A Relation Analysis of Markov Decision Process Frameworks

Tien Mai[*]
SIS, Singapore Management University
atmai@smu.edu.sg

Patrick Jaillet
EECS, Massachusetts Institute of Technology
Jaillet@mit.edu

Abstract

We study the relation between different Markov Decision Process (MDP) frameworks in the machine learning and econometrics literatures, including the standard MDP, the entropy and general regularized MDP, and stochastic MDP, where the latter is based on the assumption that the reward function is stochastic and follows a given distribution. We show that the entropy-regularized MDP is equivalent to a stochastic MDP model, and is strictly subsumed by the general regularized MDP. Moreover, we propose a distributional stochastic MDP framework by assuming that the distribution of the reward function is ambiguous. We further show that the distributional stochastic MDP is equivalent to the regularized MDP, in the sense that they always yield the same optimal policies. We also provide a connection between stochastic/regularized MDP and constrained MDP. Our work gives a unified view on several important MDP frameworks, which would lead new ways to interpret the (entropy/general) regularized MDP frameworks through the lens of stochastic rewards and vice-versa. Given the recent popularity of regularized MDP in (deep) reinforcement learning, our work brings new understandings of how such algorithmic schemes work and suggest ideas to develop new ones.

1 Introduction

Markov decision processes (MDPs) are appealing for various problems involving sequential decision-making. Beside the standard MDP framework [26], which is already popular in the machine learning, operations research and econometrics communities, the literature has seen several variant MDP frameworks proposed over the recent decades, with various motivations. For example, relative entropy regularizers have been added to the reward function to improve exploration and compositionality in reinforcement learning (RL), e.g., Soft-Q learning [12, 28, 15] and Sotf-Actor-Critic [16, 17]. Kullback-Leibler (KL) divergence regularizers have been used to penalize the "*divergence*" between two conservative policies, with the motivation of preventing earlier convergence to sub-optimal policies in a policy iteration algorithm. This idea have been implemented in many state-of-the-art RL algorithms, e.g., Dynamic Policy Programming (DPP) [6], Maximum A Posteriori Policy Optimization (MPO) [2, 1, 22]. [21, 8] add a Tsallis entropy to the Markov problem with the motivation of having sparse and non-deterministic policies, and a theory a general regularized MDP framework has been studied in [14]. The use of entropy-regularizers in MDP is also advantageous in removing ambiguity in some imitation learning/inverse reinforcement learning (IRL) applications [32, 18, 13, 11].

In the econometrics community, the stochastic MDP framework (a.k.a dynamic discrete choice modeling) have been proposed and intensively studied since 1980s [27, 20] (see [3] for a survey). This framework is based on the assumption that the reward function is no-longer deterministic but stochastic and follows a given distribution, and decisions are made according to each realization

---

[*]Corresponding author

of the random rewards. The main motivation is to overcome the deterministic notation of optimal policies resulting from the standard MDP by allowing the agent to make different decisions when the rewards vary, thus improve exploration. Such stochastic Markov models, under some settings, yield randomized/soft policies, which is helpful in casting some behavioral learning procedures as probabilistic inference problems [20]. From this viewpoint, the stochastic MDP framework shares some common motivations with the entropy-regularized MDP one [32].

In this paper, we make the first attempt to thoroughly study the relation between aforementioned MDP frameworks, aiming at providing new ways to view each MDP modeling scheme through different angles and perspectives. We consider a broad class of MDP frameworks, including the standard MDP, entropy-regularized MDP [32], general regularized MDP [14], stochastic MDP with i.i.d Gumbel distributed rewards (a.k.a logit dynamic discrete choice), and general stochastic MDP [27]. We summarize our contributions as follows.

(i) We propose a new MDP framework, called *distributional stochastic MDP*, by assuming that the distribution of the reward function in stochastic MDPs is not given certainly, but belongs to an ambiguity set of distributions. We then show that the corresponding Markov problem is more general and more tractable, as compared to the standard stochastic MDP, in the sense an optimal policy can be found by value iteration and convex optimization, under some distributional settings with marginal information.

(ii) We show that, interestingly, there is a nested relation between these MDP frameworks, that is, (i) the entropy-regularized MDP and stochastic MDP with i.i.d Gumbel distributed rewards are "*equivalent*" and strictly covered by the stochastic MDP, and (ii) the general regularized MDP and the distributional stochastic MDP are equivalent and strictly cover other frameworks. Here, "*equivalence*" means that the uses of these MDP frameworks are equivalent, in both RL and IRL perspectives. There are many interesting points that can be seen from this nested relationship, namely, we can view (entropy) regularized MDP through the lens of stochastic-reward MDP and vice-versa. This also indicates the robustness of regularized MDP from a viewpoint of stochastic-reward with ambiguous distribution. The results also lead to new regularizers with interpretations from stochastic-reward perspectives.

(iii) We also make a connection between regularized/stochastic MDP and constrained MDP [4] by showing that any aforementioned MDP model can be *equivalently* represented by a constrained MDP model. That is, adding regularizers or random noises to the reward function is *equivalent* to imposing some constraints on the desired optimal policies. We also provide discussions on how this constrained-MDP viewpoint would be beneficial to develop generalized RL algorithms.

The equivalence between entropy-regularized MDP and logit dynamic discrete choice in the finite-horizon case was recognized previously by [9] and in our work we show it holds for the infinite case as well. Indeed, we cover a much broader class of MDP frameworks. Our work generally gives a unified picture of several useful MDP frameworks in the literature, thus brings new understandings of how they work and suggests ideas to develop new algorithms.

The paper is structures as follows. In Section 2 we revisit some well-known MDP frameworks in the literature. Section 3 introduces the distributional stochastic MDP framework. Section 4 presents the nested relationship between regularized and stochastic MDP. We show a connection between constrained MDP and other frameworks in Section 5 and finally, Section 6 concludes. Proofs and some relevant discussions are given in the supplementary material. We use $|\mathcal{S}|$ to denote the cardinality of set $\mathcal{S}$. Boldface characters represent matrices (or vectors) or a collection of values/functions.

2 Preliminaries

Consider a infinite-horizon Markov decision process (MDP) for an agent with finite states and actions, defined as $(\mathcal{S}, \mathcal{A}, \mathbf{\Pi}, \mathbf{q}, \mathbf{r}, \gamma)$, where $\mathcal{S}$ is a set of states $\mathcal{S} = \{1, 2, \ldots, |\mathcal{S}|\}$, $\mathcal{A}$ is a finite set of actions, $\mathbf{\Pi} = \{\boldsymbol{\pi}^0, \boldsymbol{\pi}^1, \ldots\}$ is a policy function where $\pi^t(a_t|s_t)$ is the probability of making action $a_t \in \mathcal{A}$ in state $s_t \in \mathcal{A}$ at time $t \in \{0, \ldots, \infty\}$, $\mathbf{q} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ is a transition probability function, i.e., $q(s'_t|a_t, s_t)$ is the probability of moving to state $s'_t \in \mathcal{S}$ from $s_t \in \mathcal{S}$ by performing action $a_t \in \mathcal{A}$ at time t, $\mathbf{r}$ is a vector of reward functions such that $r(a_t|s_t)$ is a reward function associated with state $s_t \in \mathcal{S}$ and action $a_t \in \mathcal{A}$ at time t, and $\gamma \in [0, 1]$ is a discount factor.

In the following we review different MDP frameworks, including the standard one, regularized MDP [32], stochastic MDP [20, 27] (a.k.a dynamic discrete choice). We also briefly discuss how optimal policies can be computed. They are important ingredients to study the relation.

Standard MDPs. In a standard MDP model, the aim is to find an optimal sequential policy $\mathbf{\Pi}$ that maximizes the expected discounted reward

$$\sup_{\substack{\boldsymbol{\pi}^t \in \Delta^\pi \\ t=0,\dots,\infty}} \left\{ \mathbb{E}_{\tau \sim (\mathbf{\Pi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} r(a_t|s_t) \right] \right\}, \tag{1}$$

where τ is a trajectory $\tau = (s_0, a_0, \dots, s_\infty, a_\infty)$ and Δ^π is the set of all possible policy functions, i.e., $\Delta^\pi = \{\boldsymbol{\pi}|\ \pi(a_t|s_t) \geq 0,\ \sum_{a_t} \pi(a_t|s_t) = 1\}$. Here, $\gamma^{[t]}$ refers to *"γ to the power of t"* to distinguish it from a superscript t. It is well-known the problem admits a deterministic notion of optimality, i.e., optimal policies are deterministic.

Regularized MDPs. In a general regularized MDP model [14], a set of regularizers $\boldsymbol{\phi}$ are added to the rewards; $\boldsymbol{\phi} = \{\phi_s(\boldsymbol{\pi}_s),\ s \in \mathcal{S}\}$, where $\boldsymbol{\pi}_s = \{\pi(a|s), \forall a \in \mathcal{A}\}$, for any $s \in \mathcal{S}$. The Markov decision problem can be stated as

$$\sup_{\substack{\boldsymbol{\pi}^t \in \Delta^\pi \\ t=0,\dots}} \left\{ \mathbb{E}_{\tau \sim (\mathbf{\Pi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) + \phi_{s_t}(\boldsymbol{\pi}_{s_t}^t) \Big) \right] \right\}. \tag{2}$$

It is typically assumed that $\phi_s(\boldsymbol{\pi}_s)$ is (strictly) concave. To compute an optimal policy, we define, for any $s \in \mathcal{S}$, the value functions

$$V^{\text{RE},\boldsymbol{\phi}}(s) = \sup_{\substack{\boldsymbol{\pi}^t \in \Delta^\pi \\ t=0,\dots}} \mathbb{E}_{\tau \sim (\mathbf{\Pi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) + \phi_{s_t}(\boldsymbol{\pi}_{s_t}^t) \Big) \Bigg| s_0 = s \right]; \tag{3}$$

and the Bellman update $\mathcal{T}^{\text{RE},\boldsymbol{\phi}}[V](s) = \sup_{\boldsymbol{\pi}_s} \mathbb{E}_{\boldsymbol{\pi}_s}[r(a|s) + \gamma \sum_{s'} q(s'|s,a)V(s')] + \phi_s(\boldsymbol{\pi}_s)$. Then one can show that $\mathcal{T}^{\text{RE},\boldsymbol{\phi}}[V]$ is a contraction [14], guaranteeing that the Bellman equation always have a unique fixed point solution. Furthermore, $V^{\text{RE},\boldsymbol{\phi}}$ defined above is the unique solution to $\mathcal{T}^{\text{RE},\boldsymbol{\phi}}[V] = V$. Moreover, if we denote by $\phi_s^*(\mathbf{w}_s)$ the Legendre-Fenchel transform (or convex conjugate) of $-\phi_s(\boldsymbol{\pi}_s)$, i.e.,

$$\phi_s^*(\mathbf{w}_s) = \sup_{\boldsymbol{\pi}_s} \left\{ \mathbf{w}_s^{\text{T}} \boldsymbol{\pi}_s + \phi_s(\boldsymbol{\pi}_s) \right\}, \forall \mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}, \tag{4}$$

then the value iteration can be updated as $\mathcal{T}^{\text{RE},\boldsymbol{\phi}}[V](s) = \phi_s^*(\mathbf{w}_s)$, where $\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}$ with elements $w_{sa} = r(a|s) + \sum_{s'} q(s'|s,a)V(s')$. If $\phi_s(\boldsymbol{\pi}_s) = -\sum_{a \in \mathcal{A}} \pi(a|s) \ln \pi(a|s)$ (relative entropy), the model becomes the entropy-regularized one. We call this MDP framework as **Entropy-Regularized MDP** (ER-MDP).

Stochastic MDPs (a.k.a. Dynamic Discrete Choice Modeling). This framework has been developed and widely studied in the econometrics community for the problem of learning/predicting people behavior from observations (a.k.a IRL) [e.g., 27]. In this framework, it is assumed that the reward function is stochastic and can be written as $\widetilde{r}(a|s) = r(a|s) + \epsilon(a|s)$, where $r(a|s)$ is deterministic and $\epsilon(a|s)$ is a random noise. Let $\boldsymbol{\epsilon}_s = \{\epsilon(a|s)|\ a \in \mathcal{A}\}$. We further assume that the distributions of $\boldsymbol{\epsilon}_s, \forall s \in \mathcal{S}$, are continuous and independent over states. We define a sequence of decision rule $\boldsymbol{\Psi} = \{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \dots\}$, where $\boldsymbol{\psi}^t$ is a set of decision rules at time t, i.e., $\boldsymbol{\psi}^t = \{\psi(s_t|\boldsymbol{\epsilon}_{s_t}),\ s_t \in \mathcal{S}\}$, where $\psi(s_t|\boldsymbol{\epsilon}_{s_t}) : \mathbb{R}^{|\mathcal{A}|} \to \mathcal{A}$ is a deterministic policy at state s_t under a realization of the random noise $\boldsymbol{\epsilon}_{s_t}$. The aim is to select a decision rule to maximize the expected discounted stochastic reward

$$\sup_{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \dots} \mathbb{E}_{\boldsymbol{\epsilon}} \left\{ \left[\mathbb{E}_{\tau \sim (\boldsymbol{\Psi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) + \epsilon(a_t|s_t) \Big) \right] \right] \right\}. \tag{5}$$

To solve the above Markov problem, we also define the value function

$$V^{\text{S}}(s) = \sup_{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \dots} \mathbb{E}_{\boldsymbol{\epsilon}} \left\{ \left[\mathbb{E}_{\tau \sim (\boldsymbol{\Psi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) + \epsilon(a_t|s_t) \Big) \Bigg| s = s_0 \right] \right] \right\}, \forall a \in \mathcal{S}.$$

Then [20] shows that V^{S} is a solution to the contraction mapping $\mathcal{T}^{\mathrm{S}}[V](s) = \sup_{\boldsymbol{\psi}_s}\{\mathbb{E}_{\boldsymbol{\epsilon}_s}[z(s,a|V)]\}$ where $z(s,a|V) = r(a|s) + \epsilon(a|s) + \mathbb{E}[V(s')]$, and there is an memory-less optimal decision rule specified as $\psi(s|\boldsymbol{\epsilon}_s) = \arg\max_{a\in\mathcal{A}}\{z(s,a,V^{\mathrm{S}})\}$. Basically, a set of decision rule $\{\boldsymbol{\psi}_s, s\in\mathcal{S}\}$ can be viewed as a policy where the probability of making action a at state s is also the density of making action a, i.e., $\pi(a|s) = \mathbb{P}_{\boldsymbol{\epsilon}}(\phi(s|\boldsymbol{\epsilon}_s)=a)$. In particular, if $\epsilon(a|s)$ are i.i.d and follow the Gumbel distribution (or Extreme Value Type I distribution), we obtain a MDP framework so-called *the logit-based dynamic discrete choice model* [20]. We call this model as **Stochastic EV MDP** (or EV-MDP for short).

3   Distributional Stochastic MDPs

The Bellman update in a stochastic MDP model involves an expectation of the form $\mathbb{E}_{\boldsymbol{\epsilon}_s}[\max_{a\in\mathcal{A}}\{w_{sa}+\epsilon(a|s))\}]$, which is typically difficult to evaluate even when the distribution of $\boldsymbol{\epsilon}_s$ is well specified [30]. To over come this, inspired by the persistent model studied in [7, 25], we introduce a new MDP framework, called *Distributional Stochastic MDP* (DS-MDP), by extending the S-MDP framework by assuming that the distribution of the random terms $\boldsymbol{\epsilon}$ is not given certainly, but belongs to a set of distributions Ξ. This new setting makes the corresponding Markov problem easier to solve by convex optimization, under some classes of distributions with marginal information. We will show later this new framework is strongly related to the regularized MDP one.

Under the ambiguity setting, we define a Markov problem where the objective is to select a sequence of distributions $\mathbf{D} = \{\mathbf{d}^0, \mathbf{d}^1, \ldots\}$ and a sequence of decision rules $\boldsymbol{\Psi} = \{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots\}$ in such a way that the following expected reward function is maximized.

$$\sup_{\substack{\mathbf{d}^t\in\Xi \\ t=0,1,\ldots}} \left\{ \sup_{\boldsymbol{\psi}^0,\boldsymbol{\psi}^1,\ldots} \left\{ \mathbb{E}_{\substack{\boldsymbol{\epsilon}^t\sim\mathbf{d}^t \\ t=0,\ldots}} \left[\mathbb{E}_{\tau\sim(\boldsymbol{\psi},\mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]}\Big(r(a_t|s_t) + \epsilon(a|s) \Big) \right] \right] \right\} \right\}, \tag{6}$$

where $\boldsymbol{\epsilon}$ the random noise vector at time t. That is, at time step t, we select a distribution $\mathbf{d}^t \in \Xi$ describing how the random noises $\boldsymbol{\epsilon}$ vary and a decision rule $\boldsymbol{\psi}^t$ to select action according to each realization of the random noise vector. To make the Markov problem tractable, we assume that Ξ is state-wise decomposable, i.e., $\Xi = \otimes_{s\in\mathcal{S}}\Xi_s$, where Ξ_s is an ambiguity set for the distributions of the noise vector $\boldsymbol{\epsilon}_s$ at state s. This assumption allows to solve the Markov problem by value iteration or policy iteration, similarly to other MDP frameworks.

We define the value function (expected accumulated discounted reward) $V^{\mathrm{DS},\Xi} : \mathcal{S} \to \mathbb{R}$ as

$$V^{\mathrm{DS},\Xi}(s) = \sup_{\substack{\mathbf{d}^t\in\Xi \\ t=0,\ldots}} \sup_{\boldsymbol{\psi}^0,\boldsymbol{\psi}^1,\ldots} \left\{ \mathbb{E}_{\substack{\boldsymbol{\epsilon}^t\sim\mathbf{d}^t \\ t=0,\ldots}} \left[\mathbb{E}_{\tau\sim(\boldsymbol{\Psi},\mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]}\Big(r(a_t|s_t) + \epsilon(a_t|s_t) \Big) \,\Big|\, s_0 = s \right] \right] \right\}, \tag{7}$$

and the mapping

$$\mathcal{T}^{\mathrm{DS},\Xi}[V](s) = \sup_{d_s\in\Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s\sim d_s} \left[\max_{a\in\mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma\sum_{s'} q(s'|s,a)V(s') \right\} \right] \right\}. \tag{8}$$

The following theorem shows that basic theoretical properties that hold for other MDP frameworks also hold for the DS-MDP.

Theorem 1 *For any nonempty set of distributions Ξ, we have the following properties*

  (i) *$\mathcal{T}^{\mathrm{DS},\Xi}[V]$ is monotonically increasing and translation invariant, i.e., $\mathcal{T}^{\mathrm{DS},\Xi}[V] \geq \mathcal{T}^{\mathrm{DS},\Xi}[V']$ for any $V, V' \in \mathbb{R}^{|\mathcal{S}|}, V \geq V'$, and $\mathcal{T}^{\mathrm{DS},\Xi}[V+\alpha\boldsymbol{e}] = \mathcal{T}^{\mathrm{DS},\Xi}[V] + \alpha$, where α is a scalar and $\boldsymbol{e}$ is a unit vector of size $|\mathcal{S}|$.*

  (ii) *$\mathcal{T}^{\mathrm{DS},\Xi}[V]$ is a contraction mapping of parameter γ w.r.t. the infinity norm, i.e., for any $V, V' \in \mathbb{R}^{|\mathcal{S}|}$ we have $||\mathcal{T}^{\mathrm{DS},\Xi}[V] - \mathcal{T}^{\mathrm{DS},\Xi}[V']||_\infty \leq \gamma||V - V'||_\infty$.*

  (iii) *If $\sup_{d_s\in\Xi}\mathbb{E}_{\boldsymbol{\epsilon}_s\sim d_s}[\max_a\{r(a|s)+\epsilon(a|s)\}]$ are bounded from above for all $s\in\mathcal{S}$, then $V^{\mathrm{DS},\Xi}$ is the unique fixed point solution to the system $\mathcal{T}^{\mathrm{DS},\Xi}[V] = V$.*

4

(iv) The densities of the optimal decision rules (i.e. optimal policy) $\boldsymbol{\psi}_s$, $\forall s \in \mathcal{S}$, can be computed as

$$\pi^{\mathrm{DS},*}(a|s) = \partial \varphi_s(\boldsymbol{w}_s)/\partial w_{sa}, \ \forall a \in \mathcal{A}, s \in \mathcal{S}$$

where $\boldsymbol{w}_s$ is a vector of size $|\mathcal{A}|$ with entries $w_{sa} = r(a|s) + \gamma \sum_{s'} q(s'|s,a) V^{\mathrm{DS},\Xi}(s')$ and

$$\varphi_s(\boldsymbol{w}_s) = \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s}[\max_a \{w_{sa} + \epsilon(a|s)\}] \right\}. \tag{9}$$

Theorem 1-(i) is quite obvious to validate. The contraction property (ii) can be done by looking at the structure of the mapping (8) and (iii) can be validated by extending (8) recursively. The assumption in (iii) generally holds if the noise $\boldsymbol{\epsilon}_s$ has a finite first moment for any distribution $d_s \in \Xi$ for any $s \in \mathcal{S}$, i.e., $\sup_{d_s \in \Xi} \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s}[\boldsymbol{\epsilon}_s] < \infty$. Given (ii) and (iii), one can find an optimal decision rule by solving the Bellman equation (8). Hence, we can make use of the properties of the semi semi-parametric choice model propose by [25] to obtain the optimal decision rule (iv). Note that in the S-MDP framework, the optimal policy is defined similarly, but for only one distribution d_s. Let $d_s(\mathbf{w}_s)$ denote the distribution or a limit of a sequence of distributions (in case that the set Ξ_s is not compact) that attains the optimal value $\varphi_s(\mathbf{w}_s)$, the optimal policy can be expressed as $\pi^{\mathrm{DS},*}(a|s) = \mathbb{P}_{\boldsymbol{\epsilon}_s \sim d_s(\mathbf{w}_s)} \left[a = \mathrm{argmax}_{a \in \mathcal{A}} \{w_{sa} + \epsilon(a|s)\} \right]$. It is also interesting to look at a relevant version of the DS-MDP, where we take an *inf* over the distribution set Ξ, instead of *sup* as in (6), i.e., $\inf_{\mathbf{d}^0, \mathbf{d}^1, \ldots} \sup_{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots} \{\cdot\}$. In this case, an optimal decision rule is more difficult to identify, even for one-stage problems, thus less interesting [25].

4 Nested Relation between Stochastic and Regularized MDP Frameworks

In this section we analyze the relation between the stochastic and regularized MDP frameworks mentioned above. First, we summarize all the MDP frameworks considered in Tab. 1. The constrained MDP one on the last row will be discussed in the next section.

MDP frameworks	Abbreviation
Standard MDP	MDP
Stochastic EV MDP	EV-MDP
Stochastic MDP	S-MDP
Distributional Stochastic MDP	DS-MDP
Regularized MDP	R-MDP
Entropy Regularized MDP	ER-MDP
Constrained MDP	CT-MDP

Table 1: Summary of MDP frameworks
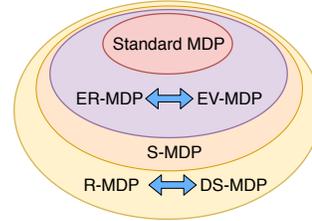
Figure 1: Nested relation between MDP frameworks

In general, we will validate the diagram in Fig. 1, which says that (i) the R-MDP are equivalent to DS-MDPs and strictly covers other frameworks, (ii) the ER-MDP is equivalent to the S-MDP with Gumbel distributed random noises and both are strictly covered by the S-MDP. This diagram shows a broad and insightful connection/relation between regularized and stochastic MDP frameworks, thus provides new ways to interpret each MDP framework through others. To facilitate our exposition, we need the following definitions.

Definition 1 (Equivalent MDP models) *Two MDP models X and Y are said to be "equivalent", denoted by $X \equiv Y$, if and only if they are specified by the same tuple $(\mathcal{S}, \mathcal{A}, \boldsymbol{q}, \gamma)$ and there are a constant vector $\boldsymbol{\alpha}$ of size $(|\mathcal{S}| \times |\mathcal{A}|)$ such that for any reward functions $\boldsymbol{r}^X, \boldsymbol{r}^Y \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, $\boldsymbol{r}^X - \boldsymbol{r}^Y = \boldsymbol{\alpha}$, model X with rewards $\boldsymbol{r}^X$ and model Y with rewards $\boldsymbol{r}^Y$ always yield the same the value function and a common stationary optimal policy $\{\pi^*(a|s), \ \forall a, s\}$.*

The equivalence property between two MDP models implies that the two models are identical to use, in both perspectives of RL and IRL. Under this definition, two ER-MDP models based on a relative entropy regularizer $\eta_1 \sum_a \pi(a|s) \ln(\pi(a|s))$ and a KL divergence $\eta_2 \mathrm{KL}(\boldsymbol{\pi}_s || \bar{\boldsymbol{\pi}}_s)$ are equivalent if $\eta_1 = \eta_2$, but not equivalent if $\eta_1 \neq \eta_2$.

We can view a MDP framework as a set of MDP models, e.g., the S-MDP framework contains S-MDP models, each is specified by a distribution set $\{d_s, \ s \in \mathcal{S}\}$ of the random noises $\{\boldsymbol{\epsilon}_s, \ s \in \mathcal{S}\}$. We define "equivalent" MDP frameworks as follows.

Definition 2 (Equivalent MDP frameworks) *Two MDP frameworks $\mathcal{X}$ and $\mathcal{Y}$ are said to be "equivalent" if and only if for any MDP model $X \in \mathcal{X}$, there is a MDP model $Y \in \mathcal{Y}$ such that $X \equiv Y$, and vice-versa.*

In other words, the "*equivalence*" between two MDP frameworks implies that they share the same set of equivalent MDP models, and the uses of these MDP frameworks are indeed identical. We also define when a framework is strictly subsumed by another one.

Definition 3 (Subsumed MDP framework) *Given two MDP frameworks $\mathcal{X}$ and $\mathcal{Y}$, we say "$\mathcal{X}$ strictly subsumes $\mathcal{Y}$" if and only if for any MDP model $Y \in \mathcal{Y}$ there is always $X \in \mathcal{X}$ such that $X \equiv Y$, but there is $X \in \mathcal{X}$ such that there is no $Y \in \mathcal{Y}$ such that X and Y are equivalent.*

So if a MDP framework $\mathcal{X}$ subsumes another $\mathcal{Y}$, it means $\mathcal{X}$ is more flexible to use. It is clear that the R-MDP subsumes the EV-MDP and standard MDP, and the DS-MDP subsumes the S-MDP and EV-MDP ones (but not sure the subsumption is strict). The question now is what is the relation between two modeling approaches: *stochastic* and *regularized*. We start investigating this by the following proposition.

Proposition 1 *The ER-MDP and EV-MDP frameworks are equivalent and strictly subsume the standard unconstrained MDP one.*

The equivalence between ER-MDP and EV-MDP in the finite-horizon case has been shown in [9] and can be validated easily by looking at the formulations used to compute optimal policies in previous studies [32, 20]. The claim that ER-MDP and EV-MDP strictly subsumes the standard one is also easy to verify, using the fact the standard MDP typically yields deterministic policies while the other frameworks give randomized and soft optimal policies.

There is a well-known similarity between Shannon's entropy [29] and the multinomial logit model [5]. Proposition 1 extends this similarity to a multi-stage decision-making setting. It also indicates a connection between EV-MDPs and the maximum causal entropy principle [32], which implies that an optimal decision rule given by an EV-MDP model will maximize a causal entropy function and satisfy a prediction log-loss guarantee.

An application of the ER-MDP framework is to add a KL divergence $\phi_s(\boldsymbol{\pi}_s) = -\eta \text{KL}(\boldsymbol{\pi}_s || \overline{\boldsymbol{\pi}}_s) = \eta \sum_a \pi(a|s) \ln \frac{\pi(a|s)}{\overline{\pi}(a|s)}$ to control the divergence between a desired optimal policy and a given policy [15, 2]. According to Theorem 1, this ER-MDP model is *equivalent* to a stochastic MDP model with rewards $r'(a|s) = r(a|s) - \eta \overline{\pi}(a|s)$ and Gumbel distributed i.i.d random noises. In other words, by adding stochastic noises to the reward function, we can penalize the divergence between two policies.

We now move to the S-MDP framework and have Proposition 2. It is obvious to see that the EV-MDP is just special case of the S-MDP when the random noises follow the Gumbel distribution. Thus S-MDP subsumes both the EV-MDP and ER-MDP. However, it is not straightforward to see that the S-MDP *strictly* subsumes the other frameworks, as one can show that an ER- or EV-MDP model can yield any optimal optimal policy in Δ^π when varying the reward function. We do this by giving a counter example where the random noises in the S-MDP model are no-longer i.i.d.

Proposition 2 *The S-MDP strictly subsumes the EV-MDP and ER-MDP frameworks.*

We now look at broader frameworks. The R-MDP is the most general frameworks among regularized ones, and the DS-MDP covers all other stochastic MDP frameworks. Interestingly, Theorem 2 below tells use that the two general frameworks are actually equivalent. The proof makes use of Theorem 1 and some results in [10]. The equivalence provides a new way to view regularized MDPs through the lens of stochastic MDPs. That is, adding (concave) regularizers to the MDP model is equivalent to assuming that the reward function contains some random noises of ambiguous distribution. This also shows the robustness of the R-MDP framework from a stochastic-reward viewpoint.

Theorem 2 *The R-MDP and DS-MDP frameworks are equivalent and strictly subsume the S-MDP framework if $|\mathcal{A}| \geq 3$.*

We also see that, given an ambiguity set of continuous distributions Ξ, the DS-MDP model is equivalent to a R-MDP model with regularizers $\phi_s(\boldsymbol{\pi}_s) = -\varphi_s^*(\boldsymbol{\pi}_s)$, where $\varphi_s^*(\boldsymbol{\pi}_s)$ is the convex conjugate of function $\varphi_s(\mathbf{w}_s)$ defined in (9), i.e., $\phi_s(\boldsymbol{\pi}_s) = -\varphi_s^*(\boldsymbol{\pi}_s) = -\sup_{\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}} \{ \boldsymbol{\pi}_s^{\mathsf{T}} \mathbf{w}_s -$

$\varphi_s(\mathbf{w}_s)\}$. Under some ambiguity settings with marginal information, one can obtain a closed form to compute the convex conjugate functions $\varphi_s^*(\boldsymbol{\pi}_s)$ and we provide some examples in Proposition 3 below, which may suggest new regularizers with interpretations from a stochastic-rewards viewpoint. Proposition 3 below provides some examples of DS-MDP models and their equivalent R-MDP ones. These examples extend some results developed for the persistent model [25] to our MDP settings.

Proposition 3 *The DS-MDP models on the left column are equivalent to the corresponding R-MDP models on the right column. where $F_{as}(\cdot)$ is the marginal (continuous) cumulative distribution*

Case	DS-MDP, for any $s \in \mathcal{S}$	R-MDP, for any $s \in \mathcal{S}$
(i)	$\Xi_s = \left\{ d_s \middle| \epsilon(a\|s) \sim F_{sa}(\cdot), \forall a \in \mathcal{A} \right\}$	$\phi_s(\boldsymbol{\pi}_s) = \sum_{a \in \mathcal{A}} \int_{1-\pi(a\|s)}^1 F_{sa}^{-1}(t)dt$
(ii)	$\Xi_s = \left\{ d_s \middle| \begin{array}{l} \mathrm{dev}(\epsilon(a\|s)) = \sigma(a\|s) \\ \mathbb{E}[\epsilon(a\|s)] = 0 \end{array} \right\}$	$\phi_s(\boldsymbol{\pi}_s) = \sum_{a \in \mathcal{A}} \sigma(a\|s)\sqrt{\pi(a\|s)(1 - \pi(a\|s))}$
(iii)	$\Xi_s = \left\{ d_s \middle| \begin{array}{l} \mathrm{Cov}(\boldsymbol{\epsilon}_s) = \Sigma_s \\ \mathbb{E}[\boldsymbol{\epsilon}_s] = 0 \end{array} \right\}$	$\phi_s(\boldsymbol{\pi}_s) = \mathrm{trace}\left(\Sigma_s^{1/2} \left(\mathrm{Diag}(\boldsymbol{\pi}_s) - \boldsymbol{\pi}_s\boldsymbol{\pi}_s^{\mathrm{T}}\right)\Sigma_s^{1/2} \right)^{1/2}$

function of $\epsilon(a\|s)$, $dev(\epsilon(a\|s))$ is the marginal standard deviation of $\epsilon(a\|s)$, $trace(\mathbf{A})$ is the trace of matrix $\mathbf{A}$, $Diag(\mathbf{v})$ is the diagonal matrix with the elements of vector $\mathbf{v}$ on the main diagonal, and $Cov(\boldsymbol{\epsilon}_s)$ is the covariance matrix of the random vector $\boldsymbol{\epsilon}_s$. Here we assume that Σ_s is positive definite.

For (i) and (ii), the models are based on the the marginal distribution and marginal moment models considered [25], in which only the cumulative distribution, or the first and moments of the marginal distributions of the random noises are known. Example (iii) concerns models in which only the covariance information of the distribution in Ξ is known and the equivalent R-MDP model is derived using [24].

5 Connection to Constrained MDPs

We investigate a relation between the MDP frameworks considered above and the constrained MDP (CT-MDP) one, an extension of the standard MDP where constraints are imposed on the policy. A typically CT-MDP model has constraints on a long-term performance of the policy [4]. In this section, we only consider constraints on state policies, i.e., the policy variables at each state belong to a feasible set, and these sets are independent over states. This allows to solve the constrained MDP problem by value iteration or policy iteration as in the case of standard MDPs.

5.1 Constrained MDPs

The constrained version requires that the policy $\boldsymbol{\pi}$ belong to a set $\boldsymbol{\pi} \in \mathcal{D}^{\pi} \subset \Delta^{\pi}$. We further assume that $\mathcal{D}^{\pi}$ is state-wise decomposable, i.e., $\mathcal{D}^{\pi} = \otimes_{s \in \mathcal{S}} \mathcal{D}_s^{\pi}$, where $\mathcal{D}_s^{\pi} \subset \Delta(|\mathcal{A}|)$ is an support set for the policy $\{\pi(a|s),\ a \in \mathcal{A}\}$, for any $s \in \mathcal{S}$. We further assume that $\mathcal{D}_s^{\pi}$ is nonempty. The Markov decision problem becomes

$$\sup_{\substack{\boldsymbol{\pi}^t \in \mathcal{D}^{\pi} \\ t=0,\dots}} \left\{ \mathbb{E}_{\tau \sim (\mathbf{\Pi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} r(a_t|s_t) \right] \right\}$$

We can then define a value function $V^{\mathrm{CT}, \mathcal{D}^{\pi}}$ and Bellman equation $\mathcal{T}^{\mathrm{CT}, \mathcal{D}^{\pi}}[V]$ analogously to the standard MDP, except that the policy $\boldsymbol{\pi}$ are required to lie in the set $\mathcal{D}^{\pi}$. Note that the Bellman update in this case $\mathcal{T}^{\mathrm{CT}, \mathcal{D}^{\pi}}$ is contraction for any non-empty set $\mathcal{D}^{\pi}$.

5.2 CT-MDP and other MDP frameworks

We first note that the CT-MDP covers the standard MDP framework when $\mathcal{D}_s^{\pi}$ is the simplex in $\mathbb{R}^{|\mathcal{A}|}$. Even-though the CT-MDP looks general, as we can introduce any feasible set $\mathcal{D}_s^{\pi}$, Theorem 3 below shows that, in the context of Definitions 1 and 2, the CT-MDP framework is not subsumed or cover any stochastic or regularized MDP frameworks considered above, except the standard one. On the other hand, the R-MDP and DS-MDP frameworks are the most general among those presented in Fig. 1, but cannot cover the CT-MDP.

Theorem 3 *There is an ER-MDP model X with parameter $\eta > 0$ such that there is no CT-MDP model that is equivalent to X. On the other hand, there is also a CT-MDP model Y such that there is no R-MDP model being equivalent to Y.*

However, if we fix the reward function, then we can show that, even though the CT-MDP framework does not cover the R-MDP one, but for any R-MDP model, there is a CT-MDP model specified by convex constraints that yields an identical optimal policy.

Proposition 4 *Given a R-MDP model defined by a tuple $(\mathcal{S}, \mathcal{A}, \boldsymbol{q}, \boldsymbol{r}, \gamma)$ and the set of strictly concave functions $\{\phi_s | \ s \in \mathcal{S}\}$, there exists a set of constants $\{c_s, \ s \in \mathcal{S}\}$ such that the CT-MDP model defined by the same tuple but with a shifted reward function $\boldsymbol{r}'_s = \boldsymbol{r}_s - c_s$ and feasible sets $\mathcal{D}^\pi_s = \{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)| \ - \phi_s(\boldsymbol{\pi}_s) \leq c_s\}$ will yield an identical optimal policy.*

It is important to note that $\{c_s, \ \forall s\}$ specified in Proposition 4 depend on the reward vector $\mathbf{r}$, thus the CT-MDP model cannot yield identical optimal policies for any reward function. The reverse is also true, that is, for any CT-MDP model specified by a set of convex constraints, we always can find a R-MDP model that yield an identical optimal policy (see the supplementary material for details).

5.3 Discussion

In general, we have seen that a stochastic MDP model can be interpreted as a regularized one and vice-versa. Proposition 4 shows that any R-MDP or DS-MDP model can be always represented by a CT-MDP one. In other words, adding regularizers or random noises to the reward function is *equivalent* to constraining the desired optimal policies. Let us illustrate this by some examples.

A popular and useful regularizer in RL is a KL divergence $-\eta\mathrm{KL}(\boldsymbol{\pi}_s||\overline{\boldsymbol{\pi}}_s)$ between two consecutive policies. We have shown that this regularized model is equivalent to a stochastic model with shifted rewards and Gumbel distributed random noises. From Proposition 4 we see that these ER-MDP and EV-MDP models can be equivalently represented by a CT-MDP with constraints $\mathrm{KL}(\boldsymbol{\pi}_s||\overline{\boldsymbol{\pi}}_s) \leq c_s/\eta$, where $c_s \geq 0$ for all $s \in \mathcal{S}$. Clearly, if $c_s \to 0$ then the optimal policy should converge to $\overline{\boldsymbol{\pi}}_s$ and if we increase c_s, the feasible set $\mathcal{D}^\pi_s$ will converge to $\Delta(|\mathcal{A}|)$ (i.e. the simplex in $\mathbb{R}^{|\mathcal{A}|}$) and we will retain a deterministic policy as in the unconstrained MDP model.

Several RL and IRL algorithms make use of soft policies resulting from relative entropy regularizers $\phi_s(\boldsymbol{\pi}_s) = -\eta \sum_a \pi(a|s) \ln \pi(a|s)$ [32]. Let $\widetilde{\boldsymbol{\pi}}^*_s = \mathrm{argmax}_{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)} \phi_s(\boldsymbol{\pi}_s)$, we see that (i) $\widetilde{\boldsymbol{\pi}}^*_s = \mathbf{e}/|\mathcal{A}|$ (a random walk, $\mathbf{e}$ is a unit vector of size $|\mathcal{A}|$) and (ii) for any $\boldsymbol{\pi}_s$

$$\mathrm{KL}(\boldsymbol{\pi}_s||\widetilde{\boldsymbol{\pi}}^*_s) = \sum_{a \in \mathcal{A}} \pi(a|s) \ln \pi(a|s) - \ln |\mathcal{A}| = -\phi_s(\boldsymbol{\pi})/\eta - \ln |\mathcal{A}|.$$

Thus the constraint $-\phi_s(\boldsymbol{\pi}_s) \leq c_s$ becomes $\mathrm{KL}(\boldsymbol{\pi}_s||\widetilde{\boldsymbol{\pi}}^*_s) \leq c_s/\eta - \ln |\mathcal{A}|$. So, we see that by adding the relative entropy terms and varying the parameter η, we are controlling the KL divergence between a desired policy and the random walk vector $\widetilde{\boldsymbol{\pi}}^*_s$. Making the optimal policy closer to the random walk indeed helps improve exploration.

This analysis can be extended to a general R-MDP model with (strictly) concave regularizers $\{\phi_s, s \in \mathcal{S}\}$. Also let $\widetilde{\boldsymbol{\pi}}^*_s = \mathrm{argmax}_{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)} \phi_s(\boldsymbol{\pi}_s)$, which are always uniquely determined. We can show that if $\widetilde{\boldsymbol{\pi}}^*_s$ contain non-zero elements, then we can write $\nabla \phi_s(\widetilde{\boldsymbol{\pi}}^*_s) = \alpha \mathbf{e}$, where α is a scalar. The Bregman divergence (BD) generated by $-\phi_s(\cdot)$ between two policies becomes

$$\mathrm{BD}^{-\phi_s}(\boldsymbol{\pi}_s||\boldsymbol{\pi}^*_s) = -\phi_s(\boldsymbol{\pi}_s) + \phi_s(\widetilde{\boldsymbol{\pi}}^*_s) + \alpha - \alpha \mathbf{e}^\mathsf{T}\widetilde{\boldsymbol{\pi}}^*_s.$$

Thus, the constraints $-\phi_s(\boldsymbol{\pi}_s) \leq c_s$ becomes $\mathrm{BD}^{-\phi_s}(\boldsymbol{\pi}_s||\boldsymbol{\pi}^*_s) \leq c_s + \phi_s(\widetilde{\boldsymbol{\pi}}^*_s) + \alpha - \alpha \mathbf{e}^\mathsf{T}\widetilde{\boldsymbol{\pi}}^*_s$. This generally implies that using a regularized MDP model is equivalent to imposing a Bregman divergence requirement on the desired policies.

The idea of using KL or Bregman divergence to control the distance between two consecutive policies in a policy iteration is interesting and useful [6, 2, 1, 22, 21, 8]. The CT-MDP framework provides a flexible and general way to implement it. For example, one can consider other types distance measures at a greedy step, i.e., we solve a constrained Bellman update with feasible sets defined as $\mathcal{D}^\pi_s = \{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|), \ ||\boldsymbol{\pi}_s - \boldsymbol{\pi}^k_s|| \leq \xi\}$, where the norm $|| \cdot ||$ can be L1, L2 or L-infinity norm. If the norm is L1 or L-infinity, we greedy step can be efficiently performed by solving a linear program.

If the norm is L2, then the greed step is still a convex optimization problem, which is generally easy to solve. We can incorporate several previous policies to have a more robust greedy step, i.e., $\mathcal{D}_s^\pi = \{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|), \; \max_{i \in \mathcal{K}} ||\boldsymbol{\pi}_s - \boldsymbol{\pi}_s^i|| \leq \xi\}$, where $\mathcal{K}$ is a set of indexes of previous policies considered. The constraints require that a desired policy after the greedy step should not be *too far* from a set of some policies computed from previous iterations. Moreover, different types of distance measures can be combined, e.g., KL divergence or different norms, while keeping the constrained Bellman update tractable. We provide some computational complexity analyses for such constrained models in the supplementary material.

6 Conclusion

In this paper we study the relation between stochastic, regularized and constrained MDP frameworks. We first propose the DS-MDP by assuming the distribution of the random noises in the S-MDP is ambiguous and we validate some basic properties of the new framework, which allows to solve a DS-MDP problem by value iterations. We show that there is a nested relationship between stochastic and regularized MDP frameworks, which provides new ways to interpret regularized MDP frameworks through the lens of stochastic ones and vice-versa. We also show a strong connection between stochastic/regularized MDPs and constrained MDPs, and show how this finding would be useful in some well-known algorithmic schemes.

References

[1] Abbas Abdolmaleki, Jost Tobias Springenberg, Jonas Degrave, Steven Bohez, Yuval Tassa, Dan Belov, Nicolas Heess, and Martin Riedmiller. Relative entropy regularized policy iteration. *arXiv preprint arXiv:1812.02256*, 2018.

[2] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. *arXiv preprint arXiv:1806.06920*, 2018.

[3] Victor Aguirregabiria and Pedro Mira. Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67, 2010.

[4] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

[5] Alex Anas. Discrete choice theory, information theory and the multinomial logit and gravity models. *Transportation Research Part B: Methodological*, 17(1):13–23, 1983.

[6] Mohammad Gheshlaghi Azar, Vicenç Gómez, and Hilbert J Kappen. Dynamic policy programming. *Journal of Machine Learning Research*, 13(Nov):3207–3245, 2012.

[7] Dimitris Bertsimas, Karthik Natarajan, and Chung-Piaw Teo. Persistence in discrete optimization under data uncertainty. *Mathematical programming*, 108(2-3):251–274, 2006.

[8] Yinlam Chow, Ofir Nachum, and Mohammad Ghavamzadeh. Path consistency learning in tsallis entropy regularized mdps. In *International Conference on Machine Learning*, pages 979–988, 2018.

[9] Stefano Ermon, Yexiang Xue, Russell Toth, Bistra Dilkina, Richard Bernstein, Theodoros Damoulas, Patrick Clark, Steve DeGloria, Andrew Mude, Christopher Barrett, et al. Learning large-scale dynamic discrete choice models of spatio-temporal preferences with application to migratory pastoralism in east africa. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[10] Guiyun Feng, Xiaobo Li, and Zizhuo Wang. On the relation between several discrete choice models. *Operations research*, 65(6):1516–1525, 2017.

[11] Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*, 2016.

[12] Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.

[13] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.

[14] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. *arXiv preprint arXiv:1901.11275*, 2019.

[15] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1352–1361. JMLR. org, 2017.

[16] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International conference on machine learning*, 2018.

[17] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

[18] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.

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

[20] Rust John. Maximum likelihood estimation of discrete control processes. *SIAM journal on control and optimization*, 26(5):1006–1024, 1988.

[21] Kyungjae Lee, Sungjoon Choi, and Songhwai Oh. Sparse markov decision processes with causal sparse tsallis entropy regularization for reinforcement learning. *IEEE Robotics and Automation Letters*, 3(3):1466–1473, 2018.

[22] Daniel J Mankowitz, Nir Levine, Rae Jeong, Abbas Abdolmaleki, Jost Tobias Springenberg, Timothy Mann, Todd Hester, and Martin Riedmiller. Robust reinforcement learning for continuous control with model misspecification. *arXiv preprint arXiv:1906.07516*, 2019.

[23] Daniel McFadden. Econometric models of probabilistic choice. *Structural analysis of discrete data with econometric applications*, 198272, 1981.

[24] Vinit Kumar Mishra, Karthik Natarajan, Dhanesh Padmanabhan, Chung-Piaw Teo, and Xiaobo Li. On theoretical and empirical aspects of marginal distribution choice models. *Management Science*, 60(6):1511–1531, 2014.

[25] Karthik Natarajan, Miao Song, and Chung-Piaw Teo. Persistency model and its applications in choice modeling. *Management Science*, 55(3):453–469, 2009.

[26] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

[27] John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033, 1987.

[28] John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.

[29] Claude E Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948.

[30] J Michael Steele. *Probability theory and combinatorial optimization*, volume 69. Siam, 1997.

[31] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.

[32] Brian D Ziebart, J Andrew Bagnell, and Anind K Dey. Modeling interaction via the principle of maximum causal entropy. 2010.

A Proofs of the Main Results

A.1 Proof of Theorem 1

For (i), we see that, for any $V, V' \in \mathbb{R}^{|\mathcal{S}|}$ such that $V \geq V'$

$$\max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V(s') \right\}$$

$$\geq \max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V'(s') \right\}, \quad \forall s \in \mathcal{S}, \quad (10)$$

thus $\mathcal{T}^{\mathrm{DS},\Xi}[V](s) \geq \mathcal{T}^{\mathrm{DS},\Xi}[V'](s)$ for any $s \in \mathcal{S}$. The translation variance property is also easy to validate, as for any $s \in \mathcal{S}$

$$\mathcal{T}^{\mathrm{DS},\Xi}[V + \alpha \mathbf{e}](s) = \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} \left[\max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V(s') + \alpha \right\} \right] \right\}$$

$$= \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} \left[\max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V(s') \right\} \right] \right\} + \alpha$$

$$= \mathcal{T}^{\mathrm{DS},\Xi}[V](s) + \alpha.$$

For the **contraction property (ii)**, give any $V^1, V^2 \in \mathbb{R}^{|\mathcal{S}|}$ and a state $s \in \mathcal{S}$, in analogy to a standard proof, we consider two cases: $\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) \geq \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s)$ and show that $\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) \geq \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s) \leq ||V^1 - V^2||_\infty$, the other case can be done in a similar way. Now for any $\xi_1 > 0$, let us choose a distribution $\overline{d}_s^1 \in \Xi_s$ such that

$$\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) \leq \mathbb{E}_{\boldsymbol{\epsilon}_s \sim \overline{d}_s^1} \left[\max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V^1(s') \right\} \right] + \xi_1$$

For the sake of simplicity, we denote

$$z_s^1(a|\boldsymbol{\epsilon}_s) = r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V^1(s')$$

$$z_s^2(a|\boldsymbol{\epsilon}_s) = r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a)V^2(s')$$

For any fixed $\boldsymbol{\epsilon}_s$, we write

$$|\max_a\{z_s^1(a|\boldsymbol{\epsilon}_s)\} - \max_a\{z_s^2(a|\boldsymbol{\epsilon}_s)\}|$$

$$= \max \left\{ \max_a\{z_s^1(a|\boldsymbol{\epsilon}_s)\} - \max_a\{z_s^2(a|\boldsymbol{\epsilon}_s); \max_a\{z_s^2(a|\boldsymbol{\epsilon}_s)\} - \max_a\{z_s^1(a|\boldsymbol{\epsilon}_s)\} \right\}$$

$$\overset{(i)}{\leq} \max\{z_s^1(a_1|\boldsymbol{\epsilon}_s) - z_s^2(a_1|\boldsymbol{\epsilon}_s); \; z_s^2(a_2|\boldsymbol{\epsilon}_s) - z_s^1(a_2|\boldsymbol{\epsilon}_s)\}$$

$$= \gamma \max \left\{ \sum_{s'} q(s'|s,a_1)(V^1(s') - V^2(s')); \; \sum_{s'} q(s'|s,a_2)(V^2(s') - V^1(s')) \right\}$$

$$\leq \gamma ||V^1 - V^2||_\infty \quad (11)$$

where in (i), $a_1 = \mathrm{argmax}_{a \in \mathcal{A}} \{z^1(a|\boldsymbol{\epsilon}_s)\}$ and $a_2 = \mathrm{argmax}_{a \in \mathcal{A}} \{z^2(a|\boldsymbol{\epsilon}_s)\}$. Now we evaluate $\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) - \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s)$ as follows

$$\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) - \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s)$$

$$\leq \mathbb{E}_{\boldsymbol{\epsilon}_s \sim \overline{d}_s^1} \left[\max_{a \in \mathcal{A}} \left\{ z^1(a|\boldsymbol{\epsilon}_s) \right\} \right] - \mathbb{E}_{\boldsymbol{\epsilon}_s \sim \overline{d}_s^1} \left[\max_{a \in \mathcal{A}} \left\{ z^2(a|\boldsymbol{\epsilon}_s) \right\} \right] + \xi_1$$

$$\leq \int_{\boldsymbol{\epsilon}_s \sim \overline{d}_s^1} \left| \max_{a \in \mathcal{A}} \left\{ z^1(a|\boldsymbol{\epsilon}_s) \right\} - \max_{a \in \mathcal{A}} \left\{ z^2(a|\boldsymbol{\epsilon}_s) \right\} \right| f_s(\boldsymbol{\epsilon}_s) d(\boldsymbol{\epsilon}_s) + \xi - 1$$

$$\overset{(ii)}{\leq} \int_{\boldsymbol{\epsilon}_s \sim \overline{d}_s^1} \gamma ||V^1 - V^2||_\infty f_s(\boldsymbol{\epsilon}_s) d(\boldsymbol{\epsilon}_s) + \xi_1 = \gamma ||V^1 - V^2||_\infty + \xi_1,$$

12

where $f_s(\cdot)$ is the probability density function of $\boldsymbol{\epsilon}_s$ and (ii) is due to (11). Since ξ_1 can be chosen arbitrarily small, we let $\xi_1 \to 0$ and obtain $\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) - \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s) \leq ||V^1 - V^2||_\infty$. In a similar manner, we can show that if $\mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) < \mathcal{T}^{\mathrm{DS},\Xi}[V^2](s)$ then we also have $\mathcal{T}^{\mathrm{DS},\Xi}[V^2](s) - \mathcal{T}^{\mathrm{DS},\Xi}[V^1](s) \leq \gamma ||V^1 - V2||$. Thus, $||\mathcal{T}^{\mathrm{DS},\Xi}[V^2] - \mathcal{T}^{\mathrm{DS},\Xi}[V^1]||_\infty \leq \gamma ||V^1 - V^2||$, indicating the contraction property of $\mathcal{T}^{\mathrm{DS},\Xi}[V]$.

We now validate the **Markov optimality (iii)**, i.e., $V^{\mathrm{DS},\Xi}$ defined in (7) is the unique fixed point solution to the system $\mathcal{T}^{\mathrm{DS},\Xi}[V] = V$. The contraction property guarantees that the system $\mathcal{T}^{\mathrm{DS},\Xi}[V] = V$ always has a unique fixed point solution. We now prove that this a solution is also $V^{\mathrm{DS},\Xi}$ defined in (7). Let V^* be a solution to the above contraction system, then for any decision policy $\boldsymbol{\Psi} = \{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots\}$ and any sequence of distribution $\{\tilde{\mathbf{d}}^0, \tilde{\mathbf{d}}^1, \ldots\}$, where $\tilde{\mathbf{d}}^t \in \Xi$ for $t = 0, 1, \ldots$, we have

$$V^*(s) = \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} \left[\max_{a \in \mathcal{A}} \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|s,a) V(s') \right\} \right] \right\}$$

$$\geq \mathbb{E}_{\boldsymbol{\epsilon}_{s_0} \sim \tilde{d}^0_{s_0}} \left[r(a_0|s_0) + \epsilon(a_0|s_0) + \gamma \sum_{s'} q(s'|s_0, a_0) V(s') \,\Big|\, s_0 = s, a_0 \sim \boldsymbol{\psi}^0 \right]$$

$$\geq \mathbb{E}_{\tilde{\mathbf{d}}^0, \tilde{\mathbf{d}}^1} \left[r(a_0|s_0) + \epsilon(a_0|s_0) + \gamma(r(a_1|s_1) + \epsilon(a_1|s_1)) + \gamma^2 V^*(s_2) \right.$$

$$\left. \Big|\, s_0 = s, (a_0, a_1) \sim (\boldsymbol{\psi}^0, \boldsymbol{\psi}^1) \right]$$

By continuing expanding the recursion, we obtain

$$V^*(s) \geq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots, \tilde{\mathbf{d}}^n} \left[\sum_{t=0}^n \gamma^{[t]}(r(a_t|s_t) + \epsilon(a_t|s_t)) + \gamma^{[n+1]} V^*(s_{n+1}) \right.$$

$$\left. \Big|\, s_0 = s, (a_0, \ldots, a_n) \sim (\boldsymbol{\psi}^0, \ldots, \boldsymbol{\psi}^n) \right]$$

$$\geq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=0}^\infty \gamma^{[t]}(r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \boldsymbol{\Psi} \right] - \gamma^{[n+1]} ||V^*||_\infty$$

$$- \gamma^{[n+1]} \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=n+1}^\infty \gamma^{[t-n-1]}(r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \boldsymbol{\Psi} \right]$$

$$\geq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=0}^\infty \gamma^{[t]}(r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \boldsymbol{\Psi} \right] - \gamma^{[n+1]} ||V^*||_\infty - \frac{\gamma^{[n+1]}}{1-\gamma} \alpha$$

$$(12)$$

where $\alpha = \sup_{s \in \mathcal{S}} \sup_{d_s \in \Xi_s} \mathbb{E} \left[\max_a \left\{ r(a|s) + \epsilon(a|s) \right\} \right] < \infty$ which is always finite, according to our assumptions. From (12), by letting $n \to \infty$ to obtain, for any decision policy $\boldsymbol{\Psi} = \{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots\}$ and sequence of distribution $\{\tilde{\mathbf{d}}^0, \tilde{\mathbf{d}}^1, \ldots\}$,

$$V^*(s) \geq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=0}^\infty \gamma^{[t]}(r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \boldsymbol{\Psi} \right]$$

Thus,

$$V^*(s) \geq \sup_{\substack{\mathbf{d}^t \in \Xi; \boldsymbol{\psi}^t \\ t=0, \ldots}} \mathbb{E}_{\substack{\boldsymbol{\epsilon}^t \sim \mathbf{d}^t \\ t=0, \ldots}} \left[\sum_{t=0}^\infty \gamma^{[t]}(r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \boldsymbol{\Psi} \right] \quad (13)$$

Now, to prove the opposite direction of Inequality (13), the Bellman equation tells use that for any $\xi > 0$, there are distributions $\{\tilde{\mathbf{d}}^0, \tilde{\mathbf{d}}^1, \ldots\}$ and decision rule $\mathbf{\Psi} = \{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1, \ldots\}$ such that

$$V^*(s) \leq \mathbb{E}_{\boldsymbol{\epsilon}_{s_0}^0 \sim \tilde{d}_{s_0}^0} \left[r(a_0|s_0) + \epsilon(a_0|s_0) + \gamma V^*(s_1) \,\Big|\, s_0 = s, a_0 \sim \boldsymbol{\psi}^0 \right] + \xi$$

$$\leq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots, \tilde{\mathbf{d}}^n} \left[\sum_{t=0}^n \gamma^{[t]} (r(a_t|s_t) + \epsilon(a_t|s_t)) + \gamma^{[n+1]} V^*(s_{n+1}) \right.$$

$$\left. \,\Big|\, s_0 = s, (a_0, \ldots, a_n) \sim (\boldsymbol{\psi}^0, \ldots, \boldsymbol{\psi}^n) \right] + \xi \sum_{t=0}^n \gamma^{[t]}$$

$$\leq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=0}^\infty \gamma^{[t]} (r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \mathbf{\Psi} \right] + \gamma^{[n+1]} \|V^*\|_\infty$$

$$+ \gamma^{[n+1]} \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=n+1}^\infty \gamma^{t-n-1} (r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \mathbf{\Psi} \right] + \xi \sum_{t=0}^n \gamma^{[t]}$$

$$\leq \mathbb{E}_{\tilde{\mathbf{d}}^0, \ldots} \left[\sum_{t=0}^\infty \gamma^{[t]} (r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \mathbf{\Psi} \right] + \gamma^{[n+1]} \|V^*\|_\infty +$$

$$+ \frac{\gamma^{[n+1]} \alpha}{1 - \gamma} + \xi \sum_{t=0}^n \gamma^{[t]}$$

So let $n \to \infty$ we have

$$V^*(s) \leq \sup_{\substack{\mathbf{d}^t \in \Xi; \boldsymbol{\psi}^t \\ t=0, \ldots}} \mathbb{E}_{\substack{\boldsymbol{\epsilon}^t \sim \mathbf{d}^t \\ t=0, \ldots}} \left[\sum_{t=0}^\infty \gamma^{[t]} (r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \mathbf{\Psi} \right] + \frac{\xi}{1 - \gamma},$$

then by letting $\xi \to 0$ and combining with (13) we have

$$V^*(s) = \sup_{\substack{\mathbf{d}^t \in \Xi; \boldsymbol{\psi}^t \\ t=0, \ldots}} \mathbb{E}_{\substack{\boldsymbol{\epsilon}^t \sim \mathbf{d}^t \\ t=0, \ldots}} \left[\sum_{t=0}^\infty \gamma^{[t]} (r(a_t|s_t) + \epsilon(a_t|s_t)) \,\Big|\, s_0 = s, (a_0, \ldots) \sim \Psi \right] = V^{\mathrm{DS},\Xi}(s), \ \forall s \in \mathcal{S},$$

which is the desired result.

For (iv), from the above results, we see that an optimal decision rule can be computed as

$$\psi^*(s|\boldsymbol{\epsilon}_s) = \mathrm{argmax}_a \left\{ r(a|s) + \epsilon(a|s) + \gamma \sum_{s'} q(s'|a, s) V^{\mathrm{DS},\Xi}(s') \right\}$$

and assume that d_s^* is the probability distribution of $\boldsymbol{\epsilon}_s$ that attains the maximization in the Bellman update $\mathcal{T}^{\mathrm{DS},\Xi}[V(s)$, i.e.,

$$d_s^* = \mathrm{argmax}_{d_s \in \Xi_s} \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} \left[\max_a \left\{ w_{sa} + \epsilon(a|s) \right\} \right],$$

where $w_{sa} = r(a|s) + \gamma \sum_{s'} q(s'|a, s) V^{\mathrm{DS},\Xi}(s')$ for any $a \in \mathcal{A}, s \in \mathcal{S}$. The density of the optimal decision rule $\boldsymbol{\psi}_s$ (i.e. optimal policy) becomes

$$\pi^{\mathrm{DS},*}(a|s) = \mathbb{P}\left[\psi^*(s|\boldsymbol{\epsilon}_s) = a \,\Big|\, \boldsymbol{\epsilon} \sim d_s^* \right]$$

$$\stackrel{(i)}{=} \frac{\partial \varphi_s(\mathbf{w}_s)}{\partial w_{sa}},$$

where $\varphi_s(\mathbf{w}_s) = \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s^*}[\max_a \{w_{sa} + \epsilon(a|s)\}]$ and (i) is due to [23]. Note that

$$\varphi_s(\mathbf{w}_s) = \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s}[\max_a \{w_{sa} + \epsilon(a|s)\}] \right\}.$$

We obtain the desired results and complete the proof.

A.2 Proof of Proposition 1

Consider an ER-MDPs defined by the following maximization of expected discounted regularized reward

$$\sup_{\boldsymbol{\pi}^0,\dots} \left\{ \mathbb{E}_{\tau \sim (\boldsymbol{\Pi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) - \eta \ln(\pi^t(a_t|s_t)) \Big) \right] \right\}, \tag{14}$$

where $\eta \geq 0$. According to [31, 15], the following policy is optimal to the ER-MDP problem under reward function $\mathbf{r}$

$$\pi^{\text{EV},\eta}(a|s) = \frac{\exp\left(h(a|s)/\eta\right)}{\sum_{a' \in \mathcal{A}} \exp\left(h(a'|s)/\eta\right)}, \quad \forall s \in \mathcal{S}, a \in \mathcal{A} \tag{15}$$

where $h(a|s) = r(a|s) + \gamma \sum_{s'} q(s'|s,a) V^{\text{EV}}(s')$, V^{EV} is the unique fixed point solution to the system of equations

$$V(s) = \eta \ln \left(\sum_{a \in \mathcal{A}} \exp \left(\frac{1}{\eta} r(a|s) + \gamma \frac{1}{\eta} \sum_{s'} q(s'|s,a) V(s') \right) \right).$$

Now look at the an EV-MDP defined as follows

$$\sup_{\boldsymbol{\psi}^0, \boldsymbol{\psi}^1,\dots} \mathbb{E}_{\boldsymbol{\epsilon}} \left\{ \left[\mathbb{E}_{\tau \sim (\boldsymbol{\Psi}, \mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} \Big(r(a_t|s_t) + \eta \epsilon(a_t|s_t) \Big) \right] \right] \right\}, \tag{16}$$

where $\epsilon(a|s)$, $\forall a, s$, are i.i.d and follow the Gumbel distribution, then [27] show that the MDP problem (16) yield a value function and a set of optimal choice probabilities (i.e. policy) that are exactly equal to those given in (15). This indicates the equivalence between EV-MDPs and ER-MDPs. We also note here that if $\eta = 0$, which describes the case that the entropy terms $\pi(a|s) \ln \pi(a|s)$ or the noise term $\epsilon(a|s)$ are neglected from the MDPs, then the EV-MDP and ER-MDP become the standard unconstrained MDP. On the other hand, it is not difficult to give an example of MDP such that any optimal policy to the standard MDP need to satisfy $\pi^*(a|s)$ only takes a value of 0 or 1, but any optimal policy to the ER-MDP problem satisfies $\pi^{*,\text{EV}}(a|s) \in (0,1)$ for some pairs of $(a, s) \in \mathcal{A} \times \mathcal{S}$. This validates the claim that ER-MDPs and ER-MDPs strictly subsumes the standard MDPs. We complete the proof.

A.3 Proof of Theorem 2

To prove the equivalence, let us consider a R-MDP model specified by a tuple $(\mathcal{S}, \mathcal{A}, \boldsymbol{\pi}, \mathbf{q}, \mathbf{r}, \gamma)$ and by the set of concave function $\boldsymbol{\phi} = \{\phi_s, \ s \in \mathcal{S}\}$, and a DS-MDP model specified by the same tuple and ambiguity sets $\{\Xi_s, s \in \mathcal{S}\}$. Theorem 1 and [14] tell us that these two models yield the following value functions that are unique solutions to the following contraction systems

$$\mathcal{T}^{\text{RE},\boldsymbol{\phi}}[V](s) = \sup_{\boldsymbol{\pi}_s} \left\{ \mathbf{w}_s^{\text{T}} \boldsymbol{\pi}_s + \phi_s(\boldsymbol{\pi}_s) \right\} \tag{17}$$

$$\mathcal{T}^{\text{DS},\Xi}[V](s) = \sup_{d_s \in \Xi_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} [\max_a \{w_{sa} + \epsilon(a|s)\}] \right\}, \tag{18}$$

where $\mathbf{w}_s$ is a vector of size $|\mathcal{A}|$ with entries $w_{sa} = r(a|s) + \gamma \sum_{s'} q(s'|s,a) V(s')$. Using Theorem 1 in [10], we have that for any strictly convex function $\widetilde{\phi}_s(\cdot)$, there is always a distribution set $\widetilde{\Xi}_s$ such that $\mathcal{T}^{\text{RE},\widetilde{\phi}}[V](s) = \mathcal{T}^{\text{DS},\widetilde{\Xi}}[V](s)$ for any $\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}$, and vise-versa. As a result, for any set of convex functions $\widetilde{\phi} = \{\widetilde{\phi}_s, s \in \mathcal{S}\}$ there always a set $\widetilde{\Xi} = \{\widetilde{\Xi}_s, s \in \mathcal{S}\}$ such that $\mathcal{T}^{\text{RE},\phi}[V] = \mathcal{T}^{\text{DS},\Xi}[V]$ for any $V \in \mathbb{R}^{|\mathcal{S}|}$, and vice-versa. So, for any R-MDP model, there is a DS-MDP model such that the two contraction systems yield the same value functions, i.e., $V^{\text{RE},\widetilde{\phi}} = V^{\text{DS},\widetilde{\Xi}}$, and vice-versa. On other hand, the optimal policies given by the two MDP models are

$$\boldsymbol{\pi}_s^{\text{RE},\widetilde{\phi}} = \text{argmax}_{\boldsymbol{\pi}_s} \left\{ (\mathbf{w}_s^{\text{RE}})^{\text{T}} \boldsymbol{\pi}_s + \widetilde{\phi}_s(\boldsymbol{\pi}_s); \right\} \tag{19}$$

$$\boldsymbol{\pi}_s^{\text{DS},\widetilde{\Xi}} = \nabla_{\mathbf{w}} \left\{ \sup_{d_s \in \widetilde{\Xi}_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} [\max_a \{w_{sa}^{\text{DS}} + \epsilon(a|s)\}] \right\} \right\} \tag{20}$$

where $\mathbf{w}_s^{\text{RE}}$ and $\mathbf{w}_s^{\text{DS}}$ are vectors of size $|\mathcal{A}|$ with entries $w_{sa}^{\text{RE}} = r(a|s) + \gamma \sum_{s'} q(s'|s,a) V^{\text{RE},\widetilde{\boldsymbol{\phi}}}(s')$ and $w_{sa}^{\text{RE}} = r(a|s) + \gamma \sum_{s'} q(s'|s,a) V^{\text{DS},\widetilde{\boldsymbol{\phi}}}(s')$ for all $a \in \mathcal{A}$. Since $V^{\text{RE},\widetilde{\boldsymbol{\phi}}} = V^{\text{DS},\widetilde{\Xi}}$, $\mathbf{w}_s^{\text{RE}} = \mathbf{w}_s^{\text{DS}}$ for all $s \in \mathcal{S}$ and any reward function $\mathbf{r}$. Moreover, from the way we select $\widetilde{\phi}_s$ and $\widetilde{\Xi}_s$ we have

$$\varphi_s(\mathbf{w}_s) = \max_{\boldsymbol{\pi}_s} \left\{ (\mathbf{w}_s)^{\text{T}} \boldsymbol{\pi}_s + \widetilde{\phi}_s(\boldsymbol{\pi}_s) \right\} = \sup_{d_s \in \widetilde{\Xi}_s} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} [\max_a \{w_{sa} + \epsilon(a|s)\}] \right\}, \; \forall \mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}.$$

From [10] we also have

$$\nabla_{\mathbf{w}_s} \varphi_s(\mathbf{w}_s) = \text{argmax}_{\boldsymbol{\pi}_s} \left\{ (\mathbf{w}_s)^{\text{T}} \boldsymbol{\pi}_s + \widetilde{\phi}_s(\boldsymbol{\pi}_s) \right\}.$$

Thus, $\boldsymbol{\pi}_s^{\text{RE},\widetilde{\boldsymbol{\phi}}} = \boldsymbol{\pi}_s^{\text{DS},\widetilde{\Xi}}$ for any reward function $\mathbf{r}$. So, in summary, for any R-MDP model specified by a set $\widetilde{\boldsymbol{\phi}}$, there is always a DS-MDP model specified by a distributions set $\widetilde{\Xi}$ such that $\boldsymbol{\pi}_s^{\text{RE},\widetilde{\boldsymbol{\phi}}} = \boldsymbol{\pi}_s^{\text{DS},\widetilde{\Xi}}$ for any s, and vice-versa. In other words, the R-MDP and DS-MDP frameworks are equivalent.

To prove the claim that the R-MDP framework strictly subsumes the S-MDP one, we just use a result from [10] (for one-step decision problems) saying that if $|\mathcal{A}| \geq 3$ there exist a convex function $\phi_s : \mathbb{R}^{|\mathcal{A}|} \to \mathbb{R}$ such that there is no distribution d_s such that the following holds for any $\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}$

$$\text{argmax}_{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)} \left\{ (\mathbf{w}_s^{\text{T}} \boldsymbol{\pi}_s + \phi_s(\boldsymbol{\pi}_s)) \right\} = \nabla_{\mathbf{w}} \left\{ \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_s \sim d_s} [\max_a \{w_{sa} + \epsilon(a|s)\}] \right\} \right\}.$$

We now complete the proof.

A.4 Proof of Proposition 2

It is clear that the S-MDPs covers the ER-MDPs, thus also covers the EV-MDP framework. To show that there is a S-MDP model that is not equivalent to any EV-MDP model, we just take a small MDP example of three states $\mathcal{S} = \{s_0, s_1, s_2\}$ and assume that there are two possible action $\mathcal{A} = \{a_1, a_2\}$. The transition probabilities are given as $q(s_1|a_1, s_0) = 1$, $q(s_2|a_2, s_0) = 1$, $q(s_1|a_1, s_1) = q(s_1|a_2, s_1) = 1$ and $q(s_2|a_2, s_2) = q(s_2|a_1, s_2)$. In other words, we can move from state s_0 to s_1 with probability 1 by making action a_1 and move from s_0 to s_2 by making action a_2. From states s_1 or s_2, we have a probability of 1 to stay at the same state, under any action. We further assume that $r(a|S_1) = r(a|s_2) = 0$ for all $a \in \mathcal{A}$. It is easy to see that an optimal policy under an EV-MDP or ER-MDP model with parameter $\eta > 0$ is given as

$$\pi^{\text{EV},\eta}(a_1|s_0) = \frac{e^{r(a_1|s_0)/\eta}}{e^{r(a_1|s_0)/\eta} + e^{r(a_2|s_0)/\eta}}; \; \pi^{\text{EV},\eta}(a_2|s_0) = \frac{e^{r(a_2|s_0)/\eta}}{e^{r(a_1|s_0)/\eta} + e^{r(a_2|s_0)/\eta}}; \quad (21)$$

which leads to

$$\frac{\pi^{\text{EV},\eta}(a_1|s_0)}{\pi^{\text{EV},\eta}(a_2|s_0)} = \exp\left(r(a_1|s_0)/\eta - r(a_2|s_0)/\eta\right). \quad (22)$$

Now, to show that the S-MDP framework strictly subsumes the EV-MDP one, we consider the above example and use a S-MDP model with $\epsilon(a|s) = 0$ for all $(a,s) \in \mathcal{A} \times \mathcal{S}$ except $\epsilon(a_1|s_0)$ that follows a uniform distribution in $[0,1]$. Assume that the S-MDP with reward $\mathbf{r}' = \mathbf{r} + \boldsymbol{\alpha}$ will yield similar optimal policies for any $\mathbf{r} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$. Clearly, the S-MDP model has the following optimal policy

$$\pi^{\text{S},\boldsymbol{\epsilon}}(a_1|s_0) = \mathbb{P}\left[r(a_1|s_0) + \alpha_{s_0 a_1} + \epsilon(a_1|s_0) \geq r(a_2|s_0) + \alpha_{s_0 a_2}\right]$$

and the ration between the two action properties under the S-MDP model becomes

$$\frac{\pi^{\text{S},\boldsymbol{\epsilon}}(a_1|s_0)}{\pi^{\text{S},\boldsymbol{\epsilon}}(a_2|s_0)} \overset{(i)}{=} \frac{\mathbb{P}\left[\epsilon(a_1|s_0) \geq r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0}\right]}{\mathbb{P}\left[\epsilon(a_1|s_0) < r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0}\right]}$$

$$= \begin{cases} 0 & \text{if } r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0} \leq 0 \\ \infty & \text{if } r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0} \geq 1 \\ \frac{r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0}}{1 - (r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0})} & \text{if } r(a_2|s_0) - r(a_1|s_0) + \beta_{s_0} \in (0, 1), \end{cases} \quad (23)$$

where in (i), $\beta_{s_0} = \alpha_{s_0 a_2} - \alpha_{s_0 a_1}$. Now we note that the probabilities given in (21) is a unique solution to the EV-MDP model. Moreover, it is clear that the two ratios in (15) and (23) can not be equal for any reward function $\mathbf{r}$. This implies that the EV-MDP framework does not cover the above S-MDP model. Thus, the S-MDP framework strictly subsumes the ER-MDP and ER-MDP ones. We complete the proof.

A.5 Proof of Theorem 3

We first note that, in a CT-MDP model, the value function is defined as

$$V^{\mathrm{CT}, \mathcal{D}^\pi}(s) = \sup_{\pi \in \mathcal{D}^\pi} \left\{ \mathbb{E}_{\tau \sim (\pi, \mathbf{q})} \left[\sum_{t=0}^\infty \gamma^{[t]} r(a_t | s_t) \middle| s = s_0 \right] \right\},$$

and the Bellman equation

$$\mathcal{T}^{\mathrm{CT}, \mathcal{D}^\pi}[V](s) = \sup_{\pi_s \in \mathcal{D}_s^\pi} \left\{ \mathbb{E}_{\pi_s} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V(s') \right] \right\}. \tag{24}$$

Then we see that $\mathcal{T}^{\mathrm{CT}, \mathcal{D}^\pi}[V]$ is a contraction of parameter γ and $V^{\mathrm{CT}, \mathcal{D}^\pi}$ is the unique solution to the system $\mathcal{T}^{\mathrm{CT}, \mathcal{D}^\pi}[V] = V$. Here the contraction property holds for any nonempty set $\mathcal{D}_s^\pi$, not necessarily convex or compact (see Section B.1 below for a more detailed discussion about this). To compute an optimal policy, we need to solve the Bellman equation (24) to obtain $V^{\mathrm{CT}, \mathcal{D}^\pi}$ and solve the following problem to get an optimal policy

$$\pi_s^* = \mathrm{argmax}_{\pi_s \in \mathcal{D}_s^\pi} \left\{ \mathbb{E}_{\pi_s} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V^{\mathrm{CT}, \mathcal{D}^\pi}(s') \right] \right\}.$$

Now, we prove the results by contradiction. For the sake of simplicity, we just use a small MDPs of three states $\mathcal{S} = \{s_0, s_1, s_2\}$ and two actions $\mathcal{A} = \{a_1, a_2\}$ as in the proof of Proposition 2. Now, consider an ER-MDP model based on this set of states and actions, we see that the optimal policy for a given reward function $\mathbf{r}$ becomes $\pi(a_1|s_0) = \exp(r(a_1|s_0))/R$ and $\pi(a_2|s_0) = \exp(r(a_2|s_0))/R$, where $R = \exp(r(a_1|s_0)) + \exp(r(a_2|s_0))$. Clearly, π_{s_0} will span all over the simple $\Delta(2)$ when the reward $\mathbf{r}$ varies in $\mathbb{R}^2$. So if there is a CT-MDP model that is equivalent to the above ER-MDP model, than $\mathcal{D}_{s_0}^\pi$ has to cover all the simplex $\Delta(2)$ as well, i.e., the CT-MDP model becomes a standard unconstrained one, thus yielding deterministic optimal policies (i.e. the probability of making an action only takes a value of 0 or 1).

For the second statement, we consider a CT-MDP model with a singleton feasible set that contains only one deterministic solution $\widetilde{\pi}_{s_0} = [1, 0]$. Now, assume that there is a R-MDP model specified by a set of concave function ϕ that is equivalent to the above CT-MDP model. That is, for any reward function $\mathbf{r}$, the R-MDP model always yields an optimal policy such that $\pi(a_1|s_0) = 1$ and $\pi(a_1|s_0) = 0$. We will prove that this cannot be true for any reward function $\mathbf{r}$. To this end, we look at the Bellman equation at state s_0,

$$\pi_{s_0}^* = \mathrm{argmax}_{\pi_{s_0}} \left\{ \sum_{a \in \mathcal{A}} r(a|s_0) \pi(a|s_0) + \phi(\pi_{s_0}) \right\}, \tag{25}$$

with a note that $\phi_{s_0}(\pi_{s_0})$ is bounded in $\Delta(2)$, i.e., there exists $U^\phi, L^\phi \in \mathbb{R}$ such that $L^\phi \leq \phi(\pi_{s_0}) \leq U^\phi$, $\forall \pi_{s_0} \in \Delta(2)$. So if we set $r(a_1|s_0) = 0$ and $r(a_1|s_0) = |U^\phi - L^\phi| + 1$, then we have $\phi([1, 0]) - \phi([0, 1]) < U^\phi - L^\phi < r(a_2|s_0)$, or $r(a_1|s_0) + \phi_{s_0}([1, 0]) < r(a_2|s_0) + \phi_{s_0}([0, 1])$, implying that $\widetilde{\pi}_{s_0} = [1, 0]$ is not an optimal solution to (25). This contradicts our initial assumption on the equivalence between the R-MDP and CT-MDP models. We complete the proof.

A.6 Proof of Proposition 4

For any vector $\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}$, let π_s^* be an optimal solution to the problem $\sup_{\pi_s} \{\mathbf{w}_s^\mathrm{T} \pi_s + \phi_s(\pi_s)\}$. Moreover, we choose $c_s = -\phi_s(\pi_s^*)$ and consider the two following problems

$$(\text{P1}) \qquad \sup_{\pi_s \in \Delta(|\mathcal{A}|)} \{\mathbf{w}_s^\mathrm{T} \pi_s + \phi_s(\pi_s)\}$$

$$(\text{P2}) \qquad \sup_{\pi_s \in \Delta(|\mathcal{A}|)} \left\{ \mathbf{w}_s^\mathrm{T} \pi_s - c_s \middle| - \phi_s(\pi_s) \leq c_s \right\}$$

We will show that any optimal solution to (P2) is also optimal to (P1). Clearly, if π_s^* is optimal to (P1) than it is feasible to (P2). Now let π_s^{**} is optimal to (P2), from the feasibility of π_s^* to (P2) we have

$$\mathbf{w}_s^\mathrm{T} \pi_s^* \leq \mathbf{w}_s^\mathrm{T} \pi_s^{**}, \tag{26}$$

which leads to

$$\mathbf{w}_s^{\mathrm{T}}\boldsymbol{\pi}_s^* + \phi_s(\boldsymbol{\pi}_s^*) \overset{(i)}{=} \mathbf{w}_s^{\mathrm{T}}\boldsymbol{\pi}_s^* - c_s \overset{(ii)}{\leq} \mathbf{w}_s^{\mathrm{T}}\boldsymbol{\pi}_s^{**} + \phi_s(\boldsymbol{\pi}_s^{**}),$$

where (i) is because $c_s = -\phi_s(\boldsymbol{\pi}_s^{\mathrm{RE},\boldsymbol{\phi}}) = -\phi_s(\boldsymbol{\pi}_s^*)$ and (ii) is due to (26) and the fact that $\boldsymbol{\pi}_s^{**}$ is feasible to (P2), thus $\phi_s(\boldsymbol{\pi}_s^{**}) \geq -c_s$. Since $\boldsymbol{\pi}_s^*$ is optimal to (P1), $\boldsymbol{\pi}_s^{**}$ is also optimal to (P1). Moreover, (P1) has a strictly concave objective function, it always yields a unique optimal solution $\boldsymbol{\pi}_s^*$. Thus for any $\boldsymbol{\pi}_s^{**}$ optimal to (P2), we have $\boldsymbol{\pi}_s^{**} = \boldsymbol{\pi}_s^*$.

Now, let $\boldsymbol{\pi}^*$ is an optimal solution to a R-MDP problem with $\boldsymbol{\phi} = \{\phi_s, s \in \mathcal{S}\}$ and define $c_s = -\phi_s(\boldsymbol{\pi}_s^*)$ and $\mathcal{D}_s^{\pi} = \{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)| - \phi_s(\boldsymbol{\pi}_s) \leq c_s\}$. From the above results, we see that the two following systems always yield the same optimal solution and optimal value.

$$\sup_{\boldsymbol{\pi}_s} \left\{ \mathbb{E}_{\boldsymbol{\pi}_s} \left[r(a|s) + \gamma \sum_{s'} q(s'|a, s) V^{\mathrm{RE},\boldsymbol{\phi}}(s') \right] + \phi_s(\boldsymbol{\pi}_s) \right\}$$

$$\max_{\boldsymbol{\pi}_s \in \mathcal{D}_s^{\pi}} \left\{ \mathbb{E}_{\boldsymbol{\pi}_s} \left[r(a|s) + \gamma \sum_{s'} q(s'|a, s) V^{\mathrm{RE},\boldsymbol{\phi}}(s') \right] \right\}.$$

Therefore, $\mathcal{T}^{\mathrm{RE},\boldsymbol{\phi}}[V^{\mathrm{RE},\boldsymbol{\phi}}](s) = \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^{\mathrm{RE},\boldsymbol{\phi}}](s)$. Since $V^{\mathrm{RE},\boldsymbol{\phi}}$ is a solution to the system $\mathcal{T}^{\mathrm{RE},\boldsymbol{\phi}}[V] = V$, we have

$$\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^{\mathrm{RE},\boldsymbol{\phi}}](s) = V^{\mathrm{RE},\boldsymbol{\phi}}(s).$$

Hence, $V^{\mathrm{RE},\boldsymbol{\phi}}$ is also a fixed point solution to $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V] = V$, i.e., $V^{\mathrm{RE},\boldsymbol{\phi}} = V^{\mathrm{CT},\mathcal{D}^{\pi}}$. As a result, the two models yield the same optimal policy. We complete the proof.

B Relevant Discussions

B.1 Contraction Property of $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}$

We revisit the contraction property of $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V]$ under arbitrarily non-empty feasible sets $\{\mathcal{D}_s^{\pi}, s \in \mathcal{S}\}$. The proof is similar to the unconstrained model but we provide it here for the sake of self-contained. For any $V^1, V^2 \in \mathbb{R}^{|\mathcal{A}|}$, for each state $s \in \mathcal{S}$, we consider two cases: $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1](s) \geq \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2](s)$ or $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1](s) < \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2](s)$. The following evaluations are for the formal case and the latter can be done analogously. For any $\xi > 0$, let $\overline{\boldsymbol{\pi}}_s^1 \in \mathcal{D}_s^{\pi}$ a solution such that

$$\sup_{\boldsymbol{\pi}_s \in \mathcal{D}_s^{\pi}} \left\{ \mathbb{E}_{\boldsymbol{\pi}_s} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V^1(s') \right] \right\} \leq \mathbb{E}_{\overline{\boldsymbol{\pi}}_s^1} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V^1(s') \right] + \xi,$$

and we write, for any $s \in \mathcal{S}$,

$$|\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1](s) - \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2](s)| = \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1](s) - \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2](s)$$

$$\leq \mathbb{E}_{\overline{\boldsymbol{\pi}}_s^1} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V^1(s') \right] - \mathbb{E}_{\overline{\boldsymbol{\pi}}_s^1} \left[r(a|s) + \gamma \sum_{s'} q(s'|s, a) V^2(s') \right] + \xi$$

$$= \gamma \sum_{a \in \mathcal{A}, s' \in \mathcal{S}} \overline{\pi}^1(a|s) q(s'|s, a)(V^1(s') - V^2(s')) + \xi$$

$$\leq \gamma ||V^1 - V^2||_{\infty} + \xi.$$

Thus, $||\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1] - \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2]||_{\infty} \leq \leq \gamma ||V^1 - V^2||_{\infty} + \xi$. Since ξ can be arbitrarily small, we let $\xi \to 0$ to have $||\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1] - \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2]||_{\infty} \leq \leq \gamma ||V^1 - V^2||_{\infty}$. The case that $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^1](s) < \mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V^2](s)$ can be done similarly, which validates the contraction property of $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V]$. It is also clear that the contraction property holds for any non-empty set $\mathcal{D}^{\pi}$.

Now we validate the result that if V^* is the unique solution of the contraction mapping $\mathcal{T}^{\mathrm{CT},\mathcal{D}^{\pi}}[V] = V$, then

$$V^*(s) = V^{\mathrm{CT},\mathcal{D}^{\pi}}(s) = \sup_{\boldsymbol{\pi}^t \in \mathcal{D}^{\pi}; t=0,\dots} \left\{ \mathbb{E}_{\tau \sim (\boldsymbol{\Pi},\mathbf{q})} \left[\sum_{t=0}^{\infty} \gamma^{[t]} r(a_t|s_t) \,\middle|\, s = s_0 \right] \right\}.$$

This can be done similarly as in the unconstrained MDP case, thus we skip the rest.

B.2 From CT-MDPs to R-MDPs

The following proposition shows that a CT-MDP model specified by a set of convex constraints can be equivalently represented by a R-MDP model. The proposition is easy to validate, using Lagrange duality.

Proposition 5 *Given any CT-MDP model $\mathbf{X}$ defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathbf{q}, T = \infty, \mathbf{r}, \gamma)$ and feasible sets $\mathcal{D}_s^\pi = \{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)|\ l_s^k(\boldsymbol{\pi}_s) \leq c_s^k, k = 1, ..., K_s\}$, where $l_s^k(\boldsymbol{\pi}_s)$ are strictly convex in $\Delta(|\mathcal{A}|)$. Assuming also that the Slater conditions hold for all $s \in \mathcal{S}$, i.e., $c_s^k > \inf_{\boldsymbol{\pi}_s} l_s^k(\boldsymbol{\pi}_s)$, $\forall k \in [K], s \in \mathcal{S}$, then there are a set of constants $0 \leq \overline{\lambda}_{sk} < \infty$, $\forall s, k$ such that the R-MDP model defined by the same tuple with the set of regularizers $\{\phi_s(\boldsymbol{\pi}_s) = -\sum_k \overline{\lambda}_{sk}(l_s^k(\boldsymbol{\pi}_s) - c_s), \forall s \in \mathcal{S}\}$ yields an identical optimal policy.*

Proof. We write the Bellman update of the CT-MDP model as follows

$$g_s(\mathbf{w}_s) = \underset{\boldsymbol{\pi}_s}{\text{maximize}} \quad \mathbf{w}_s^{\mathrm{T}} \boldsymbol{\pi}_s \tag{27}$$
$$\text{subject to} \quad l_s^k(\boldsymbol{\pi}_s) \leq c_s^k,\ k = 1, \ldots, K_s$$
$$\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|).$$

The *Lagrange dual function* is defined as

$$\mathcal{L}_s(\boldsymbol{\lambda}|\mathbf{w}_s) = \max_{\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)} \left\{ \mathbf{w}_s^{\mathrm{T}} \boldsymbol{\pi}_s - \sum_{k=1}^{K_s} \lambda_k(l_s^k(\boldsymbol{\pi}_s) - c_s^k) \right\}.$$

Since the *Slater conditions* hold, we always have

$$g_s(\mathbf{w}_s) = \min_{\boldsymbol{\lambda} \geq 0} \{\mathcal{L}_s(\boldsymbol{\lambda}|\mathbf{w}_s)\}.$$

Now, let $\overline{\boldsymbol{\lambda}}_s$ be an optimal solution to $\{\mathcal{L}_s(\boldsymbol{\lambda}|\mathbf{w}_s)\}$. The KKT conditions imply that

$$g_s(\mathbf{w}_s) = \mathcal{L}_s(\overline{\boldsymbol{\lambda}}_s|\mathbf{w}_s)$$
$$\overline{\lambda}_{sk}(l_s^k(\boldsymbol{\pi}_s) - c_s^k) = 0,\ \forall k.$$

So, if we define $\phi_s(\boldsymbol{\pi}_s) = -\sum_{k=1}^{K_s} \overline{\lambda}_{sk}(l_s^k(\boldsymbol{\pi}_s) - c_s)$, then the problem $\max_{\boldsymbol{\pi}_s} \{\mathbf{w}_s^{\mathrm{T}} \boldsymbol{\pi}_s + \phi_s(\boldsymbol{\pi}_s)\}$ will yield the same optimal value and optimal solutions as (27). Then, in analogy to the proof of Proposition 4, we can also show that the R-MDP model specified by $\boldsymbol{\phi} = \{\phi_s, s \in \mathcal{S}\}$ and the CT-MDP will yield the same optimal value and optimal policy. ∎

B.3 Constrained MDPs and Some Computational Complexity Remarks

In this section, we look at some CT-MDP models and investigate their computational complexities. An mentioned, constraints based on a KL divergence would be useful to control the distance between a desired policy and a given policy. The Bellman update can be written as

$$\mathcal{T}^{\mathrm{CT}, \mathcal{D}^\pi}[V](s) = \underset{\boldsymbol{\pi}_s}{\text{maximize}} \quad \mathbf{w}_s^{\mathrm{T}} \boldsymbol{\pi}_s \tag{28}$$
$$\text{subject to} \quad \mathrm{KL}(\boldsymbol{\pi}_s | \overline{\boldsymbol{\pi}}_s) \leq c_s$$
$$\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|),$$

where $\mathbf{w}_s \in \mathbb{R}^{|\mathcal{A}|}$ with entries $w_{sa} = r(a|s) + \gamma \sum_{s'} q(s'|a, s) V(s')$, for all $a \in \mathcal{A}$. The following proposition shows how to solve (28) efficiently (in polynomial time).

Proposition 6 *Problem* (28) *is equivalent to*

$$-\min_{y \in \mathbb{R}_+} \left\{ c_s y + y \log \left(\sum_a \overline{\boldsymbol{\pi}}(a|s) \exp \left(\frac{w_{sa}}{y} \right) \right) \right\},$$

which can be solved by bisection and the complexity to get an ϵ-optimal solution is $\mathcal{O}(|\mathcal{A}| \ln(1/\epsilon))$.

Proof. We convert the problem into a minimization problem and the proof is similar to Lemma 4.1 of [19]. ∎

If the distance is restricted by a L1 norm, i.e., $\mathcal{D}_s^\pi = \{\boldsymbol{\pi}_s | \, ||\boldsymbol{\pi}_s - \overline{\boldsymbol{\pi}}_s||_1 \leq c_s\}$, we have the following Bellman update

$$\mathcal{T}^{\text{CT},\mathcal{D}^\pi}[V](s) = \underset{\boldsymbol{\pi}_s}{\text{maximize}} \quad \mathbf{w}_s^\text{T}\boldsymbol{\pi}_s \tag{29}$$
$$\text{subject to} \quad ||\boldsymbol{\pi}_s - \overline{\boldsymbol{\pi}}_s)||_1 \leq c_s,$$
$$\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)$$

We also have the proposition.

Proposition 7 *The objective value of* (29) *is equal to*

$$\boldsymbol{w}_s^T\overline{\boldsymbol{\pi}}_s + \frac{c_s}{2} \min_{\boldsymbol{\mu}_s \in \mathbb{R}_+^{|\mathcal{A}|}} \left\{ \max_a\{w_{sa} + \mu_{sa}\} - \min_a\{w_{sa} + \mu_{sa}\} \right\},$$

which can be solved exactly in $\mathcal{O}(|\mathcal{A}| \ln |\mathcal{A}|)$.

Proof. Similar to the proof of Lemma 4.3 in [19]. Note that in this case we can solve the Bellman equation exactly. ∎

Proposition 8 below examines the constrained problem under a L2 norm

$$\mathcal{T}^{\text{CT},\mathcal{D}^\pi}[V](s) = \underset{\boldsymbol{\pi}_s}{\text{maximize}} \quad \mathbf{w}_s^\text{T}\boldsymbol{\pi}_s \tag{30}$$
$$\text{subject to} \quad \sum_a \frac{(\pi(a|s) - \overline{\pi}(a|s))^2}{\overline{\pi}(a|s)} \leq c_s$$
$$\boldsymbol{\pi}_s \in \Delta(|\mathcal{A}|)$$

Proposition 8 *The optimal value of* (30) *is equal to*

$$\min_{\boldsymbol{\mu}_s \in \mathbb{R}_+^{|\mathcal{A}|}} \left\{ \sum_a \overline{\pi}(a|s)(w_{sa} + \mu_{sa}) + \sqrt{c_s \sum_a \overline{\pi}(a|s)(w_{sa} + \mu_{sa})^2} \right\},$$

which can be solved exactly in $\mathcal{O}(|\mathcal{A}| \ln(|\mathcal{A}|))$.

Proof. We convert the problem into a minimization one and prove the results similarly to Lemma 4.2 in [19]. ∎

In summary, we show that, beside the widely-used KL divergence, other types of constraints can be used to control the divergence between policies, noting that Bellman updates under L1 or L2 norms can be solved exactly, and the complexity can be bounded by $\mathcal{O}(|\mathcal{A}| \ln |\mathcal{A}|)$. In this aspect, L1 and L2 norms are more advantageous to use, as compared to the KL divergence. However, the use of such norm-based constraints need to be justified by numerical experiments. This is however out-of-scope of the paper and we keep the idea of future work.