigma_\tau - \sigma'_\tau\|_1 \leq \rho$. Let us also note the corresponding terminal OSS (because trajectories beyond iteration $N_{\tau+1}$ do not go further) $\sigma_{\tau+1} = T(\sigma_\tau, \beta_\tau)$ and $\sigma'_{\tau+1} = T(\sigma'_\tau, \beta'_\tau)$. Now, we show that the second trajectory should not have happened, *i.e.*, $\overline{V}(\sigma'_\tau) - \underline{V}(\sigma'_\tau) \leq \text{thr}(\tau)$.

Combining the previous lemmas,

$$\begin{aligned} \overline{V}(\sigma'_\tau) &\leq \overline{V}(\sigma_\tau) + \lambda_\tau \|\sigma_\tau - \sigma'_\tau\|_1 && \text{(By Lipschitz-Continuity)} \\ &\leq \max_{\beta_\tau^1} \overline{W}_\tau(\sigma_\tau, \beta_\tau^1) + \lambda_\tau \|\sigma_\tau - \sigma'_\tau\|_1 && \text{(Lemma 12)} \end{aligned}$$

$$\begin{aligned} &\leq \overline{W}_\tau(\sigma_\tau, \beta_\tau^1) + \lambda_\tau \|\sigma_\tau - \sigma'_\tau\|_1 && \text{(Lemma 11)} \\ &= \overline{W}_\tau(\sigma_\tau, \beta_\tau^1) + \lambda_\tau \rho. \end{aligned}$$

Symmetrically, we also have

$$\underline{V}(\sigma'_\tau) \geq \underline{W}_\tau(\sigma_\tau, \beta_\tau^2) - \lambda_\tau \rho.$$

Hence,

$$\begin{aligned} \overline{V}(\sigma'_\tau) - \underline{V}(\sigma'_\tau) &\leq \left(\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) + \lambda_\tau \rho \right) - \left(\underline{W}_\tau(\sigma_\tau, \beta_\tau^2) - \lambda_\tau \rho \right) \\ &= \left(\overline{W}_\tau(\sigma_\tau, \beta_\tau^1) - \underline{W}_\tau(\sigma_\tau, \beta_\tau^2) \right) + 2\lambda_\tau \rho \\ &\leq \gamma \text{thr}(\tau + 1) + 2\lambda_\tau \rho && \text{(Lemma 10)} \\ &= \gamma \left(\gamma^{-(\tau+1)} \epsilon - \sum_{i=1}^{\tau+1} 2\rho \lambda_{\tau+1-i} \gamma^{-i} \right) + 2\lambda_\tau \rho \\ &= \gamma^{-\tau} \epsilon - \sum_{i=1}^{\tau+1} 2\rho \lambda_{\tau+1-i} \gamma^{-i+1} + 2\lambda_\tau \rho \\ &= \gamma^{-\tau} \epsilon - \sum_{j=0}^{\tau} 2\rho \lambda_{\tau-j} \gamma^{-j} + 2\lambda_\tau \rho \\ &= \gamma^{-\tau} \epsilon - \cancel{2\rho \lambda_{\tau-0} \gamma^{-0}} - \sum_{j=1}^{\tau} 2\rho \lambda_{\tau-j} \gamma^{-j} + \cancel{2\lambda_\tau \rho} = \text{thr}(\tau). \end{aligned}$$

Therefore, σ'_τ should not have been expanded. This shows that the algorithm will generate only a finite number of trajectories of length τ . \square

E.2.2 Handling Infinite Horizons

Proposition 17 (originally stated on page 17) *When $\gamma < 1$, the length of trajectories is upper bounded by $T_{\max} \stackrel{\text{def}}{=} \left\lceil \log_\gamma \frac{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}}{W - \frac{2\rho\lambda^\infty}{1-\gamma}} \right\rceil$, where λ^∞ is a depth-independent Lipschitz constant and $W \stackrel{\text{def}}{=} \|\overline{V}^{(0)} - \underline{V}^{(0)}\|_\infty$ is the maximum width between initializations.*

Proof (detailed version) Since W is the largest possible width, any trajectory stops in the worst case at depth τ such that

$$\begin{aligned} \text{thr}(\tau) &< W \\ \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \frac{\gamma^{-\tau} - 1}{1 - \gamma} &< W && \text{(from Eq. (E24))} \\ \gamma^{-\tau} \epsilon - 2\rho\lambda^\infty \frac{\gamma^{-\tau}}{1 - \gamma} - 2\rho\lambda^\infty \frac{-1}{1 - \gamma} &< W \\ \underbrace{\gamma^{-\tau} \left(\epsilon - \frac{2\rho\lambda^\infty}{1 - \gamma} \right)}_{>0 \text{ (Thm. 16)}} &< W - \frac{2\rho\lambda^\infty}{1 - \gamma} \end{aligned}$$

$$\begin{aligned} \gamma^{-\tau} &< \frac{W - \frac{2\rho\lambda^\infty}{1-\gamma}}{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}} \\ \exp(-\tau \ln(\gamma)) &< \exp\left(\ln\left(\frac{W - \frac{2\rho\lambda^\infty}{1-\gamma}}{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}}\right)\right) \\ -\tau \ln(\gamma) &< \ln\left(\frac{W - \frac{2\rho\lambda^\infty}{1-\gamma}}{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}}\right) \\ \tau \ln(\gamma) &> \ln\left(\frac{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}}{W - \frac{2\rho\lambda^\infty}{1-\gamma}}\right) \\ \tau &< \log_\gamma\left(\frac{\epsilon - \frac{2\rho\lambda^\infty}{1-\gamma}}{W - \frac{2\rho\lambda^\infty}{1-\gamma}}\right). \end{aligned}$$

□

Even if the problem horizon is infinite, trajectories will thus have bounded length. Then, everything beyond this *effective* horizon will rely on the upper- and lower-bound initializations and the corresponding strategies.

Bibliography

- [1] Åström, K.J.: Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications* **10**(1), 174–205 (1965)
- [2] Basilio, N., De Nittis, G., Gatti, N.: A security game combining patrolling and alarm-triggered responses under spatial and detection uncertainties. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence* (2016)
- [3] Basu, A., Stettner, L.: Finite- and infinite-horizon Shapley games with nonsymmetric partial observation. *SIAM Journal on Control and Optimization* **53**(6), 3584–3619 (2015)
- [4] Bernstein, D.S., Givan, R., Immerman, N., Zilberstein, S.: The complexity of decentralized control of markov decision processes. *Mathematics of operations research* **27**(4), 819–840 (2002)
- [5] Bono, G., Dibangoye, J., Matignon, L., Pereyron, F., Simonin, O.: Cooperative multi-agent policy gradient. In: *Proceedings of the Twenty-Eight European Conference on Machine Learning* (2018)

- [6] Bošanský, B., Kiekintveld, C., Lisý, V., Pěchouček, M.: An exact double-oracle algorithm for zero-sum extensive-form games with imperfect information. *Journal of Artificial Intelligence Research* **51**, 829–866 (2014). <https://doi.org/10.1613/jair.4477>
- [7] Brown, N., Sandholm, T.: Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* **359**(6374), 418–424 (2018)
- [8] Brown, N., Kroer, C., Sandholm, T.: Dynamic thresholding and pruning for regret minimization. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31 (2017)
- [9] Brown, N., Bakhtin, A., Lerer, A., Gong, Q.: Combining deep reinforcement learning and search for imperfect-information games. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, pp. 17057–17069 (2020)
- [10] Burch, N., Johanson, M., Bowling, M.: Solving imperfect information games using decomposition. In: *Twenty-eighth AAAI Conference on Artificial Intelligence* (2014)
- [11] Burch, N., Moravcik, M., Schmid, M.: Revisiting CFR+ and alternating updates. *Journal of Artificial Intelligence Research* **64**, 429–443 (2019)
- [12] Chatterjee, K., Doyen, L.: Partial-observation stochastic games: How to win when belief fails. *ACM Transactions on Computational Logic* **15**(2), 16 (2014)
- [13] Cole, H.L., Kocherlakota, N.: Dynamic games with hidden actions and hidden states. *Journal of Economic Theory* **98**(1), 114–126 (2001)
- [14] Dibangoye, J., Amato, C., Buffet, O., Charpillet, F.: Optimally solving Dec-POMDPs as continuous-state MDPs. *Journal of Artificial Intelligence Research* **55**, 443–497 (2016)
- [15] Ghosh, M.K., McDonald, D.R., Sinha, S.: Zero-sum stochastic games with partial information. *Journal of Optimization Theory and Applications* **121**(1), 99–118 (2004)
- [16] Hansen, E.A., Zhou, R.: Anytime heuristic search **28**, 267–297 (2007). <https://doi.org/10.1613/jair.2096>
- [17] Hansen, E.A., Bernstein, D.S., Zilberstein, S.: Dynamic programming for partially observable stochastic games. In: *AAAI*, vol. 4, pp. 709–715 (2004)
- [18] Hoda, S., Gilpin, A., Peña, J., Sandholm, T.: Smoothing techniques

- for computing Nash equilibria of sequential games. *Mathematics of Operations Research* **35**(2), 494–512 (2010)
- [19] Horák, K., Bošanský, B., Kovařík, V., Kiekintveld, C.: Solving zero-sum one-sided partially observable stochastic games. *Artificial Intelligence* **316**, 103838 (2023)
- [20] Horák, K.: Scalable algorithms for solving stochastic games with limited partial observability. PhD thesis, Czech Technical University in Prague, Faculty of Electrical Engineering (2019)
- [21] Horák, K., Bošanský, B.: Solving partially observable stochastic games with public observations. In: *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pp. 2029–2036 (2019)
- [22] Horák, K., Bošanský, B., Pěchouček, M.: Heuristic search value iteration for one-sided partially observable stochastic games. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pp. 558–564 (2017)
- [23] Koller, D., Megiddo, N., von Stengel, B.: Fast algorithms for finding randomized strategies in game trees. In: *Proceedings of the 26th ACM Symposium on the Theory of Computing (STOC'94)*, pp. 750–759 (1994)
- [24] Koller, D., Megiddo, N., von Stengel, B.: Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior* **14**(51), 220–246 (1996)
- [25] Korf, R.E.: Real-time heuristic search. *Artificial Intelligence* **42**(2), 189–211 (1990). [https://doi.org/10.1016/0004-3702\(90\)90054-4](https://doi.org/10.1016/0004-3702(90)90054-4)
- [26] Kovařík, V., Seitz, D., Lisý, V., Rudolf, J., Sun, S., Ha, K.: Value functions for depth-limited solving in imperfect-information games. *Corr abs/1906.06412* (2019) [arXiv:1906.06412](https://arxiv.org/abs/1906.06412)
- [27] Kovařík, V., Seitz, D., Lisý, V., Rudolf, J., Sun, S., Ha, K.: Value functions for depth-limited solving in zero-sum imperfect-information games. *Artificial Intelligence*, 103805 (2022). <https://doi.org/10.1016/j.artint.2022.103805>
- [28] Kovařík, V., Schmid, M., Burch, N., Bowling, M., Lisý, V.: Rethinking formal models of partially observable multiagent decision making. *CoRR abs/1906.11110* (2019) [arXiv:1906.11110](https://arxiv.org/abs/1906.11110)
- [29] Kroer, C., Waugh, K., Kılınc-Karzan, F., Sandholm, T.: Faster algorithms for extensive-form game solving via improved smoothing functions. *Mathematical Programming* **179**, 385–417 (2020). <https://doi.org/10.1007/>

s10107-018-1336-7

- [30] Kuhn, H.W.: Simplified two-person Poker. In: Kuhn, H.W., Tucker, A.W. (eds.) *Contributions to the Theory of Games*, vol. 1 (1950)
- [31] Lanctot, M., Waugh, K., Zinkevich, M., Bowling, M.: Monte Carlo sampling for regret minimization in extensive games. *Advances in Neural Information Processing Systems* **22** (2009)
- [32] Lanctot, M., Lockhart, E., Lespiau, J.-B., Zambaldi, V., Upadhyay, S., Pérolat, J., Srinivasan, S., Timbers, F., Tuyls, K., Omidshafiei, S., Hennes, D., Morrill, D., Muller, P., Ewalds, T., Faulkner, R., Kramár, J., Vylder, B.D., Saeta, B., Bradbury, J., Ding, D., Borgeaud, S., Lai, M., Schrittwieser, J., Anthony, T., Hughes, E., Danihelka, I., Ryan-Davis, J.: OpenSpiel: A framework for reinforcement learning in games. *CoRR* **abs/1908.09453** (2019) [arXiv:1908.09453](https://arxiv.org/abs/1908.09453) [cs.LG]
- [33] Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M., Bowling, M.: DeepStack: Expert-level artificial intelligence in heads-up no-limit Poker. *Science* **356**(6337), 508–513 (2017)
- [34] Oliehoek, F., Vlassis, N.: Dec-POMDPs and extensive form games: equivalence of models and algorithms. Technical Report IAS-UVA-06-02, Intelligent Systems Laboratory Amsterdam, University of Amsterdam (2006)
- [35] Oliehoek, F.A.: Sufficient plan-time statistics for decentralized pomdps. In: *Twenty-Third International Joint Conference on Artificial Intelligence* (2013)
- [36] Schmid, M.: Search in imperfect information games. PhD thesis, Charles University - Univerzita Karlova, Prague (2021). <https://arxiv.org/pdf/2111.05884.pdf>
- [37] Schmid, M., Moravcik, M., Burch, N., Kadlec, R., Davidson, J., Waugh, K., Bard, N., Timbers, F., Lanctot, M., Holland, Z., Davoodi, E., Christianson, A., Bowling, M.: Player of games. *CoRR* **abs/2112.03178** (2021) [arXiv:2112.03178](https://arxiv.org/abs/2112.03178)
- [38] Smith, T.: Probabilistic planning for robotic exploration. PhD thesis, The Robotics Institute, Carnegie Mellon University (2007)
- [39] Smith, T., Simmons, R.G.: Point-based POMDP algorithms: Improved analysis and implementation. In: *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, pp. 542–549 (2005)

- [40] Sokota, S., D’Orazio, R., Ling, C.K., Wu, D.J., Kolter, J.Z., Brown, N.: Abstracting imperfect information away from two-player zero-sum games. arXiv preprint arXiv:2301.09159 (2023)
- [41] Szer, D., Charpillet, F., Zilberstein, S.: MAA*: A heuristic search algorithm for solving decentralized POMDPs. In: Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence, pp. 576–583 (2005)
- [42] Tammelin, O.: Solving large imperfect information games using CFR+. CoRR (2014) [arXiv:1407.5042](https://arxiv.org/abs/1407.5042)
- [43] von Neumann, J.: Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* **100** (1928)
- [44] von Stengel, B.: Efficient computation of behavior strategies. *Games and Economic Behavior* **14**(50), 220–246 (1996)
- [45] Wiggers, A.: Structure in the value function of two-player zero-sum games of incomplete information. Master’s thesis, University of Amsterdam (2015)
- [46] Wiggers, A., Oliehoek, F., Roijers, D.: Structure in the value function of two-player zero-sum games of incomplete information. In: Proceedings of the Twenty-Second European Conference on Artificial Intelligence, pp. 1628–1629 (2016a). <https://doi.org/10.3233/978-1-61499-672-9-1628>
- [47] Wiggers, A., Oliehoek, F., Roijers, D.: Structure in the value function of two-player zero-sum games of incomplete information. *Computing Research Repository* **abs/1606.06888** (2016b) [arXiv:1606.06888](https://arxiv.org/abs/1606.06888)
- [48] Xie, Y., Dibangoye, J., Buffet, O.: Optimally solving two-agent decentralized pomdps under one-sided information sharing. In: International Conference on Machine Learning, pp. 10473–10482 (2020). PMLR
- [49] Zinkevich, M., Johanson, M., Bowling, M., Piccione, C.: Regret minimization in games with incomplete information. In: Advances in Neural Information Processing Systems 20 (2007)