On the Existence of Uniformly Optimal Policies for MDPs under Epistemic Uncertainty

Axel Benyamine[1,2*] **Julien Grand-Clément**[3] **Marek Petrik**[4]
Michael I. Jordan[2,5] **Alain Durmus**[1]

[1]CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris
[2]Inria Paris, École Normale Supérieure, PSL Research University
[3]ISOM Department, HEC Paris [4]University of New Hampshire
[5]University of California, Berkeley

Abstract

We investigate the existence of uniformly optimal stationary policies in Markov Decision Processes (MDPs) under epistemic (parameter) uncertainty that are optimal for any initial distribution. Our study is conducted within two complementary frameworks. For robust MDPs, we establish that (s- and sa-) rectangular models are equivalent to the existence of uniformly optimal stationary policies and provide several interesting counterexamples, including a simple r-rectangular instance with only history-dependent optimal policies. For ambiguity-averse MDPs where the uncertain transition kernel is a random variable, and the return is evaluated via a law-invariant risk measure, we show that uniformly optimal stationary policies can exist only for the essential supremum or essential infimum (under mild monotone continuity assumptions), thus excluding other risk measures such as CVaR, VaR, and the Entropic Risk Measure. We conclude our work by relating dynamic programming and the existence of uniformly optimal policies. One of the main takeaways of our impossibility results is the need for new models to go beyond uniformly optimal stationary policies.

1 Introduction

In this paper, we study the structural properties of Markov decision processes (MDPs) and their optimal policies [21]. A particularly interesting feature of the MDP framework for sequential decision-making is that optimal policies enjoy remarkably simple structural properties. First, optimal policies may be chosen to be stationary, that is, the action depends only on the current state and not on the full history. Second, they may also be taken to be deterministic, meaning that the same action is always selected whenever the system visits the same state. Finally, there exists a stationary deterministic policy that is optimal for every initial distribution. Such a policy is known as *uniformly optimal*. These structural properties simplify both the learning and implementation of policies, while also supporting interpretability in high-stakes applications such as vehicle routing [19] and healthcare [8, 25].

The existence of uniformly optimal stationary deterministic policies is well understood only in *nominal MDPs*, which assume a known and fixed MDP model. In practice, the MDP model is subject to estimation errors, often referred to as *epistemic uncertainty*, which must be accounted for in any assertion of optimality. Many frameworks have been proposed to study this issue. Robust MDPs (RMDPs) optimize for the worst-case parameters within an uncertainty set [28]. There exist several different models of uncertainty sets in the literature, e.g., s-rectangularity [28], sa-rectangularity [13, 20], r-rectangularity [9], and k-rectangularity [18]. Ambiguity-averse MDPs treat

*axel.benyamine@polytechnique.edu

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

the uncertain kernel as a random variable, so that the return of a policy is itself a random variable, and a risk measure is used to aggregate the random return into a scalar objective [1]. It has been shown that ambiguity-averse MDPs encompass many MDP variants, such as robust MDPs, multi-model MDPs [25], and percentile optimization [4, 5].

In this paper, we establish key structural properties of robust and ambiguity-averse MDPs that allow for stationary and deterministic optimal policies. Specifically, our main contributions are as follows.

1) For *Robust MDPs*, where the return is evaluated as the worst case over an uncertainty set, we fully characterize the class of uncertainty sets that guarantee the existence of uniformly optimal stationary policies. Specifically, we show that s-rectangularity is equivalent to the existence of uniformly optimal stationary (randomized) policies, while sa-rectangularity is equivalent to the existence of uniformly optimal stationary *deterministic* policies. We complement these characterizations with some instructive examples and counterexamples: in particular, we exhibit a simple r-rectangular instance where optimal policies must be non-stationary, and introduce the classes of component-wise and trajectory-wise s-rectangularity.

2) For *Ambiguity-Averse MDPs*, where the return is evaluated using a risk measure, we prove that any law-invariant risk measure that guarantees uniformly optimal stationary policies must be either the essential supremum or the essential infimum (i.e., pure optimism or pure pessimism) under mild continuity assumptions. Notably, this result explains why many standard criteria considered in the literature, such as multi-model or soft-robust MDPs, do not admit optimal stationary policies.

Overall, our results provide a complete understanding of the class of models that guarantee the existence of uniformly optimal stationary policies. Our work also clarifies the relation between this latter property and the dynamic programming principle. Our main conclusion is that going beyond these models (rectangular uncertainty sets with full pessimism/optimism) requires new techniques and likely entails giving up uniformly optimal stationary policies. One must then allow either for optimal stationary policies that depend on the initial distribution, or for uniformly optimal non-stationary policies, for instance based on sufficient statistics through state augmentation.

Notation. Random variables are represented by uppercase letters with a tilde, such as $\tilde{X}$ or $\tilde{S}$. We write L_c^∞ for the set of bounded random variables in $\mathbb{R}$ with convex support. The probability simplex over $\mathcal{S}$ is denoted by $\Delta(\mathcal{S})$. For vectors $v, w \in \mathbb{R}^{\mathcal{S}}$, the relation $v \leqslant w$ is interpreted component-wise, i.e., $v(s) \leqslant w(s)$ for every $s \in \mathcal{S}$. We treat vectors in $\mathbb{R}^{\mathcal{S}}$ interchangeably with functions $\mathcal{S} \to \mathbb{R}$.

2 Preliminaries on Nominal MDPs

We begin by briefly reviewing standard (nominal) Markov decision processes (MDPs) to introduce our notation and the key concepts used throughout the paper. We refer the reader to [21] for more details.

A (discounted) Markov Decision Process (MDP) is a tuple $\mathsf{M}_P = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where $\mathcal{S}$ and $\mathcal{A}$ are finite state and action sets, $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ is a reward function, $P \in \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ is a transition kernel and $\gamma \in [0, 1)$ is a discount factor. A history-dependent policy is a map $\pi : \sqcup_{n\in\mathbb{N}}(\mathcal{S} \times \mathcal{A})^n \times \mathcal{S} \to \Delta(\mathcal{A})$. The set Π_{H} denotes all history-dependent policies. A *stationary* policy only chooses actions based on the current state and can be identified as a function from $\mathcal{S}$ to $\Delta(\mathcal{A})$. A deterministic policy returns Dirac distributions over $\mathcal{A}$. The sets of stationary and stationary deterministic policies are denoted by Π_{S} and Π_{SD}, respectively.

The *value function* $V^{\pi,P} \in \mathbb{R}^{\mathcal{S}}$ of a policy $\pi \in \Pi_{\mathsf{H}}$ is defined as the expected discounted return from each state $s \in \mathcal{S}$: $V^{\pi,P}(s) = \mathbb{E}_s^{\pi,P}[\sum_{t=0}^{\infty} \gamma^t r(\tilde{S}_t, \tilde{A}_t, \tilde{S}_{t+1})]$ with $\mathbb{E}_s^{\pi,P}$ denoting the expectation over the trajectories $(\tilde{S}_t, \tilde{A}_t)_{t\in\mathbb{N}}$ induced by π, P when starting in state $s \in \mathcal{S}$, and $(\tilde{S}_t, \tilde{A}_t)$ is the (random) state-action pair visited at period t. Given an initial distribution $\mu \in \Delta(\mathcal{S})$, an *optimal policy* is a maximizer of the function $\pi \mapsto \mu^\top V^{\pi,P}$ over $\pi \in \Pi_{\mathsf{H}}$. In this work, we study the following stronger notion of *uniform optimality*.

Definition 2.1. A policy $\pi^\star \in \Pi_{\mathsf{H}}$ is *uniformly optimal* if it is optimal for each choice of $\mu \in \Delta(\mathcal{S})$, i.e., if $\forall\, \mu \in \Delta(\mathcal{S}),\ \mu^\top V^{\pi^\star,P} = \sup_{\pi\in\Pi_{\mathsf{H}}} \mu^\top V^{\pi,P}$.

In this paper, we study whether comparable uniformly optimal policies exist in the more general settings of robust and ambiguity-averse MDPs. This question is motivated by a remarkable property

of nominal MDPs: such policies are always guaranteed to exist (see, e.g., Theorem 6.2.10 in [21]). This property is important in practice: uniform optimality lets a decision maker deploy the same policy for any initial distribution $\mu \in \Delta(\mathcal{S})$; stationarity greatly simplifies deployment and analysis compared with history-dependent rules; and determinism is particularly valuable in sensitive domains such as healthcare, where transparency and interpretability are essential.

3 Uniformly Optimal Policies in Robust MDPs

In many real-life applications, the transition function P is not known exactly and can only be estimated approximately. To compute policies that are robust to small changes in the transition parameters, *Robust MDPs* (RMDPs) assume that the unknown P belongs to a known uncertainty set $\mathcal{P} \subseteq \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$. Formally, a robust MDP is a tuple $\mathsf{M}_{\mathcal{P}} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where $\mathcal{S}, \mathcal{A}, r, \gamma$ are defined as in nominal MDPs and where $\mathcal{P}$ is the uncertainty set. Once the initial distribution $\mu \in \Delta(\mathcal{S})$ is given, the goal of the decision maker is to compute a policy that maximizes the worst-case return $\pi \mapsto \inf_{P\in\mathcal{P}} \mu^\top V^{\pi,P}$ over $\pi \in \Pi_{\mathsf{H}}$.

As in MDPs, we define *uniformly optimal* policies in robust MDPs as follows.

Definition 3.1 (Uniformly Optimal policy for RMDPs)**.** Given a robust MDP $\mathsf{M}_{\mathcal{P}}$, a policy $\pi^\star \in \Pi_{\mathsf{H}}$ is *uniformly optimal* if for all $\mu \in \Delta(\mathcal{S})$ we have $\inf_{P\in\mathcal{P}} \mu^\top V^{\pi^\star,P} = \sup_{\pi\in\Pi_{\mathsf{H}}} \inf_{P\in\mathcal{P}} \mu^\top V^{\pi,P}$.

In robust MDPs, it is known that uniformly optimal *stationary* and *deterministic* policies are guaranteed to exist under the following structural assumptions [28].

Definition 3.2 (Rectangularity)**.** An uncertainty set $\mathcal{P} \subseteq \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ is:

- *s-rectangular* if it can be written as a Cartesian product across states, *i.e.*, $\mathcal{P} = \bigtimes_{s\in\mathcal{S}} \mathcal{P}_s$, where $\mathcal{P}_s$ corresponds to the marginal-set for s associated to $\mathcal{P}$:

$$\mathcal{P}_s := \{(p_a)_{a\in\mathcal{A}} \in \Delta(\mathcal{S})^{\mathcal{A}} \ : \ \exists\, P \in \mathcal{P}, (P_{sa})_{a\in\mathcal{A}} = (p_a)_{a\in\mathcal{A}}\} \subseteq \Delta(\mathcal{S})^{\mathcal{A}} \,. \tag{1}$$

- *sa-rectangular* if it can be written as a Cartesian product across all state-action pairs, *i.e.*, $\mathcal{P} = \bigtimes_{(s,a)\in\mathcal{S}\times\mathcal{A}} \mathcal{P}_{s,a}$, where $\mathcal{P}_{s,a}$ corresponds to the marginal-set for s, a associated to $\mathcal{P}$:

$$\mathcal{P}_{s,a} := \{p \in \Delta(\mathcal{S}) \ : \ \exists\, P \in \mathcal{P}, P_{sa} = p\} \subseteq \Delta(\mathcal{S}) \,. \tag{2}$$

Equivalence of rectangularity and uniform optimality. Rectangularity plays a central role in robust MDPs. For convex compact uncertainty sets, the s-rectangularity property is sufficient for the existence of uniformly optimal stationary policies [28], while sa-rectangularity additionally implies that these uniformly optimal stationary policies can be chosen deterministic [13, 20]. There have been efforts to define non-rectangular models of uncertainty that allow for tractable optimal policies, e.g., k-rectangular models [18], r-rectangular models [9], or (ξ, η)-rectangular uncertainty [12], but it is unknown if these sets would yield uniformly optimal stationary policies. Surprisingly, our main theorems below show that s-rectangularity and sa-rectangularity are the *only* models of uncertainty that guarantee the existence of uniformly optimal *stationary* and *stationary deterministic* policies, respectively.

Theorem 3.3. *The following assertions are equivalent for any compact convex $\mathcal{P}$ and $|\mathcal{A}| \geqslant 2$:*

(i) For any choice of $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}, \gamma \in [0,1)$, there exists a uniformly optimal stationary policy $\pi^\star$:
$\forall\, r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}, \gamma \in [0,1), \exists\, \pi^\star \in \Pi_{\mathsf{S}}, \forall\, \mu \in \Delta(\mathcal{S}), \quad \inf_{P\in\mathcal{P}} \mu^\top V^{\pi^\star,P} = \sup_{\pi\in\Pi_{\mathsf{H}}} \inf_{P\in\mathcal{P}} \mu^\top V^{\pi,P}$.

(ii) $\mathcal{P}$ is s-rectangular.

Theorem 3.4. *The following assertions are equivalent for any compact convex $\mathcal{P}$ and $|\mathcal{A}| \geqslant 2$:*

(i) For any choice of $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and $\gamma \in [0,1)$, there exists a uniformly optimal stationary deterministic policy $\pi^\star$:
$\forall\, r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}, \gamma \in [0,1), \exists\, \pi^\star \in \Pi_{\mathsf{SD}}, \forall\, \mu \in \Delta(\mathcal{S}), \quad \inf_{P\in\mathcal{P}} \mu^\top V^{\pi^\star,P} = \sup_{\pi\in\Pi_{\mathsf{H}}} \inf_{P\in\mathcal{P}} \mu^\top V^{\pi,P}$.

(ii) $\mathcal{P}$ is sa-rectangular.

Note that in both theorems $V^{\pi,P}$ is an implicit function of the MDP M_P and therefore of r and γ.

It is worth emphasizing that the first characterizations in Theorems 3.3 and 3.4 must hold for every choice of reward function and discount factor. This requirement is natural, since for particular rewards

or discount factors the problem may become artificially trivial, making the existence of uniformly optimal stationary policies immediate, for instance when $r = 0$. From a practical standpoint, our goal is to address the full class of admissible parameters rather than a few special cases.

A broader question is to characterize the triples (uncertainty set, reward function, discount factor) for which uniformly optimal stationary policies exist. Some results are available when more structure is imposed on the rewards: when $r(s, a, s')$ only depends on (s, a) and not on s', several classes of non-s-rectangular uncertainty sets admit uniformly optimal stationary policies, such as r-rectangular sets, sr-rectangular sets and several others (see Section 4 in [11] and Section 3 in [16]). Restricting to less general reward models thus allows for many classes of uncertainty sets that are *sufficient* for the existence of uniformly optimal stationary policies, but it becomes difficult to obtain *necessary* conditions, which is the main goal of this work.

We provide some intuition for our proofs, as we believe that they shed new light on some important properties of robust MDPs. For the sake of conciseness, we focus only on the proof of Theorem 3.3 here, and we refer to Appendix A for both detailed proofs. To prove Theorem 3.3, we only need to prove that the first assertion implies the second (the converse is already known from seminal papers, e.g., [28]). To do so, we prove that for *any* non-s-rectangular uncertainty set $\mathcal{P}$, we can find a choice of rewards $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and discount factor $\gamma \in [0, 1)$ such that no uniformly optimal stationary policy exists. Although Propositions 1 and 2 in [28] show two particular examples of non-s-rectangular $\mathcal{P}$ for which no uniformly optimal stationary policy exists, the main difficulty in our work is to prove that this remains true for *any* non-s-rectangular uncertainty set. We proceed by contradiction.

To this end, we choose $\gamma = 0$ and use a structural result [11, Theorem 2], which states that for convex compact $\mathcal{P}$, s-rectangularity is equivalent to the following Simultaneous Solvability Property:

$$\forall\, U \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}},\ \bigcap_{s\in\mathcal{S}} \operatorname{argmin}_{P\in\mathcal{P}} \sum_{a\in\mathcal{A}} P_{sa}^{\top} U_{sa} \neq \emptyset\ . \tag{3}$$

Thus, for $\mathcal{P}$ non-s-rectangular, there exists $U \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ such that $\bigcap_{s\in\mathcal{S}} \operatorname{argmin}_{P\in\mathcal{P}} \sum_{a\in\mathcal{A}} P_{sa}^{\top} U_{sa} = \emptyset$. Intuitively, this means that the state-wise worst case over $\mathcal{P}$ cannot be attained simultaneously across all states. For a given $\pi \in \Pi_{\mathsf{S}}$ and $\bar{\mu} \in \Delta(\mathcal{S})$ the uniform initial distribution, we compare two "worst-case functions" of π: $F_U(\pi)$ is a worst case over $\mathcal{P}$, defined as $F_U(\pi) = \min_{P\in\mathcal{P}} \sum_{(s,a)} \bar{\mu}(s)\pi(s, a)P_{sa}^{\top}U_{sa}$, and G_U^{s} is a worst case over $\bigtimes_s \mathcal{P}_s$ (the s-rectangular extension of $\mathcal{P}$) defined as $G_U^{\mathrm{s}}(\pi) = \sum_{s\in\mathcal{S}} \bar{\mu}(s) \min_{P\in\mathcal{P}} \sum_{a\in\mathcal{A}} \pi(s, a)P_{sa}^{\top}U_{sa}$. We always have $F_U \geqslant G_U^{\mathrm{s}}$, and U is chosen so that $F_U(\bar{\pi}) > G_U^{\mathrm{s}}(\bar{\pi})$ for the uniform policy $\bar{\pi}$. We use this strict inequality to show that $F_U - G_U^{\mathrm{s}}$ is non-constant. Then, since F_U, G_U^{s} are two concave functions with non-constant difference, a separation argument based on Fenchel conjugacy (see Lemma A.1) yields $c \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}}$ such that no policy maximizes both $F_U + \langle c, \cdot\rangle$ and $G_U^{\mathrm{s}} + \langle c, \cdot\rangle$ over Π_{S}. Now we choose the instantaneous rewards as $r(s, a, s') = U_{sas'} + c_{sa}/\bar{\mu}(s)$. In this case $F_U(\pi)$ is the worst-case return $\min_{P\in\mathcal{P}} \bar{\mu}^{\top} V^{\pi,P}$, while $G_U^{\mathrm{s}}(\pi)$ is the average of $\min_{P\in\mathcal{P}} V^{\pi,P}(s)$ across $s \in \mathcal{S}$. However, if a uniformly optimal $\pi^{\star} \in \Pi_{\mathsf{S}}$ exists, then it must maximize both F_U and G_U^{s} by optimality for $\bar{\mu}$ and for every Dirac initial distribution δ_s, which would yield a contradiction. Therefore, there does not exist a uniformly optimal policy in Π_{S}.

It is worth noting that we use $\gamma = 0$ in our proofs of Theorems 3.3 and 3.4, for which the sets of returns achievable by policies in Π_{H} and Π_{S} coincide. Therefore, we actually showed the stronger statements that s-rectangularity is equivalent to the existence of uniformly optimal policies (possibly history-dependent), and sa-rectangularity is equivalent to the existence of uniformly optimal *deterministic* policies (possibly history-dependent). The following corollary summarizes this discussion.

Corollary 3.5. *Let $\mathcal{P}$ be compact convex and $|\mathcal{A}| \geqslant 2$. The following assertions are equivalent:*

(i) For any choice of $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}, \gamma \in [0, 1)$, there exists a uniformly optimal (resp., uniformly optimal deterministic) policy.

(ii) $\mathcal{P}$ is s-rectangular (resp., sa-rectangular).

Beyond uniform optimality in RMDPs. In light of Theorems 3.3 and 3.4, a natural research question is to characterize the uncertainty sets that guarantee the existence of optimal stationary (randomized or deterministic) policies *that may depend on the initial distribution* (i.e., optimal but not uniformly optimal). As a first step, we now analyze three such generalizations of rectangularity.

While r-rectangularity has been studied in the literature, the component-wise and trajectory-wise s-rectangular classes are introduced in this paper, and their analysis is one of our contributions.

A first candidate is the class of **r-rectangular** uncertainty sets [8, 9] (also named *d-rectangularity* in the finite horizon case [17]) that generalizes sa-rectangularity. An uncertainty set $\mathcal{P}$ is r-rectangular if for any $P \in \mathcal{P}$ and for any s and a, the probabilities P_{sa} are convex combinations of the same underlying vectors $w^1, ..., w^d \in \Delta(\mathcal{S})$ chosen in a Cartesian product. A formal definition is provided in Section B.1. Under the assumptions that the rewards $r(s, a, s')$ depend only on (s, a) and are independent of s', [9] show that uniformly optimal stationary deterministic policies always exist for r-rectangular uncertainty sets. Interestingly, we show that this is false in full generality in the next proposition, even in very simple MDP instances.

Proposition 3.6. *There exists a compact convex r-rectangular uncertainty set $\mathcal{P}$ with $|\mathcal{S}| = |\mathcal{A}| = 2$ and rank $d = 1$, for which for some choice of rewards $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and initial distribution $\mu \in \Delta(\mathcal{S})$, an optimal policy must be non-stationary for any choice of $\gamma \in (0, 1)$.*

The proof of Proposition 3.6 is based on a very simple counterexample with two states and $d = 1$ (see Figure 1a). The intuition is that the next state in $\{s_1, s_2\}$ is sampled with distribution $(p, 1 - p)$ at every period, and the adversary chooses a fixed value for $p \in [0, 1]$. A non-stationary decision maker can "learn" the value of p over time and therefore can strictly improve upon stationary policies. It is worth emphasizing that existing algorithms for solving r-rectangular models involve (to the best of our knowledge) optimizing over stationary policies [2, 9, 17]. Yet, our counterexample shows that this may be strictly suboptimal, even in very simple instances. Given the use of r-rectangular models for real-life high-stakes healthcare applications [8, 10], this calls for new methods to infer optimal policies that depend on the initial distribution of the problem at hand.

A second natural candidate we investigate is the class of **component-wise s-rectangular** uncertainty sets, i.e., instances where we can partition the set of states into subsets where the transitions are s-rectangular, while the full uncertainty set may not be s-rectangular; see Appendix B.2 for a formal definition. An example is given in Figure 1b–1c, with $m = 2$, $\mathcal{S}_1 = \{a, b, c, d, e\}$ and $\mathcal{S}_2 = \{f, g, h\}$, parametrized by $p, p' \in [0, 1]$. The uncertainty set obtained by transitions out of $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2$ is not rectangular because the same parameter $p' \in [0, 1]$ influences the transition out of both states b and f, but the set of feasible transitions out of states in $\mathcal{S}_1$ or out of states in $\mathcal{S}_2$ is s-rectangular. Component-wise s-rectangular uncertainty sets have the nice property that if we restrict to deterministic initial distributions, we can recover uniformly optimal stationary policies. However, perhaps surprisingly, non-stationary optimal policies may still be needed if we allow for randomized initial distributions that can start across different components of the partition of $\mathcal{S}$ (this happens in Figure 1b–1c if one starts in a, b or f with probability $1/3$ and for a certain choice of rewards, see Appendix B.2).

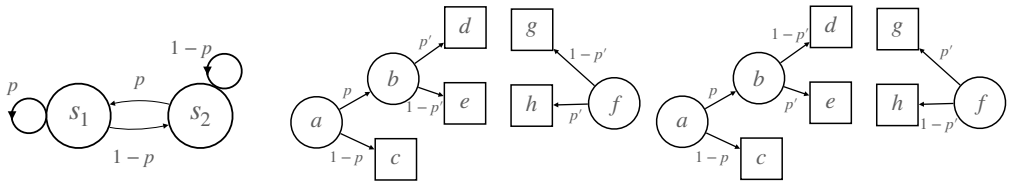
Proposition 3.7. *(i) If $\mathcal{P}$ is a compact convex component-wise s-rectangular uncertainty set, then, for any $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and $\gamma \in [0, 1)$ there exists $\pi^\star \in \Pi_{\mathsf{S}}$ such that for any $s \in \mathcal{S}$, $\inf_{P\in\mathcal{P}} V^{\pi^\star,P}(s) = \sup_{\pi\in\Pi_{\mathsf{H}}} \inf_{P\in\mathcal{P}} V^{\pi,P}(s)$.*

(ii) There exists a compact convex component-wise s-rectangular uncertainty set $\mathcal{P}$ such that for some $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}, \gamma \in [0, 1)$ and $\mu \in \Delta(\mathcal{S})$, optimal policies must be non-stationary.

The last class of uncertainty sets we consider is that of **trajectory-wise s-rectangular** uncertainty sets, that is, sets for which the transitions from states visited along any feasible trajectory remain s-rectangular. We show that this class may also require history-dependent policies, even when the initial distribution is deterministic. Due to space limitations, we defer the formal statement and its proof to Appendix B.3.

4 Uniformly Optimal Policies for Ambiguity-averse MDPs

Robust MDPs are known to be overly conservative in many settings. Ambiguity-averse MDPs [1] aim to mitigate this conservativeness by replacing the worst-case transition kernel of robust MDPs with an uncertain transition kernel $\tilde{P}$, which is unknown but constant throughout the execution. Ambiguity-averse MDPs replace the worst-case objective of robust MDPs with a more general risk-based objective using a risk measure $\rho : L_c^\infty \to \mathbb{R}$. A risk measure generalizes the expectation operator by penalizing the uncertainty of the random variable [7]. In particular, when $\rho = \operatorname{ess\,inf}$, this framework is equivalent to robust MDPs, and when $\rho = \operatorname{ess\,sup}$, it is equivalent to optimistic



(a) Rank 1 r-rectangular MDP with only two states

(b) Component-wise s-rectangular MDP: transitions for action a_1

(c) Component-wise s-rectangular MDP: transitions for action a_2

Figure 1: An r-rectangular RMDP (Figure 1a) and a component-wise s-rectangular RMDP (Figure 1b–1c) with two actions $\{a_1, a_2\}$. The transition probabilities are indicated above each arc; square states are absorbing with reward zero. The choice of rewards for the counterexamples of Proposition 3.6 and Proposition 3.7 are given in Appendix B

MDPs. This section answers analogous questions to Section 3; it characterizes the precise conditions on ρ that allow for uniformly optimal stationary policies.

Definition 4.1. An ambiguity-averse MDP is a tuple $\mathsf{M}_\nu = (\mathcal{S}, \mathcal{A}, \nu, r, \gamma)$, where $\mathcal{S}, \mathcal{A}, r, \gamma$ are defined as in nominal MDPs, and ν is a probability measure over the set of possible transition kernels. We denote by $\mathcal{M}_c$ the set of ambiguity-averse MDPs M_ν where ν has convex support in $\Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$.

Since the transition kernel $\tilde{P} \sim \nu$ is a random variable, the return $\mu^\top V^{\pi,\tilde{P}}$ of a policy $\pi \in \Pi_{\mathsf{H}}$ is itself a random variable (that is $\sigma(\tilde{P})$-measurable). The decision maker evaluates it through a risk measure $\rho : L_c^\infty \to \mathbb{R}$ that maps bounded random variables to scalars, and solves:

$$\sup_{\pi\in\Pi_{\mathsf{H}}} \rho\big(\mu^\top V^{\pi,\tilde{P}}\big) \,. \tag{4}$$

We can now define uniform optimality exactly as in the nominal and robust cases.

Definition 4.2. A policy $\pi^\star \in \Pi_{\mathsf{H}}$ is *uniformly optimal* if for any $\mu \in \Delta(\mathcal{S})$, we have $\rho(\mu^\top V^{\pi^\star,\tilde{P}}) = \sup_{\pi\in\Pi_{\mathsf{H}}} \rho(\mu^\top V^{\pi,\tilde{P}})$.

The following product structures play the role of rectangularity for ambiguity-averse MDPs.

Definition 4.3 (Product Structures)**.** A probability measure ν over transition kernels $\Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ has **s-product structure** if the random variables $(\tilde{P}_s)_{s\in\mathcal{S}}$ are all independent for $\tilde{P} \sim \nu$; and **sa-product structure** if the random variables $(\tilde{P}_{sa})_{(s,a)\in\mathcal{S}\times\mathcal{A}}$ are all independent for $\tilde{P} \sim \nu$.
Note that if ν has an s-product (resp., sa-product) structure, then its support (as a subset of $\Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$) is s-rectangular (resp., sa-rectangular).

For specific choices of risk measures ρ, ambiguity-averse MDPs recover many models of MDPs capturing uncertainty on the transition kernel. Robust MDPs correspond to the essential infimum [28], optimistic MDPs to the essential supremum, multi-model MDPs to $\rho = \mathbb{E}^\nu$ [25], and percentile optimization to the Value-at-Risk, i.e., the upper $(1-\alpha)$-quantile [4, 5].

Remark 4.4 (Comparison with risk-averse MDPs)**.** In ambiguity-averse MDPs, the kernel $\tilde{P}$ is random but the objective for each P is the expected discounted return $\mathbb{E}_\mu^{\pi,\tilde{P}}[\sum_{t=0}^{+\infty} \gamma^t r(\tilde{S}_t, \tilde{A}_t, \tilde{S}_{t+1})]$, evaluated through ρ over the randomness of $\tilde{P}$. In contrast, risk-averse MDPs [23] assume a fixed kernel P and apply ρ directly to $\sum_{t=0}^{+\infty} \gamma^t r(\tilde{S}_t, \tilde{A}_t, \tilde{S}_{t+1})$, the random return along the trajectory. We refer to [1] for further discussion.

Properties of risk measures. To fully specify an instance of ambiguity-averse MDPs, we must choose a risk measure $\rho : L_c^\infty \to \mathbb{R}$ to evaluate the objective in (4). Several properties of risk measures have been extensively studied in the literature. We review the most relevant ones and refer the reader to [7] for an introduction to risk measures and to [1] for their use in MDPs.

Definition 4.5. Let $\rho : L_c^\infty \to \mathbb{R}$ be a risk measure. We say that ρ is **law-invariant** if $\rho(\tilde{X}) = \rho(\tilde{Y})$ whenever $\tilde{X}$ and $\tilde{Y}$ have the same distribution; **monotone** if $\rho(\tilde{X}) \geqslant \rho(\tilde{Y})$ whenever $\tilde{X} \geqslant \tilde{Y}$ a.s.; **translation-invariant** if $\rho(\tilde{X} + c) = \rho(\tilde{X}) + c$ for each $c \in \mathbb{R}$; and **non-expansive** if $|\rho(\tilde{X}) - \rho(\tilde{Y})| \leqslant \|\tilde{X} - \tilde{Y}\|_\infty$ for all $\tilde{X}, \tilde{Y} \in L_c^\infty$.

Note that monotonicity and translation-invariance imply non-expansiveness: for each $\tilde{X}, \tilde{Y}$, let $\delta = \|\tilde{X} - \tilde{Y}\|_\infty$, so $\tilde{Y} - \delta \leqslant \tilde{X} \leqslant \tilde{Y} + \delta$ a.s., hence $\rho(\tilde{Y}) - \delta \leqslant \rho(\tilde{X}) \leqslant \rho(\tilde{Y}) + \delta$. For our purposes we also require a mild continuity notion.

Definition 4.6 (Left and right continuity at constants)**.** A risk measure ρ is left-continuous (resp., right-continuous) at constants if, for each constant $c \in \mathbb{R}$ and sequence $\tilde{X}_n \uparrow c$ a.s. (resp., $\tilde{X}_n \downarrow c$ a.s.) we have $\rho(\tilde{X}_n) \to \rho(c)$.

Law-invariance, monotonicity, translation-invariance and both left and right continuity at constants are satisfied by the vast majority of risk measures in the MDP literature, including the expectation $\mathbb{E}$, the Value-at-Risk VaR_α, the Conditional Value-at-Risk CVaR_α, the Entropic Risk Measure ERM_β, and the Entropic Value-at-Risk EVaR_α. Notable risk measures that are not simultaneously left *and* right continuous are the essential infimum and the essential supremum, defined for $\tilde{X} \in L_c^\infty$ as:

$$\operatorname{ess\,sup}(\tilde{X}) \coloneqq \inf\{c \in \mathbb{R} \mid \tilde{X} \leqslant c \text{ a.s.}\}, \qquad \text{and} \qquad \operatorname{ess\,inf}(\tilde{X}) \coloneqq \sup\{c \in \mathbb{R} \mid \tilde{X} \geqslant c \text{ a.s.}\}\,.$$

Both $\operatorname{ess\,sup}$ and $\operatorname{ess\,inf}$ are law-invariant, monotone, and translation-invariant, but $\operatorname{ess\,sup}$ is only left-continuous at constants (not right-continuous), while $\operatorname{ess\,inf}$ is only right-continuous at constants (not left-continuous). This asymmetry plays an important role in our characterization theorem below.

Remark 4.7. Left or right continuity at constants is weaker than many common assumptions on risk measures. In particular, they are implied by continuity in distribution or almost sure continuity, since they only require convergence along monotone sequences converging to a constant (i.e., a Dirac mass). It is also worth relating these notions to the Fatou property [6, 7], a common regularity assumption in risk theory: a law-invariant, monotone risk measure satisfies the Fatou property if it is lower semicontinuous under almost sure convergence. The essential supremum satisfies the Fatou property [14], while the essential infimum satisfies the reverse Fatou property.

To analyze the existence of uniformly optimal stationary or stationary deterministic policies for all MDP instances, we formulate the following formal conditions on the risk measure ρ:

Condition 1. (Existence of Uniformly Optimal Stationary Policies). We say that a risk measure ρ satisfies this condition if there exists a uniformly optimal stationary policy for each $\mathsf{M}_\nu \in \mathcal{M}_c$ where ν has s-product structure.

Condition 2. (Existence of Uniformly Optimal Stationary Deterministic Policies). We say that a risk measure ρ satisfies this condition if there exists a uniformly optimal stationary deterministic policy for each $\mathsf{M}_\nu \in \mathcal{M}_c$ where ν has sa-product structure.

Before diving into our technical results, it is worth contemplating the implications of Conditions 1 and 2. Recall that the uncertain transition kernel $\tilde{P}$ is drawn *once* from ν at time $t = 0$ and then remains fixed over time. Therefore, in principle a history-dependent policy $\pi \in \Pi_{\mathsf{H}}$ can use the observed trajectory $(s_0, a_0, s_1, \dots)$ to acquire some information regarding the realized value of $\tilde{P}$ and adapt its action choices accordingly to increase its objective value. In contrast, a stationary policy $\pi \in \Pi_{\mathsf{S}}$ must play the same mixed actions in $\Delta(\mathcal{A})$ over time and cannot learn the realization of $\tilde{P}$. Requiring a stationary policy to be uniformly optimal means that ρ is insensitive to this "information advantage" of history-dependent policies over stationary policies, and it is perhaps surprising in the first place that Conditions 1 and 2 may be true for some choices of ρ.

In particular, the fact that $\rho = \operatorname{ess\,inf}$ satisfies Condition 2 follows from [13, 20], and the fact that $\rho = \operatorname{ess\,inf}$ satisfies Condition 1 follows from [28] (recall that the product structure of ν yields rectangular uncertainty sets in $\Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$). The same arguments show the essential supremum ($\operatorname{ess\,sup}$) also satisfies both Conditions 1 and 2. The main goal of this section is to inquire whether other risk measures could satisfy these two conditions. The next lemma shows that Conditions 1 and 2 imply that ρ "commutes" with the pointwise maximum.

Lemma 4.8. *Let ρ be a risk measure that is law-invariant, monotone and non-expansive.*

(i) If ρ satisfies Condition 1, then, for any $m \in \mathbb{N}$ and $(\tilde{X}_i)_{i\in[m]} \in (L_c^\infty)^m$ with convex joint support, there exists $\lambda^\star \in \Delta(m)$ such that $\rho\big(\sum_{i=1}^m \lambda_i^\star \tilde{X}_i\big) = \rho\big(\max_{i\in[m]}(\tilde{X}_i)\big)$.

(ii) If ρ satisfies Condition 2, then, for any $m \in \mathbb{N}$ and independent $(\tilde{X}_i)_{i\in[m]} \in (L_c^\infty)^m$, we have

$$\max_{i\in[m]} \rho\big(\tilde{X}_i\big) = \rho\big(\max_{i\in[m]}(\tilde{X}_i)\big)\,.$$

Lemma 4.8 is proved by constructing an "explore then exploit" ambiguity-averse MDP example, presented in Figure 2. The MDP example is built such that the agent can only choose an action a_i in $\{a_1, \ldots, a_m\}$ from state Start, which has the effect of giving a random return of $\tilde{X}_i$ (via transitioning to Win or Lose and returning to Start). A history-dependent policy can first "explore" to get empirical estimates of these unknown realizations of the random variables $\tilde{X}_1, \ldots, \tilde{X}_m$, and then "exploit" by committing to the empirically optimal action. Asymptotically, this dynamic exploration yields a random return of $\max\{\tilde{X}_i\}$. Conversely, a stationary policy is forced to fix $\lambda \in \Delta(m)$ (the probability to play each action a_i in state Start) and obtain a random return of $\sum_{i=1}^{m} \lambda_i \tilde{X}_i$.

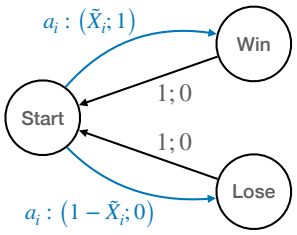


Figure 2: MDP instance for proving Lemma 4.8. Each arc is labeled by the pair (transition; reward) for each action a_i.

Condition 1 implies that at least one stationary policy must perform at least as well as any history-dependent exploration scheme, yielding (i). To obtain (ii) we use the fact that uniformly optimal stationary policies can be chosen *deterministic*, i.e., $\lambda_i = 1$ for some $i \in [m]$.

Under mild continuity assumptions, Lemma 4.8 allows us to completely characterize the risk measures that satisfy Condition 1 or Condition 2, as we show in the next theorem.

Theorem 4.9. *Let ρ be a risk measure that is law-invariant, monotone, non-expansive, normalized (i.e., $\rho(c) = c$ for any $c \in \mathbb{R}$) and either left- or right-continuous at constants. The following assertions are equivalent:*

(i) ρ satisfies Condition 1, *(ii) ρ satisfies Condition 2, and* *(iii) $\rho \in \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$.*

The proof of Theorem 4.9 relies on applying Lemma 4.8 to an increasing number of independent copies $\tilde{X}_1, ..., \tilde{X}_m$ of a random variable $\tilde{X}$. This yields that $\rho(\max_{i\in[m]} \tilde{X}_i) = \rho(\tilde{X})$ while, as $m \to +\infty$, $\max_{i\in[m]} \tilde{X}_i$ converges almost surely to the constant random variable $\operatorname{ess\,sup}(\tilde{X}) \in \mathbb{R}$, so that (ii) and left-continuity readily imply that $\rho(\tilde{X}) = \operatorname{ess\,sup}(\tilde{X})$ for each $\tilde{X} \in L_c^\infty$. The proofs for right-continuity are less straightforward but rely on similar ideas, using random variables $\tilde{Y}_i^{(m)}$ for which $\max_{i\in[m]} \tilde{Y}_i^{(m)}$ is distributed as $\tilde{X}$, in which case $\tilde{Y}_i^{(m)}$ decreases to $\operatorname{ess\,inf}(\tilde{X})$ as $m \to \infty$.

We emphasize that Theorem 4.9 rules out the existence of uniformly optimal stationary policies for all the usual choices of risk measures beyond ess sup and ess inf, e.g., for expected utilities or distorted means [7, 26], or for classical ambiguity-averse models such as percentile optimization (ρ is the value-at-risk [5]) and multi-model MDPs (ρ is the expectation over some kernels [3, 25]). In particular, for multi-model MDPs with infinite horizon, this question was unanswered. It is also interesting to note that a single counterexample (the MDP from Figure 2) can disprove the existence of uniformly optimal stationary policies across *all* risk measures simultaneously (except $\rho \in \{\operatorname{ess\,sup}, \operatorname{ess\,inf}\}$).

On the product-form assumption. Conditions 1 and 2 only involve distributions ν with s- or sa-product structure, which excludes correlated uncertainty across states and actions. This is deliberate: product structures are "favorable" assumptions, in the sense that they aim at allowing as many risk measures ρ as possible, and yet only ess inf and ess sup remain possible under them. Indeed, under any weaker structural assumption on ν (e.g., block-wise or weakly correlated structures), the implication "uniformly optimal stationary policies exist $\Rightarrow \rho \in \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$" would still hold, since the resulting class of instances contains the s-product case and Lemma 4.8 (i) would still apply; however, the converse would fail. For $\rho = \operatorname{ess\,inf}$, the objective only depends on the support of ν, and Section 3 shows that uniformly optimal stationary policies exist for all r and γ only when this support is s-rectangular. Thus, weakening the product assumption cannot yield additional admissible risk measures, and may even rule out these two extreme cases.

Stronger impossibility results. The "explore then exploit" counterexample of Figure 2 reveals that the impossibility results of Theorem 4.9 above remain valid under two weaker assumptions:

(i) *From stationary to Markovian policies.* The set of *Markovian policies* Π_{M} is the set of policies that choose their actions based only on the current observed state and time period. For our counterexample in Figure 2 (where only the decisions at Start have consequences), the return of any Markovian policy can be achieved by a stationary policy when starting at Start. Indeed, any Markovian policy

Table 1: Overview of our results for the existence of uniformly optimal policies for robust MDPs and ambiguity-averse MDPs. ✓means "existence for every instance of the class, for all r, γ"; $\times$ means "non-existence for some particular instance". The component-wise s-rectangular model and the trajectory-wise s-rectangular model are defined formally in Section B. Robust MDPs: $\mathcal{P}$ compact convex; the first two columns are characterizations. Ambiguity-averse MDPs: ρ as in Theorem 4.9, ν with s-product (sa-product for the first column).

	Uniformly optimal stationary		Optimal stationary for each		
Model	deterministic	randomized	$\mu \in \Delta(\mathcal{S})$	deterministic μ	Results
Robust MDPs					
sa-rectangular $\mathcal{P}$	✓	✓	✓	✓	Thm. 3.4
s-rectangular $\mathcal{P}$	$\times$	✓	✓	✓	Thm. 3.3
component-wise s-rec. $\mathcal{P}$	$\times$	$\times$	$\times$	✓	Prop. 3.7
trajectory-wise s-rec. $\mathcal{P}$	$\times$	$\times$	$\times$	$\times$	Prop. B.4
r-rectangular $\mathcal{P}$	$\times$	$\times$	$\times$	$\times$	Prop. 3.6
Ambiguity-averse MDPs					
$\rho \in \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$	✓	✓	✓	✓	Thm. 4.9
$\rho \notin \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$	$\times$	$\times$	$\times$	$\times$	Thm. 4.9

can be characterized by a sequence $(\lambda_t)_{t\in\mathbb{N}} \in \Delta(m)^{\mathbb{N}}$ of decision rules at each period and its return is also achieved by the stationary policy that corresponds to $\lambda = \sum_{k\in\mathbb{N}} \gamma^{2k}(1-\gamma^2)\lambda_{2k} \in \Delta(m)$.

(ii) *From uniform to μ-dependent optimality.* In the "explore then exploit" MDP of Figure 2, even if we were to start in Win or Lose, we would first obtain no reward at period $t = 0$, then transition to Start after one period. Therefore, Theorem 4.9 still holds if we only require the existence of *an* initial distribution μ for which there exists *an* optimal stationary policy (possibly dependent on μ).

To summarize this discussion, even the existence of a Markovian optimal policy for *some* initial distribution forces pure optimism or pure pessimism. This is stated rigorously in the next corollary.

Corollary 4.10. *Let ρ satisfy the assumptions of Theorem 4.9. If for any $\mathsf{M}_\nu \in \mathcal{M}_c$ where ν has s-product structure, there exists an initial distribution $\mu \in \Delta(\mathcal{S})$ for which an optimal Markovian policy exists (possibly dependent on μ), then, ρ is the essential supremum or the essential infimum.*

5 Discussion

Connection to dynamic programming. Intuitively, the dynamic programming equations imply that the optimal value function is the fixed point of a monotone contracting operator, so that optimal value functions dominate the value functions of any other policies component-wise, which implies in particular that we can find optimal policies that are optimal starting from any initial distribution. Under certain classical assumptions, our results highlight interesting connections between dynamic programming and the existence of uniformly optimal stationary policies. For RMDPs, it is shown in [11] that dynamic programming principles for evaluating the worst-case return of a policy as a fixed point of some contracting operators are equivalent to s- and sa-rectangularity, which, by our Theorems 3.3 and 3.4, are also equivalent to the existence of uniformly optimal policies with desirable properties (stationary and deterministic). For ambiguity-averse MDPs, the authors in [1] study the dynamic programming equations and completely characterize when (4) can be solved by computing the fixed point of a contracting (ambiguity-averse) Bellman operator. In particular, Theorem 4.4 in [1] shows that among law-invariant, monotone risk measures, dynamic programming holds if and only if ρ is the essential infimum or essential supremum. This mirrors our Theorem 4.9, albeit we use a slightly different set of assumptions (we need left- or right-continuity, as well as non-expansiveness).

Overall, our main results show that existence of the dynamic programming equations is equivalent to the existence of uniformly optimal stationary policies for RMDPs (a condition on $\mathcal{P}$) and ambiguity-averse MDPs (a condition on ρ). We refer to Appendix D for the rigorous statements. Our results additionally identify when algorithm design can restrict the optimization problem to stationary policies: they specify when stationary-policy optimization is theoretically justified, and when algorithms built on these assumptions can fail.

Results for robust vs. ambiguity-averse MDPs. There is an asymmetry in our main results from Sections 3 and 4. For RMDPs, we can in fact characterize the existence of uniformly optimal policies, even without assuming stationarity (see Corollary 3.5), and the distinction between randomized and deterministic uniformly optimal policies yields the distinction between s- and sa-rectangularity. In contrast, for ambiguity-averse MDPs, our characterization results rely more on the "adaptive" nature of optimal policies (history-dependent, Markovian or stationary) rather than their uniform optimality (Corollary 4.10), and our conclusions are the same whether assuming determinism or not.

Beyond uniformly optimal stationary policies. Our work fully settles the questions of the existence of uniformly optimal stationary policies for MDPs under epistemic uncertainty. Table 1 summarizes our results. If we insist on requiring uniformly optimal stationary policies, then we must use s- or sa-rectangular uncertainty sets (for robust MDPs), and we must use the essential supremum or infimum as a risk measure (for ambiguity-averse MDPs). From these strong characterization results, one can relax the uniformity or stationarity requirement, or restrict the instance class:

- *Non-uniformly optimal stationary policies.* We can search for stationary optimal policies that may depend on the initial distribution. We are not aware of any "non-trivial" models of uncertainty sets or risk measures for which robust MDPs and ambiguity-averse MDPs would have this property (beyond rectangular models of uncertainty and $\rho \in \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$). We may also accept uniform optimality only across deterministic initial distributions (for instance); the class of component-wise s-rectangular models has this favorable property, and it would be interesting to find other such models.
- *Uniformly optimal non-stationary policies.* We can optimize for uniformly optimal non-stationary policies. It would be interesting to find models where the non-stationarity is "limited," e.g., where the decisions may depend only on a small number of sufficient statistics (rather than the full history).
- *Smaller instance universe.* We may also decide to work on a strict subset of MDP instances, e.g., we may want to consider rewards that only depend on the current state-action pair but are independent of the next state; in this case, many uncertainty sets yield uniformly optimal stationary deterministic policies (including r-rectangularity, see Section 4.1.2 in [11]). Understanding what other types of constraints yield the same properties is an interesting next step.

Limitations. Our assumptions in this paper are standard. For RMDPs, convexity and compactness are natural when uncertainty sets are built from confidence regions or distance-based neighborhoods around empirical transition estimates. For risk measures, monotonicity and translation-invariance yield the well-studied class of *monetary* risk measures. Non-law-invariant risk measures have been studied in sequential decision-making (e.g., nested risk measures [23]) but are more difficult to interpret. We derive all our results for the case of *static* kernels, where the transition probabilities remain fixed over time. Studying RMDPs with *time-varying* kernels [11, 27] and ambiguity-averse models with resampled kernels (Definition 3.2 in [1]) is an interesting next step. As mentioned in the discussion following Theorem 3.4, characterizing the triples (uncertainty set, reward function, discount factor) for which similar optimal policies exist is an important direction for future work. Finally, we only consider finite state and action spaces. Our impossibility results carry over to any broader class of models that contains finite state spaces as a special case. Analogous analysis for robust MDPs with countably infinite or continuous state spaces remains relatively scarce (continuous state spaces were only very recently considered [15]) and falls beyond the scope of this paper.

6 Conclusion

This paper characterizes the conditions under which uniformly optimal stationary policies exist under epistemic uncertainty. For robust MDPs, s-rectangularity is both sufficient *and necessary* for their existence, while sa-rectangularity plays the same role for stationary deterministic ones. For ambiguity-averse MDPs with static kernels, we provide an impossibility result: under mild continuity assumptions, any law-invariant, monotone and non-expansive risk measure compatible with the existence of uniformly optimal stationary policies must coincide with the essential supremum or essential infimum. Overall, we highlight that the existence of uniformly optimal stationary policies is a defining feature of robust and optimistic MDPs with rectangular uncertainty sets, which naturally suggests some new research directions, including designing models of uncertainty and objectives that allow for optimal stationary policies that depend on where the process starts or that would allow for simple history-dependent policies (e.g., relying on a small number of sufficient statistics).

Acknowledgements

Axel Benyamine, Michael I. Jordan and Alain Durmus are funded by the European Union (ERC-2022-SYG-OCEAN-101071601). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Julien Grand-Clément and Alain Durmus were supported by Hi! Paris and Agence Nationale de la Recherche (Grant 11-LABX-0047). Alain Durmus a reçu un financement de la Fondation de l'École polytechnique dans le cadre de sa campagne Servir la science. Alain Durmus is supported by the France 2030 program with the reference ANR-25-PEIA-0001 (THEOREM project).

References

[1] A. Benyamine, J. Grand-Clément, M. Petrik, M. I. Jordan, and A. Durmus. Dynamic programming for epistemic uncertainty in Markov decision processes. In *International Conference on Machine Learning*, 2026.

[2] J. Blanchet, M. Lu, T. Zhang, and H. Zhong. Double pessimism is provably efficient for distributionally robust offline reinforcement learning: Generic algorithm and robust partial coverage. In *Advances in Neural Information Processing Systems*, 2023.

[3] P. Buchholz and D. Scheftelowitsch. Computation of weighted sums of rewards for concurrent MDPs. *Mathematical Methods of Operations Research*, 89(1):1–42, 2019.

[4] C. Cousins, E. Lobo, M. Petrik, and Y. Zick. Percentile criterion optimization in offline reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023.

[5] E. Delage and S. Mannor. Percentile optimization for Markov decision processes with parameter uncertainty. *Operations Research*, 58(1):203–213, 2010.

[6] F. Delbaen. Coherent risk measures on general probability spaces. In *Advances in Finance and Stochastics: Essays in Honour of Dieter Sondermann*, pages 1–37. Springer, 2002.

[7] H. Föllmer and A. Schied. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, 4th edition, 2016.

[8] J. Goh, M. Bayati, S. A. Zenios, S. Singh, and D. Moore. Data uncertainty in Markov chains: Application to cost-effectiveness analyses of medical innovations. *Operations Research*, 66(3): 697–715, 2018.

[9] V. Goyal and J. Grand-Clément. Robust Markov decision processes: Beyond rectangularity. *Mathematics of Operations Research*, 48(1):203–226, 2023.

[10] J. Grand-Clément, C. W. Chan, V. Goyal, and G. Escobar. Robustness of proactive intensive care unit transfer policies. *Operations Research*, 71(5):1653–1688, 2023.

[11] J. Grand-Clément, N. Si, and S. Wang. Tractable robust Markov decision processes. *arXiv preprint arXiv:2411.08435*, 2024.

[12] Y. Hu, H. Ma, N. Li, and B. Dai. Efficient duple perturbation robustness in low-rank MDPs. In *Annual Learning for Dynamics & Control Conference*, 2025.

[13] G. N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2): 257–280, 2005.

[14] S. Kusuoka. On law invariant coherent risk measures. In *Advances in Mathematical Economics*, volume 3, pages 83–95. Springer, 2001.

[15] M. Li, Y. Hu, D. Kuhn, and Y. Li. Robust Markov decision processes on continuous state spaces. *arXiv preprint arXiv:2605.28706*, 2026.

[16] Y. Li and A. Shapiro. Rectangularity and duality of distributionally robust Markov decision processes. *Mathematical Programming*, 2025.

[17] X. Ma, Z. Liang, J. Blanchet, M. Liu, L. Xia, J. Zhang, Q. Zhao, and Z. Zhou. Distributionally robust offline reinforcement learning with linear function approximation. *arXiv preprint arXiv:2209.06620*, 2022.

[18] S. Mannor, O. Mebel, and H. Xu. Robust MDPs with k-rectangular uncertainty. *Mathematics of Operations Research*, 41(4):1484–1509, 2016.

[19] F. Miao, S. Han, A. M. Hendawi, M. E. Khalefa, J. A. Stankovic, and G. J. Pappas. Data-driven distributionally robust vehicle balancing using dynamic region partitions. In *International Conference on Cyber-Physical Systems*, 2017.

[20] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[21] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.

[22] R. T. Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, 1970.

[23] A. Ruszczyński. Risk-averse dynamic programming for Markov decision processes. *Mathematical Programming*, 125(2):235–261, 2010.

[24] M. Shaked and J. G. Shanthikumar. *Stochastic Orders*. Springer, 2007.

[25] L. N. Steimle, D. L. Kaufman, and B. T. Denton. Multi-model Markov decision processes. *IISE Transactions*, 53(10):1124–1139, 2021.

[26] S. Wang. Premium calculation by transforming the layer premium density. *ASTIN Bulletin: The Journal of the IAA*, 26(1):71–92, 1996.

[27] S. Wang, N. Si, J. Blanchet, and Z. Zhou. On the foundation of distributionally robust reinforcement learning. *arXiv preprint arXiv:2311.09018*, 2023.

[28] W. Wiesemann, D. Kuhn, and B. Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

A Proofs of Theorems 3.3, 3.4, and Corollary 3.5

Proof of Theorems 3.3 and 3.4. The implications from rectangularity to the existence of uniformly optimal policies are known: the s-rectangular case follows from [28], and the sa-rectangular deterministic case follows from [13, 20]. We prove both converse implications simultaneously by contraposition.

Let $|\mathcal{A}| \geqslant 2$, and let $\mathcal{P}$ be compact and convex. The case $|\mathcal{S}| = 1$ is immediate (any singleton is s-rectangular, and the set of policies becomes a singleton itself), so we assume in the following $|\mathcal{S}| \geqslant 2$. Fix $q \in \{\mathrm{s}, \mathrm{sa}\}$, assume that $\mathcal{P}$ is not q-rectangular, and take $\gamma = 0$.

Step 1: choice of U from SSP. By the Simultaneous Solvability Property [11, Theorems 2 and 4], there exists $U^0 \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ such that the corresponding intersection of minimizers is empty:

$$\begin{cases} \bigcap_{s\in\mathcal{S}} \operatorname{argmin}_{P\in\mathcal{P}} \sum_{a\in\mathcal{A}} P_{s,a}^\top U_{s,a}^0 = \emptyset, & q = \mathrm{s}, \\ \bigcap_{(s,a)\in\mathcal{S}\times\mathcal{A}} \operatorname{argmin}_{P\in\mathcal{P}} P_{s,a}^\top U_{s,a}^0 = \emptyset, & q = \mathrm{sa}. \end{cases} \tag{5}$$

Let $\bar{\mu} = \frac{1}{|\mathcal{S}|}(1, \dots, 1)$ be the uniform initial distribution. For $U \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$, define the following functions $F_U : \Pi_{\mathsf{S}} \to \mathbb{R}$ and $G_U^q : \Pi_{\mathsf{S}} \to \mathbb{R}$:

$$F_U(\pi) = \min_{P\in\mathcal{P}} \sum_{s\in\mathcal{S}} \bar{\mu}(s) \sum_{a\in\mathcal{A}} \pi(s,a) P_{s,a}^\top U_{sa}, \tag{6}$$

$$G_U^q(\pi) = \begin{cases} \sum_{s\in\mathcal{S}} \bar{\mu}(s) \min_{P\in\mathcal{P}} \sum_{a\in\mathcal{A}} \pi(s,a) P_{s,a}^\top U_{sa}, & q = \mathrm{s}, \\ \sum_{s\in\mathcal{S}} \bar{\mu}(s) \sum_{a\in\mathcal{A}} \pi(s,a) \min_{P\in\mathcal{P}} P_{s,a}^\top U_{sa}, & q = \mathrm{sa}. \end{cases} \tag{7}$$

Note that $F_U(\pi)$ is $\min_{P\in\mathcal{P}} \bar{\mu}^\top V^{\pi,P}$, while $G_U^q(\pi)$ is $\min_{P\in\hat{\mathcal{P}}} \bar{\mu}^\top V^{\pi,P}$ for $\hat{\mathcal{P}}$ the q-rectangular extension of $\mathcal{P}$. In particular,

$$F_U(\pi) - G_U^q(\pi) \geqslant 0, \qquad \forall\, U \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}},\ \forall\, \pi \in \Pi_{\mathsf{S}},$$

and, for $\bar{\pi}$ the uniform policy, Equation (5) gives

$$F_{U^0}(\bar{\pi}) - G_{U^0}^q(\bar{\pi}) > 0.$$

We now prove that U can be chosen so that $F_U - G_U^q$ is not constant on Π_{S}. We inspect two cases:

- If this already holds for U^0, take $U = U^0$.
- Otherwise, $F_{U^0} - G_{U^0}^q$ is a strictly positive constant (it is strictly larger than 0 by evaluating at $\bar{\pi}$). Fix $a_1 \in \mathcal{A}$ and define

  $$U_{sa_1} = U_{s,a_1}^0, \qquad U_{sa} = 0 \quad \text{for } a \neq a_1.$$

  For the deterministic policy $\pi^1 : s \mapsto a_1$,

  $$(F_U - G_U^q)(\pi^1) = (F_{U^0} - G_{U^0}^q)(\pi^1) > 0,$$

  whereas $(F_U - G_U^q)(\pi^2) = 0$ for any deterministic policy $\pi^2 : s \mapsto a_2$ with $a_2 \neq a_1$.

Hence, in both cases, $F_U - G_U^q$ is not constant on Π_{S}.

Step 2: choice of rewards. We use the following separation lemma.

Lemma A.1. *Let $\mathcal{X} \subset \mathbb{R}^d$ be a convex compact set. Let f, g be two continuous concave functions. If $f - g$ is not constant, then there exists a linear function $\ell : \mathcal{X} \to \mathbb{R}$ such that*

$$(\operatorname{argmax}_{x\in\mathcal{X}}\{g(x) + \ell(x)\}) \cap (\operatorname{argmax}_{x\in\mathcal{X}}\{f(x) + \ell(x)\}) = \emptyset.$$

We apply this lemma to $\mathcal{X} = \Pi_{\mathsf{S}} \subset \mathbb{R}^{\mathcal{S}\times\mathcal{A}}, f = F_U, g = G_U^q$. Note that F_U, G_U^q are concave continuous functions as pointwise minima of linear functions, and we have chosen U in Step 1 precisely so that $F_U - G_U^q$ is not constant. Therefore, there exists $c \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}}$ such that

$$\operatorname{argmax}_{\pi\in\Pi_{\mathsf{S}}}\{G_U^q(\pi) + \langle c, \pi\rangle\} \cap \operatorname{argmax}_{\pi\in\Pi_{\mathsf{S}}}\{F_U(\pi) + \langle c, \pi\rangle\} = \emptyset. \tag{8}$$

We pick the rewards as follows:

$$r(s,a,s') = U_{sas'} + \frac{c_{sa}}{\bar{\mu}(s)}, \qquad \forall\,(s,a,s') \in \mathcal{S}\times\mathcal{A}\times\mathcal{S}.$$

For every $\pi \in \Pi_{\mathsf{S}}$, this gives

$$F_U(\pi) + \langle c, \pi\rangle = \inf_{P\in\mathcal{P}} \bar{\mu}^\top V^{\pi,P}, \tag{9}$$

$$G_U^q(\pi) + \langle c, \pi\rangle = \begin{cases} \displaystyle\sum_{s\in\mathcal{S}} \bar{\mu}(s) \inf_{P\in\mathcal{P}} \delta_s^\top V^{\pi,P}, & q = \mathrm{s}, \\ \displaystyle\sum_{s\in\mathcal{S}} \bar{\mu}(s) \sum_{a\in\mathcal{A}} \pi(s,a) \inf_{P\in\mathcal{P}} \delta_s^\top V^{\pi^a,P}, & q = \mathrm{sa}, \end{cases} \tag{10}$$

where π^a denotes the deterministic policy that always selects action a.

Step 3: the two conclusions. We now conclude separately for Theorem 3.3 and Theorem 3.4.

Conclusion for Theorem 3.3. Take $q = \mathrm{s}$ and suppose that a uniformly optimal stationary policy $\pi^\star \in \Pi_{\mathsf{S}}$ exists. Its optimality for $\bar{\mu}$ and (9) imply

$$\pi^\star \in \operatorname{argmax}_{\pi\in\Pi_{\mathsf{S}}}\{F_U(\pi) + \langle c, \pi\rangle\}.$$

Uniform optimality also implies that $\pi^\star$ maximizes $\inf_{P\in\mathcal{P}} \delta_s^\top V^{\pi,P}$ for every $s \in \mathcal{S}$. By the first case of (10),

$$\pi^\star \in \operatorname{argmax}_{\pi\in\Pi_{\mathsf{S}}}\{G_U^{\mathrm{s}}(\pi) + \langle c, \pi\rangle\},$$

contradicting (8). Thus, if $\mathcal{P}$ is not s-rectangular, there is a choice of r and γ for which no uniformly optimal stationary policy exists. This proves the converse implication in Theorem 3.3.

Conclusion for Theorem 3.4. Take $q = \mathrm{sa}$ and suppose that a uniformly optimal stationary deterministic policy $\pi^\star \in \Pi_{\mathsf{SD}}$ exists. As above, optimality for $\bar{\mu}$ implies that $\pi^\star$ maximizes $F_U + \langle c, \cdot\rangle$. Moreover, uniform optimality and $\gamma = 0$ imply, for every $s \in \mathcal{S}$,

$$\pi^\star(s) \in \operatorname{argmax}_{a\in\mathcal{A}} \inf_{P\in\mathcal{P}} \delta_s^\top V^{\pi^a,P}.$$

The second case of (10) therefore yields

$$\pi^\star \in \operatorname{argmax}_{\pi\in\Pi_{\mathsf{S}}}\{G_U^{\mathrm{sa}}(\pi) + \langle c, \pi\rangle\},$$

again contradicting (8). Thus, if $\mathcal{P}$ is not sa-rectangular, there is a choice of r and γ for which no uniformly optimal stationary deterministic policy exists. This proves the converse implication in Theorem 3.4. □

We now prove the separation lemma used in the proof above. The intuition behind this lemma comes from Theorem 24.9 in [22], which states that a closed proper convex function is determined up to an additive constant by its subdifferential mapping. In the proof below, the convex functions at hand are the conjugates ϕ^* and ψ^* of $-f$ and $-g$ (restricted to $\mathcal{X}$), whose subdifferentials at c are exactly the sets of maximizers of $f + c^\top x$ and $g + c^\top x$ over $\mathcal{X}$: if these subdifferentials always intersect, then ϕ^* and ψ^*, hence f and g, must differ by a constant.

Proof of Lemma A.1. Let

$$\phi = -f + \iota_{\mathcal{X}}$$
$$\psi = -g + \iota_{\mathcal{X}}$$

with $\iota_{\mathcal{X}}$ the indicator of $\mathcal{X}$: $\iota_{\mathcal{X}}(x) = 0$ if $x \in \mathcal{X}$ and $\iota_{\mathcal{X}}(x) = +\infty$ otherwise.

Then ϕ, ψ are proper lower semicontinuous convex functions with conjugates finite everywhere: for $c \in \mathbb{R}^d$ we have

$$\phi^*(c) = \max_{x \in \mathcal{X}} f(x) + c^\top x$$
$$\psi^*(c) = \max_{x \in \mathcal{X}} g(x) + c^\top x$$

and, from Fenchel equality, for ϕ^*, ψ^* the conjugate functions of ϕ, ψ we have

$$\partial\phi^*(c) = \arg\max_{x \in \mathcal{X}} f(x) + c^\top x$$
$$\partial\psi^*(c) = \arg\max_{x \in \mathcal{X}} g(x) + c^\top x$$

Therefore, we want to show that we can find c such that $\partial\phi^*(c) \cap \partial\psi^*(c) = \emptyset$. Assume on the contrary that $\partial\phi^*(c) \cap \partial\psi^*(c) \neq \emptyset$ for all $c \in \mathbb{R}^d$. Since ϕ^*, ψ^* are finite convex Lipschitz functions, the Rademacher theorem implies that they are differentiable almost everywhere, and at every point $c \in \mathbb{R}^d$ where both are differentiable, we must have $\nabla\phi^*(c) = \nabla\psi^*(c)$. Thus $\phi^* - \psi^*$ is a Lipschitz function which has gradient zero almost everywhere. Therefore, there exists some constant $\alpha \in \mathbb{R}$ such that $\phi^* = \psi^* + \alpha$. Taking the conjugates on both sides of this identity gives $\phi = \psi - \alpha$, which in turn yields that $-f = -g - \alpha$, i.e., that $f - g$ is a constant function on $\mathcal{X}$, which contradicts the assumption of the lemma. □

Proof of Corollary 3.5. The fact that (ii) implies (i) is well known [13, 20, 28]. To prove that (i) implies (ii), we just need to show that if $\mathcal{P}$ is not s-rectangular (resp. sa-rectangular), we can construct $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and $\gamma \in [0, 1)$ such that no policy (resp. deterministic policy) is uniformly optimal across all $\mu \in \Delta(\mathcal{S})$.

We use the fact that the proofs of Theorems 3.3 and 3.4 are carried out with $\gamma = 0$. In this case the discounted return depends only on the one-step transition, so the sets of achievable returns coincide. Indeed, for $\gamma = 0$ we have, for any policy $\pi \in \Pi_{\mathsf{H}}$ and any $\mathcal{P}$,

$$V^{\pi,\mathcal{P}}(s) = \inf_{P \in \mathcal{P}} \mathbb{E}_s^{\pi,P}\left[r(\tilde{S}_0, \tilde{A}_0, \tilde{S}_1)\right] = \inf_{P \in \mathcal{P}} \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in \mathcal{S}} P_{s,a,s'} r(s, a, s') \ ,$$

which depends only on the one-step action distribution at s. Hence for any history-dependent $\pi \in \Pi_{\mathsf{H}}$ there exists a stationary $\bar{\pi} \in \Pi_{\mathsf{S}}$ with

$$\bar{\pi}(s, \cdot) = \pi((s), \cdot) \ ,$$

where (s) is the history so far. This policy yields

$$V^{\pi,\mathcal{P}}(s) = V^{\bar{\pi},\mathcal{P}}(s) \ , \qquad \forall s \in \mathcal{S} \ ,$$

and the same reasoning applies to deterministic policies.

Finally, when $\gamma = 0$:

$$\{V^{\pi,\mathcal{P}} : \pi \in \Pi_{\mathsf{H}}\} = \{V^{\pi,\mathcal{P}} : \pi \in \Pi_{\mathsf{S}}\}, \text{ and } \qquad \{V^{\pi,\mathcal{P}} : \pi \in \Pi_{\mathsf{HD}}\} = \{V^{\pi,\mathcal{P}} : \pi \in \Pi_{\mathsf{SD}}\} \ . \quad (11)$$

Assume $\mathcal{P}$ is not s-rectangular (resp. not sa-rectangular). By the contrapositive direction of Theorem 3.3 (resp. Theorem 3.4), there exists a reward r and a discount factor γ (indeed one may take $\gamma = 0$) for which no uniformly optimal stationary policy exists (resp. no uniformly optimal stationary deterministic policy exists). If a uniformly optimal history-dependent policy (resp. history-dependent deterministic policy) existed for that specific $(r, \gamma = 0)$, then by Equation (11) there would also be a uniformly optimal stationary (resp. stationary deterministic) policy, a contradiction. Hence no uniformly optimal history-dependent policy (resp. history-dependent deterministic policy) exists for that $(r, \gamma = 0)$, which concludes the proof. □

B Beyond s- and sa-rectangularity

We collect the definitions and proofs for the three non-s-rectangular and non-sa-rectangular families discussed in Section 3. The goal is to show that each of these classes can break stationary optimality, even though they preserve interesting structures.

B.1 r-rectangularity

We begin with r-rectangularity, which generalizes sa-rectangularity by allowing transitions to be convex combinations of shared factors and refer to [9] for more discussion on r-rectangular models.

Definition B.1 (r-rectangularity)**.** An uncertainty set $\mathcal{P} \subset \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ is *r-rectangular of rank* d if there exist sets $\mathcal{W}_1, \ldots, \mathcal{W}_d \subset \Delta(\mathcal{S})$ and $(u_{s,a})_{(s,a)\in\mathcal{S}\times\mathcal{A}} \in \Delta(d)^{\mathcal{S}\times\mathcal{A}}$ such that,

$$\mathcal{P} = \left\{ \left(\sum_{i=1}^{d} u_{sa}^i w^i \right)_{(s,a)\in\mathcal{S}\times\mathcal{A}} : \forall i \in [d], \quad w^i \in \mathcal{W}_i \right\} ,$$

i.e., $P \in \mathcal{P}$ if and only if there exists $(w^1, \ldots, w^d) \in \mathcal{W}_1 \times \cdots \times \mathcal{W}_d$ such that $P = \left(\sum_{i=1}^{d} u_{sa}^i w^i\right)_{(s,a)\in\mathcal{S}\times\mathcal{A}}$.

We now show that this structure is insufficient to guarantee stationary optimality in general, by constructing a minimal counterexample.

Proof of Proposition 3.6. Consider the r-rectangular MDP from Figure 1a with two states $\mathcal{S} = \{s_1, s_2\}$ and $d = 1$. In this case, the next state is sampled i.i.d. at every period following the distribution $w = (p, 1-p) \in \Delta(\mathcal{S})$, with $p \in [0,1]$ the probability that the next state is s_1. Consider the case of two actions: $\mathcal{A} = \{a_1, a_2\}$, and $p \in [0,1]$. Assume that we start in state $s_1 : \mu = \delta_{s_1}$. The rewards are given as

$$\begin{aligned} r(s_1, a_1, s_1) = 0, r(s_1, a_1, s_2) = -1 \\ r(s_1, a_2, s_1) = -1, r(s_1, a_2, s_2) = 1 \\ r(s_2, a_1, s_1) = -1, r(s_2, a_1, s_2) = 0 \\ r(s_2, a_2, s_1) = 1, r(s_2, a_2, s_2) = -1 \, . \end{aligned}$$

The intuition for the decision maker facing this choice of rewards is the following. In s_1, we would choose a_1 if we think we are going to stay in s_1, otherwise we would choose a_2. Similarly, in s_2, we would choose a_1 if we think we are going to stay in s_2, otherwise we would choose a_2. The adversary chooses $p \in [0,1]$, which controls whether the next state is in s_1 (with probability p) or s_2 (with probability $1-p$). It is straightforward to show that the best stationary policy is to choose $\pi(s_1, a_1) = 2/3, \pi(s_2, a_1) = 2/3$ with a worst-case return of $-\frac{1}{3}\frac{1}{1-\gamma}$.

Now consider the following history-dependent policy: in the first two periods we match the best stationary policies: at $t = 0$ play $\pi(s_1, a_1) = 2/3$, at $t = 1$ play optimal stationary again, then for $t \geqslant 2$: if $\tilde{S}_1$ (state visited at period 1) is equal to s_1 then play forever $\pi(s_1, a_1) = 1, \pi(s_2, a_2) = 1$ (that is, we bet that $p \geqslant 1/2$, and that we will often "stay" if in s_1 and "move away" if we are in s_2) otherwise $\tilde{S}_1 = s_2$ and play forever $\pi(s_1, a_2) = 1, \pi(s_2, a_1) = 1$ (that is, we bet that $p \leqslant 1/2$, and that we will often "move away" if in s_1 and "stay" if in s_2).

For this history-dependent policy, we can compute the worst-case return as $-\frac{1}{3} - \frac{\gamma}{3} - \frac{\gamma^2}{4(1-\gamma)}$, which is larger than the best worst-case return for a stationary policy computed above. This shows that non-stationary policies strictly outperform stationary ones in this small example. □

B.2 Component-wise s-rectangularity

We next consider component-wise s-rectangularity, where the uncertainty set is s-rectangular within each closed component of the state space but may be non-rectangular globally.

Definition B.2 (Component-wise s-rectangularity)**.** An uncertainty set $\mathcal{P} \subset \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ is component-wise s-rectangular if we can partition $\mathcal{S}$ as $\mathcal{S} = \cup_{i=1}^{m} \mathcal{S}_i$ where, for each $i \in [m]$:

1. $\mathcal{S}_i \subset \mathcal{S}$ is a closed subset of states under any transition probabilities in $\mathcal{P}$, i.e.,

$$\forall\, j \in [m] \setminus \{i\}, \forall (s, s') \in \mathcal{S}_i \times \mathcal{S}_j, \forall P \in \mathcal{P}, \forall a \in \mathcal{A}, P_{s,a,s'} = 0 \, .$$

2. The set $\mathcal{P}_{\mathcal{S}_i} := \{(P_{sa})_{s\in\mathcal{S}_i, a\in\mathcal{A}} \mid P \in \mathcal{P}\}$ of transitions out of $\mathcal{S}_i$ is s-rectangular.

The next proof illustrates that, despite s-rectangularity within each component, global non-rectangularity can still force non-stationarity for optimal policies.

Proof of Proposition 3.7. (i) Let $\mathcal{P}$ be a convex compact component-wise s-rectangular uncertainty set. Let $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and $\gamma \in [0,1)$. Let $s \in \mathcal{S}$ and $\mu = \delta_s$. From the component-wise s-rectangularity assumption, we know that $s \in \mathcal{S}_i$ for some $i \in [m]$ and some subset $\mathcal{S}_i \subset \mathcal{S}$ that is closed under any $P \in \mathcal{P}$. Let $\pi \in \Pi_{\mathsf{H}}$. Then we note that the function $P \mapsto V^{\pi,P}(s)$ only depends on the $\mathcal{S}_i$-th components of P, since $\mathcal{S}_i$ is closed under any $P \in \mathcal{P}$ (we recall that a set of states is closed for a Markov chain if a Markov chain starting in this set stays in this set with probability 1). Since $\mathcal{P}_{\mathcal{S}_i}$ (as defined in Definition B.2) is s-rectangular, we get that there exists an optimal stationary policy $\pi_{\mathcal{S}_i} : \mathcal{S} \to \Delta(\mathcal{A})$ that remains optimal starting from anywhere in $\mathcal{S}_i$, i.e. for any $\mu \in \{\delta_s \mid s \in \mathcal{S}_i\}$. Now this policy $\pi_{\mathcal{S}_i}$ must also assign actions to states outside of $\mathcal{S}_i$ but these actions are irrelevant if we start in $\mathcal{S}_i$ (since we never reach the states outside of $\mathcal{S}_i$ where we would have to choose these actions). Therefore we can restrict $\pi_{\mathcal{S}_i}$ to a map $\mathcal{S}_i \to \Delta(\mathcal{A})$ instead of a map $\mathcal{S} \to \Delta(\mathcal{A})$.

Now consider the policy $\pi^\star$ such that for each $i \in [m]$ and $s \in \mathcal{S}_i$ we have $\pi^\star(s) = \pi_{\mathcal{S}_i}(s)$, with $\pi_{\mathcal{S}_i}$ constructed in the previous paragraph. By construction this policy is an optimal policy for each choice of deterministic initial distribution $\mu \in \{\delta_s \mid s \in \mathcal{S}\}$, which concludes the proof of this statement.

(ii) Consider the instance in Figure 1b–1c with the following rewards:

$$\begin{aligned} r(b,a_1,d) &= 1, & r(b,a_1,e) &= -1 \\ r(f,a_1,g) &= -1, & r(f,a_1,h) &= -1 \\ r(f,a_2,g) &= -1, & r(f,a_2,h) &= 1\ . \end{aligned}$$

All the remaining rewards are 0. The rewards are summarized in Table 2. We parameterize policies π by $\pi = (\alpha, \beta_1, \beta_2, \beta_3, \beta_4) \in [0,1]^5$, with

- α the probability to choose a_1 in state a.
- β_1 the probability to choose a_1 in state b, if the history is (a, a_1, b).
- β_2 the probability to choose a_1 in state b, if the history is (a, a_2, b).
- β_3 the probability to choose a_1 in state b, if the history is (b).
- β_4 the probability to choose a_1 in state f, if the history is (f).

Note that we can always choose $\beta_1 = \beta_2 = \beta \in [0,1]$. Here we write $R_\mu(\pi, p, p') \in \mathbb{R}$ for the scalar value $\mu^\top V^{\pi,P}$ with P the transition probabilities induced by the choice of $p, p' \in [0,1]$.

The case $\mu_a = 1$. In this case we have

$$\begin{aligned} R_\mu(\pi,p,p') &= \gamma \cdot p \cdot \Big(\alpha \cdot (\beta_1(2p'-1) + (1-\beta_1)\cdot 0) \\ &\qquad + (1-\alpha)\cdot(\beta_2(2p'-1) + (1-\beta_2)\cdot 0)\Big) \\ &= \gamma p(2p'-1)\cdot\beta\ , \end{aligned}$$

for $\beta = \beta_1 = \beta_2$.

Minimizing over p, p' we get that for $\mu_a = 1$ we have that the worst-case transitions and optimal policies are

$$p' = 0, p = 1, \beta = 0, \sup_\pi \inf_{p,p'} R_\mu(\pi,p,p') = 0\ .$$

The case $\mu_b = 1$. In this case we have

$$R_\mu(\pi,p,p') = \beta_3(2p'-1)$$

so that we have that the worst-case transitions and optimal policies are

$$p' = 0, p \in [0,1], \beta_3 = 0, \sup_\pi \inf_{p,p'} R(\pi,p,p') = 0\ .$$

Table 2: Non-zero rewards in the MDP instance from Figure 1b–1c

edge	reward	proba
(b, a_1, d)	1	p'
(b, a_1, e)	-1	$1-p'$
(f, a_1, g)	-1	$1-p'$
(f, a_1, h)	-1	p'
(f, a_2, g)	-1	p'
(f, a_2, h)	1	$1-p'$

Table 3: Optimal policies for various choices of initial distributions in Figure 1b–1c with rewards as in Table 2. "/" means any value in $[0, 1]$.

μ	p	p'	β	β_3	β_4
$\mu_a = 1$	1	0	0	/	/
$\mu_b = 1$	/	0	/	0	/
$\mu_f = 1$	/	1	/	/	/
$\mu_a = \mu_b = \mu_f = 1/3$	/	/	0	1	0

The case $\mu_f = 1$. In this case we have

$$R_\mu(\pi, p, p') = -\beta_4 + (1-\beta_4)\cdot(1-2p') \ .$$

The worst case is attained for $p' = 1$ which always yields a worst-case return of -1. Therefore the worst-case transitions and optimal policies are

$$p' = 1, p \in [0,1], \beta_4 \in [0,1], \sup_\pi \inf_{p,p'} R_\mu(\pi, p, p') = -1 \ .$$

The case $\mu_a = \mu_b = \mu_f = 1/3$. Consider the initial distribution $\mu \in \Delta(\mathcal{S})$ with $\mu_a = \mu_b = \mu_f = 1/3$. Then the discounted return of a policy π is

$$\begin{aligned} 3 \cdot R_\mu(\pi, p, p') &= R_{\delta_a}(\pi, p, p') + R_{\delta_b}(\pi, p, p') + R_{\delta_f}(\pi, p, p') \\ &= \gamma p \cdot (2p'-1)\beta + \beta_3(2p'-1) - \beta_4 + (2p'-1)(\beta_4 - 1) \ . \end{aligned}$$

Note that this function is bi-affine in (p, p') so the minimizers over (p, p') are always attained in $\{(0,0),(1,0),(1,1),(0,1)\}$, for which we get

$$\begin{aligned} 3 \cdot R_\mu(\pi, 0, 1) &= -1 + \beta_3 \\ 3 \cdot R_\mu(\pi, 1, 0) &= +1 - \gamma\beta - \beta_3 - 2\beta_4 \\ 3 \cdot R_\mu(\pi, 0, 0) &= +1 - \beta_3 - 2\beta_4 \\ 3 \cdot R_\mu(\pi, 1, 1) &= -1 + \beta_3 + \gamma\beta \ . \end{aligned}$$

Note that $R_\mu(\pi, 0, 0) \geqslant R_\mu(\pi, 1, 0)$ and $R_\mu(\pi, 1, 1) \geqslant R_\mu(\pi, 0, 1)$ so the worst-case return of a policy π is

$$\inf_{p,p'} R_\mu(\pi, p, p') = \frac{1}{3}\min\{\beta_3 - 1, 1 - \gamma\beta - \beta_3 - 2\beta_4\} \ . \tag{12}$$

Now we note that for any π we have $R_\mu(\pi, 0, 1) = \frac{1}{3}(\beta_3 - 1) \leqslant 0$ therefore $\sup_\pi \inf_P R(\pi, P) \leqslant 0$. So any policy π that achieves $\inf_P R_\mu(\pi, P) = 0$ is optimal.

If (12) is equal to 0 then we must have $\beta_3 = 1$. In this case we must have $1 - \beta_3 - \gamma\beta - 2\beta_4 = -\gamma\beta - 2\beta_4 = 0$ so that $\beta = \beta_4 = 0$. Therefore, the set of optimal policies is $\{(\alpha, 0, 1, 0) \mid \alpha \in [0,1]\}$, containing only Markovian policies, and no stationary policies can be optimal. This is true for any $\gamma \in (0, 1)$. We summarize the analysis in the next table (Table 3).

What makes this example work is the following (see Table 2):

(a) The decision maker is aiming for a policy π that achieves a worst case of 0

(b) By choosing $\beta_4 = 0$, the decision maker forces the adversary to hedge at $h = f$: p' gives -1, $1 - p'$ gives $+1$

(c) By choosing $\beta_3 = 1$, the decision maker forces the adversary to hedge at $h = b$, with opposite odds compared to starting in $h = f$: p' gives 1, $1 - p'$ gives -1

(d) By choosing $\beta = 0$, the decision maker forces a payoff of 0 (which is the best one can hope for), whereas if they had chosen $\beta = 1$ this would break the exact tensions between odds at $h = f$ and $h = b$. In particular if we choose $\beta = 1$ we would have that the adversary is hedging again at b, with p' giving 1 and $1 - p'$ giving -1. So the adversary has more incentives to choose $p' = 0$, and the decision maker would obtain a negative total return.

□

B.3 Trajectory-wise s-rectangularity

We conclude with trajectory-wise s-rectangularity, which weakens s-rectangularity only along feasible trajectories. An uncertainty set $\mathcal{P}$ is trajectory-wise s-rectangular if the set of feasible transition probabilities for the states visited along any feasible trajectory over $\mathcal{S}$ induced by a kernel in $\mathcal{P}$ is s-rectangular. The definition below formalizes this restriction and the proposition shows that history dependence can still be necessary.

Definition B.3 (Trajectory-wise s-rectangularity)**.** A set $\mathcal{P} \subset \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ is *trajectory-wise s-rectangular* if for any $t \in \mathbb{N}$, any trajectory $h = (s_0, a_0, \dots, s_t)$ with $\prod_{\tau=0}^{t-1} P_{s_\tau a_\tau s_{\tau+1}} > 0$ for some $P \in \mathcal{P}$, the restricted uncertainty set $\mathcal{P}_{\mathcal{S}_h} \subseteq \Delta(\mathcal{S})^{\mathcal{S}_h\times\mathcal{A}}$ over the visited states $\mathcal{S}_h = \{s_0, \dots, s_t\}$ is s-rectangular.

Even this trajectory-restricted notion of rectangularity is not enough to recover stationary optimality, as shown by the following construction.

Proposition B.4. *There exists a compact convex trajectory-wise s-rectangular uncertainty set $\mathcal{P} \subset \Delta(\mathcal{S})^{\mathcal{S}\times\mathcal{A}}$ such that for some choice of rewards $r \in \mathbb{R}^{\mathcal{S}\times\mathcal{A}\times\mathcal{S}}$ and deterministic $\mu \in \Delta(\mathcal{S})$, optimal policies must be non-stationary for any choice of $\gamma \in (0, 1)$.*

Proof of Proposition B.4. We build upon the instance in Figure 1b–1c. From the proof of the second item of Proposition 3.7, we know that if we start with probability $1/3$ in a, b, or f, then an optimal policy must be non-stationary. Now consider the following trajectory-wise instance: the sets of states and actions are in Figure 1b–1c, but there is an additional start state that transitions to each state in $\{a, b, f\}$ with probability $1/3$ (for both actions a_1 and a_2). The rewards are the same as for the proof of Proposition 3.7. Assume that we deterministically start in state start. Then the return is the same (up to a multiplicative scalar $\gamma \in (0, 1)$) as in the counterexample for the second item of Proposition 3.7, and therefore an optimal policy must be chosen non-stationary for this trajectory-wise s-rectangular uncertainty set. □

C Proofs for Section 4

This appendix presents the proofs for the ambiguity-averse setting. We begin with a short lemma relating two notions of monotonicity used in the proof of Theorem 4.9. We then prove Lemma 4.8 via an explore-then-exploit construction. Finally, we prove the impossibility theorem (Theorem 4.9) and Corollary 4.10 on Markovian optimal policies.

Lemma C.1. *Let ρ be a law-invariant risk measure. Then the following are equivalent:*

1. *ρ is monotone: $\tilde{X} \geqslant \tilde{Y}$ a.s. $\implies \rho(\tilde{X}) \geqslant \rho(\tilde{Y})$.*
2. *ρ is monotone w.r.t. the usual stochastic order:*

$$\Big[\forall t \in \mathbb{R}, \quad \mathbb{P}(\tilde{X} \geqslant t) \geqslant \mathbb{P}(\tilde{Y} \geqslant t)\Big] \implies \rho(\tilde{X}) \geqslant \rho(\tilde{Y}) \ .$$

Proof. From Theorem 1.A.1 in [24], $\tilde{X}$ dominates $\tilde{Y}$ in the usual stochastic order (which we denote by $\tilde{X} \geqslant_{st} \tilde{Y}$) if and only if there exist two random variables $\tilde{X}', \tilde{Y}'$ defined on the same probability space such that $\tilde{X}, \tilde{X}'$ have same law, $\tilde{Y}, \tilde{Y}'$ have same law and $\tilde{X}' \geqslant \tilde{Y}'$ (to construct such random variables it suffices to evaluate the same uniform random variable in the right-continuous inverses of both cumulative distribution functions). Evaluating in ρ concludes the proof.

□

Note that, under law-invariance, a direct consequence of Lemma C.1 is that the increase (resp. decrease) needed to apply left-continuity (resp. right-continuity) at constants can be chosen in terms of the usual stochastic order. In other terms, if ρ is law-invariant and left-continuous (resp. right-continuous) at constants and $(\tilde{Y}_n)_{n\in\mathbb{N}}$ is increasing (resp. decreasing) in the usual stochastic order and converges to a constant c, then $\rho(\tilde{Y}_n) \to \rho(c)$.

Proof of Lemma 4.8. Let ρ be a risk measure that is law-invariant on $L_c^\infty(\mathbb{R})$, monotone and non-expansive.

Let $\tilde{X}'_1, \dots, \tilde{X}'_m \in L_c^\infty$ whose joint law has convex support in $\mathbb{R}^m$. There exist $a > 0$ and $b \in \mathbb{R}$ such that $\tilde{X}_i := (\tilde{X}'_i - b)/a \in [0,1]$ a.s. for all $i \in [m]$. We compare the best stationary (or stationary deterministic) policy with a simple explore-then-exploit policy in the MDP of Figure 3.

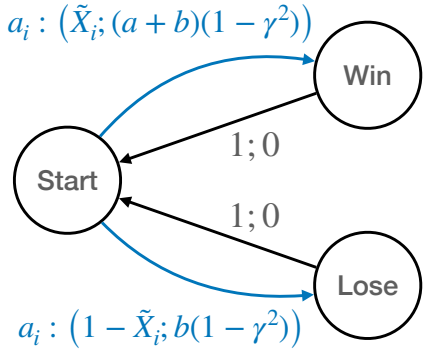


Figure 3: MDP instance for proving Lemma 4.8. Each arc is labeled by the pair of transition and reward for each action a_i.

MDP construction. Consider the ambiguity-averse MDP M with $\mathcal{S} = \{\mathsf{Start}, \mathsf{Win}, \mathsf{Lose}\}$ and $\mathcal{A} = \{a_1, \dots, a_m\}$. From Start, taking action a_i transitions to Win with probability $\tilde{X}_i$ and to Lose with probability $1 - \tilde{X}_i$. Both Win and Lose deterministically transition back to Start in one step. Rewards are given only on the transition out of Start: taking a_i yields reward $(a+b)(1-\gamma^2)$ with probability $\tilde{X}_i$ and $b(1-\gamma^2)$ otherwise; all other transitions carry zero reward. Thus every reward is collected at even times $2t$ (the system returns to Start at odd times), and the odd time steps contribute no reward. This lets us write the total discounted return as a sum over even indices only.

Explore-then-exploit policy. Define the policy π_n as follows:

(i) for each i, play a_i during epochs $t = (i-1)n, \dots, in-1$ (i.e., at times $2t$)
(ii) compute an empirical estimate $\hat{X}_i$, then commit forever to the action with the highest estimate.

When computing the highest estimate, ties are broken by a fixed rule $f : \mathbb{R}^m \to \{1, \dots, m\}$ with $f((x_i)_{i\in[m]}) \in \arg\max_i x_i$.

Using notation $\tilde{X} := (\tilde{X}_i)_{i\in[m]}$ and $\hat{X} := (\hat{X}_i)_{i\in[m]}$, we write the discounted sums over even times and the random return for policy π_n is

$$\begin{aligned}
\tilde{R}_n &= \sum_{t=0}^{\infty} \gamma^{2t} b(1-\gamma^2) + \mathbb{E}\Bigg[a(1-\gamma^2) \sum_{i=1}^{m} \sum_{t=(i-1)n}^{in-1} \gamma^{2t} \mathbb{1}_{\mathsf{Start},a_i}(\tilde{S}_{2t+1} = \mathsf{Win}) \\
&\qquad + a(1-\gamma^2) \sum_{i=1}^{m} \mathbb{1}\big(f(\hat{X}) = i\big) \sum_{t=mn}^{\infty} \gamma^{2t} \mathbb{1}_{\mathsf{Start},a_i}(\tilde{S}_{2t+1} = \mathsf{Win}) \,\Bigg|\, \tilde{X} \Bigg] \\
&\geqslant b + 0 + a(1-\gamma^2) \sum_{i=1}^{m} \sum_{t=mn}^{\infty} \gamma^{2t} \mathbb{E}\Big[\mathbb{1}\big(f(\hat{X}) = i\big) \mathbb{1}_{\mathsf{Start},a_i}(\tilde{S}_{2t+1} = \mathsf{Win}) \,\Big|\, \tilde{X} \Big] \\
&= b + a\gamma^{2mn} \sum_{i=1}^{m} \mathbb{P}\Big(f(\hat{X}) = i \,\Big|\, \tilde{X} \Big)\, \tilde{X}_i \,,
\end{aligned}$$

where the last equality comes from the fact that the two indicators are independent random variables: $\mathbb{1}(f(\hat{X}) = i)$ depends on what happened during the exploration phase (epochs $t = 0, \dots, mn-1$) while $\mathbb{1}_{\mathsf{Start},a_i}(\tilde{S}_{2t+1} = \mathsf{Win})$ depends on what happens during the exploitation phase (at epoch $t \geqslant mn$).

Let σ be the (random) permutation such that $\tilde{X}_{\sigma(1)} \geqslant \cdots \geqslant \tilde{X}_{\sigma(m)}$. Then

$$\tilde{R}_n \geqslant b + a\gamma^{2mn} \sum_{i=1}^{m} \mathbb{P}\Big(\sigma(i) = f(\hat{X}) \;\Big|\; \tilde{X}\Big)\, \tilde{X}_{\sigma(i)} \,.$$

Define gaps $\widetilde{\Delta}_j := \tilde{X}_{\sigma(1)} - \tilde{X}_{\sigma(j)}$ with $\widetilde{\Delta}_{m+1} := +\infty$, so that

$$1 = \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \,.$$

Hence

$$\begin{aligned}
\tilde{R}_n &\geqslant b + a\gamma^{2mn} \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \sum_{i=1}^{m} \mathbb{P}\left(f(\hat{X}) = \sigma(i) \mid \tilde{X}\right) \tilde{X}_{\sigma(i)} \\
&\geqslant b + a\gamma^{2mn} \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \sum_{i=1}^{j} \mathbb{P}\left(f(\hat{X}) = \sigma(i) \mid \tilde{X}\right) \tilde{X}_{\sigma(i)} \\
&\geqslant b + a\gamma^{2mn} \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \sum_{i=1}^{j} \mathbb{P}\left(f(\hat{X}) = \sigma(i) \mid \tilde{X}\right) (\tilde{X}_{\sigma(1)} - \varepsilon)^+ \\
&\geqslant b + a\gamma^{2mn} (\tilde{X}_{\sigma(1)} - \varepsilon)^+ \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \Big(1 - \sum_{i=j+1}^{m} \mathbb{P}\left(f(\hat{X}) = \sigma(i) \mid \tilde{X}\right)\Big) \,,
\end{aligned}$$

where $x^+ := \max(x, 0)$ and we used that, on $\{\widetilde{\Delta}_j \leqslant \varepsilon\}$, $\tilde{X}_{\sigma(i)} \geqslant (\tilde{X}_{\sigma(1)} - \varepsilon)^+$ for $i \leqslant j$.

By Hoeffding's inequality, for any $i \in \{j+1, \ldots, m\}$,

$$\begin{aligned}
\mathbb{P}\left(f(\hat{X}) = \sigma(i) \mid \tilde{X}\right) &\leqslant \mathbb{P}\left(\forall\, k \leqslant m,\ \hat{X}_{\sigma(i)} \geqslant \hat{X}_{\sigma(k)} \mid \tilde{X}\right) \\
&\leqslant \mathbb{P}\left(\hat{X}_{\sigma(i)} \geqslant \hat{X}_{\sigma(1)} \mid \tilde{X}\right) \\
&\leqslant \exp\left(-\frac{n\widetilde{\Delta}_{j+1}^2}{2}\right) \,,
\end{aligned}$$

so

$$\tilde{R}_n \geqslant b + a\gamma^{2mn} (\tilde{X}_{\sigma(1)} - \varepsilon)^+ \sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon) \left(1 - (m-j) e^{-n\varepsilon^2/2}\right) \,.$$

We now choose for the rest of the proof $\gamma = e^{-1/n^2}$ and $\varepsilon = n^{-1/3}$, and define

$$\begin{aligned}
\tilde{Z}_n = b + a e^{-2m/n} (\tilde{X}_{\sigma(1)} - n^{-1/3})^+ \times \\
\sum_{j=1}^{m} \mathbb{1}(\widetilde{\Delta}_j \leqslant n^{-1/3},\ \widetilde{\Delta}_{j+1} > n^{-1/3}) \left(1 - (m-j) e^{-n^{1/3}/2}\right) \,.
\end{aligned}$$

Because $\sqcup_{j=1}^{m} \{\widetilde{\Delta}_j \leqslant \varepsilon,\ \widetilde{\Delta}_{j+1} > \varepsilon\} = \Omega$ is a finite disjoint union, $\tilde{X}_{\sigma(1)} \in [0,1]$, and $|x - (x - \varepsilon)^+| \leqslant \varepsilon$, we have

$$\delta_n := \|a\tilde{X}_{\sigma(1)} + b - \tilde{Z}_n\|_\infty \leqslant a \max_{j \in [m]} \Big(n^{-1/3} + \big|1 - e^{-2m/n}\big(1 - (m-j)e^{-n^{1/3}/2}\big)\big|\Big) \to 0 \,.$$

Then, since $\tilde{R}_n \geqslant \tilde{Z}_n \geqslant a\tilde{X}_{\sigma(1)} + b - \delta_n = \max_{i \in [m]} \tilde{X}'_i - \delta_n$ a.s., and both $\tilde{R}_n = V^{\pi_n, \tilde{P}}(\mathsf{Start})$ (a continuous function of the kernel) and $\max_i \tilde{X}'_i - \delta_n$ (a continuous function of $(\tilde{X}'_i)_i$) belong to L_c^∞, monotonicity and non-expansiveness give

$$\rho(\tilde{R}_n) \geqslant \rho\big(\max_{i \in [m]} \tilde{X}'_i\big) - \delta_n \to \rho\big(\max_{i \in [m]} \tilde{X}'_i\big) \,. \tag{13}$$

Conclusion if ρ satisfies Condition 1. Because ρ satisfies Condition 1, every policy π_n is dominated by a stationary policy. Stationary policies can be represented by a decision rule $\lambda \in \Delta(m)$ that achieves a return

$$\rho\left(\sum_{t=0}^{\infty}\gamma^{2t}\sum_{i=1}^{m}\lambda_i\left(b(1-\gamma^2)+a(1-\gamma^2)\tilde{X}_i\right)\right)=\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)\ ,$$

so Condition 1 implies that for any n,

$$\sup_{\lambda\in\Delta(m)}\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)\geqslant\rho(\tilde{R}_n)\ . \tag{14}$$

Then, from Equations (13) and (14), we get

$$\sup_{\lambda\in\Delta(m)}\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)\geqslant\rho\left(\max_{i\in[m]}\tilde{X}'_i\right)\ ,$$

and, by monotonicity, this inequality is an equality

$$\sup_{\lambda\in\Delta(m)}\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)=\rho\left(\max_{i\in[m]}\tilde{X}'_i\right)\ . \tag{15}$$

Finally,

$$\left|\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)-\rho\left(\sum_{i=1}^{m}\lambda'_i\tilde{X}'_i\right)\right|\leqslant\left\|\sum_{i=1}^{m}(\lambda_i-\lambda'_i)\tilde{X}'_i\right\|_\infty\ ,$$

so $\lambda\mapsto\rho\left(\sum_{i=1}^{m}\lambda_i\tilde{X}'_i\right)$ is continuous and attains its supremum at some $\lambda^\star\in\Delta(m)$, which combined with Equation (15) concludes the proof.

Conclusion if ρ satisfies Condition 2. If $\tilde{X}'_1,\dots,\tilde{X}'_m$ are independent, because ρ satisfies Condition 2, every policy π_n is dominated by a stationary deterministic policy that always chooses one action a_i at Start.

Thus, if the $\tilde{X}'_i$ are independent, Condition 2 implies that for any n,

$$\max_{i\in[m]}\rho\left(\tilde{X}'_i\right)\geqslant\rho(\tilde{R}_n)\ , \tag{16}$$

and from Equations (13) and (16), we get

$$\max_{i\in[m]}\rho\left(\tilde{X}'_i\right)\geqslant\rho\left(\max_{i\in[m]}\tilde{X}'_i\right)\ .$$

By monotonicity this inequality is an equality

$$\max_{i\in[m]}\rho\left(\tilde{X}'_i\right)=\rho\left(\max_{i\in[m]}\tilde{X}'_i\right)\ .$$

□

Remark C.2. The statement for Condition 2 can also be obtained using only two actions and an induction argument, which we do not present since the m-action version is required anyway when assuming Condition 1.

We now leverage the properties established above to characterize the only risk measures compatible with the existence of uniformly optimal stationary (or stationary deterministic) policies. In the proof below, when ρ satisfies Condition 1 (Parts III–IV), we use Lemma 4.8–(i) only for *independent* random variables. Concretely, we rely on the fact that for any $m\in\mathbb{N}$ and any independent $(\tilde{X}_i)_{i\in[m]}\in(L_c^\infty)^m$, there exists $\lambda^\star\in\Delta(m)$ such that

$$\rho\left(\sum_{i=1}^{m}\lambda_i^\star\tilde{X}_i\right)=\rho\left(\max_{i\in[m]}\tilde{X}_i\right)\ .$$

Because independence is assumed, the same identity also follows from Lemma 4.8–(ii) when ρ satisfies Condition 2. Hence the arguments in Parts III–IV remain valid under Condition 2. We nevertheless include the deterministic-case proofs separately in Parts I–II since they are simpler and more transparent.

Proof of Theorem 4.9. As mentioned in Section 4, the fact that (iii) implies both (i) and (ii) is well known [13, 20, 28]. We now prove that (i) implies (iii) and that (ii) implies (iii).

Throughout, assume ρ is law-invariant, monotone, non-expansive, and normalized (i.e., $\rho(c) = c$ for all $c \in \mathbb{R}$).

Part I: left-continuity and Condition 2 implies $\rho = \operatorname{ess\,sup}$. Assume Condition 2 and left-continuity at constants.

Let $\tilde{X} \in L_c^\infty(\mathbb{R})$ be arbitrary. Let $(\tilde{X}_i)_{i\geq 1}$ be i.i.d. copies of $\tilde{X}$, and define

$$\tilde{Y}_n := \max_{1\leqslant i\leqslant n} \tilde{X}_i \ .$$

By Lemma 4.8–(ii), we have for every $n \geqslant 2$:

$$\rho(\tilde{Y}_n) = \rho(\tilde{X}) \ .$$

Let $M := \operatorname{ess\,sup}(\tilde{X})$. Since $\tilde{X}$ is bounded and M is the supremum of its support, $\tilde{Y}_n \uparrow M$ almost surely as $n \to \infty$ (it can be proven easily using the convergence in distribution of $\tilde{Y}_n$ to M). By left-continuity at constants and normalization,

$$\rho(\tilde{Y}_n) \to \rho(M) = M \ .$$

Since $\rho(\tilde{Y}_n) = \rho(\tilde{X})$ for all n, we conclude $\rho(\tilde{X}) = M = \operatorname{ess\,sup}(\tilde{X})$.

Part II: right-continuity and Condition 2 implies $\rho = \operatorname{ess\,inf}$. Assume Condition 2 and right-continuity at constants.

Let $\tilde{X} \in L_c^\infty(\mathbb{R})$ be arbitrary, and denote $m := \operatorname{ess\,inf}(\tilde{X})$ and $M := \operatorname{ess\,sup}(\tilde{X})$, and let F be the CDF of $\tilde{X}$. For each $n \in \mathbb{N}^\star$, we define $\tilde{Z}_i^{(n)}$ as i.i.d. random variables with CDF $F_n : x \mapsto F(x)^{1/n}$ (it exists as $F^{1/n}$ is still a non-decreasing right-continuous function with limits 0 and 1 at $\pm\infty$).

By construction, $\max_{i\leqslant n} \tilde{Z}_i^{(n)}$ has the same distribution as $\tilde{X}$, so by Lemma 4.8–(ii) and law-invariance, we have for every $n \geqslant 2$:

$$\rho(\tilde{Z}_1^{(n)}) = \max_{i\leqslant n} \rho(\tilde{Z}_i^{(n)}) = \rho(\max_{i\leqslant n} \tilde{Z}_i^{(n)}) = \rho(\tilde{X}) \ ,$$

where the first equality comes from the fact that the $\tilde{Z}_i^{(n)}$ are i.i.d. and law-invariance.

Then, the sequence $(\tilde{Z}_1^{(n)})_{n\geqslant 1}$ converges in distribution to m, which is a constant. We also have that the sequence $(\tilde{Z}_1^{(n)})_{n\geqslant 1}$ is decreasing in the usual stochastic dominance. By right-continuity at constants and normalization,

$$\rho(\tilde{Z}_1^{(n)}) \to \rho(m) = m \ .$$

Since $\rho(\tilde{Z}_1^{(n)}) = \rho(\tilde{X})$ for all n, we conclude $\rho(\tilde{X}) = m = \operatorname{ess\,inf}(\tilde{X})$.

Part III: left-continuity and Condition 1 implies $\rho = \operatorname{ess\,sup}$. Assume that ρ satisfies Condition 1 and is left-continuous at constants.

Let $\tilde{X} \in L_c^\infty(\mathbb{R})$ be arbitrary, and denote $m := \operatorname{ess\,inf}(\tilde{X})$ and $M := \operatorname{ess\,sup}(\tilde{X})$. If $M = m$, monotonicity directly gives $\rho(\tilde{X}) = M$, so we can assume $M > m$. Let $\varepsilon \in (0, M - m)$. Because $M - \varepsilon < M$, we have $p = \mathbb{P}(\tilde{X} > M - \varepsilon) > 0$.

Let $\tilde{V}_i \sim (M - m - \varepsilon) \cdot \mathrm{Beta}(1, \frac{M-m-\varepsilon}{p\varepsilon})$ an i.i.d. sequence of random variables.

By construction $\mathbb{E}[\tilde{V}_1] = (M - m - \varepsilon)\frac{1}{1+\frac{M-m-\varepsilon}{p\varepsilon}} \leqslant p\varepsilon$.

By Lemma 4.8–(i), there exists $\lambda^{(n)} \in \Delta(n)$ such that,

$$\rho\left(m + \sum_{i=1}^{n} \lambda_i^{(n)} \tilde{V}_i\right) = \sup_{\lambda \in \Delta(n)} \rho\left(m + \sum_{i=1}^{n} \lambda_i \tilde{V}_i\right) = \rho\left(m + \max_{i \in [n]} \tilde{V}_i\right) ,$$

and we define $\tilde{S}_n = \sum_{i=1}^{n} \lambda_i^{(n)} \tilde{V}_i$.

Then, we have,

$$\begin{aligned} m + \tilde{S}_n &\leqslant m + \varepsilon + (M - m - \varepsilon)\mathbb{1}_{\tilde{S}_n > \varepsilon} \\ \tilde{X} &\geqslant m + (M - m - \varepsilon)\mathbb{1}_{\tilde{X} > M - \varepsilon} . \end{aligned}$$

We see that the indicator on the first line is a Bernoulli random variable with parameter

$$\mathbb{P}(\tilde{S}_n > \varepsilon) \leqslant \frac{\mathbb{E}[\tilde{S}_n]}{\varepsilon} = \frac{\mathbb{E}[\tilde{V}_1]}{\varepsilon} \leqslant p ,$$

while the indicator on the second line is a Bernoulli random variable with parameter p.

Therefore, we have the following usual stochastic dominance:

$$m + \tilde{S}_n \leqslant_{st} \varepsilon + \tilde{X} ,$$

which by monotonicity of ρ and Lemma C.1 gives

$$\rho(m + \tilde{S}_n) \leqslant \rho(\tilde{X} + \varepsilon) .$$

By non-expansiveness we also have $\rho(\tilde{X} + \varepsilon) \leqslant \rho(\tilde{X}) + \varepsilon$, so

$$\rho(m + \tilde{S}_n) \leqslant \rho(\tilde{X}) + \varepsilon .$$

Then, by construction, for any $n \in \mathbb{N}$,

$$\rho(m + \tilde{S}_n) = \rho(m + \max_{i \in [n]} \tilde{V}_i) \to m + M - m - \varepsilon = M - \varepsilon ,$$

by left-continuity and normalization.

Finally, we have proved that for any $\varepsilon > 0$,

$$\rho(\tilde{X}) \geqslant M - 2\varepsilon ,$$

so $\rho(\tilde{X}) \geqslant M$ and by monotonicity this inequality is an equality.

Part IV: right-continuity and Condition 1 implies $\rho = \operatorname{ess\,inf}$. Assume that ρ satisfies Condition 1 and is right-continuous at constants.

Let $\tilde{X} \in L_c^\infty(\mathbb{R})$ be arbitrary, and denote $m := \operatorname{ess\,inf}(\tilde{X})$ and $M := \operatorname{ess\,sup}(\tilde{X})$. We denote by F the CDF of $\tilde{X} - m$.

For each $n \in \mathbb{N}^\star$, we define $(\tilde{Y}_i^{(n)})_{i \in [n]}$ as i.i.d. random variables with CDF $F_n : x \mapsto F(x)^{1/n}$ (it exists as $F^{1/n}$ is still a non-decreasing right-continuous function with limits 0 and 1 at $\pm\infty$). In particular, we have $F_n(0^-) = 0$ and $F_n((M-m)^+) = 1$ so $\tilde{Y}_1^{(n)} \in [0, M - m]$ almost surely.

As in the previous part, we take $\lambda^{(n)}$ the optimal weights from Lemma 4.8–(i) and $\tilde{S}_n = \sum_{i=1}^{n} \lambda_i^{(n)} \tilde{Y}_i^{(n)}$. In particular, we have for any n,

$$\rho(\tilde{X}) = \rho(m + \max_{i \in [n]} \tilde{Y}_i^{(n)}) = \rho(m + \tilde{S}_n) .$$

Then, we have

$$\mathbb{E}[\tilde{S}_n] = \mathbb{E}[\tilde{Y}_1^{(n)}] = \int_0^{M-m} (1 - F(x)^{1/n}) dx \to 0 ,$$

where the limit comes from the Dominated Convergence Theorem ($F(x)^{1/n} \to 1$ for all $x \in (0, M - m]$).

Let now $\varepsilon > 0$, we have

$$p_n = \mathbb{P}(\tilde{S}_n > \varepsilon) \leqslant \frac{\mathbb{E}[\tilde{S}_n]}{\varepsilon} = \frac{\mathbb{E}[\tilde{Y}_1^{(n)}]}{\varepsilon} \to 0 \ ,$$

and we define $q_n = \sup_{k \geqslant n} p_k$ a non-increasing sequence such that for any n, $\mathbb{P}(\tilde{S}_n > \varepsilon) \leqslant q_n$. Because $p_n \to 0$, we also have $q_n \to 0$.

For each n such that $\frac{(M-m)q_n}{\varepsilon} \leqslant 1$, we define $\tilde{V}_n \in [0, M-m]$ a random variable with CDF $x \in [0, M-m] \mapsto (1 - \frac{(M-m)q_n}{\varepsilon}) + x\frac{q_n}{\varepsilon}$ (i.e., $\tilde{V}_n$ is a mixture of δ_0 and Unif$([0, M-m])$).

In particular, $\tilde{V}_n \downarrow 0$ in distribution (the decreasing part is with respect to the usual stochastic dominance) so $\rho(\tilde{V}_n) \to 0$ by right-continuity at a constant and normalization.

Finally,

$$\tilde{S}_n \leqslant \varepsilon + (M - m - \varepsilon)\mathbb{1}_{\tilde{S}_n > \varepsilon}$$
$$\tilde{V}_n \geqslant (M - m - \varepsilon)\mathbb{1}_{\tilde{V}_n > M-m-\varepsilon} \ ,$$

so because $\mathbb{P}(\tilde{V}_n > M - m - \varepsilon) = q_n \geqslant p_n$, we have the following usual stochastic dominance

$$m + \tilde{S}_n \leqslant_{st} \varepsilon + m + \tilde{V}_n \ ,$$

and by monotonicity,

$$\rho(\tilde{X}) = \rho(m + \tilde{S}_n) \leqslant \rho(m + \tilde{V}_n + \varepsilon) \to m + \varepsilon \ .$$

This holds for any $\varepsilon > 0$ so $\rho(\tilde{X}) \leqslant m$ and by monotonicity this is an equality.

□

We now show that the existence of a Markovian optimal policy for *some* initial distribution already forces ρ to be ess sup or ess inf.

Proof of Corollary 4.10. Let ρ be law-invariant, monotone, non-expansive, normalized, and either left or right-continuous at constants. Assume that for any ambiguity-averse MDP $\mathsf{M}_\nu \in \mathcal{M}_c$ with s-product structure, there exists *some* initial distribution μ for which an optimal policy can be chosen Markovian (possibly depending on μ). Note that this assumption is strictly weaker than uniform optimality: we do *not* require a single policy to be optimal for all μ.

Let $m \in \mathbb{N}$ and let $\tilde{X}'_1, \dots, \tilde{X}'_m \in L_c^\infty$ be independent. For each $n \geqslant 1$, consider the explore-then-exploit MDP of Figure 3, built from $(\tilde{X}'_i)_{i \in [m]}$ as in the proof of Lemma 4.8, with discount factor $\gamma = \gamma_n := e^{-1/n^2}$. Let $K := \max(|b|, |a+b|)$, so that $\tilde{X}'_i \in [b, a+b] \subseteq [-K, K]$ almost surely.

Step 1: Markovian optimality to stationary optimality. This MDP has a single decision state Start, and every visit to Start is followed by a deterministic return to Start in one step. Any Markovian policy $\pi \in \Pi_{\mathsf{M}}$ is therefore fully described by a sequence of mixing weights $(\lambda_t)_{t \geqslant 0} \in \Delta(m)^{\mathbb{N}}$ at Start. When starting at Start, the state Start is visited at even times $2k$, so the return equals

$$\rho\left(\sum_{k=0}^{\infty} \gamma^{2k}(1-\gamma^2) \sum_{i=1}^{m} \lambda_{2k}^{(i)} \tilde{X}'_i\right) = \rho\left(\sum_{i=1}^{m} \bar{\lambda}^{(i)} \tilde{X}'_i\right), \qquad \bar{\lambda} := \sum_{k=0}^{\infty} \gamma^{2k}(1-\gamma^2)\lambda_{2k} \in \Delta(m) \ ,$$

which is exactly the return of the stationary policy that plays action a_i with probability $\bar{\lambda}^{(i)}$ at Start. Hence, for this MDP:

$$\sup_{\pi \in \Pi_{\mathsf{M}}} \rho\left(V^{\pi, \tilde{P}}(\mathsf{Start})\right) = \sup_{\lambda \in \Delta(m)} \rho\left(\sum_{i=1}^{m} \lambda_i \tilde{X}'_i\right) = \sup_{\pi \in \Pi_{\mathsf{S}}} \rho\left(V^{\pi, \tilde{P}}(\mathsf{Start})\right) \ .$$

More generally, let $\mu \in \Delta(\mathcal{S})$ and $c_\mu := \mu_{\mathsf{Start}} + \gamma(\mu_{\mathsf{Win}} + \mu_{\mathsf{Lose}}) \in [\gamma, 1]$. When starting from Win or Lose, the state Start is visited at odd times $2k+1$, where π uses the decision rules λ_{2k+1}. Therefore

$$\mu^\top V^{\pi, \tilde{P}} = \mu_{\mathsf{Start}} \sum_{i=1}^{m} \bar{\lambda}^{(i)} \tilde{X}'_i + \gamma(\mu_{\mathsf{Win}} + \mu_{\mathsf{Lose}}) \sum_{i=1}^{m} \bar{\lambda}^{o,(i)} \tilde{X}'_i = c_\mu \sum_{i=1}^{m} \hat{\lambda}^{(i)} \tilde{X}'_i \ ,$$

where $\bar{\lambda}^o := \sum_{k=0}^{\infty} \gamma^{2k}(1-\gamma^2)\lambda_{2k+1} \in \Delta(m)$ and $\hat{\lambda} := \big(\mu_{\mathsf{Start}}\bar{\lambda} + \gamma(\mu_{\mathsf{Win}} + \mu_{\mathsf{Lose}})\bar{\lambda}^o\big)/c_\mu \in \Delta(m)$. Since $|\sum_i \hat{\lambda}^{(i)}\tilde{X}'_i| \leqslant K$, we get $\|\mu^\top V^{\pi,\tilde{P}} - \sum_i \hat{\lambda}^{(i)}\tilde{X}'_i\|_\infty \leqslant (1-c_\mu)K \leqslant (1-\gamma)K$, and non-expansiveness yields, for any $\mu \in \Delta(\mathcal{S})$,

$$\sup_{\pi\in\Pi_{\mathsf{M}}} \rho\left(\mu^\top V^{\pi,\tilde{P}}\right) \leqslant \sup_{\lambda\in\Delta(m)} \rho\left(\sum_{i=1}^m \lambda_i \tilde{X}'_i\right) + (1-\gamma)K \ . \tag{17}$$

Step 2: shift from any μ to Start. Let π_n be the explore-then-exploit policy of the proof of Lemma 4.8. Because the MDP transitions deterministically from Win and Lose to Start in one step, the return for policy π_n from any μ decomposes into an even-time contribution (starting at Start) and an odd time-step contribution (starting at Win or Lose):

$$\mu^\top V^{\pi_n,\tilde{P}} = \mu_{\mathsf{Start}} V^{\pi_n,\tilde{P}}(\mathsf{Start}) + \gamma(\mu_{\mathsf{Win}} + \mu_{\mathsf{Lose}}) V^{\pi_n,\tilde{P}}(\mathsf{Start}) = c_\mu V^{\pi_n,\tilde{P}}(\mathsf{Start}) \ ,$$

where $V^{\pi_n,\tilde{P}}(\mathsf{Start})$ is the return when starting from Start, that exactly corresponds to $\tilde{R}_n$ in the proof of Lemma 4.8. Since $V^{\pi_n,\tilde{P}}(\mathsf{Start}) \in [b, a+b]$ almost surely and $c_\mu \in [\gamma, 1]$, we obtain that for any μ

$$\left\|\mu^\top V^{\pi_n,\tilde{P}} - V^{\pi_n,\tilde{P}}(\mathsf{Start})\right\|_\infty \leqslant (1-\gamma)K \ .$$

In particular, plugging $\gamma_n = e^{-1/n^2}$, we have from the proof of Lemma 4.8 (see (13)) that $\rho(V^{\pi_n,\tilde{P}}(\mathsf{Start})) \geqslant \rho(\max_{i\in[m]} \tilde{X}'_i) - \delta_n$ with $\delta_n \to 0$, so by non-expansiveness:

$$\rho(\mu^\top V^{\pi_n,\tilde{P}}) \ \geqslant \ \rho(V^{\pi_n,\tilde{P}}(\mathsf{Start})) - (1-e^{-1/n^2})K \ \geqslant \ \rho\big(\max_{i\in[m]} \tilde{X}'_i\big) - \delta_n - (1-e^{-1/n^2})K \ .$$

Note that this intermediary result holds for any $n \geqslant 1$ and $\mu \in \Delta(\mathcal{S})$.

By assumption, we also have that for any n there exists μ_n such that a Markovian policy is optimal when starting from μ_n. Thus,

$$\sup_{\pi\in\Pi_{\mathsf{M}}} \rho\left(\mu_n^\top V^{\pi,\tilde{P}}\right) \geqslant \rho(\mu_n^\top V^{\pi_n,\tilde{P}}) \geqslant \rho\big(\max_{i\in[m]} \tilde{X}'_i\big) - \delta_n - (1-e^{-1/n^2})K \ . \tag{18}$$

Then, adding (17) applied with $\mu = \mu_n$ and $\gamma = \gamma_n$ gives

$$\begin{aligned}
\sup_{\lambda\in\Delta(m)} \rho\left(\sum_{i=1}^m \lambda_i \tilde{X}'_i\right) &\geqslant \sup_{\pi\in\Pi_{\mathsf{M}}} \rho\left(\mu_n^\top V^{\pi,\tilde{P}}\right) - (1-e^{-1/n^2})K \\
&\geqslant \rho\big(\max_{i\in[m]} \tilde{X}'_i\big) - \delta_n - 2(1-e^{-1/n^2})K \\
&\to_n \rho\left(\max_{i\in[m]}\{\tilde{X}'_i\}\right)
\end{aligned}$$

so $\sup_{\lambda\in\Delta(m)} \rho\left(\sum_{i=1}^m \lambda_i \tilde{X}'_i\right) \geqslant \rho\left(\max_{i\in[m]}\{\tilde{X}'_i\}\right)$, and by monotonicity, this inequality is an equality.

As in Lemma 4.8 the supremum over λ is attained by continuity of $\lambda \mapsto \rho\left(\sum_{i=1}^m \lambda_i \tilde{X}'_i\right)$.

Step 3: conclude as in Theorem 4.9. The remainder of the proof is identical to the one of Theorem 4.9: left-continuity at constants forces $\rho = \operatorname{ess\,sup}$, and right-continuity at constants forces $\rho = \operatorname{ess\,inf}$. □

D Connection with Dynamic Programming

This appendix revisits dynamic programming in the nominal, robust, and ambiguity-averse settings, and connects the fixed-point principles in the literature to our results on uniformly optimal stationary policies. In the nominal case, Bellman evaluation and optimality always hold. In the robust and ambiguity-averse cases, by contrast, these principles can fail unless additional structure is imposed.

D.1 Dynamic Programming for Nominal MDPs

We recall the standard Bellman operators for a nominal MDP with kernel P and reward r. For a stationary policy $\pi \in \Pi_{\mathsf{S}}$, define the Bellman operator

$$(T^{\pi,P}v)(s) := \sum_{a\in\mathcal{A}} \pi(s,a) \sum_{s'\in\mathcal{S}} P_{s,a,s'}\big(r(s,a,s') + \gamma v(s')\big), \qquad v \in \mathbb{R}^{\mathcal{S}} .$$

Similarly define the optimal Bellman operator

$$(T^{P}v)(s) := \sup_{\pi\in\Pi_{\mathsf{S}}} (T^{\pi,P}v)(s) = \max_{a\in\mathcal{A}} \sum_{s'\in\mathcal{S}} P_{s,a,s'}\big(r(s,a,s') + \gamma v(s')\big) .$$

For nominal MDPs, the following properties hold:

- **Bellman Evaluation Principle.** For any stationary policy $\pi \in \Pi_{\mathsf{S}}$, the value function $V^{\pi,P}$ is the unique fixed point of $T^{\pi,P}$.
- **Bellman Optimality Principle.** The optimal value function $V^{P} = (\sup_{\pi\in\Pi_{\mathsf{H}}} V^{\pi,P}(s))_{s\in\mathcal{S}}$ is the unique fixed point of T^{P}.

These principles not only certify optimality but also guarantee algorithmic tractability: $T^{\pi,P}$ and T^{P} are monotone γ-contractions in the sup norm, which underpins value iteration, policy iteration, and linear programming [21].

D.2 Dynamic Programming for Robust MDPs

In robust MDPs, the Bellman principles do *not* hold automatically. The obstruction is the interaction between the infimum over $\mathcal{P}$ and the recursive structure of the value function. This is precisely where rectangularity becomes essential.

For an uncertainty set $\mathcal{P}$, define the robust value of a policy $\pi \in \Pi_{\mathsf{H}}$ by $V^{\pi,\mathcal{P}} = (\inf_{P\in\mathcal{P}} V^{\pi,P}(s))_{s\in\mathcal{S}}$, and, for $\pi \in \Pi_{\mathsf{S}}$, the robust Bellman operators

$$(T^{\pi,\mathcal{P}}v)(s) := \inf_{P\in\mathcal{P}} (T^{\pi,P}v)(s) , \qquad (T^{\mathcal{P}}v)(s) := \sup_{\pi\in\Pi_{\mathsf{S}}} (T^{\pi,\mathcal{P}}v)(s) .$$

We can now adapt the nominal Bellman principles to the robust case:

- **Robust Bellman Evaluation Principle.** For any stationary policy $\pi \in \Pi_{\mathsf{S}}$, the robust value function $V^{\pi,\mathcal{P}}$ is the unique fixed point of $T^{\pi,\mathcal{P}}$.
- **Robust Bellman Optimality Principle.** The optimal robust value function $V^{\star,\mathcal{P}} = (\sup_{\pi\in\Pi_{\mathsf{H}}} V^{\pi,\mathcal{P}}(s))_{s\in\mathcal{S}}$ is the unique fixed point of $T^{\mathcal{P}}$.

The robust Bellman evaluation principle holds for all rewards and discounts *if and only if* $\mathcal{P}$ is s-rectangular. This is the main tractability result in [11]. Combining it with our uniform optimality characterization (Theorem 3.3) yields the following equivalence.

Proposition D.1 (Equivalence for robust evaluation)**.** *Let $\mathcal{P}$ be a compact convex uncertainty set and $|\mathcal{A}| \geqslant 2$. The following assertions are equivalent:*

(i) For every choice of (r,γ), the robust Bellman evaluation principle holds.

(ii) For every choice of (r,γ), there exists a uniformly optimal stationary policy.

(iii) $\mathcal{P}$ is s-rectangular.

Proof. By Theorems 1 and 2 in [11], (i) and (iii) are equivalent. By our Theorem 3.3, (ii) and (iii) are equivalent. □

Comment on sa-rectangularity. Using Theorems 3 and 5 in [11], a parallel statement can be made for sa-rectangularity and stationary *deterministic* optimality by replacing T^{π} with another contracting operator.

D.3 Dynamic Programming for Ambiguity-Averse MDPs

Fix a risk measure ρ and an ambiguity-averse MDP M_ν where ν has s-product structure. For a policy $\pi \in \Pi_{\mathsf{H}}$, define the ambiguity-averse value function $V^{\pi,\nu} = (\rho(V^{\pi,\tilde{P}}(s)))_{s\in\mathcal{S}}$ and, for $\pi \in \Pi_{\mathsf{S}}$, the operators

$$(T^{\pi,\nu}v)(s) := \rho\left((T^{\pi,\tilde{P}}v)(s)\right), \qquad (T^{\nu}v)(s) := \sup_{\pi\in\Pi_{\mathsf{S}}} (T^{\pi,\nu}v)(s)\ .$$

If ρ is monotone and non-expansive, then $T^{\pi,\nu}$ is a monotone γ-contraction in the L^∞ norm. Hence for each stationary policy $\pi \in \Pi_{\mathsf{S}}$, $T^{\pi,\nu}$ and T^{ν} admit unique fixed points, and policy and value iteration algorithms converge exactly as in the nominal case.

We now present the ambiguity-averse analogues of Bellman principles:

• **Ambiguity-Averse Bellman Evaluation Principle.** For any stationary policy $\pi \in \Pi_{\mathsf{S}}$, the ambiguity-averse value function $V^{\pi,\nu}$ is the unique fixed point of $T^{\pi,\nu}$.

• **Ambiguity-Averse Bellman Optimality Principle.** The optimal ambiguity-averse value function $V^{\nu} = (\sup_{\pi\in\Pi_{\mathsf{H}}} V^{\pi,\nu}(s))_{s\in\mathcal{S}}$ is the unique fixed point of T^{ν}.

Theorem 4.4 in [1] proves, for law-invariant monotone risk measures and static kernels with s-product structure, that the evaluation and optimality principles hold for *all* such MDPs if and only if $\rho \in \{\operatorname{ess\,inf}, \operatorname{ess\,sup}\}$. Combining this with our Theorem 4.9 yields the following equivalence.

Proposition D.2 (Equivalence for ambiguity-averse DP)**.** *Let ρ be law-invariant, monotone, non-expansive, normalized (i.e., $\rho(c) = c$ for any $c \in \mathbb{R}$), and either left- or right-continuous at constants. The following assertions are equivalent:*

(i) For every $\mathsf{M}_\nu \in \mathcal{M}_c$ with s-product structure, the ambiguity-averse Bellman evaluation principle holds.

(ii) For every $\mathsf{M}_\nu \in \mathcal{M}_c$ with s-product structure, the ambiguity-averse Bellman optimality principle holds.

(iii) For every $\mathsf{M}_\nu \in \mathcal{M}_c$ with s-product structure, there exists a uniformly optimal stationary policy.

(iv) ρ is either $\operatorname{ess\,inf}$ *or* $\operatorname{ess\,sup}$.

Proof. By Theorem 4.4 in [1], (i), (ii) and (iv) are equivalent. By our Theorem 4.9, (iii) and (iv) are equivalent. □