

---

# Budgeted Reinforcement Learning in Continuous State Space

---

**Nicolas Carrara\***

SequeL team, INRIA Lille – Nord Europe<sup>†</sup>  
nicolas.carrara@inria.fr

**Edouard Leurent\***

SequeL team, INRIA Lille – Nord Europe<sup>†</sup>  
Renault Group, France  
edouard.leurent@inria.fr

**Romain Laroche**

Microsoft Research, Montreal, Canada  
romain.laroche@microsoft.com

**Tanguy Urvoy**

Orange Labs, Lannion, France  
tanguy.urvoy@orange.com

**Odalric-Ambrym Maillard**

SequeL team, INRIA Lille – Nord Europe  
odalric.maillard@inria.fr

**Olivier Pietquin**

Google Research - Brain Team  
SequeL team, INRIA Lille – Nord Europe<sup>†</sup>  
pietquin@google.com

## Abstract

A Budgeted Markov Decision Process (BMDP) is an extension of a Markov Decision Process to critical applications requiring safety constraints. It relies on a notion of risk implemented in the shape of a cost signal constrained to lie below an – adjustable – threshold. So far, BMDPs could only be solved in the case of finite state spaces with known dynamics. This work extends the state-of-the-art to continuous spaces environments and unknown dynamics. We show that the solution to a BMDP is a fixed point of a novel Budgeted Bellman Optimality operator. This observation allows us to introduce natural extensions of Deep Reinforcement Learning algorithms to address large-scale BMDPs. We validate our approach on two simulated applications: spoken dialogue and autonomous driving.

## 1 Introduction

Reinforcement Learning (RL) is a general framework for decision-making under uncertainty. It frames the learning objective as the optimal control of a Markov Decision Process  $(\mathcal{S}, \mathcal{A}, P, R_r, \gamma)$  with measurable state space  $\mathcal{S}$ , discrete actions  $\mathcal{A}$ , unknown rewards  $R_r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ , and unknown dynamics  $P \in \mathcal{M}(\mathcal{S})^{\mathcal{S} \times \mathcal{A}}$ , where  $\mathcal{M}(\mathcal{X})$  denotes the probability measures over a set  $\mathcal{X}$ . Formally, we seek a policy  $\pi \in \mathcal{M}(\mathcal{A})^{\mathcal{S}}$  that maximises in expectation the  $\gamma$ -discounted return of rewards  $G_r^\pi = \sum_{t=0}^{\infty} \gamma^t R_r(s_t, a_t)$ .

However, this modelling assumption comes at a price: no control is given over the spread of the performance distribution (Dann *et al.*, 2019). In many critical real-world applications where failures may turn out very costly, this is an issue as most decision-makers would rather give away some amount of expected optimality to increase the performances in the lower-tail of the distribution. This has led to the development of several risk-averse variants where the optimisation criteria include

---

\*Both authors contributed equally.

<sup>†</sup>Univ. Lille, CNRS, Centrale Lille, INRIA UMR 9189 - CRISTAL, Lille, France

other statistics of the performance, such as the worst-case realisation (Iyengar, 2005; Nilim and El Ghaoui, 2005; Wiesemann *et al.*, 2013), the variance-penalised expectation (García and Fernández, 2015; Tamar *et al.*, 2012), the Value-At-Risk (VaR) (Mausser and Rosen, 2003; Luenberger, 2013), or the Conditional Value-At-Risk (CVaR) (Chow *et al.*, 2015, 2018).

Reinforcement Learning also assumes that the performance can be described by a single reward function  $R_r$ . Conversely, real problems typically involve many aspects, some of which can be contradictory (Liu *et al.*, 2014). For instance, a self-driving car needs to balance between progressing quickly on the road and avoiding collisions. When aggregating several objectives in a single scalar signal, as often in Multi-Objectives RL (Roijers *et al.*, 2013), no control is given over their relative ratios, as high rewards can compensate high penalties. For instance, if a weighted sum is used to balance velocity  $v$  and crashes  $c$ , then for any given choice of weights  $\omega$  the optimality equation  $\omega_v \mathbb{E}[\sum \gamma^t v_t] + \omega_c \mathbb{E}[\sum \gamma^t c_t] = G_r^* = \max_{\pi} G_r^{\pi}$  is the equation of a line in  $(\mathbb{E}[\sum \gamma^t v_t], \mathbb{E}[\sum \gamma^t c_t])$ , and the automotive company cannot control where its optimal policy  $\pi^*$  lies on that line.

Both of these concerns can be addressed in the *Constrained Markov Decision Process* (CMDP) setting (Beutler and Ross, 1985; Altman, 1999). In this multi-objective formulation, task completion and safety are considered separately. We equip the MDP with a cost signal  $R_c \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$  and a cost budget  $\beta \in \mathbb{R}$ . Similarly to  $G_r^{\pi}$ , we define the return of costs  $G_c^{\pi} = \sum_{t=0}^{\infty} \gamma^t R_c(s_t, a_t)$  and the new cost-constrained objective:

$$\max_{\pi \in \mathcal{M}(\mathcal{A})^{\mathcal{S}}} \mathbb{E}[G_r^{\pi} | s_0 = s] \quad \text{s.t.} \quad \mathbb{E}[G_c^{\pi} | s_0 = s] \leq \beta \quad (1)$$

This constrained framework allows for better control of the performance-safety tradeoff. However, it suffers from a major limitation: the budget has to be chosen before training, and cannot be changed afterwards.

To address this concern, the *Budgeted Markov Decision Process* (BMDP) was introduced in (Boutilier and Lu, 2016) as an extension of CMDPs to enable the online control over the budget  $\beta$  within an interval  $\mathcal{B} \subset \mathbb{R}$  of admissible budgets. Instead of fixing the budget prior to training, the objective is now to find a generic optimal policy  $\pi^*$  that takes  $\beta$  as input so as to solve the corresponding CMDP (Eq. (1)) for all  $\beta \in \mathcal{B}$ . This gives the system designer the ability to move the optimal policy  $\pi^*$  in real-time along the Pareto-optimal curve of the different reward-cost trade-offs.

Our first contribution is to re-frame the original BMDP formulation in the context of continuous states and infinite discounted horizon. We then propose a novel Budgeted Bellman Optimality Operator and prove the optimal value function to be a fixed point of this operator. Second, we use this operator in BFTQ, a batch Reinforcement Learning algorithm, for solving BMDPs online by interaction with an environment, through function approximation and a tailored exploration procedure. Third, we scale this algorithm to large problems by providing an efficient implementation of the Budgeted Bellman Optimality Operator based on convex programming, and by leveraging tools from Deep Reinforcement Learning such as Deep Neural Networks and synchronous parallel computing. Finally, we validate our approach in two environments that display a clear trade-off between rewards and costs: a spoken dialogue system and a problem of behaviour planning for autonomous driving. The proofs of our main results are provided in Appendix A.

## 2 Budgeted Dynamic Programming

We work in the space of budgeted policies, where  $\pi$  both depends on  $\beta$  and also outputs a next budget  $\beta_a$ . Hence, the budget  $\beta$  is neither fixed nor constant as in the CMDP setting but instead evolves as part of the dynamics.

We cast the BMDP problem as a multi-objective MDP problem (Roijers *et al.*, 2013) by considering *augmented* state and action spaces  $\bar{\mathcal{S}} = \mathcal{S} \times \mathcal{B}$  and  $\bar{\mathcal{A}} = \mathcal{A} \times \mathcal{B}$ , and equip them with the augmented dynamics  $\bar{P} \in \mathcal{M}(\bar{\mathcal{S}})^{\bar{\mathcal{S}} \times \bar{\mathcal{A}}}$  defined as:

$$\bar{P}(\bar{s}' | \bar{s}, \bar{a}) = \bar{P}((s', \beta') | (s, \beta), (a, \beta_a)) \stackrel{\text{def}}{=} P(s' | s, a) \delta(\beta' - \beta_a), \quad (2)$$

where  $\delta$  is the Dirac indicator distribution.

In other words, in these augmented dynamics, the output budget  $\beta_a$  returned at time  $t$  by a budgeted policy  $\pi \in \Pi = \mathcal{M}(\bar{\mathcal{A}})^{\bar{\mathcal{S}}}$  will be used to condition the policy at the next timestep  $t + 1$ .

We stack the rewards and cost functions in a single *vectorial* signal  $R \in (\mathbb{R}^2)^{\bar{\mathcal{S}} \times \bar{\mathcal{A}}}$ . Given an augmented transition  $(\bar{s}, \bar{a}) = ((s, \beta), (a, \beta_a))$ , we define:

$$R(\bar{s}, \bar{a}) \stackrel{\text{def}}{=} \begin{bmatrix} R_r(s, a) \\ R_c(s, a) \end{bmatrix} \in \mathbb{R}^2. \quad (3)$$

Likewise, the return  $G^\pi = (G_r^\pi, G_c^\pi)$  of a budgeted policy  $\pi \in \Pi$  refers to:  $G^\pi \stackrel{\text{def}}{=} \sum_{t=0}^{\infty} \gamma^t R(\bar{s}_t, \bar{a}_t)$ , and the value functions  $V^\pi, Q^\pi$  of a budgeted policy  $\pi \in \Pi$  are defined as:

$$V^\pi(\bar{s}) = (V_r^\pi, V_c^\pi) \stackrel{\text{def}}{=} \mathbb{E}[G^\pi \mid \bar{s}_0 = \bar{s}] \quad Q^\pi(\bar{s}, \bar{a}) = (Q_r^\pi, Q_c^\pi) \stackrel{\text{def}}{=} \mathbb{E}[G^\pi \mid \bar{s}_0 = \bar{s}, \bar{a}_0 = \bar{a}]. \quad (4)$$

We restrict  $\bar{\mathcal{S}}$  to feasible budgets only:  $\bar{\mathcal{S}}_f \stackrel{\text{def}}{=} \{(s, \beta) \in \bar{\mathcal{S}} : \exists \pi \in \Pi, V_c^\pi(s) \geq \beta\}$  that we assume is non-empty for the BMDP to admit a solution. We still write  $\bar{\mathcal{S}}$  in place of  $\bar{\mathcal{S}}_f$  for brevity of notations.

**Proposition 1** (Budgeted Bellman Expectation). *The value functions  $V^\pi$  and  $Q^\pi$  verify:*

$$V^\pi(\bar{s}) = \sum_{\bar{a} \in \bar{\mathcal{A}}} \pi(\bar{a} \mid \bar{s}) Q^\pi(\bar{s}, \bar{a}) \quad Q^\pi(\bar{s}, \bar{a}) = R(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \bar{\mathcal{S}}} \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) V^\pi(\bar{s}') \quad (5)$$

Moreover, consider the Budgeted Bellman Expectation operator  $\mathcal{T}^\pi: \forall Q \in (\mathbb{R}^2)^{\bar{\mathcal{S}} \times \bar{\mathcal{A}}}, \bar{s} \in \bar{\mathcal{S}}, \bar{a} \in \bar{\mathcal{A}}$ ,

$$\mathcal{T}^\pi Q(\bar{s}, \bar{a}) \stackrel{\text{def}}{=} R(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \bar{\mathcal{S}}} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) \pi(\bar{a}' \mid \bar{s}') Q(\bar{s}', \bar{a}') \quad (6)$$

Then  $\mathcal{T}^\pi$  is a  $\gamma$ -contraction and  $Q^\pi$  is its unique fixed point.

**Definition 1** (Budgeted Optimality). *We now come to the definition of budgeted optimality. We want an optimal budgeted policy to: (i) respect the cost budget  $\beta$ , (ii) maximise the  $\gamma$ -discounted return of rewards  $G_r$ , (iii) in case of tie, minimise the  $\gamma$ -discounted return of costs  $G_c$ . To that end, we define for all  $\bar{s} \in \bar{\mathcal{S}}$ :*

(i) *Admissible policies  $\Pi_a$ :*

$$\Pi_a(\bar{s}) \stackrel{\text{def}}{=} \{\pi \in \Pi : V_c^\pi(\bar{s}) \leq \beta\} \text{ where } \bar{s} = (s, \beta) \quad (7)$$

(ii) *Optimal value function for rewards  $V_r^*$  and candidate policies  $\Pi_r$ :*

$$V_r^*(\bar{s}) \stackrel{\text{def}}{=} \max_{\pi \in \Pi_a(\bar{s})} V_r^\pi(\bar{s}) \quad \Pi_r(\bar{s}) \stackrel{\text{def}}{=} \arg \max_{\pi \in \Pi_a(\bar{s})} V_r^\pi(\bar{s}) \quad (8)$$

(iii) *Optimal value function for costs  $V_c^*$  and optimal policies  $\Pi^*$ :*

$$V_c^*(\bar{s}) \stackrel{\text{def}}{=} \min_{\pi \in \Pi_r(\bar{s})} V_c^\pi(\bar{s}), \quad \Pi^*(\bar{s}) \stackrel{\text{def}}{=} \arg \min_{\pi \in \Pi_r(\bar{s})} V_c^\pi(\bar{s}) \quad (9)$$

We define the budgeted action-value function  $Q^*$  similarly:

$$Q_r^*(\bar{s}, \bar{a}) \stackrel{\text{def}}{=} \max_{\pi \in \Pi_a(\bar{s})} Q_r^\pi(\bar{s}, \bar{a}) \quad Q_c^*(\bar{s}, \bar{a}) \stackrel{\text{def}}{=} \min_{\pi \in \Pi_r(\bar{s})} Q_c^\pi(\bar{s}, \bar{a}) \quad (10)$$

and denote  $V^* = (V_r^*, V_c^*)$ ,  $Q^* = (Q_r^*, Q_c^*)$ .

**Theorem 1** (Budgeted Bellman Optimality). *The optimal budgeted action-value function  $Q^*$  verifies:*

$$Q^*(\bar{s}, \bar{a}) = \mathcal{T}Q^*(\bar{s}, \bar{a}) \stackrel{\text{def}}{=} R(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \bar{\mathcal{S}}} \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) \sum_{\bar{a}' \in \bar{\mathcal{A}}} \pi_{\text{greedy}}(\bar{a}' \mid \bar{s}'; Q^*) Q^*(\bar{s}', \bar{a}'), \quad (11)$$

where the greedy policy  $\pi_{\text{greedy}}$  is defined by:  $\forall \bar{s} = (s, \beta) \in \bar{\mathcal{S}}, \bar{a} \in \bar{\mathcal{A}}, \forall Q \in (\mathbb{R}^2)^{\bar{\mathcal{A}} \times \bar{\mathcal{S}}}$ ,

$$\pi_{\text{greedy}}(\bar{a} \mid \bar{s}; Q) \in \arg \min_{\rho \in \Pi_r^Q} \mathbb{E}_{\bar{a} \sim \rho} Q_c(\bar{s}, \bar{a}), \quad (12a)$$

$$\text{where } \Pi_r^Q \stackrel{\text{def}}{=} \arg \max_{\rho \in \mathcal{M}(\bar{\mathcal{A}})} \mathbb{E}_{\bar{a} \sim \rho} Q_r(\bar{s}, \bar{a}) \quad (12b)$$

$$\text{s.t. } \mathbb{E}_{\bar{a} \sim \rho} Q_c(\bar{s}, \bar{a}) \leq \beta. \quad (12c)$$

**Remark 1** (Appearance of the greedy policy). *In classical Reinforcement Learning, the greedy policy takes a simple form  $\pi_{\text{greedy}}(s; Q^*) = \arg \max_{a \in \mathcal{A}} Q^*(s, a)$ , and the term  $\pi_{\text{greedy}}(a' | s'; Q^*) Q^*(s', a')$  in (11) conveniently simplifies to  $\max_{a' \in \mathcal{A}} Q^*(s', a')$ . Unfortunately, in a budgeted setting the greedy policy requires solving the nested constrained optimisation program (12) at each state and budget in order to apply this Budgeted Bellman Optimality operator.*

**Proposition 2** (Optimality of the greedy policy). *The greedy policy  $\pi_{\text{greedy}}(\cdot; Q^*)$  is uniformly optimal: for all  $\bar{s} \in \bar{\mathcal{S}}$ ,  $\pi_{\text{greedy}}(\cdot; Q^*) \in \Pi^*(\bar{s})$ . In particular,  $V^{\pi_{\text{greedy}}(\cdot; Q^*)} = V^*$  and  $Q^{\pi_{\text{greedy}}(\cdot; Q^*)} = Q^*$ .*

**Budgeted Value Iteration** The Budgeted Bellman Optimality equation is a fixed-point equation, which motivates the introduction of a fixed-point iteration procedure. We introduce Algorithm 1, a Dynamic Programming algorithm for solving known BMDPs. If it were to converge to a unique fixed point, this algorithm would provide a way to compute  $Q^*$  and recover the associated optimal budgeted policy  $\pi_{\text{greedy}}(\cdot; Q^*)$ .

**Theorem 2** (Non-contractivity of  $\mathcal{T}$ ). *For any BMDP  $(\mathcal{S}, \mathcal{A}, P, R_r, R_c, \gamma)$  with  $|\mathcal{A}| \geq 2$ ,  $\mathcal{T}$  is not a contraction. Precisely:  $\forall \varepsilon > 0, \exists Q^1, Q^2 \in (\mathbb{R}^2)^{\bar{\mathcal{S}}\mathcal{A}} : \|\mathcal{T}Q^1 - \mathcal{T}Q^2\|_\infty \geq \frac{1}{\varepsilon} \|Q^1 - Q^2\|_\infty$ .*

Unfortunately, as  $\mathcal{T}$  is not a contraction, we can guarantee neither the convergence of Algorithm 1 nor the unicity of its fixed points. Despite those theoretical limitations, we empirically observed the convergence to a fixed point in our experiments (Section 5). We conjecture a possible explanation:

**Remark 2** (Contractivity of  $\mathcal{T}$  on smooth  $Q$ -functions). *We conjecture that  $\mathcal{T}$  is a contraction when restricted to the subset  $\mathcal{L}_\gamma$  of  $Q$ -functions such that " $Q_r$  is  $L$ -Lipschitz with respect to  $Q_c$ ", with  $L < \frac{1}{\gamma} - 1$ . We lengthily discuss some intuition on why that should be the case in Appendix A.5.*

### 3 Budgeted Reinforcement Learning

In this section, we consider BMDPs with unknown parameters that must be solved by interaction with an environment.

#### 3.1 Budgeted Fitted-Q

When the BMDP is unknown, we need to adapt Algorithm 1 to work with a batch of samples  $\mathcal{D} = \{(\bar{s}_i, \bar{a}_i, r_i, \bar{s}'_i)\}_{i \in [0, N]}$  collected by interaction with the environment. Applying  $\mathcal{T}$  in (11) would require computing an expectation  $\mathbb{E}_{\bar{s}' \sim \bar{P}}$  over next states  $\bar{s}'$  and hence an access to the model  $\bar{P}$ . We instead use  $\hat{\mathcal{T}}$ , a sampling operator, in which this expectation is replaced by:

$$\hat{\mathcal{T}}Q(\bar{s}_i, \bar{a}_i, r_i, \bar{s}'_i) \stackrel{\text{def}}{=} r_i + \gamma \sum_{\bar{a}'_i \in \mathcal{A}_i} \pi_{\text{greedy}}(\bar{a}'_i | \bar{s}'_i; Q) Q(\bar{s}'_i, \bar{a}'_i).$$

We introduce in Algorithm 2 the *Budgeted-Fitted-Q* (BFTQ) algorithm, an extension of the *Fitted-Q* (FTQ) algorithm (Ernst et al., 2005; Riedmiller, 2005) adapted to solve unknown BMDPs. Because we work with continuous state space  $\mathcal{S}$  and budget space  $\mathcal{B}$ , we need to employ function-approximation in order to generalise to nearby states and budgets. Precisely, given a parametrized model  $Q_\theta$ , we seek to minimise a regression loss  $\mathcal{L}(Q_\theta, Q_{\text{target}}; \mathcal{D}) = \sum_{\mathcal{D}} \|Q_\theta(\bar{s}, \bar{a}) - Q_{\text{target}}(\bar{s}, \bar{a}, r, \bar{s}')\|_2^2$ . Any model can be used, such as linear models, regression trees, or neural networks.

---

#### Algorithm 1: Budgeted Value Iteration

---

**Data:**  $P, R_r, R_c$

**Result:**  $Q^*$

- 1  $Q_0 \leftarrow 0$
  - 2 **repeat**
  - 3    $Q_{k+1} \leftarrow \mathcal{T}Q_k$
  - 4 **until** convergence
- 

---

#### Algorithm 2: Budgeted Fitted-Q

---

**Data:**  $\mathcal{D}$

**Result:**  $Q^*$

- 1  $Q_{\theta_0} \leftarrow 0$
  - 2 **repeat**
  - 3    $\theta_{k+1} \leftarrow \arg \min_{\theta} \mathcal{L}(Q_\theta, \hat{\mathcal{T}}Q_{\theta_k}; \mathcal{D})$
  - 4 **until** convergence
-

### 3.2 Risk-sensitive exploration

In order to run Algorithm 2, we must first gather a batch of samples  $\mathcal{D}$ . Ideally we would need samples from the asymptotic state-budget distribution  $\lim_{t \rightarrow \infty} \mathbb{P}(\bar{s}_t)$  induced by an optimal policy  $\pi^*$  given an initial distribution  $\mathbb{P}(\bar{s}_0)$ , but as we are actually building this policy, it is not possible. Following the same idea of  $\varepsilon$ -greedy exploration for FTQ (Ernst *et al.*, 2005; Riedmiller, 2005), we introduce an algorithm for risk-sensitive exploration. We follow an exploration policy: a mixture between a random budgeted policy  $\pi_{\text{rand}}$  and the current greedy policy  $\pi_{\text{greedy}}$ . The batch  $\mathcal{D}$  is split into several mini-batches generated sequentially, and  $\pi_{\text{greedy}}$  is updated by running Algorithm 2 on  $\mathcal{D}$  upon mini-batch completion.  $\pi_{\text{rand}}$  is designed to obtain trajectories that only explore feasible budgets: we impose that the joint distribution  $\mathbb{P}(a, \beta_a | s, \beta)$  verifies  $\mathbb{E}[\beta_a] \leq \beta$ . This condition defines a probability simplex  $\Delta_{\bar{\mathcal{A}}}$  from which we sample uniformly. Finally, when interacting with an environment the initial state  $s_0$  is usually sampled from a starting distribution  $\mathbb{P}(s_0)$ . In the budgeted setting, we also need to sample the initial budget  $\beta_0$ . Importantly, we pick a uniform distribution  $\mathbb{P}(\beta_0) = \mathcal{U}(\mathcal{B})$  so that the entire range of risk-level is explored, and not only reward-seeking behaviours as would be the case with a traditional risk-neutral  $\varepsilon$ -greedy strategy. The pseudo-code of our exploration procedure is shown in Algorithm 4 in Appendix B.

## 4 A Scalable Implementation

In this section, we introduce an implementation of the BFTQ algorithm designed to operate efficiently and handle large batches of experiences  $\mathcal{D}$ .

### 4.1 How to compute the greedy policy?

As stated in Remark 1, computing the greedy policy  $\pi_{\text{greedy}}$  in (11) is not trivial since it requires solving the nested constrained optimisation program (12). However, it can be solved efficiently by exploiting the *structure* of the set of solutions with respect to  $\beta$ , that is, concave and increasing.

**Proposition 3** (Equality of  $\pi_{\text{greedy}}$  and  $\pi_{\text{hull}}$ ). *Algorithm 1 and Algorithm 2 can be run by replacing  $\pi_{\text{greedy}}$  in the equation (11) of  $\mathcal{T}$  with  $\pi_{\text{hull}}$  as described in Algorithm 3.*

$$\pi_{\text{greedy}}(\bar{a} | \bar{s}; Q) = \pi_{\text{hull}}(\bar{a} | \bar{s}; Q)$$

---

**Algorithm 3:** Convex hull policy  $\pi_{\text{hull}}(\bar{a} | \bar{s}; Q)$

---

**Data:**  $\bar{s} = (s, \beta), Q$

- 1  $Q^+ \leftarrow \{Q_c > \min\{Q_c(\bar{s}, \bar{a}) \text{ s.t. } \bar{a} \in \arg \max_{\bar{a}} Q_r(\bar{s}, \bar{a})\}\}$  // dominated points
- 2  $\mathcal{F} \leftarrow \text{top frontier of } \text{convex\_hull}(Q(\bar{s}, \bar{\mathcal{A}}) \setminus Q^+)$  // candidate mixtures
- 3  $\mathcal{F}_Q \leftarrow \mathcal{F} \cap Q(\bar{s}, \bar{\mathcal{A}})$
- 4 **for** points  $q = Q(\bar{s}, \bar{a}) \in \mathcal{F}_Q$  **in clockwise order do**
- 5     **if** find two successive points  $((q_c^1, q_r^1), (q_c^2, q_r^2))$  of  $\mathcal{F}_Q$  such that  $q_c^1 \leq \beta < q_c^2$  **then**
- 6          $p \leftarrow (\beta - q_c^1) / (q_c^2 - q_c^1)$
- 7         **return** the mixture  $(1 - p)\delta(\bar{a} - \bar{a}^1) + p\delta(\bar{a} - \bar{a}^2)$
- 8 **end**
- 9 **else return**  $\delta(\bar{a} - \arg \max_{\bar{a}} Q_r(\bar{s}, \bar{a}))$  // budget  $\beta$  always respected

---

The computation of  $\pi_{\text{hull}}$  in Algorithm 3 is illustrated in Figure 1.

### 4.2 Function approximation

Neural networks are well suited to model Q-functions in Reinforcement Learning algorithms (Riedmiller, 2005; Mnih *et al.*, 2015). We approximate  $Q = (Q_r, Q_c)$  using one single neural network. Thus, the two components are jointly optimised which accelerates convergence and fosters learning of useful shared representations. Moreover, as in (Mnih *et al.*, 2015) we are dealing with a finite (categorical) action space  $\mathcal{A}$ , instead of including the action in the input we add the output of the Q-function for each action to the last layer. Again, it provides a faster convergence toward useful shared representations and it only requires one forward pass to evaluate all action values. Finally,




Figure 1: Representation of  $\pi_{\text{hull}}$ . When the budget lies between  $Q(\bar{s}, \bar{a}_1)$  and  $Q(\bar{s}, \bar{a}_2)$ , two points of the top frontier of the convex hull, then the policy is a mixture of these two points.

beside the state  $s$  there is one more input to a budgeted  $Q$ -function: the budget  $\beta_a$ . This budget is a scalar value whereas the state  $s$  is a vector of potentially large size. To avoid a weak influence of  $\beta$  compared to  $s$  in the prediction, we include an additional encoder for the budget, whose width and depth may depend on the application. A straightforward choice is a single layer with the same width as the state. The overall architecture is shown in Figure 7 in Appendix C.

### 4.3 Parallel computing

In a simulated environment, a first process that can be distributed is the collection of samples in the exploration procedure of Algorithm 4, as  $\pi_{\text{greedy}}$  stays constant within each mini-batch which avoids the need of synchronisation between workers. Second, the main bottleneck of BFTQ is the computation of the target  $\mathcal{T}Q$ . Indeed, when computing  $\pi_{\text{hull}}$  we must perform at each epoch a Graham-scan of complexity  $\mathcal{O}(|\mathcal{A}||\tilde{\mathcal{B}}| \log |\mathcal{A}\tilde{\mathcal{B}}|)$  per sample in  $\mathcal{D}$  to compute the convex hulls of  $Q$  (where  $\tilde{\mathcal{B}}$  is a finite discretisation of  $\mathcal{B}$ ). The resulting total time-complexity is  $\mathcal{O}(\frac{|\mathcal{D}||\mathcal{A}||\tilde{\mathcal{B}}|}{1-\gamma} \log |\mathcal{A}||\tilde{\mathcal{B}}|)$ . This operation can easily be distributed over several CPUs provided that we first evaluate the model  $Q(s', \mathcal{A}\tilde{\mathcal{B}})$  for each sample  $s' \in \mathcal{D}$ , which can be done in a single forward pass. By using multiprocessing in the computations of  $\pi_{\text{hull}}$ , we enjoy a linear speedup. The full description of our scalable implementation of BFTQ is recalled in Algorithm 5 in Appendix C.

## 5 Experiments

There are two hypotheses we want to validate.

**Exploration strategies** We claimed in Section 3.2 that a risk-sensitive exploration was required in the setting of BMDPs. We test this hypotheses by confronting our strategy to a classical risk-neutral strategy. The latter is chosen to be a  $\varepsilon$ -greedy policy slowly transitioning from a random to a greedy policy<sup>3</sup> that aims to maximise  $\mathbb{E}_\pi G_r^\pi$  regardless of  $\mathbb{E}_\pi G_c^\pi$ . The quality of the resulting batches  $\mathcal{D}$  is assessed by training a BFTQ policy and comparing the resulting performance.

**Budgeted algorithms** We compare our scalable BFTQ algorithm described in Section 4 to an FTQ( $\lambda$ ) baseline. This baseline consists in approximating the BMDP by a finite set of CMDPs problems. We solve each of these CMDP using the standard technique of Lagrangian Relaxation: the cost constraint is converted to a soft penalty weighted by a Lagrangian multiplier  $\lambda$  in a surrogate reward function:  $\max_\pi \mathbb{E}_\pi [G_r^\pi - \lambda G_c^\pi]$ . The resulting MDP can be solved by any RL algorithm, and we chose FTQ for being closest to BFTQ. In our experiments, a single training of BFTQ corresponds to 10 trainings of FTQ( $\lambda$ ) policies. Each run was repeated  $N_{\text{seeds}}$  times. Parameters of the algorithms can be found in Appendix E.3.2

<sup>3</sup>We train this greedy policy using FTQ.

## 5.1 Environments

We evaluate our method on three different environments involving reward-cost trade-offs. Their parameters can be found in Appendix E.3.1

**Corridors** This simple environment is only meant to highlight clearly the specificity of exploration in a budgeted setting. It is a continuous gridworld with Gaussian perturbations, consisting in a maze composed of two corridors: a risky one with high rewards and costs, and a safe one with low rewards and no cost. In both corridors the outermost cell is the one yielding the most reward, which motivates a deep exploration.

**Spoken dialogue system** Our second application is a dialogue-based slot-filling simulation that has already benefited from batch RL optimisation in the past (Li *et al.*, 2009; Chandramohan *et al.*, 2010; Pietquin *et al.*, 2011). The system fills in a form of slot-values by interacting a user through speech, before sending them a response. For example, in a restaurant reservation domain, it may ask for three slots: the area of the restaurant, the price-range and the food type. The user could respectively provide those three slot-values : Cambridge, Cheap and Indian-food. In this application, we do not focus on how to extract such information from the user utterances, we rather focus on decision-making for filling in the form. To that end, the system can choose among a set of generic actions. As in (Carrara *et al.*, 2018), there are two ways of asking for a slot value: a slot value can be either be provided with an utterance, which may cause speech recognition errors with some probability, or by requiring the user to fill-in the slots by using a numeric pad. In this case, there are no recognition errors but a counterpart risk of hang-up: we assume that manually filling a key-value form is time-consuming and annoying. The environment yields a reward if all slots are filled without errors, and a constraint if the user hang-ups. Thus, there is a clear trade-off between using utterances and potentially committing a mistake, or using the numeric pad and risking a premature hang-up.

**Autonomous driving** In our third application, we use the `highway-env` environment (Leurent *et al.*, 2018) for simulated highway driving and behavioural decision-making. We define a task that displays a clear trade-off between safety and efficiency. The agent controls a vehicle with a finite set of manoeuvres implemented by low-lever controllers:  $\mathcal{A} = \{\text{no-op, right-lane, left-lane, faster, slower}\}$ . It is driving on a two-lane road populated with other traffic participants: the vehicles in front of the agent drive slowly, and there are incoming vehicles on the opposite lane. Their behaviours are randomised, which introduces some uncertainty with respect to their possible future trajectories. The task consists in driving as fast as possible, which is modelled by a reward proportional to the velocity:  $R_r(s_t, a_t) \propto v_t$ . This motivates the agent to try and overtake its preceding vehicles by driving fast on the opposite lane. This optimal but overly aggressive behaviour can be tempered through a cost function that embodies a safety objective:  $R_c(s_t, a_t)$  is set to  $1/H$  whenever the ego-vehicle is driving on the opposite lane, where  $H$  is the episode horizon. Thus, the constrained signal  $G_c^\pi$  is the maximum proportion of time that the agent is allowed to drive on the wrong side of the road.

## 5.2 Results

In the following figures, each patch represents the mean and 95% confidence interval over  $N_{\text{seeds}}$  seeds of the means of  $(G_r^\pi, G_c^\pi)$  over  $N_{\text{trajs}}$  trajectories. That way, we display the variation related to learning (and batches) rather than the variation in the execution of the policies.

We first bring to light the role of risk-sensitive exploration in the corridors environment: Figure 2 shows the set of trajectories collected by each exploration strategy<sup>4</sup>, and the resulting performance of a budgeted policy trained on each batch. The trajectories (orange) in the risk-neutral batch are concentrated along the risky corridor (red) and ignore the safe corridor (green), which results in bad performances in the low-risk regime. Conversely, trajectories in the risk-sensitive batch (blue) are well distributed among both corridors and the corresponding budgeted policy achieves good performance across the whole spectrum of risk budgets.

In a second experiment displayed in Figure 3, we compare the performance of FTQ( $\lambda$ ) to that of BFTQ in the dialogue and autonomous driving tasks. For each algorithm, we plot the reward-cost

<sup>4</sup>Animations are available in Appendix E.1



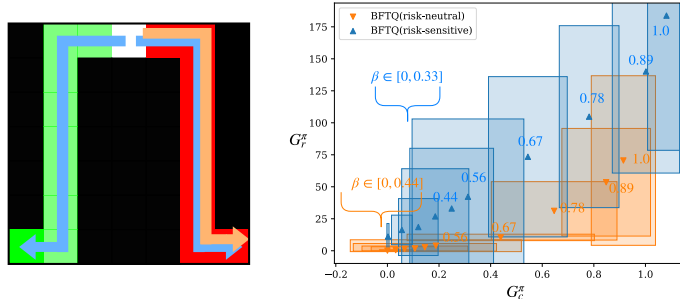


Figure 2: Trajectories (left) and performances (right) of two exploration strategies in the `corridors` environment.

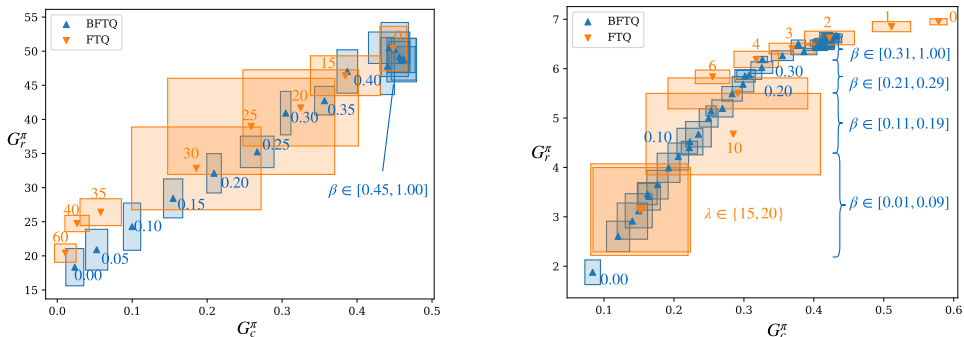


Figure 3: Performance comparison of  $\text{FTQ}(\lambda)$  and BFTQ on slot-filling (left) and highway-env(right)

trade-off curve. In both cases, BFTQ performs almost as well as  $\text{FTQ}(\lambda)$  despite only requiring a single model. All budgets are well-respected on slot-filling, but on highway-env we can observe an underestimation of  $Q_c$ , since e.g.  $\mathbb{E}[G_c|\beta = 0] \simeq 0.1$ . This underestimation can be a consequence of two approximations: the use of the sampling operator  $\hat{\mathcal{T}}$  instead of the true environmental operator  $\mathcal{T}$ , and the use of the neural network function approximation  $Q_\theta$  instead of  $Q$ . Still, BFTQ provides a better control on the expected cost of the policy, than  $\text{FTQ}(\lambda)$ . In addition, BFTQ behaves more consistently than  $\text{FTQ}(\lambda)$  overall, as shown by its lower extra-seed variance. Examples of policy execution can be found in Appendix E.2.

## 6 Discussion

Algorithm 2 is an algorithm for solving large unknown BMDPs with continuous states. To the best of our knowledge, there is no algorithm in the current literature that combines all those features.

Algorithms have been proposed for CMDPs, which are less flexible sub-problems of the more general BMDP. When the environment parameters ( $P$ ,  $R_r$ ,  $R_c$ ) are known but not tractable, solutions relying on function approximation (Undurti *et al.*, 2011) or approximate linear programming (Poupart *et al.*, 2015) have been proposed. For unknown environments, online algorithms (Geibel and Wysotzki, 2005; Abe and others, 2010; Chow *et al.*, 2018; Achiam *et al.*, 2017) and a batch algorithm (Thomas *et al.*, 2015; Petrik *et al.*, 2016; Larocche and Trichelair, 2019; Le *et al.*, 2019) can solve large unknown CMDPs. Nevertheless, these approaches are limited in that the constraints thresholds are fixed prior to training and cannot be updated in real-time at policy execution to select the desired level of risk.

To our knowledge, there were only two ways of solving a BMDP. The first one is to approximate it with a finite set of CMDPs (e.g. see our  $\text{FTQ}(\lambda)$  baseline). The solutions of these CMDPs take the form of mixtures between two deterministic policies (Theorem 4.4, Beutler and Ross, 1985). To



obtain these policies, one needs to evaluate their expected cost by interacting with the environment<sup>5</sup>. Our solution not only requires one single model but also avoids any supplementary interaction.

The only other existing BMDP algorithm, and closest work to ours, is the Dynamic Programming algorithm proposed by Boutilier and Lu (2016). However, their work was established for finite state spaces only, and their solution relies heavily on this property. For instance, they enumerate and sort the next states  $s' \in \mathcal{S}$  by their expected value-by-cost, which could not be performed in a continuous state space  $\mathcal{S}$ . Moreover, they rely on the knowledge of the model  $(P, R_r, R_c)$ , and do not address the question of learning from interaction data.

## 7 Conclusion

The BMDP framework is a principled framework for safe decision making under uncertainty, which could be beneficial to the diffusion of Reinforcement Learning in industrial applications. However, BMDPs could so far only be solved in finite state spaces which limits their interest in many use-cases. We extend their definition to continuous states by introducing of a novel Dynamic Programming operator, that we build upon to propose a Reinforcement Learning algorithm. In order to scale to large problems, we provide an efficient implementation that exploits the structure of the value function and leverages tools from Deep Distributed Reinforcement Learning. We show that on two practical tasks our solution performs similarly to a baseline Lagrangian relaxation method while only requiring a single model to train, and relying on an interpretable  $\beta$  instead of the tedious tuning of the penalty  $\lambda$ .

## Acknowledgments

This work has been supported by CPER Nord-Pas de Calais/FEDER DATA Advanced data science and technologies 2015-2020, the French Ministry of Higher Education and Research, INRIA, and the French Agence Nationale de la Recherche (ANR). We thank Guillaume Gautier, Fabrice Clerot, Xuedong Shang for the helpful discussions and valuable insights.

## References

- Naoki Abe et al. Optimizing debt collections using constrained reinforcement learning. In *Special Interest Group on Knowledge Discovery and Data Mining (SIGKDD)*, 2010.
- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Eitan Altman. *Constrained Markov Decision Processes*. CRC Press, 1999.
- Frederick J. Beutler and Keith W. Ross. Optimal policies for controlled markov chains with a constraint. In *Journal of Mathematical Analysis and Applications*, 1985.
- Craig Boutilier and Tyler Lu. Budget allocation using weakly coupled, constrained markov decision processes. In *Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Nicolas Carrara, Romain Laroche, Jean-Léon Bouraoui, Tanguy Urvoy, and Olivier Pietquin. Safe transfer learning for dialogue applications. In *International Conference on Statistical Language and Speech Processing (SLSP)*, 2018.
- Senthilkumar Chandramohan, Matthieu Geist, and Olivier Pietquin. Optimizing spoken dialogue management with fitted value iteration. In *Conference of the International Speech Communication Association (InterSpeech)*, 2010.
- Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-Sensitive and Robust Decision-Making: a CVaR Optimization Approach. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. In *Journal of Machine Learning Research (JMLR)*, 2018.

---

<sup>5</sup>More details are provided in Appendix D

- Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-Based Batch Mode Reinforcement Learning. In *Journal of Machine Learning Research (JMLR)*, 2005.
- Javier García and Fernando Fernández. A Comprehensive Survey on Safe Reinforcement Learning . In *Journal of Machine Learning Research (JMLR)*, 2015.
- Peter Geibel and Fritz Wysotzki. Risk-sensitive reinforcement learning applied to control under constraints. In *Journal of Artificial Intelligence Research (JAIR)*, 2005.
- Garud N. Iyengar. Robust Dynamic Programming . In *Mathematics of Operations Research*, 2005.
- Hatim Khouzaimi, Romain Laroche, and Fabrice. Lefevre. Optimising turn-taking strategies with reinforcement learning. . In *Special Interest Group on Discourse and Dialogue (SIGDIAL)*, 2015.
- Romain Laroche and Rémi Trichelair, Paul and Tachet des Combes. Safe policy improvement with baseline bootstrapping. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Hoang M. Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Edouard Leurent, Yann Blanco, Denis Efimov, and Odalric-Ambrym Maillard. Approximate Robust Control of Uncertain Dynamical Systems . In *Neural Information Processing Systems (NeurIPS), Workshop on Machine Learning for Intelligent Transportation Systems*, 2018.
- Lihong Li, Jason D. Williams, and Suhrud Balakrishnan. Reinforcement learning for dialog management using least-squares policy iteration and fast feature selection. In *Conference of the International Speech Communication Association (InterSpeech)*, 2009.
- Chunming Liu, Xin Xu, and Dewen Hu. Multiobjective Reinforcement Learning: A Comprehensive Overview. In *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2014.
- David G. Luenberger. *Investment science*. Oxford University Press, Incorporated, 2013.
- H. Mausser and D. Rosen. Beyond VaR: from measuring risk to managing risk. In *Proceedings of the IEEE Conference on Computational Intelligence for Financial Engineering*, 2003.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-mare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Arnab Nilim and Laurent El Ghaoui. Robust Control of Markov Decision Processes with Uncertain Transition Matrices . In *Operations Research*, 2005.
- Mohammad Petrik, Marek Ghavamzadeh, , and Yinlam Chow. Safe policy improvement by minimizing robust baseline regret. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Olivier Pietquin, Matthieu Geist, Senthilkumar Chandramohan, and Hervé Frezza-Buet. Sample-efficient batch reinforcement learning for dialogue management optimization. *ACM Transactions on Speech and Language Processing (TSLP)*, 7(3):7, 2011.
- Pascal Poupart, Aarti Malhotra, Pei Pei, Kee-Eung Kim, Bongseok Goh, and Michael Bowling. Approximate linear programming for constrained partially observable markov decision processes. In *Proceedings of the Association for the Advancement of Artificial Intelligence Conference (AAAI)*, 2015.
- Martin Riedmiller. Neural fitted Q iteration - First experiences with a data efficient neural Reinforcement Learning method. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2005.
- Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. In *Journal of Artificial Intelligence Research (JAIR)*, 2013.
- Aviv Tamar, Dotan Di Castro , and Shie Mannor. Policy Gradients with Variance Related Risk Criteria . In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.

- Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High confidence policy improvement. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- Aditya Undurti, Alborz Geramifard, and Jonathan P. How. Function approximation for continuous constrained mdps. In *Tech Report*, 2011.
- Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust markov decision processes. In *Mathematics of Operations Research*, 2013.

# Appendices

**Outline** This paper gathers all the supplementary material and goes as follows: Appendix A details all the proofs of the main results. Appendix B and Appendix C recall respectively the scalable BFTQ algorithm and the risk-sensitive exploration procedure. Appendix D describes a naive alternative to BFTQ based on Lagrangian Relaxation. The Appendix E assembles all the assets for visualising and reproducing the experiments, including visualisations of policy executions, algorithms and environment parameters, and instructions for executing the attached source code. Finally we fill the Machine Learning Reproducibility Checklist and we justify each statement in Appendix F.

## A Proofs of Main Results

### A.1 Proposition 1

*Proof.* This proof is the same as that in classical multi-objective MDPs.

$$\begin{aligned}
V^\pi(\bar{s}) &\stackrel{\text{def}}{=} \mathbb{E}[G^\pi \mid \bar{s}_0 = \bar{s}] \\
&= \sum_{\bar{a} \in \bar{\mathcal{A}}} \mathbb{P}(\bar{a}_0 = \bar{a} \mid \bar{s}_0 = \bar{s}) \mathbb{E}[G^\pi \mid \bar{s}_0 = \bar{s}, \bar{a}_0 = \bar{a}] \\
&= \sum_{\bar{a} \in \bar{\mathcal{A}}} \pi(\bar{a} \mid \bar{s}) Q^\pi(\bar{s}, \bar{a}) \\
Q^\pi(\bar{s}, \bar{a}) &\stackrel{\text{def}}{=} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R(\bar{s}_t, \bar{a}_t) \mid \bar{s}_0 = \bar{s}, \bar{a}_0 = \bar{a}\right] \\
&= R(\bar{s}, \bar{a}) + \sum_{\bar{s}' \in \bar{\mathcal{S}}} \mathbb{P}(\bar{s}_1 = \bar{s}' \mid \bar{s}_0 = \bar{s}, \bar{a}_0 = \bar{a}) \cdot \mathbb{E}\left[\sum_{t=1}^{\infty} \gamma^t R(\bar{s}_t, \bar{a}_t) \mid \bar{s}_1 = \bar{s}'\right] \\
&= R(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \bar{\mathcal{S}}} \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R(\bar{s}_t, \bar{a}_t) \mid \bar{s}_0 = \bar{s}'\right] \\
&= R(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \bar{\mathcal{S}}} \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) V^\pi(\bar{s}')
\end{aligned}$$

**Contraction of  $\mathcal{T}^\pi$ :** Let  $\pi \in \Pi, Q_1, Q_2 \in (\mathbb{R}^2)^{\bar{\mathcal{S}}\bar{\mathcal{A}}}$ .

$$\begin{aligned}
\forall \bar{s} \in \bar{\mathcal{S}}, \bar{a} \in \bar{\mathcal{A}}, \quad |\mathcal{T}^\pi Q_1(\bar{s}, \bar{a}) - \mathcal{T}^\pi Q_2(\bar{s}, \bar{a})| &= \left| \gamma \mathbb{E}_{\substack{\bar{s}' \sim \bar{P}(\bar{s}' \mid \bar{s}, \bar{a}) \\ \bar{a}' \sim \pi(\bar{a}' \mid \bar{s}')}} Q_1(\bar{s}', \bar{a}') - Q_2(\bar{s}', \bar{a}') \right| \\
&\leq \gamma \|Q_1 - Q_2\|_\infty
\end{aligned}$$

Hence,  $\|\mathcal{T}^\pi Q_1 - \mathcal{T}^\pi Q_2\|_\infty \leq \gamma \|Q_1 - Q_2\|_\infty$

According to the Banach fixed point theorem,  $\mathcal{T}^\pi$  admits a unique fixed point. It can be easily verified that  $Q^\pi$  is indeed this fixed point by combining the two Bellman Expectation equations (5). □

### A.2 Theorem 1

*Proof.* Let  $\bar{s}, \bar{a} \in \bar{\mathcal{A}} \times \bar{\mathcal{S}}$ . For this proof, we consider potentially non-stationary policies  $\pi = (\rho, \pi')$ , with  $\rho \in \mathcal{M}(\bar{\mathcal{A}}), \pi' \in \mathcal{M}(\bar{\mathcal{A}})^{\mathbb{N}}$ . The results will apply to the particular case of stationary optimal policies, when they exist.

$$Q_r^*(\bar{s}, \bar{a}) = \max_{\rho, \pi'} Q_r^{\rho, \pi'}(\bar{s}', \bar{a}') \quad (13)$$

$$= \max_{\rho, \pi'} R_r(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \mathcal{S}} P(\bar{s}' | \bar{s}, \bar{a}) V_r^{\rho, \pi'}(\bar{s}') \quad (14)$$

$$= R_r(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \mathcal{S}} P(\bar{s}' | \bar{s}, \bar{a}) \max_{\rho, \pi'} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \rho(\bar{a}' | \bar{s}') Q_r^{\pi'}(\bar{s}', \bar{a}') \quad (15)$$

$$= R_r(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \mathcal{S}} P(\bar{s}' | \bar{s}, \bar{a}) \max_{\rho} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \rho(\bar{a}' | \bar{s}') \max_{\pi' \in \Pi_a(\bar{s}')} Q_r^{\pi'}(\bar{s}', \bar{a}') \quad (16)$$

$$= R_r(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \mathcal{S}} P(\bar{s}' | \bar{s}, \bar{a}) \max_{\rho} \mathbb{E}_{\bar{a}' \sim \rho} Q_r^*(\bar{s}', \bar{a}') \quad (17)$$

where  $\pi = (\rho, \pi') \in \Pi_a(\bar{s})$  and  $\pi' \in \Pi_a(\bar{s}')$ .

This follows from:

(13). Definition of  $Q^*$ .

(14). Bellman Expectation expansion from Proposition 1.

(15). Marginalisation on  $\bar{a}'$ .

(16). • Trivially  $\max_{\pi' \in \Pi_a(\bar{s}')} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \cdot \leq \sum_{\bar{a}' \in \bar{\mathcal{A}}} \max_{\pi' \in \Pi_a(\bar{s})} \cdot$   
 • Let  $\bar{\pi} \in \arg \max_{\pi' \in \Pi_a(\bar{s}')} Q_r^{\pi'}(\bar{s}', \bar{a}')$ , then:

$$\begin{aligned} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \rho(\bar{a}' | \bar{s}') \max_{\pi' \in \Pi_a(\bar{s}')} Q_r^{\pi'}(\bar{s}', \bar{a}') &= \sum_{\bar{a}' \in \bar{\mathcal{A}}} \rho(\bar{a}' | \bar{s}') Q_r^{\bar{\pi}}(\bar{s}', \bar{a}') \\ &\leq \max_{\pi' \in \Pi_a(\bar{s}')} \sum_{\bar{a}' \in \bar{\mathcal{A}}} \rho(\bar{a}' | \bar{s}') Q_r^{\pi'}(\bar{s}', \bar{a}') \end{aligned}$$

(17). Definition of  $Q^*$ .

Moreover, the condition  $\pi = (\rho, \pi') \in \Pi_a(\bar{s})$  gives

$$\mathbb{E}_{\bar{a}' \sim \rho} Q_c^*(\bar{s}, \bar{a}) = \mathbb{E}_{\bar{a}' \sim \rho} Q_c^{\pi'}(\bar{s}, \bar{a}) = V_c^{\pi}(\bar{s}) \leq \beta$$

Consequently,  $\pi_{\text{greedy}}(\cdot; Q^*)$  belongs to the arg max of (17), and in particular:

$$Q_r^*(\bar{s}, \bar{a}) = r(\bar{s}, \bar{a}) + \gamma \sum_{\bar{s}' \in \mathcal{S}} P(\bar{s}' | \bar{s}, \bar{a}) \mathbb{E}_{\bar{a}' \sim \pi_{\text{greedy}}(\bar{s}', Q^*)} Q_r^*(\bar{s}', \bar{a}')$$

The same reasoning can be made for  $Q_c^*$  by replacing max operators by min, and  $\Pi_a$  by  $\Pi_r$ .  $\square$

### A.3 Proposition 2

*Proof.* Notice from the definitions of  $\mathcal{T}$  and  $\mathcal{T}^{\pi}$  in (11) and (6) that  $\mathcal{T}$  and  $\mathcal{T}^{\pi_{\text{greedy}}(\cdot; Q^*)}$  coincide on  $Q^*$ . Moreover, since  $Q^* = \mathcal{T}Q^*$  by Theorem 1, we have:  $\mathcal{T}^{\pi_{\text{greedy}}(\cdot; Q^*)}Q^* = \mathcal{T}Q^* = Q^*$ . Hence,  $Q^*$  is a fixed point of  $\mathcal{T}^{\pi_{\text{greedy}}(\cdot; Q^*)}$ , and by Proposition 1 it must be equal to  $Q^{\pi_{\text{greedy}}(\cdot; Q^*)}$ .

To show the same result for  $V^*$ , notice that

$$V^{\pi_{\text{greedy}}(Q^*)}(\bar{s}) = \mathbb{E}_{\bar{a} \sim \pi_{\text{greedy}}(Q^*)} Q^{\pi_{\text{greedy}}(Q^*)}(\bar{s}, \bar{a}) = \mathbb{E}_{\bar{a} \sim \pi_{\text{greedy}}(Q^*)} Q^*(\bar{s}, \bar{a})$$

By applying the definitions of  $Q^*$  and  $\pi_{\text{greedy}}$ , we recover the definition of  $V^*$ .  $\square$




Figure 4: Representation of  $Q_\varepsilon^1$  (blue) and  $Q_\varepsilon^2$  (yellow)

#### A.4 Theorem 2

*Proof.* In the trivial case  $|\mathcal{A}| = 1$ , there exists only one policy  $\pi$  and  $\mathcal{T} = \mathcal{T}^\pi$ , which is a contraction by Proposition 1.

In the general case  $|\mathcal{A}| \geq 2$ , we can build the following counter-example:

Let  $(\mathcal{S}, \mathcal{A}, P, R_r, R_c)$  be a BMDP. For any  $\varepsilon > 0$ , we define  $Q_\varepsilon^1$  and  $Q_\varepsilon^2$  as:

$$Q_\varepsilon^1(\bar{s}, \bar{a}) = \begin{cases} (0, 0), & \text{if } a = a_0 \\ \left(\frac{1}{\gamma}, \varepsilon\right), & \text{if } a \neq a_0 \end{cases}$$

$$Q_\varepsilon^2(\bar{s}, \bar{a}) = \begin{cases} (0, \varepsilon), & \text{if } a = a_0 \\ \left(\frac{1}{\gamma}, 2\varepsilon\right), & \text{if } a \neq a_0 \end{cases}$$

Then,  $\|Q_1 - Q_2\|_\infty = \varepsilon$ .  $Q_\varepsilon^1$  and  $Q_\varepsilon^2$  are represented in Figure 4.

But for  $\bar{a} = (a, \beta_a)$  with  $\beta_a = \varepsilon$ , we have:

$$\begin{aligned} \|\mathcal{T}Q_\varepsilon^1(\bar{s}, \bar{a}) - \mathcal{T}Q_\varepsilon^2(\bar{s}, \bar{a})\|_\infty &= \gamma \left\| \mathbb{E}_{\bar{s}' \sim \bar{P}(\bar{s}'|\bar{s}, \bar{a})} \mathbb{E}_{\bar{a}' \sim \pi_{\text{greedy}}(Q_\varepsilon^1)} Q_\varepsilon^1(\bar{s}', \bar{a}') - \mathbb{E}_{\bar{a}' \sim \pi_{\text{greedy}}(Q_\varepsilon^2)} Q_\varepsilon^2(\bar{s}', \bar{a}') \right\|_\infty \\ &= \gamma \left\| \mathbb{E}_{\bar{s}' \sim \bar{P}(\bar{s}'|\bar{s}, \bar{a})} \left(\frac{1}{\gamma}, \varepsilon\right) - (0, \varepsilon) \right\|_\infty \\ &= \gamma \frac{1}{\gamma} = 1 \end{aligned}$$

Hence,

$$\|\mathcal{T}Q_\varepsilon^1 - \mathcal{T}Q_\varepsilon^2\|_\infty \geq 1 = \frac{1}{\varepsilon} \|Q_1 - Q_2\|_\infty$$

In particular, there does not exist  $L > 0$  such that:

$$\forall Q_1, Q_2 \in (\mathbb{R}^2)^{\bar{\mathcal{S}}\bar{\mathcal{A}}}, \|\mathcal{T}Q^1 - \mathcal{T}Q^2\|_\infty \leq L \|Q^1 - Q^2\|_\infty$$

In other words,  $\mathcal{T}$  is not a contraction for  $\|\cdot\|_\infty$ .  $\square$

## A.5 Remark 2

*Proof.* We now study the contractivity of  $\mathcal{T}$  when restricted to the functions of  $\mathcal{L}_\gamma$  defined as follows:

$$\mathcal{L}_\gamma = \left\{ Q \in (\mathbb{R}^2)^{\overline{\mathcal{S}\mathcal{A}}} \text{ s.t. } \exists L < \frac{1}{\gamma} - 1 : \forall \bar{s} \in \overline{\mathcal{S}}, \bar{a}_1, \bar{a}_2 \in \overline{\mathcal{A}}, \right. \\ \left. |Q_r(\bar{s}, \bar{a}_1) - Q_r(\bar{s}, \bar{a}_2)| \leq L |Q_c(\bar{s}, \bar{a}_1) - Q_c(\bar{s}, \bar{a}_2)| \right\} \quad (18)$$

That is, for all state  $\bar{s}$ , the set  $Q(\bar{s}, \overline{\mathcal{A}})$  plot in the  $(Q_c, Q_r)$  plane must be the *graph* of a  $L$ -Lipschitz function, with  $L < 1/\gamma - 1$ .

We impose such structure for the following reason: the counter-example presented above prevented contraction because it was a pathological case in which the slope of  $Q$  can be arbitrary large. As a consequence, when solving  $Q_r^*$  such that  $Q_c^* = \beta$ , a vertical slice of a  $\|\cdot\|_\infty$  ball around  $Q_1$  (which must contain  $Q_2$ ) can be arbitrary large as well.

This sketch of proof makes use of insights detailed in the proof of Proposition 3, which we recommend the reader to consult first.

We denote  $\mathcal{B}(Q, R)$  the ball of centre  $Q$  and radius  $R$  for the  $\|\cdot\|_\infty$ -norm:

$$\mathcal{B}(Q, R) = \{Q' \in (\mathbb{R}^2)^{\overline{\mathcal{S}\mathcal{A}}} : \|Q - Q'\|_\infty \leq R\}$$

We give the three main steps required to show that  $\mathcal{T}$  restricted to  $\mathcal{L}_\gamma$  is a contraction. Given  $Q^1, Q^2 \in \mathcal{L}_\gamma$ , show that:

1.  $Q^2 \in \mathcal{B}(Q^1, R) \implies \mathcal{F}^2 \in \mathcal{B}(\mathcal{F}^1, R), \forall \bar{s} \in \overline{\mathcal{S}}$ , where  $\mathcal{F}$  is the top frontier of the convex hull of undominated points, as defined in Appendix A.6.
2.  $Q \in \mathcal{L}_\gamma \implies \mathcal{F}$  is the graph of a  $L$ -Lipschitz function,  $\forall \bar{s} \in \overline{\mathcal{S}}$ .
3. taking the slice  $Q_c = \beta$  of a ball  $\mathcal{B}(\mathcal{F}, R)$  with  $\mathcal{F}$   $L$ -Lipschitz results in an interval on  $Q_r$  of range at most  $(L + 1)R$

These three steps will allow us to control  $Q_r^{2*} - Q_r^{1*}$  as a function of  $R = \|Q^2 - Q^1\|_\infty$ .

**Step 1:** we want to show that if  $Q^1$  and  $Q^2$  are close, then  $\mathcal{F}^1$  and  $\mathcal{F}^2$  are close as well in the following sense:

$$\mathcal{F}^2 \in \mathcal{B}(\mathcal{F}^1, R) \iff d(\mathcal{F}^1, \mathcal{F}^2) \leq R \iff \max_{q^2 \in \mathcal{F}^2} \min_{q^1 \in \mathcal{F}^1} \|q^2 - q^1\|_\infty \leq R \quad (19)$$

Assume  $Q^2 \in \mathcal{B}(Q^1, R)$ . We start by showing this result for  $\mathcal{C}^2(Q^{1-})$  and  $\mathcal{C}^2(Q^{2-})$  as defined in Appendix A.6:

Let  $\bar{s} \in \overline{\mathcal{S}}$  and  $q^2 \in \mathcal{C}^2(Q^{2-}), \exists \lambda \in [0, 1], \bar{a}_1, \bar{a}_2 \in \overline{\mathcal{A}} : q^2 = (1 - \lambda)Q^2(\bar{s}, \bar{a}_1) + \lambda Q^2(\bar{s}, \bar{a}_2)$ . Define  $q^1 = (1 - \lambda)Q^1(\bar{s}, \bar{a}_1) + \lambda Q^1(\bar{s}, \bar{a}_2)$ . Then

$$\begin{aligned} \|q^2 - q^1\|_\infty &= \|(1 - \lambda)(Q^2(\bar{s}, \bar{a}_1) - Q^1(\bar{s}, \bar{a}_1)) + \lambda(Q^2(\bar{s}, \bar{a}_2) - Q^1(\bar{s}, \bar{a}_2))\|_\infty \\ &\leq (1 - \lambda)\|Q^2(\bar{s}, \bar{a}_1) - Q^1(\bar{s}, \bar{a}_1)\|_\infty + \lambda\|Q^2(\bar{s}, \bar{a}_2) - Q^1(\bar{s}, \bar{a}_2)\|_\infty \\ &\leq (1 - \lambda)R + \lambda R = R \end{aligned}$$

It remains to show that when taking the top frontiers of the convex sets  $\mathcal{C}^2(Q^{1-})$  and  $\mathcal{C}^2(Q^{2-})$ , they remain at a distance of at most  $R$ .

This is illustrated in Figure 5: given a function  $Q^1$ , we show the locus  $\mathcal{B}(Q_1, R)$  of  $Q^2$ . We then draw  $\mathcal{F}^1$  the top frontier of the convex hull of  $Q^1$  and alongside the locus of all possible  $\mathcal{F}^2$ , which belong to a ball  $\mathcal{B}(\mathcal{F}^1, R)$ .

**Step 2:** We want to show that if  $Q \in \mathcal{L}_\gamma$ ,  $\mathcal{F}$  is the graph of an  $L$ -Lipschitz function:

$$\forall q^1, q^2 \in \mathcal{F}, |q_r^2 - q_r^1| \leq |q_c^2 - q_c^1| \quad (20)$$

Let  $Q \in \mathcal{L}_\gamma$  and  $\bar{s} \in \overline{\mathcal{S}}$ ,  $\mathcal{F}$  the corresponding top frontier of convex hull. For all  $q^1, q^2 \in \mathcal{F}, \exists \lambda, \mu \in [0, 1], q^{11}, q^{12}, q^{21}, q^{22} \in Q(\bar{s}, \overline{\mathcal{A}})$  such that  $q^1 = (1 - \lambda)q^{11} + \lambda q^{12}$  and  $q^2 = (1 - \mu)q^{21} + \mu q^{22}$ .






Figure 5: We represent the range of possible solutions  $Q_r^{2*}$  for any  $Q^2 \in \mathcal{B}(Q^1)$ , given  $Q_1 \in \mathcal{L}_\lambda$

Without loss of generality, we can assume  $q_c^{11} \leq q_c^{12}$  and  $q_c^{21} \leq q_c^{22}$ . We also consider the worst case in terms of maximum  $q_r$  deviation:  $q_c^{12} \leq q_c^{21}$ . Then the maximum increment  $q_r^2 - q_r^1$  is:

$$\begin{aligned}
\|q_r^2 - q_r^1\| &\leq \|q_r^{12} - q_r^1\| + \|q_r^{21} - q_r^{12}\| + \|q_r^2 - q_r^{21}\| \\
&= (1 - \lambda)\|q_r^{12} - q_r^1\| + \|q_r^{21} - q_r^{12}\| + \mu\|q_r^{22} - q_r^{21}\| \\
&\leq (1 - \lambda)L\|q_c^{12} - q_c^1\| + L\|q_c^{21} - q_c^{12}\| + \mu L\|q_c^{22} - q_c^{21}\| \\
&= L\|q_c^{12} - q_c^1\| + L\|q_c^{21} - q_c^{12}\| + L\|q_c^2 - q_c^{21}\| \\
&= L\|q_c^2 - q_c^1\|
\end{aligned}$$

This can also be seen in Figure 5: the maximum slope of the  $\mathcal{F}^1$  is lower than the maximum slope between two points of  $Q^1$ .

**Step 3:** Let  $\mathcal{F}_1$  be a L-Lipschitz set as defined in (20), and consider a ball  $\mathcal{B}(\mathcal{F}_1, R)$  around it as defined in (19).

We want to bound the optimal reward value  $Q_r^{2*}$  under constraint  $Q_c^{2*} = \beta$  (regular case in Appendix A.6 where the constraint is saturated), for any  $\mathcal{F}^2 \in \mathcal{B}(\mathcal{F}_1, R)$ . This quantity is represented as a red double-ended arrow in Figure 5.

Because we are only interested in what happens locally at  $Q_c = \beta$ , we can zoom in on Figure 5 and only consider a thin  $\varepsilon$ -section around  $\beta$ . In the limit  $\varepsilon \rightarrow 0$ , this section becomes the tangent to  $\mathcal{F}^1$  at  $Q_c^1 = \beta$ . It is represented in Figure 6, from which we derive a geometrical proof:




Figure 6: We represent a section  $[\beta - \varepsilon, \beta + \varepsilon]$  of  $\mathcal{F}^1$  and  $\mathcal{B}(\mathcal{F}^1, R)$ . We want to bound the range of  $Q_r^{2*}$ .

$$\begin{aligned}
\Delta Q_r^{2*} &= b + c \\
&\leq La + c && (\mathcal{F}^1 \text{ } L\text{-Lipschitz}) \\
&= 2LR + 2R = 2R(L + 1)
\end{aligned}$$

Hence,

$$|Q_r^{2*} - Q_r^{1*}| \leq \frac{\Delta Q_r^{2*}}{2} = R(L + 1)$$

and  $Q_c^{1*} = Q_c^{2*} = \beta$ . Consequently,  $\|Q^{2*} - Q^{1*}\|_\infty \leq (L + 1)R$

For completeness, the edge case in Appendix A.6 should be considered as well.

### Wrapping it up:

We've shown that for any  $Q^1, Q^2 \in \mathcal{L}_\gamma$ , and all  $\bar{s} \in \bar{\mathcal{S}}, \mathcal{F}^2 \in \mathcal{B}(\mathcal{F}^1, \|Q^2 - Q^1\|_\infty)$  and  $\mathcal{F}^1$  is the graph of a  $L$ -Lipschitz function with  $L < 1/\gamma - 1$ . Moreover, the solutions of  $\pi_{\text{greedy}}(Q^1)$  and  $\pi_{\text{greedy}}(Q^2)$  at  $\bar{s}$  are such that  $\|Q^{2*} - Q^{1*}\|_\infty \leq (L + 1)\|Q^2 - Q^1\|_\infty$ .

Hence, for all  $\bar{a}$ ,

$$\begin{aligned}
\|\mathcal{T}Q^1(\bar{s}, \bar{a}) - \mathcal{T}Q^2(\bar{s}, \bar{a})\|_\infty &= \gamma \left\| \mathbb{E}_{\bar{s}' \sim \bar{P}(\bar{s}, \bar{a})} \mathbb{E}_{\bar{a}' \sim \pi_{\text{greedy}}(Q^1)} Q^1(\bar{s}', \bar{a}') - \mathbb{E}_{\bar{a}' \sim \pi_{\text{greedy}}(Q^2)} Q^2(\bar{s}', \bar{a}') \right\|_\infty \\
&= \gamma \|Q^{2*} - Q^{1*}\|_\infty \\
&\leq \gamma(L + 1)\|Q^2 - Q^1\|_\infty
\end{aligned}$$

Taking the sup on  $\bar{\mathcal{S}\mathcal{A}}$ ,

$$\|\mathcal{T}Q^1 - \mathcal{T}Q^2\|_\infty \leq \gamma(L + 1)\|Q^1 - Q^2\|_\infty$$

with  $\gamma(L + 1) < 1$ . As a conclusion,  $\mathcal{T}$  is a  $\gamma(L + 1)$ -contraction on  $\mathcal{L}_\gamma$ .  $\square$

## A.6 Proposition 3

**Definition 2.** Let  $A$  be a set, and  $f$  a function defined on  $A$ . We define:

- Convex hull of  $A$ :  $\mathcal{C}(A) = \{\sum_{i=1}^p \lambda_i a_i : a_i \in A, \lambda_i \in \mathbb{R}^+, \sum_{i=1}^p \lambda_i = 1, p \in \mathbb{N}\}$
- Convex edges of  $A$ :  $\mathcal{C}^2(A) = \{\lambda a_1 + (1 - \lambda)a_2 : a_1, a_2 \in A, \lambda \in [0, 1]\}$
- Dirac distributions of  $A$ :  $\delta(A) = \{\delta(a - a_0) : a_0 \in A\}$
- Image of  $A$  by  $f$ :  $f(A) = \{f(a) : a \in A\}$

*Proof.* Let  $\bar{s} = (s, \beta) \in \bar{\mathcal{S}}$  and  $Q \in (\mathbb{R}^2)^{\bar{\mathcal{S}\mathcal{A}}}$ . We recall the definition of  $\pi_{\text{greedy}}$ :

$$\pi_{\text{greedy}}(\bar{a}|\bar{s}; Q) \in \arg \min_{\rho \in \Pi_r^Q} \mathbb{E}_{\bar{a} \sim \rho} Q_c(\bar{s}, \bar{a}) \quad (12a)$$

$$\text{where } \Pi_r^Q = \arg \max_{\rho \in \mathcal{M}(\bar{\mathcal{A}})} \mathbb{E}_{\bar{a} \sim \rho} Q_r(\bar{s}, \bar{a}) \quad (12b)$$

$$\text{s.t. } \mathbb{E}_{\bar{a} \sim \rho} Q_c(\bar{s}, \bar{a}) \leq \beta \quad (12c)$$

Note that any policy in the arg min in (12a) is suitable to compute  $\mathcal{T}$ . We first reduce the set of candidate optimal policies. Consider the problem described in (12b),(12c): it can be seen as a single-step CMDP problem with reward  $R_r = Q_r$  and cost  $R_c = Q_c$ . By (Theorem 4.4 Beutler and Ross, 1985), we know that the solutions are mixtures of two deterministic policies. Hence, we can replace  $\mathcal{M}(\bar{\mathcal{A}})$  by  $\mathcal{C}^2(\delta(\bar{\mathcal{A}}))$  in (12b).

Moreover, remark that:

$$\begin{aligned} \{\mathbb{E}_{\bar{a} \sim \rho} Q(\bar{s}, \bar{a}) : \rho \in \mathcal{C}^2(\delta(\bar{\mathcal{A}}))\} &= \{\mathbb{E}_{\bar{a} \sim \rho} Q(\bar{s}, \bar{a}) : \rho = (1 - \lambda)\delta(\bar{a} - \bar{a}_1) + \lambda\delta(\bar{a} - \bar{a}_2), \bar{a}_1, \bar{a}_2 \in \bar{\mathcal{A}}, \lambda \in [0, 1]\} \\ &= \{(1 - \lambda)Q(\bar{s}, \bar{a}_1) + \lambda Q(\bar{s}, \bar{a}_2), \bar{a}_1, \bar{a}_2 \in \bar{\mathcal{A}}, \lambda \in [0, 1]\} \\ &= \mathcal{C}^2(Q(\bar{s}, \bar{\mathcal{A}})) \end{aligned}$$

Hence, the problem (12b), (12c) has become:

$$\tilde{\Pi}_r^Q = \arg \max_{(q_r, q_c) \in \mathcal{C}^2(Q(\bar{s}, \bar{\mathcal{A}}))} q_r \quad \text{s.t.} \quad q_c \leq \beta$$

and the solution of  $\pi_{\text{greedy}}$  is  $q^* = \arg \min_{q \in \tilde{\Pi}_r^Q} q_c$ .

The original problem in the space of actions  $\bar{\mathcal{A}}$  is now expressed in the space of values  $Q(\bar{s}, \bar{\mathcal{A}})$  (which is why we use = instead of  $\in$  before  $\arg \min$  here).

We further restrict the search space of  $q^*$  following two observations:

1.  $q^*$  belongs to the *undominated* points  $\mathcal{C}^2(Q^-)$ :

$$Q^+ = \{(q_c, q_r) : q_c > q_c^\pm = \min_{q^+} q_c^+ \text{ s.t. } q^+ \in \arg \max_{q \in Q(\bar{s}, \bar{\mathcal{A}})} q_r\} \quad (22)$$

$$Q^- = Q(\bar{s}, \bar{\mathcal{A}}) \setminus Q^+ \quad (23)$$

Denote  $q^* = (1 - \lambda)q^1 + \lambda q^2$ , with  $q^1, q^2 \in Q(\bar{s}, \bar{\mathcal{A}})$ . There are three possible cases:

- (a)  $q^1, q^2 \notin Q^-$ . Then  $q_c^* = (1 - \lambda)q_c^1 + \lambda q_c^2 > q_c^\pm$ . But then  $q_c^\pm < q_c^* \leq \beta$  so  $q^\pm \in \tilde{\Pi}_r^Q$  with a strictly lower  $q_c$  than  $q^*$ , which contradicts the  $\arg \min$ .
- (b)  $q^1 \in Q^-, q^2 \notin Q^-$ . But then consider the mixture  $q^\top = (1 - \lambda)q^1 + \lambda q^\pm$ . Since  $q_r^\pm \geq q_r^2$  and  $q_r^\pm < q_r^1$ , we also have  $q_r^\top \geq q_r^*$  and  $q_c^\top < q_c^*$ , which also contradicts the  $\arg \min$ .
- (c)  $q^1, q^2 \in Q^-$  is the only remaining possibility.

2.  $q^*$  belongs to the *top frontier*  $\mathcal{F}$ :

$$\mathcal{F}_Q = \{q \in \mathcal{C}^2(Q^-) : \nexists q' \in \mathcal{C}^2(Q^-) : q_c = q'_c \text{ and } q_r < q'_r\}$$

Trivially, otherwise  $q'$  would be a better candidate than  $q^*$ .

Let us characterise this frontier  $\mathcal{F}$ . It is both:

1. the *graph of a non-decreasing function*:  $\forall q^1, q^2 \in \mathcal{F}$  such that  $q_c^1 \leq q_c^2$  then  $q_r^1 \leq q_r^2$ .  
By contradiction, if we had  $q_r^1 > q_r^2$ , we could define  $q^\top = (1 - \lambda)q^1 + \lambda q^2$  where  $q^\pm$  is the dominant point as defined in (22). By choosing  $\lambda = (q_c^2 - q_c^1)/(q_c^\pm - q_c^1)$  such that  $q_c^\top = q_c^2$ , then since  $q_r^\pm \geq q_r^1 > q_r^2$  we also have  $q_r^\top > q_r^2$  which contradicts  $q^2 \in \mathcal{F}$ .
2. the *graph of a concave function*:  $\forall q^1, q^2, q^3 \in \mathcal{F}$  such that  $q_c^1 \leq q_c^2 \leq q_c^3$  with  $\lambda$  such that  $q_c^2 = (1 - \lambda)q_c^1 + \lambda q_c^3$ , then  $q_r^2 \geq (1 - \lambda)q_r^1 + \lambda q_r^3$ .  
Trivially, otherwise the point  $q^\top = (1 - \lambda)q^1 + \lambda q^3$  would verify  $q_c^\top = q_c^2$  and  $q_r^\top > q_r^2$ , which would contradict  $q^2 \in \mathcal{F}$ .

We denote  $\mathcal{F}_Q = \mathcal{F} \cap Q$ . Clearly,  $q^* \in \mathcal{C}^2(\mathcal{F}_Q)$ : let  $q^1, q^2 \in Q^-$  such that  $q^* = (1 - \lambda)q^1 + \lambda q^2$ . First,  $q^1, q^2 \in Q^- \subset \mathcal{C}^2(Q^-)$ . Then, by contradiction, if there existed  $q^{1'}$  or  $q^{2'}$  with equal  $q_c$  and strictly higher  $q_r$ , again we could build an admissible mixture  $q^\top = (1 - \lambda)q^{1'} + \lambda q^{2'}$  strictly better than  $q^*$ .

$q^*$  can be written as  $q^* = (1 - \lambda)q^1 + \lambda q^2$  with  $q^1, q^2 \in \mathcal{F}_Q$  and, without loss of generality,  $q_c^1 \leq q_c^2$ .

**Regular case:** there exists  $q^0 \in \mathcal{F}_Q$  such that  $q_c^0 \geq \beta$ .

Then  $q^1$  and  $q^2$  must flank the budget:  $q_c^1 \leq \beta \leq q_c^2$ . Indeed, by contradiction, if  $q_c^2 \geq q_c^1 > \beta$  then  $q_c^* > \beta$  which contradicts  $\Pi_r^Q$ . Conversely, if  $q_c^1 \leq q_c^2 < \beta$  then  $q_c^* < \beta \leq q_c^0$ , which would make  $q^*$

a worse candidate than  $q^\top = (1 - \lambda)q^* + \lambda q^0$  when  $\lambda$  is chosen such that  $q_c^\top = \beta$ , and contradict  $\Pi_r^Q$  again.

Because  $\mathcal{F}$  is the graph of a non-decreasing function,  $\lambda$  should be as high as possible, as long as the budget  $q^* \leq \beta$  is respected. We reach the highest  $q_r^*$  when  $q_c^* = \beta$ , that is:  $\lambda = (\beta - q_c^1)/(q_c^2 - q_c^1)$ .

It remains to show that  $q^1$  and  $q^2$  are two successive points in  $\mathcal{F}_Q$ :  $\nexists q \in \mathcal{F}_Q \setminus \{q^1, q^2\} : q_c^1 \leq q_c \leq q_c^2$ . Otherwise, as  $\mathcal{F}$  is the graph of a concave function, we would have  $q_r \geq (1 - \mu)q_r^1 + \mu q_r^2$ .  $q_r$  cannot be strictly greater than  $(1 - \mu)q_r^1 + \mu q_r^2$  which would contradict  $q^*$ , but it can still be equal, which means the tree points  $q, q^1, q^2$  are aligned. In fact, every points aligned with  $q^1$  and  $q^2$  can also be used to construct mixtures resulting in  $q^*$ , but among these solutions we can still choose  $q^1$  and  $q^2$  as the two points in  $\mathcal{F}_Q$  closest to  $q^*$ .

**Edge case:**  $\forall q \in \mathcal{F}_Q, q_c < \beta$ . Then  $q^* = \arg \max_{q \in \mathcal{F}} q_r = q^\pm = \arg \max_{q \in Q^-} q_r$

□

## B Risk-Sensitive Exploration

We recall the Risk-Sensitive Exploration algorithm in Algorithm 4

---

### Algorithm 4: Risk-sensitive exploration

---

**Data:** An environment, a BFTQ solver,  $W$  CPU workers

**Result:** A batch of transitions  $\mathcal{D}$

```

1  $\mathcal{D} \leftarrow \{\}$ 
2 for each intermediate batch do
3   split episodes between  $W$  workers
4   for each episode in batch do // run this loop on each worker in parallel
5     sample initial budget  $\beta \sim \mathcal{U}(\mathcal{B})$ .
6     while episode not done do
7       update  $\varepsilon$  from schedule.
8       sample  $z \sim \mathcal{U}([0, 1])$ .
9       if  $z < \varepsilon$  then sample  $(a, \beta_a) \sim \mathcal{U}(\Delta_{\mathcal{AB}})$ . // Explore
10      else sample  $(a, \beta_a) \sim \pi_{\text{greedy}}(a, \beta_a | s, \beta; Q^*)$ . // Exploit
11      append transition  $(s, \beta, a, \beta_a, R, C, s')$  to batch  $\mathcal{D}$ .
12      step episode budget  $\beta \leftarrow \beta_a$ 
13    end
14  end
15   $\pi_{\text{greedy}}(\cdot \sim; Q^*) \leftarrow \text{BFTQ}(\mathcal{D})$ .
16 end
17 return the batch of transitions  $\mathcal{D}$ 

```

---

## C Scalable Implementation of BFTQ

We recall the scalable version of BFTQ in Algorithm 5 and the architecture of the neural network Figure 7.

## D The Lagrangian Relaxation Baseline

As explained on Figure 8, the optimal deterministic policy can be obtained by a line-search on the Lagrange multiplier values  $\lambda$ .

Then, according to Beutler and Ross (1985, Theorem 4.4), the optimal policy is a randomised mixture of two deterministic policies: the safest deterministic policy that violates the constraint  $\pi_{\lambda-}$  and the riskier of the feasible ones  $\pi_{\lambda+}$ .

Fitted-Q (FTQ) (Ernst et al., 2005; Riedmiller, 2005) can be easily adapted for continuous states CMDP and BMDP through this methodology, but given the high variance it requires a lot of sim-




Figure 7: Neural Network for  $Q$ -functions approximation when  $S = \mathbb{R}^2$  and  $|\mathcal{A}| = 2$ .

---

**Algorithm 5:** Scalable BFTQ

---

**Data:**  $\mathcal{D}, \tilde{\mathcal{B}}$  a finite subset of  $\mathcal{B}$ ,  $\gamma$ , a model  $Q \in (\mathbb{R}^2)^{S\mathcal{A}}$ , a regression algorithm `fit`, a set of CPU workers  $W$

**Result:**  $Q^*$

```

1  $Q \leftarrow 0$ 
2  $X \leftarrow \{s_i, a_i, \beta_{a_i}\}_{i \in [0, |\mathcal{D}|]}$ 
3  $S' \leftarrow \{s'_i\}_{i \in [0, |\mathcal{D}|]}$ 
4 repeat
5   Evaluate  $Q(S', \mathcal{A}, \tilde{\mathcal{B}})$  in a single forward pass
6   Split  $\mathcal{D}$  among workers:  $\mathcal{D} = \cup_{w \in W} \mathcal{D}_w$ 
7   for  $w \in W$  do // Run in parallel
8     for  $(\cdot, \cdot, \beta_{a_i}, R_{r_i}, R_{c_i}, s'_i) \in \mathcal{D}$  do
9        $\mathcal{P} \leftarrow \{(Q_c(s'_i, \mathcal{A}, \tilde{\mathcal{B}}), Q_r(s'_i, \mathcal{A}, \tilde{\mathcal{B}}))\}$ 
10       $\mathcal{P}.\text{prune}()$  // Remove all dominated points
11       $\mathcal{H} \leftarrow \text{convex\_hull}(\mathcal{P}).\text{vertices}()$  // in cw order
12       $k \leftarrow \min\{k : \beta_i \geq q_c \text{ with } (q_c, q_r) = \mathcal{H}[k]\}$ 
13       $q_c^2, q_r^2, q_c^1, q_r^1 \leftarrow \mathcal{H}[k], \mathcal{H}[k-1]$ 
14       $p \leftarrow (\beta_{a_i} - q_a^1) / (q_c^2 - q_c^1)$ 
15       $Y_c^{w,i} \leftarrow R_{c_i} + \gamma((1-p)q_c^1 + pq_c^2)$ 
16       $Y_r^{w,i} \leftarrow R_{r_i} + \gamma((1-p)q_r^1 + pq_r^2)$ 
17    end
18  end
19  Join the results:  $Y \leftarrow \cup_{w \in W} (Y_c^w, Y_r^w)$ 
20   $Q \leftarrow \text{fit}(X, Y)$ 
21 until convergence

```

---




Figure 8: Calibration of a penalty multiplier according to the budget  $\beta$ . The optimal multiplier  $\lambda_{\text{avg}}^*$  is the smallest one to satisfy the budget constraint on average. Safer policies can also be selected according to the largest deviation from this mean cost.

ulations to get a proper estimate of the calibration curve. Our purpose is to avoid this calibration phase.

## E Experiments

### E.1 Examples of different exploration strategies

We compare two approaches for constructing a batch of samples. The animations from the html page [exploration.html](#) display the trajectories collected in each intermediate sub-batch. The first row corresponds to a classical risk neutral epsilon-greedy exploration policy while the second row showcases a risk-sensitive exploration strategy introduced in the paper. Each animation corresponds to a different seed.

### E.2 Examples of BFTQ policies executions

We display the evolution in the budgeted policy behaviour with respect to the budget on different environments. The policies have been learnt with a risk-sensitive exploration.

**Highway-Env** On the `highway-env`, the budgeted agents display a wide variety of behaviours. Animations are displayed on the html page [highway-env.html](#). When  $\beta = 1$ , the ego-vehicle drives in a very aggressive style: it immediately switches to the opposite lane and drives as fast as possible to pass slower vehicles, swiftly changing lanes to avoid incoming traffic. On the contrary when  $\beta = 0$ , the ego-vehicle is conservative: it stays on its lane and drives at a low velocity. With intermediate budgets such as  $\beta = 0.2$ , the agent sometimes decides to overtake its front vehicle but promptly steers back to its original lane afterwards.

### Slot-filling

**Remark on the slot-filling environment** When receiving an utterance, the system can either understand it ( $\mu = \mu_u$ ) or misunderstand it ( $\mu = \mu_m$ ) with a fixed probability called the sentence error rate  $ser$ . Then, the speech recognition score is simulated (Khouzaimi *et al.*, 2015):  $srs = (1 + \exp(-x))^{-1}$  with  $x \sim N(\mu, \sigma)$ . It's the confidence score of the natural language understanding module about the last utterance. Note that here are no recognition errors ( $ser = 0$  and  $srs = 1$ ) when the user provides information using the numeric pad.

In Table 1, we display two dialogues done with the same BFTQ policy on `slot-filling`. The policy is given two budgets to respect in expectation,  $\beta = 0$  and  $\beta = 0.5$ . For  $\beta = 0$ , one can see that the system never uses the `ask_num_pad` action. Instead, it uses `ask_oral`, an action subject to recognition errors. The system keeps asking for the same slot 2, because it has the lowest speech

turn	$\beta = 0$	$\beta = 0.5$
turn 0	valid slots : [0, 0, 0] srs : [ None None None ] system says ASK_ORAL(1) user says INFORM	valid slots : [0, 0, 0] srs : [ None None None ] system says ASK_ORAL(2) user says INFORM
turn 1	valid slots : [0, 0, 0] srs : [ None 0.48 None ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 0, 1] srs : [ None None 0.56 ] system says ASK_ORAL(0) user says INFORM
turn 2	valid slots : [0, 0, 0] srs : [ None 0.48 0.22 ] system says ASK_ORAL(0) user says INFORM	valid slots : [0, 0, 1] srs : [ 0.30 None 0.56 ] system says ASK_ORAL(1) user says INFORM
turn 3	valid slots : [0, 0, 0] srs : [ 0.62 0.48 0.22 ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 0, 1] srs : [ 0.30 0.54 0.56 ] system says ASK_ORAL(0) user says INFORM
turn 4	valid slots : [0, 0, 0] srs : [ 0.62 0.48 0.66 ] system says ASK_ORAL(1) user says INFORM	valid slots : [0, 0, 1] srs : [ 0.68 0.54 0.56 ] system says ASK_NUM_PAD(1) user says INFORM
turn 5	valid slots : [0, 1, 0] srs : [ 0.62 0.56 0.66 ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 1, 1] srs : [ 0.68 1.00 0.56 ] system says SUMMARIZE_AND_INFORM user says DENY_SUMMARIZE
turn 6	valid slots : [0, 1, 0] srs : [ 0.62 0.56 0.14 ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 1, 1] srs : [ 0.68 1.00 0.56 ] system says ASK_NUM_PAD(2) user says INFORM
turn 7	valid slots : [0, 1, 1] srs : [ 0.62 0.56 0.30 ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 1, 1] srs : [ 0.68 1.00 1.00 ] system says SUMMARIZE_AND_INFORM user says DENY_SUMMARIZE
turn 8	valid slots : [0, 1, 1] srs : [ 0.62 0.56 0.49 ] system says ASK_ORAL(2) user says INFORM	valid slots : [0, 1, 1] srs : [ 0.68 1.00 1.00 ] system says ASK_NUM_PAD(0) user hangs up !
turn 9	valid slots : [0, 1, 1] srs : [ 0.62 0.56 0.65 ] system says SUMMARIZE_AND_INFORM max size reached !	

Table 1: Two dialogues generated by a safe policy ( $\beta = 0$ ) on the left and a risky one ( $\beta = 0.5$ ) on the right.

recognition score. It eventually summarises the form to the user, but then reaches the maximum dialogue length and thus faces a dialogue failure. For  $\beta = 0.5$ , the system first asks in a safe way, with `ask_oral`. It may want to `ask_num_pad` if one of the speech recognition score is low. Then, the system proceeds to a confirmation of the slot values. If it is incorrect, the system continues the dialogue using unsafe the `ask_num_pad` action to be certain of the slot values.

**Corridors** Animations are displayed on the html page [corridors.html](#) for the `corridors` environment. When the budget is low, the agent takes the safest path on the left. When the budget increases, it gradually switches to the other lane, earning higher rewards but also costs. This gradual process could not be achieved with a deterministic policy as it would chose either one path or the other. Each animation corresponds to a different seed.



### E.3 Reproducibility

The following section displays environments and algorithms parameters and instructions to reproduce the exact same results displayed in Section 5.

#### E.3.1 Environments Parameters

All environments parameters are displayed in Table 2, Table 3 and Table 4.

**State-Space** The states  $s$  (from  $\bar{s} = (s, \beta)$ ) of the agent are described in the following:

- Corridors:  $s = (x, y)$  where  $x$  and  $y$  are the 2D coordinates of the agent.
- Slot-Filling:  $s = (\text{srs}, \text{min}, a_u, a_s, t)$  where srs is a vector of the speech recognition score for each slot, min is a one hot vector describing the minimum of the srs vector,  $a_u$  is a one hot vector of the last user dialogue act and  $a_s$  is the one hot vector of the last system dialogue act. Finally  $t \in [0, 1]$  is the fraction of the current turn with the maximum number of turns authorised.
- Highway-Env: the positions  $(x, y)$  and velocities  $(\dot{x}, \dot{y})$  of every vehicle on the road.

Parameter	Description	Value
-	Size of the environment	7 x 6
-	Standard deviation of the Gaussian noise applied to actions	(0.25,0.25)
H	Episode duration	9

Table 2: Parameters of Corridors

Parameter	Description	Value
ser	Sentence Error Rate	0.6
$\mu_m$	Gaussian mean for misunderstanding	-0.25
$\mu_u$	Gaussian mean for understanding	0.25
$\sigma$	Gaussian standard deviation	0.6
$p$	Probability of hang-up	0.25
H	Episode duration	10
-	Number of slots	3

Table 3: Parameters of Slot-Filling

Parameter	Description	Value
$N_v$	Number of vehicles	2 - 6
$\sigma_p$	Standard deviation of vehicles initial positions	100 m
$\sigma_v$	Standard deviation of vehicles initial velocities	3 m/s
H	Episode duration	15 s

Table 4: Parameters of highway-env

#### E.3.2 Algorithm parameters

All algorithm parameters are displayed in Table 5, Table 6 and Table 7.

**A note on the parameters search** We performed a shallow grid-search for the classic Neural-Network parameters. Most of the parameters don't have a strong influence on the results, however in the slot-filling environment, the choice of the regulation weight is decisive.

Parameters	BFTQ(risk-sensitive)	BFTQ(risk-neutral)
architecture	256x128x64	256x128x64
regularisation	0.001	0.001
activation	relu	relu
size beta encoder	3	3
initialisation	xavier	xavier
loss function	L2	L2
optimizer	adam	adam
learning rate	0.001	0.001
epoch (NN)	1000	5000
normalize reward	true	true
epoch (FTQ)	12	12
$\tilde{\beta}$	0:0.01:1	-
$\gamma$	1	1
$N =  \mathcal{D} $	5000	5000
$N_{\text{minibatch}}$	10	10
$N_{\text{seeds}}$	4	4
$N_{\text{test}}$	1000	1000
decay epsilon scheduling	0.001	0.001

Table 5: Algorithms parameters for Corridors

Parameters	BFTQ	FTQ
architecture	256x128x64	128x64x32
regularisation	0.0005	0.0005
activation	relu	relu
size beta encoder	50	-
initialisation	xavier	xavier
loss function	L2	L2
optimizer	adam	adam
learning rate	0.001	0.001
epoch (NN)	5000	5000
normalize reward	true	true
epoch (FTQ)	11	11
$\tilde{\beta}$	0:0.01:1	-
$\gamma$	1	1
$N =  \mathcal{D} $	5000	5000
$N_{\text{minibatch}}$	10	10
$N_{\text{seeds}}$	6	6
$N_{\text{test}}$	1000	1000
decay epsilon scheduling	0.001	0.001

Table 6: Algorithms parameters for Slot-Filling

Parameters	BFTQ	FTQ
architecture	256x128x64	128x64x32
regularisation	0.0005	0
activation	relu	relu
size beta encoder	50	-
initialisation	xavier	xavier
loss function	L2	L2
optimizer	adam	adam
learning rate	0.001	0.01
epoch (NN)	5000	400
normalize reward	true	true
epoch (FTQ)	15	15
$\tilde{\mathcal{B}}$	0:0.01:1	-
$\gamma$	0.9	0.9
$N =  \mathcal{D} $	10000	10000
$N_{\text{minibatch}}$	10	10
$N_{\text{seeds}}$	10	10
$N_{\text{test}}$	150	150
decay epsilon scheduling	0.0003	0.0003

Table 7: Algorithms parameters for Highway-Env

Listing 1: bash version

```
# Install highway-env
pip3 install --user git+https://github.com/eleurent/rl-agents
# Change python path to the path of this repository
export PYTHONPATH="code/scaling-up-brl"
# Navigate to budgeted-rl folder
cd code/scaling-up-brl/budgeted-rl/
# Run main script using any config file
# Choose the range of seeds you want to test on
python3 main/egreedy/main-egreedy.py config/slot-filling.json 0 6
python3 main/egreedy/main-egreedy.py config/corridors.json 0 4
python3 main/egreedy/main-egreedy.py config/highway-easy.json 0 10
```

Figure 9: Instructions to reproduce experiments

### E.3.3 Instructions for reproducibility

To reproduce the result displayed in Section 5, first install the following conventional libraries for python3: pycairo, numpy, scipy and pytorch. Then, execute the commands in Figure 9 on a Linux Operating System. The Graphic Processing Unit used for experiments is an NVIDIA GeForce GTX 1080 Ti and the Computational Processing Unit is an Intel Xeon E7.

## F The machine learning reproducibility checklist

For all models and algorithms presented, indicate if you include:

- A clear description of the mathematical setting, algorithm, and/or model:
  - **yes**, see Section 1, Section 2, Section 3, Appendix B and Appendix C.
- An analysis of the complexity (time, space, sample size) of any algorithm:
  - **yes**, see Section 4.3.
- A link to a downloadable source code, with specification of all dependencies, including external libraries:

- **yes**, see Appendix E.3.3 and the folder `code` in the supplementary material zip file.

For any theoretical claim, indicate if you include:

- A statement of the result:
  - **yes**, see Section 2 and Section 3.
- A clear explanation of any assumptions:
  - we make one assumption in Section 2. We assume the program is feasible for any state. If not, no algorithm would be able to solve it anyway.
- A complete proof of the claim:
  - **yes**, see Appendix A. We formulate a conjecture in Remark 2 but we provide a sketch of the proof in Appendix A.5.

For all figures and tables that present empirical results, indicate if you include:

- A complete description of the data collection process, including sample size:
  - **yes**, see Section 5 and Appendix E.3.2.
- A link to a downloadable version of the dataset or simulation environment:
  - **yes**, two environments are shipped with the supplementary material (in the `code` folder) and the third one is fetch from a public repository, see Appendix E.3.3 for details.
- An explanation of any data that were excluded, description of any pre-processing step:
  - it's **not applicable** as data comes from simulated environments, so pre-processing steps are not needed.
- An explanation of how samples were allocated for training / validation / testing:
  - it's **not applicable**. The complete dataset is used for training. There is no need for validation set. Testing is performed in the true environment as in classical online learning approaches.
- The range of hyper-parameters considered, method to select the best hyper-parameter configuration, and specification of all hyper-parameters used to generate results:
  - **yes**, see Appendix E.3.2.
- The exact number of evaluation runs:
  - **yes**, see  $N_{seeds}$  in the tables from Appendix E.3.2.
- A description of how experiments were run:
  - **yes**, see the two first paragraphs of Section 5.
- A clear definition of the specific measure or statistics used to report results:
  - **yes**, see Section 5.2.
- Clearly defined error bars:
  - **yes**, we plot 95% confidence intervals in all figures, see Section 5.2.
- A description of results with central tendency (e.g. mean) variation (e.g. stddev):
  - **yes**, we even observe less variability with our novel approach, see Section 5.2.
- A description of the computing infrastructure used:
  - The Graphic Processing Unit used for experiments is an NVIDIA GeForce GTX 1080 Ti and the Computational Processing Unit is an Intel Xeon E7.