

Policy Gradient and Actor–Critic Learning in Continuous Time and Space: Theory and Algorithms

Yanwei Jia

*Department of Industrial Engineering and Operations Research
Columbia University
New York, NY 10027, USA*

YJ2650@COLUMBIA.EDU

Xun Yu Zhou

*Department of Industrial Engineering and Operations Research &
The Data Science Institute
Columbia University
New York, NY 10027, USA*

XZ2574@COLUMBIA.EDU

Editor: Shie Mannor

Abstract

We study policy gradient (PG) for reinforcement learning in continuous time and space under the regularized exploratory formulation developed by Wang et al. (2020). We represent the gradient of the value function with respect to a given parameterized stochastic policy as the expected integration of an auxiliary running reward function that can be evaluated using samples and the current value function. This representation effectively turns PG into a policy evaluation (PE) problem, enabling us to apply the martingale approach recently developed by Jia and Zhou (2022a) for PE to solve our PG problem. Based on this analysis, we propose two types of actor-critic algorithms for RL, where we learn and update value functions and policies simultaneously and alternately. The first type is based directly on the aforementioned representation, which involves future trajectories and is offline. The second type, designed for online learning, employs the first-order condition of the policy gradient and turns it into martingale orthogonality conditions. These conditions are then incorporated using stochastic approximation when updating policies. Finally, we demonstrate the algorithms by simulations in two concrete examples.

Keywords: reinforcement learning, continuous time and space, policy gradient, policy evaluation, actor–critic algorithms, martingale

1. Introduction

The essence of reinforcement learning (RL) is “trial and error”: repeatedly trying a policy for actions, receiving and evaluating reward signals, and improving the policy. This manifests three key components of RL: 1) **exploration** with *stochastic* policies - to broaden search space via randomization; 2) **policy evaluation** - to evaluate the value function of a current policy; and 3) **policy improvement** - to improve the current policy. Numerous algorithms have been proposed in the RL literature, generally categorized into three types: critic-only, actor-only, and actor–critic. Here, an *actor* refers to a policy that governs the actions, and a *critic* refers to the value function that evaluates the performance of a policy. The *critic-only* approach learns a value function to compare the estimated outcomes of different actions

and selects the best one following the current value function. The *actor-only* approach acts directly without learning the expected outcomes of different policies. The *actor-critic* approach uses an actor simultaneously to improve the policy for generating actions given the current state of the environment and a critic to judge the selected policy and guide improving the actor. See Sutton and Barto (2018) and the references therein for extensive discussions on these methods.

All these algorithms and indeed the general RL study have been hitherto predominantly limited to *discrete-time* Markov decision processes (MDPs). From a practical point of view, however, the study on continuous-time RL with possibly continuous state and action spaces is more important. The world is inherently continuous-time, and a discrete-time dynamic is just an approximation of the reality by taking a sequence of snapshots of the world over time. As a result, in real life, examples abound in which an agent can or actually *needs* to interact with a random environment at an ultra-high frequency or outright continuously, e.g., high-frequency stock trading, autonomous driving, and robot navigation. Solving these problems in the discrete-time setting has a notorious drawback: the resulting algorithms are highly sensitive to time discretization; see Tallec et al. (2019); Yildiz et al. (2021) and the references therein.

Theoretically, it remains a largely uncharted territory to study RL in continuous time and spaces. The few existing papers on RL in the continuous setting are mostly restricted to *deterministic* systems; see, for example Baird (1993); Doya (2000); Munos (2006); Vamvoudakis and Lewis (2010); Frémaux et al. (2013); Lee and Sutton (2021); Yildiz et al. (2021); Kim et al. (2021) where there are no environmental noises. Munos and Bourgine (1997) introduce RL for diffusion-based stochastic control problems without proposing a data-driven solution. Model-based methods such as those in Basei et al. (2020); Szpruch et al. (2021) aim to estimate model coefficients by assuming their known and simple functional forms, which are still prone to model misspecification errors. The RL research for continuous-time diffusion processes with data/sample-driven solutions started only recently. Wang et al. (2020) propose an entropy-regularized stochastic relaxed control framework to study RL in continuous time and space and derive Boltzmann distributions as the generally optimal stochastic policies for exploring the environment and generating actions. In particular, when the problem is linear-quadratic (LQ), namely, when the dynamic is linear and the reward is quadratic in state and action, the optimal policy specializes to Gaussian distributions. Extensions and applications of this work include Wang and Zhou (2020); Dai et al. (2020); Guo et al. (2022); Gao et al. (2022).

While Wang et al. (2020) address the first component of RL – exploration – for the continuous setting, Jia and Zhou (2022a) investigate the second component, namely policy evaluation (PE), aiming at establishing a theoretical foundation for PE in continuous time and space. They show that PE is theoretically equivalent to maintaining the martingale condition of a specifically defined stochastic process, based on which they propose several online and offline PE algorithms. These algorithms have discrete-time counterparts, such as gradient Monte Carlo, TD(λ), and GTD, that scatter around in the MDP RL literature. Therefore, through the “martingale lens”, Jia and Zhou (2022a) not only devise new PE algorithms for the continuous case but also interpret and unify many classical algorithms initially designed for MDPs.

The current paper is a continuation of Wang et al. (2020) and Jia and Zhou (2022a), dealing with the third component of RL – policy improvement – in the continuous setting under stochastic policies and, thereby, completing the whole procedure for typical RL tasks. Note that Wang and Zhou (2020) put forth a policy improvement theorem for the special case of a continuous-time mean–variance portfolio selection problem. Furthermore, they show that defining a new policy by maximizing the Hamiltonian of the currently learned value function is proved to achieve a better objective value than the current policy. However, this method, akin to Q-learning for MDPs, has a drawback in requiring the functional form of the Hamiltonian, which in turn involves the knowledge of the environment.¹ Moreover, even if the Hamiltonian is known, maximizing a potentially complex function in high dimensions is computationally demanding or daunting.

In this paper, we take a different approach – that of *policy gradient* (PG) – which optimizes the value function over a parameterized family of policies. This approach has at least two advantages. First, selecting actions does not involve maximization, and actions are sampled from a known parametric distribution. Second, approximating policies directly facilitates more efficient learning if one has prior knowledge or intuition about the classes of potentially optimal policies (e.g., Gaussian distributions), leading to fewer parameters of the parametric family to be learned.

PG as a general sub-method of RL has a long history that can be traced back to Aleksandrov et al. (1968); Glynn (1990); Williams (1992); Barto et al. (1983); see also Bhatnagar et al. (2009) for more literature review and references therein. PG theorems specifically for MDPs are established in Sutton et al. (1999) and Marbach and Tsitsiklis (2001). Deterministic policy gradient algorithms for semi-MDPs (with discrete time and continuous action space) are developed in Silver et al. (2014) and later extended to incorporate deep neural networks in Lillicrap et al. (2015). Empirically, however, such algorithms tend to be unstable (Duan et al., 2016). Recent studies have focused on stochastic policies with possible entropy regularizers, also known as the *softmax* method; see for example Mnih et al. (2016); Schulman et al. (2017a,b); Haarnoja et al. (2018).

PG updates and improves policies along the gradient ascent direction, and is often carried out simultaneously and alternatingly with PE. The resulting algorithms for RL, therefore, are essentially actor–critic (AC) ones. Such methods have been successful in many real-world applications, notably AlphaGo (Silver et al., 2017) and dexterous hand manipulation (Haarnoja et al., 2018). But then, again, most PG and AC algorithms have been developed for discrete-time MDPs, and many of them in heuristic and *ad hoc* manners. Existing works on PG and AC in continuous time either focus on deterministic systems (Frémaux et al., 2013; Kim et al., 2021) or study specific models such as linear–quadratic ones (Wang et al., 2021). There are a few papers on applications of specifically designed continuous-time PG- and/or AC-based algorithms. Toy examples include the cart-pole swing-up problem in Doya (2000) and Half-Cheetah in Wawrzynski (2007), both involving physical laws of motion. Real-world applications include portfolio selections (Wang and Zhou, 2020), traffic control (Aragon-Gómez and Clempner, 2020), autonomous driving (Kiran et al., 2021), and biological neural networks (Frémaux et al., 2013; Zambrano et al., 2015).

1. In a more recent working paper Jia and Zhou (2022b), we develop a (little) q-learning theory to learn essentially the Hamiltonian from samples only, including a general policy improvement theorem.

In sum, it remains a significant open question to develop general continuous-time PG and AC algorithms and, more importantly, to lay an overarching theoretical underpinning for them. This paper aims to answer these questions by studying PG for a general problem in continuous time and space. Based on this, we develop model-free, data-driven AC algorithms for RL, covering both episodic and continuing, and both online and offline tasks. As its predecessors Wang et al. (2020) and Jia and Zhou (2022a), we develop theory in continuous time and discretize time only at the final algorithmic implementation stage, instead of discretizing time upfront and applying the existing MDP results. Specifically, we conduct our analysis in the stochastic relaxed control framework of Wang et al. (2020) involving distribution-valued stochastic policies. As such, it is necessary to first extend the PE theory of Jia and Zhou (2022a), including the martingale characterization and the resulting methods of the martingale loss function and the martingale orthogonality conditions, from deterministic policies to stochastic ones. This extension is technically non-trivial. Our main contributions, however, are a thorough analysis of the PG and the resulting AC algorithms. More precisely, we deduce the representation of the gradient of the current value function with respect to a parameterized (stochastic) policy. This representation turns out to have the *same form* as the value function in the PE step, effectively turning PG into an auxiliary PE problem. However, a subtle difficulty is that the corresponding “auxiliary” reward depends on the Hamiltonian and hence on the functional forms of the system dynamics. We solve this difficulty by integration by parts and Itô’s formula, transforming the representation into the expected integration of functions that can be evaluated using samples along with the current value function approximator.

The aforementioned representation is forward-looking. Namely, it is the conditional expectation of a term involving *future* states. Hence it is suitable for offline learning only. For online learning, we employ the first-order condition of the policy gradient and turn it into martingale orthogonality conditions. These conditions are then incorporated using stochastic approximation when updating policies. Finally, combining the newly developed PG methods in this paper and the PE methods in Jia and Zhou (2022a), we propose several AC algorithms for episodic and continuing/ergodic tasks.

Within the continuous-time stochastic relaxed control framework, there are several studies involving updating policies. For example, Wang and Zhou (2020) consider mean–variance portfolio selection and update policies by maximizing the Hamiltonian. As mentioned earlier, this requires knowledge about the market and hence the method is essentially model-based. Dai et al. (2020) address the time-inconsistency issue and focus on learning equilibrium policies. Guo et al. (2022) study multi-agent RL by solving an LQ mean-field game. These two papers rely on differentiating with respect to the policy hence are both model-based methods. In contrast, the present paper provides general model-free (up to the underlying dynamics being diffusion processes) AC algorithms that can be applied in all the above problems. In particular, we apply our algorithms to the mean–variance portfolio selection problem in Wang and Zhou (2020) and show they outperform significantly.

The rest of the paper proceeds as follows. In Section 2, we review Wang et al. (2020)’s entropy-regularized, exploratory formulation for RL in continuous time and space, and put forth an equivalent formulation convenient for the subsequent analysis. In Section 3, we develop a theory for PG, based on which we present general AC algorithms. Section 4 is devoted to an extension to ergodic tasks. We demonstrate our algorithms by simulation

with two concrete examples in Section 5. Finally, Section 6 concludes. In Appendix, we discuss the connection of our results with their discrete-time counterparts, present some theoretical results used in the simulation studies, and supply proofs of the results stated in the main text.²

2. Problem Formulation and Preliminaries

Throughout this paper, by convention all vectors are *column* vectors unless otherwise specified, and \mathbb{R}^k is the space of all k -dimensional vectors (hence $k \times 1$ matrices). Let A and B be two matrices of the same size. We denote by $A \circ B$ the inner product between A and B , by $|A|$ the Eculidean/Frobenius norm of A , and write $A^2 := AA^\top$, where A^\top is A 's transpose. For a positive semidefinite matrix A , we write $\sqrt{A} = UD^{1/2}V^\top$, where $A = UDV^\top$ is its singular value decomposition with U, V two orthogonal matrices and D a diagonal matrix, and $D^{1/2}$ is the diagonal matrix whose entries are the square root of those of D . We use $f = f(\cdot)$ to denote the *function* f , and $f(x)$ to denote the *function value* of f at x . For any stochastic process $X = \{X_s, s \geq 0\}$, we denote by $\{\mathcal{F}_s^X\}_{s \geq 0}$ the natural filtration generated by X . Finally, for any filtration $\mathcal{G} = \{\mathcal{G}_s\}_{s \geq 0}$ and any semi-martingale $Y = \{Y_s, s \geq 0\}$, we denote

$$L_{\mathcal{G}}^2([0, T]; Y) = \left\{ \kappa = \{\kappa_t, 0 \leq t \leq T\} : \right. \\ \left. \kappa \text{ is } \mathcal{G}_t\text{-progressively measurable and } \mathbb{E} \int_0^T |\kappa_t|^2 d\langle Y \rangle_t < \infty \right\},$$

which is a Hilbert space with the L^2 -norm $\|\kappa\|_{L^2} = \left(\mathbb{E} \int_0^T \kappa_t^2 d\langle Y \rangle_t \right)^{\frac{1}{2}}$, where $\langle \cdot \rangle$ is the quadratic variation of a given process.

Let d, n be given positive integers, $T > 0$, and $b : [0, T] \times \mathbb{R}^d \times \mathcal{A} \mapsto \mathbb{R}^d$ and $\sigma : [0, T] \times \mathbb{R}^d \times \mathcal{A} \mapsto \mathbb{R}^{d \times n}$ be given functions, where \mathcal{A} is the action set. The classical stochastic control problem is to control the *state* (or *feature*) dynamics governed by a stochastic differential equation (SDE), defined on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}^W; \{\mathcal{F}_s^W\}_{s \geq 0})$ along with a standard n -dimensional Brownian motion $W = \{W_s, s \geq 0\}$:

$$dX_s^a = b(s, X_s^a, a_s)ds + \sigma(s, X_s^a, a_s)dW_s, \quad s \in [0, T], \quad (1)$$

where a_s stands for the agent's action (control) at time s .

The goal of stochastic control is, for each initial time-state pair (t, x) of (1), to find the optimal $\{\mathcal{F}_s^W\}_{s \geq 0}$ -progressively measurable (continuous) sequence of actions $a = \{a_s, t \leq s \leq T\}$ – also called the optimal strategy – that maximizes the expected total reward:

$$\mathbb{E}^{\mathbb{P}^W} \left[\int_t^T e^{-\beta(s-t)} r(s, X_s^a, a_s) ds + e^{-\beta(T-t)} h(X_T^a) \Big| X_t^a = x \right], \quad (2)$$

where r is the (expected) running reward function, h is the (expected) lump-sum reward function applied at the end of the planning period T , and $\beta \geq 0$ is a discount factor that

2. The codes to reproduce our simulation studies are publicly available at <https://www.dropbox.com/sh/ezbuntcfje3d7kg/AABW-ndK-4j9N8E3kVTDZLz6a?dl=0>.

measures the time-value of the payoff or the impatience level of the agent. Note in the above the state process $X^a = \{X_s^a, t \leq s \leq T\}$ also depends on (t, x) . However, to ease notation, here (and similarly in the sequel) we use X^a instead of $X^{t,x,a} = \{X_s^{t,x,a}, t \leq s \leq T\}$ to denote the solution to SDE (1) with initial condition $X_t^a = x$ whenever no ambiguity may arise.

Let \mathcal{L}^a be the *infinitesimal generator* associated with the diffusion process governed by (1):

$$\mathcal{L}^a \varphi(t, x) := \frac{\partial \varphi}{\partial t}(t, x) + b(t, x, a) \circ \frac{\partial \varphi}{\partial x}(t, x) + \frac{1}{2} \sigma^2(t, x, a) \circ \frac{\partial^2 \varphi}{\partial x^2}(t, x), \quad a \in \mathcal{A},$$

where $\frac{\partial \varphi}{\partial x} \in \mathbb{R}^d$ is the gradient, and $\frac{\partial^2 \varphi}{\partial x^2} \in \mathbb{R}^{d \times d}$ is the Hessian. We make the following assumption to ensure theoretically the well-posedness of the stochastic control problem (1)–(2).

Assumption 1 *The following conditions for the state dynamics and reward functions hold true:*

- (i) b, σ, r, h are all continuous functions in their respective arguments;
- (ii) b, σ are uniformly Lipschitz continuous in x , i.e., for $\varphi \in \{b, \sigma\}$, there exists a constant $C > 0$ such that

$$|\varphi(t, x, a) - \varphi(t, x', a)| \leq C|x - x'|, \quad \forall (t, a) \in [0, T] \times \mathcal{A}, \quad \forall x, x' \in \mathbb{R}^d;$$

- (iii) b, σ have linear growth in x , i.e., for $\varphi \in \{b, \sigma\}$, there exists a constant $C > 0$ such that

$$|\varphi(t, x, a)| \leq C(1 + |x|), \quad \forall (t, x, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{A};$$

- (iv) r and h have polynomial growth in (x, a) and x respectively, i.e., there exists a constant $C > 0$ and $\mu \geq 1$ such that

$$|r(t, x, a)| \leq C(1 + |x|^\mu + |a|^\mu), \quad |h(x)| \leq C(1 + |x|^\mu), \quad \forall (t, x, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{A}.$$

Classical model-based stochastic control theory has been well developed (e.g., Fleming and Soner, 2006 and Yong and Zhou, 1999) to solve the above problem, under the premise that the functional forms of b, σ, r, h are all given and known. In the RL setting, however, the agent does not have this knowledge of the environment. Instead, what she can do is “trial and error” – to try a sequence of actions $a = \{a_s, t \leq s \leq T\}$, observe the corresponding state process $X^a = \{X_s^a, t \leq s \leq T\}$ and collect both a stream of discounted running rewards $\{e^{-\beta(s-t)} r(s, X_s^a, a_s), t \leq s \leq T\}$ and a discounted, end-of-period lump-sum reward $e^{-\beta(T-t)} h(X_T^a)$ where β is a given, known discount factor. In the offline setting, the agent can repeatedly try different sequences of actions over the same time period $[0, T]$ and record the corresponding state processes and payoffs. In the online setting, the agent updates the actions as she goes, based on all the up-to-date historical observations.

A critical question is how to *generate* these trial-and-error sequences of actions. The idea is *randomization*, namely, the agent employs a *stochastic policy*, which is a probability

distribution on the action space, to produce actions according to the current time–state pair. It is important to note that this randomization itself is *independent* of the underlying Brownian motion W , the random source of the original control problem that stands for the environmental noise. Wang et al. (2020) formulate an RL problem in continuous time and space, incorporating distribution-valued stochastic policies with an entropy regularizer to account for the tradeoff between exploration and exploitation. Specifically, assume the probability space is rich enough to support a random variable Z that is uniformly distributed on $[0, 1]$ and independent of W . We then expand the original filtered probability space to $(\Omega, \mathcal{F}, \mathbb{P}; \{\mathcal{F}_s\}_{s \geq 0})$ where $\mathcal{F}_s = \mathcal{F}_s^W \vee \sigma(Z)$ and \mathbb{P} is now the probability measure on \mathcal{F}_T .³ Let $\pi : (t, x) \in [0, T] \times \mathbb{R}^d \mapsto \pi(\cdot|t, x) \in \mathcal{P}(\mathcal{A})$ be a given (feedback) policy, where $\mathcal{P}(\mathcal{A})$ is a suitable collection of probability density functions (pdfs).⁴ At each time s , an action a_s is generated or sampled from the distribution $\pi(\cdot|s, X_s)$.

Given a stochastic policy π , an initial time–state pair (t, x) , and an $\{\mathcal{F}_s\}_{s \geq 0}$ -progressively measurable action process $a^\pi = \{a_s^\pi, t \leq s \leq T\}$ generated from π , the corresponding state process $X^\pi = \{X_s^\pi, t \leq s \leq T\}$ follows

$$dX_s^\pi = b(s, X_s^\pi, a_s^\pi)ds + \sigma(s, X_s^\pi, a_s^\pi)dW_s, \quad s \in [t, T]; \quad X_t^\pi = x \quad (3)$$

defined on $(\Omega, \mathcal{F}, \mathbb{P}; \{\mathcal{F}_s\}_{s \geq 0})$. Moreover, following Wang et al. (2020), we add a regularizer to the reward function to encourage exploration (represented by the stochastic policy), leading to

$$J(t, x; \pi) = \mathbb{E}^\mathbb{P} \left[\int_t^T e^{-\beta(s-t)} [r(s, X_s^\pi, a_s^\pi) + \gamma p(s, X_s^\pi, a_s^\pi, \pi(\cdot|s, X_s^\pi))] ds + e^{-\beta(T-t)} h(X_T^\pi) \middle| X_t^\pi = x \right], \quad (4)$$

where $\mathbb{E}^\mathbb{P}$ is the expectation with respect to (w.r.t.) both the Brownian motion and the action randomization. In the above, $p : [0, T] \times \mathbb{R}^d \times \mathcal{A} \times \mathcal{P}(\mathcal{A}) \mapsto \mathbb{R}$ is the regularizer and $\gamma \geq 0$ a weighting parameter on exploration, also known as the *temperature* parameter. Wang et al. (2020) take the differential entropy as the regularizer, which corresponds to

$$p(t, x, a, \pi(\cdot)) = -\log \pi(a).$$

Through a law of large number argument, Wang et al. (2020) show that $\{X_s^\pi, t \leq s \leq T\}$ has the same distribution as the solution to the following SDE, denoted by $\{\tilde{X}_s^\pi, t \leq s \leq T\}$:

$$dX_s = \tilde{b}(s, X_s, \pi(\cdot|s, X_s))dt + \tilde{\sigma}(s, X_s, \pi(\cdot|s, X_s))dW_s, \quad s \in [t, T]; \quad X_t = x, \quad (5)$$

-
3. Note that a single uniform random variable Z can produce many *independent* random variables having density functions. No dynamics are needed for these random variables and they are all independent of each other and of the Brownian motion. The independence means it makes no difference if these variables are given all at once at time 0 or are revealed as time evolves. We opt for the (mathematically speaking) easier construction where these are all defined using one single uniform Z . Meanwhile, \mathbb{P} is the product extension from \mathbb{P}^W ; the two probability measures coincide when restricted to \mathcal{F}_T^W .
 4. Here we assume that the action space \mathcal{A} is continuous and randomization is restricted to those distributions that have density functions. The analysis and results of this paper can be easily extended to the cases of discrete action spaces and/or randomization with probability mass functions.

where

$$\tilde{b}(s, x, \pi(\cdot)) = \int_{\mathcal{A}} b(s, x, a) \pi(a) da, \quad \tilde{\sigma}(s, x, \pi(\cdot)) = \sqrt{\int_{\mathcal{A}} \sigma^2(s, x, a) \pi(a) da}.$$

Moreover, the reward function (4) is identical to

$$\begin{aligned} J(t, x; \boldsymbol{\pi}) = & \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T e^{-\beta(s-t)} \int_{\mathcal{A}} [r(s, \tilde{X}_s^{\boldsymbol{\pi}}, a) + \gamma p(s, \tilde{X}_s^{\boldsymbol{\pi}}, a, \boldsymbol{\pi}(\cdot|s, \tilde{X}_s^{\boldsymbol{\pi}}))] \boldsymbol{\pi}(a|s, \tilde{X}_s^{\boldsymbol{\pi}}) da ds \right. \\ & \left. + e^{-\beta(T-t)} h(\tilde{X}_T^{\boldsymbol{\pi}}) \Big| \tilde{X}_t^{\boldsymbol{\pi}} = x \right]. \end{aligned} \tag{6}$$

Mathematically, (5) and (6) together form a so-called *relaxed stochastic control* problem where the effect of individually sampled actions has been averaged out (over the randomization/exploration) and, hence, one can focus on how a policy $\boldsymbol{\pi}$ impacts the *distribution* of the “averaged” state \tilde{X} ; see Wang et al. (2020).

Here, $J(t, x; \boldsymbol{\pi})$ is called the *value function* of the policy $\boldsymbol{\pi}$, and the task of RL is to find

$$J^*(t, x) = \max_{\boldsymbol{\pi} \in \boldsymbol{\Pi}} J(t, x; \boldsymbol{\pi}), \tag{7}$$

where $\boldsymbol{\Pi}$ stands for the set of admissible policies. The following gives the precise definition of admissible (feedback) policies.

Definition 1 *A policy $\boldsymbol{\pi} = \boldsymbol{\pi}(\cdot|\cdot, \cdot)$ is called admissible if*

- (i) $\boldsymbol{\pi}(\cdot|t, x) \in \mathcal{P}(\mathcal{A})$, and $\boldsymbol{\pi}(a|t, x) : (t, x, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{A} \mapsto \mathbb{R}$ is measurable;
- (ii) the SDE (5) admits a unique weak solution (in the sense of distribution) for any initial $(t, x) \in [0, T] \times \mathbb{R}^d$;
- (iii) $\int_{\mathcal{A}} |r(t, x, a) + \gamma p(t, x, a, \boldsymbol{\pi}(\cdot|t, x))| \boldsymbol{\pi}(a|t, x) da \leq C(1 + |x|^\mu)$, $\forall (t, x)$ where $C > 0$ and $\mu \geq 1$ are constants;
- (iv) $\boldsymbol{\pi}(a|t, x)$ is continuous in (t, x) and uniformly Lipschitz continuous in x in the total variation distance, i.e., for each fixed a , $\int_{\mathcal{A}} |\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t', x')| da \rightarrow 0$ as $(t', x') \rightarrow (t, x)$, and there is a constant $C > 0$ independent of (t, a) such that

$$\int_{\mathcal{A}} |\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')| da \leq C|x - x'|, \quad \forall x, x' \in \mathbb{R}^d.$$

The conditions required in the above definition, while not necessarily the weakest ones, are to theoretically guarantee the well-posedness of the control problem (5)–(6). This is implied by the following result.

Lemma 2 *Let Assumptions 1 hold and $\boldsymbol{\pi}$ be a given admissible policy. Then the SDE (5) admits a unique strong solution. Moreover, for any $\mu \geq 2$, the solution satisfies the growth condition $\mathbb{E}^{\mathbb{P}^W} \left[\max_{t \leq s \leq T} |\tilde{X}_s^{\boldsymbol{\pi}}|^\mu \Big| \tilde{X}_t^{\boldsymbol{\pi}} = x \right] \leq C(1 + |x|^\mu)$ for some constant $C = C(\mu)$. Finally, the expected payoff (6) is finite.*

We stress that the solution to (5), $\{\tilde{X}_s^\pi, t \leq s \leq T\}$, is the *average* of the sample trajectories over infinitely many randomized actions and is in itself *not* a sample trajectory *nor* observable. The stochastic relaxed control problem (5)–(6), introduced in Wang et al. (2020), just provides a framework for *theoretical* analysis. In contrast, the solution to (3), $\{X_s^\pi, t \leq s \leq T\}$, is a sample trajectory under a realization of action sequence, $\{a_s^\pi, t \leq s \leq T\}$, generated from the policy π , and can indeed be observed. Meanwhile, the difference between (3) and (1) is that actions in the former are randomized: a^π is also driven by the randomization and hence is *not* \mathcal{F}_t^W -adapted. By taking the expectation w.r.t. the action randomization, the expectation in (4) reduces to the expectation in (6). In other words, the problem (5)–(6) is mathematically equivalent to the problem (3)–(4); yet they serve different purposes in our study: the former provides a framework for theoretical analysis of the value function while the latter directly involves observable samples.

Unlike most RL problems that are formulated in an infinite planning horizon (known as *continuing tasks*), the current paper mainly focuses on a finite horizon setting (known as *episodic tasks*). Finite horizons reflect limited lifespans of real-life tasks, e.g., a trader sells a financial contract with a maturity date, a robot finishes a task before a deadline, and a game player strives to pass a checkpoint given a time limit. If we let $T \rightarrow \infty$, under suitable regularity conditions (e.g., when β is large enough) our formulation covers the discounted formulation of the continuing tasks. In addition, later we will consider an ergodic setting as an alternative formulation for continuing tasks in Section 4.

3. Theoretical Foundation of Actor–Critic Algorithms

An actor-critic (AC) algorithm consists of two parts: to estimate the value function of a given policy and to update (improve) the policy. In this section, we provide the theoretical analysis to guide devising such an algorithm through policy evaluation (PE) and policy gradient (PG).

3.1 Policy Evaluation

Jia and Zhou (2022a) take a martingale perspective to characterize PE as well as its link to solving a linear partial differential equation (PDE) numerically. However, they consider only *deterministic* policies (i.e. no randomization/exploration), without explicitly involving actions sampled from a stochastic policy. The extension to the case of stochastic policies is non-trivial and specific statements of the corresponding results are important for the subsequent PG and AC algorithm design; so we present and prove them here.

For a given stochastic policy π , $J(\cdot, \cdot; \pi)$ can be characterized by a PDE based on the celebrated Feynman–Kac formula (cf. Karatzas and Shreve, 2014), which also holds true for the relaxed control setting.

Lemma 3 *Assume there is a unique viscosity solution $v \in C([0, T] \times \mathbb{R}^d)$ to the following PDE:*

$$\int_{\mathcal{A}} [\mathcal{L}^a v(t, x) + r(t, x, a) + \gamma p(t, x, a, \pi(\cdot|t, x)) - \beta v(t, x)] \pi(a|t, x) da = 0, \quad (t, x) \in [0, T] \times \mathbb{R}^d, \quad (8)$$

with the terminal condition $v(T, x) = h(x)$, $x \in \mathbb{R}^d$, which satisfies $|v(t, x)| \leq C(1 + |x|^\mu)$ for a constant $C > 0$ and $\mu \geq 1$. Then v is the value function, that is, $v(t, x) = J(t, x; \pi)$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$.

To avoid unduly technicalities, we assume throughout this paper that the value function $J \in C^{1,2}([0, T] \times \mathbb{R}^d) \cap C([0, T] \times \mathbb{R}^d)$. There is a rich literature on conditions ensuring the unique existence and regularity of the viscosity solution to the type of equations like (8); but see Tang et al. (2021) for some latest results.

The following is the main theoretical result underpinning PE, extended from the setting of deterministic feedback policies in Jia and Zhou (2022a) to that of stochastic policies.

Theorem 4 *A function $J(\cdot, \cdot; \pi)$ is the value function associated with the policy π if and only if it satisfies terminal condition $J(T, x; \pi) = h(x)$, and for any initial $(t, x) \in [0, T] \times \mathbb{R}^d$:*

$$e^{-\beta s} J(s, \tilde{X}_s^\pi; \pi) + \int_t^s e^{-\beta s'} \int_{\mathcal{A}} [r(s', \tilde{X}_{s'}^\pi, a) + \gamma p(s', \tilde{X}_{s'}^\pi, a, \pi(\cdot | s', \tilde{X}_{s'}^\pi))] \pi(a | s', \tilde{X}_{s'}^\pi) da ds'$$

is an $(\mathcal{F}^{\tilde{X}^\pi}, \mathbb{P}^W)$ -martingale on $[t, T]$. Moreover, it is also equivalent to the martingale orthogonality condition:

$$\mathbb{E}^\mathbb{P} \int_0^T \xi_t \left[dJ(t, X_t^\pi; \pi) + r(t, X_t^\pi, a_t^\pi) dt + \gamma p(t, X_t^\pi, a_t^\pi, \pi(\cdot | t, X_t^\pi)) dt - \beta J(t, X_t^\pi; \pi) dt \right] = 0, \quad (9)$$

for any $\xi \in L^2_{\mathcal{F}^{\pi}}([0, T]; J(\cdot, X^\pi; \pi))$.

In the above theorem, ξ is called a *test function* by convention, although in general it is actually a stochastic process.

In RL, one typically employs function approximation for learning functions of interest. Specifically, for PE, one uses a family of parameterized functions $J^\theta \equiv J^\theta(\cdot, \cdot; \pi)$ on $[0, T] \times \mathbb{R}^d$ to approximate J , where $\theta \in \Theta \subseteq \mathbb{R}^{L_\theta}$, and the problem is reduced to finding the “best” (in some sense) θ . We make the following assumption on these function approximators to be used. (Henceforth we may drop π from $J^\theta(\cdot, \cdot; \pi)$ whenever no ambiguity arises.)

Assumption 2 *For all $\theta \in \Theta$, $J^\theta \in C^{1,2}([0, T] \times \mathbb{R}^d) \cap C([0, T] \times \mathbb{R}^d)$ and satisfies the polynomial growth condition in x . Moreover, $J^\theta(t, x)$ is a smooth function in θ with $\frac{\partial J^\theta}{\partial \theta}, \frac{\partial^2 J^\theta}{\partial \theta^2} \in C^{1,2}([0, T] \times \mathbb{R}^d) \cap C([0, T] \times \mathbb{R}^d)$ satisfying the polynomial growth condition in x .*

Thanks to the martingale characterization in Theorem 4, the PE algorithms developed in Jia and Zhou (2022a) can be adapted to the current setting in a straightforward manner. We now summarize them.

- (i) Minimize the martingale loss function (offline):

$$\min_{\theta \in \Theta} \mathbb{E}^\mathbb{P} \left[\int_0^T \left(e^{-\beta T} h(X_T^\pi) - e^{-\beta t} J^\theta(t, X_t^\pi) + \int_t^T e^{-\beta s} [r(s, X_s^\pi, a_s^\pi) + \gamma p(s, X_s^\pi, a_s^\pi, \pi(\cdot | s, X_s^\pi))] ds \right)^2 dt \right].$$

This objective corresponds to the gradient Monte-Carlo algorithm for discrete MDPs (Sutton and Barto, 2018).

- (ii) Solve the martingale orthogonality condition (online/offline):

$$\mathbb{E}^{\mathbb{P}} \left\{ \int_0^T \xi_t [dJ^\theta(t, X_t^\pi) + r(t, X_t^\pi, a_t^\pi) dt + \gamma p(t, X_t^\pi, a_t^\pi, \pi(\cdot|t, X_t^\pi)) dt - \beta J^\theta(t, X_t^\pi) dt] \right\} = 0.$$

This objective corresponds to various (semi-gradient) TD algorithms and their variants for MDPs (Sutton, 1988; Bradtke and Barto, 1996), depending on the choices of the test function ξ .

- (iii) Minimize a quadratic form of the martingale orthogonality condition (online/offline):

$$\min_{\theta \in \Theta} \mathbb{E}^{\mathbb{P}} \left\{ \int_0^T \xi_t [dJ^\theta(t, X_t^\pi) + r(t, X_t^\pi, a_t^\pi) dt + \gamma p(t, X_t^\pi, a_t^\pi, \pi(\cdot|t, X_t^\pi)) dt - \beta J^\theta(t, X_t^\pi) dt] \right\}^\top A \mathbb{E}^{\mathbb{P}} \left\{ \int_0^T \xi_t [dJ^\theta(t, X_t^\pi) + r(t, X_t^\pi, a_t^\pi) dt + \gamma p(t, X_t^\pi, a_t^\pi, \pi(\cdot|t, X_t^\pi)) dt - \beta J^\theta(t, X_t^\pi) dt] \right\},$$

where A is a positive definite matrix of a suitable size. Typical choices are $A = I$ or $A = (\mathbb{E}^{\mathbb{P}}[\int_0^T \xi_t \xi_t^\top dt])^{-1}$. This objective corresponds to the gradient TD algorithms and their variants for MDPs (Sutton et al., 2008, 2009; Maei et al., 2009).

In the above, the choice of the parametric family J^θ may be guided by exploiting some special structure of the underlying problem; see Wang and Zhou (2020) for an example. More general choices include linear combinations of some basis functions or neural networks. On the other hand, common choices of the test functions are $\xi_t = \frac{\partial J^\theta}{\partial \theta}(t, X_t^\pi)$ or $\xi_t = \int_0^t \lambda^{s-t} \frac{\partial J^\theta}{\partial \theta}(s, X_s^\pi) ds$. Refer to the aforementioned references for details, and in particular to Jia and Zhou (2022a) for the continuous setting. Finally, when implementing these algorithms we need to discretize time, and the convergence when the mesh size goes to zero is established in Jia and Zhou (2022a), which can be readily extended to the current setting.

3.2 Policy Gradient

Given an admissible policy, suppose we have carried out the PE step and obtained an estimate of the corresponding value function. The next step is PG, namely, to estimate the gradient of the (learned) value function w.r.t. the policy. Specifically, let π^ϕ be a parametric family of policies with the parameter $\phi \in \Phi \subset \mathbb{R}^{L_\phi}$. We aim to compute the policy gradient $g(t, x; \phi) := \frac{\partial}{\partial \phi} J(t, x; \pi^\phi) \in \mathbb{R}^{L_\phi}$ at the current time–state pair (t, x) . Here and throughout we always assume π^ϕ is an admissible policy.

Based on the PDE characterization (8) of the value function, we take the derivative in ϕ on both sides of (8), with $v(t, x)$ replaced by $J(t, x; \pi^\phi)$, to get a new PDE satisfied by

$g(t, x; \phi)$:

$$\begin{cases} \int_{\mathcal{A}} \left\{ [\mathcal{L}^a g(t, x; \phi) - \beta g(t, x; \phi) + \gamma q(t, x, a, \phi)] \pi^\phi(a|t, x) \right. \\ \left. + [\mathcal{L}^a J(t, x; \pi^\phi) + r(t, x, a) + \gamma p(t, x, a, \pi^\phi(\cdot|t, x)) - \beta J(t, x; \pi^\phi)] \frac{\partial \pi^\phi}{\partial \phi}(a|t, x) \right\} da = 0, \\ g(T, x; \phi) = 0, \end{cases} \quad (10)$$

where $q(t, x, a, \phi) = \frac{\partial}{\partial \phi} p(t, x, a, \pi^\phi(\cdot|t, x))$ that maps $[0, T] \times \mathbb{R}^d \times \mathcal{A} \times \Phi$ to \mathbb{R}^{L_ϕ} . Note that (10) is a *system* of L_ϕ equations, and $\mathcal{L}^a g$ denotes applying the operator \mathcal{L}^a to each component of the \mathbb{R}^{L_ϕ} -valued function $g(\cdot, \cdot; \phi)$.

Define

$$\begin{aligned} \check{r}(t, x, a; \phi) &= \left[\mathcal{L}^a J(t, x; \pi^\phi) + r(t, x, a) + \gamma p(t, x, a, \pi^\phi(\cdot|t, x)) - \beta J(t, x; \pi^\phi) \right] \frac{\frac{\partial \pi^\phi}{\partial \phi}(a|t, x)}{\pi^\phi(a|t, x)} \\ &\quad + \gamma q(t, x, a, \phi) \\ &= \left[\mathcal{L}^a J(t, x; \pi^\phi) + r(t, x, a) + \gamma p(t, x, a, \pi^\phi(\cdot|t, x)) - \beta J(t, x; \pi^\phi) \right] \frac{\partial}{\partial \phi} \log \pi^\phi(a|t, x) \\ &\quad + \gamma q(t, x, a, \phi), \end{aligned}$$

which is again a function that maps $[0, T] \times \mathbb{R}^d \times \mathcal{A} \times \Phi$ to \mathbb{R}^{L_ϕ} . Then (10) can be written as

$$\int_{\mathcal{A}} [\mathcal{L}^a g(t, x; \phi) - \beta g(t, x; \phi) + \check{r}(t, x, a; \phi)] \pi^\phi(a|t, x) da = 0, \quad g(T, x; \phi) = 0. \quad (11)$$

Observe that (11) has the similar form to (8). Thus a Feynman–Kac formula (similar to Lemma 3) represents g as

$$\begin{aligned} g(t, x; \phi) &= \mathbb{E}^{\mathbb{P}} \left[\int_t^T e^{-\beta(s-t)} \check{r}(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}; \phi) ds \mid X_t^{\pi^\phi} = x \right] \\ &= \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T e^{-\beta(s-t)} \int_{\mathcal{A}} \check{r}(s, \tilde{X}_s^{\pi^\phi}, a; \phi) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da ds \mid \tilde{X}_t^{\pi^\phi} = x \right]. \end{aligned} \quad (12)$$

Therefore, computing PG boils down mathematically to a PE problem with a *different* reward function. Indeed, the task here is much easier because we only need to compute the function *value*, $g(t, x; \phi)$, via (12) at *some* (t, x) along a sample trajectory, instead of learning the *entire* function $g(\cdot, \cdot; \phi)$ as in PE. However, unlike a normal PE problem, the new reward function \check{r} involves the operator \mathcal{L}^a applied to J which can not be observed nor computed without the knowledge of the environment.

The remedy to overcome this difficulty rests with Itô's lemma and martingality. We now provide an informal argument for explanation before presenting the formal result. Suppose at time t , an action a is generated from $\pi^\phi(\cdot|t, X_t)$ and applied to the system within a small time window $[t, t + \Delta t]$. Apply Itô's lemma to obtain

$$J(t + \Delta t, X_{t+\Delta t}^a; \pi^\phi) - J(t, X_t; \pi^\phi) = \int_t^{t+\Delta t} \mathcal{L}^a J(s, X_s^a; \pi^\phi) ds + \frac{\partial J}{\partial x}(s, X_s^a; \pi^\phi)^\top \sigma_s dW_s.$$

Therefore,

$$\begin{aligned}
 & \tilde{r}(s, X_s^a, a; \phi) ds \\
 & \equiv \left[\mathcal{L}^a J(s, X_s^a; \pi^\phi) + r(s, X_s^a, a) + \gamma p(s, X_s^a, a, \pi^\phi(\cdot|s, X_s^a)) - \beta J(s, X_s^a; \pi^\phi) \right] \\
 & \quad \times \frac{\partial}{\partial \phi} \log \pi^\phi(a|s, X_s^a) ds + \gamma q(s, X_s^a, a, \phi) ds \\
 & \approx \frac{\partial}{\partial \phi} \log \pi^\phi(a|s, X_s^a) \left\{ dJ(s, X_s^a; \pi^\phi) + [r(s, X_s^a, a) + \gamma p(s, X_s^a, a, \pi^\phi(\cdot|s, X_s^a)) \right. \\
 & \quad \left. - \beta J(s, X_s^a; \pi^\phi)] ds - \frac{\partial J}{\partial x}(s, X_s^a; \pi)^\top \sigma_s dW_s \right\} + \gamma q(s, X_s^a, a, \phi) ds.
 \end{aligned} \tag{13}$$

Since the stochastic integral w.r.t. the dW term above is a martingale (under suitable regularity conditions), such a term, even if unknown, does not contribute to the expectation and thus can be ignored. As a result, \tilde{r} can be incrementally estimated based on observations of samples and the learned value function.

Before stating the main result of this paper, we impose the following technical conditions on the policy approximators.

Assumption 3 $\pi^\phi(a|t, x)$ is smooth in $\phi \in \Phi$ for all (t, x, a) . Moreover,

$$\int_{\mathcal{A}} |\tilde{r}(t, x, a; \phi)| \pi^\phi(a|t, x) da \leq C(1 + |x|^\mu)$$

for all (t, x, ϕ) , where $C > 0, \mu \geq 1$ are constants. Furthermore, $\int_{\mathcal{A}} |\frac{\partial}{\partial \phi} \log \pi^\phi(a|t, x)|^2 \pi^\phi(a|t, x) da$ is continuous in (t, x) for all $\phi \in \Phi$.

Theorem 5 Given an admissible parameterized policy π^ϕ , its policy gradient $g(t, x; \phi) = \frac{\partial}{\partial \phi} J(t, x; \pi^\phi)$ admits the following representation:

$$\begin{aligned}
 g(t, x; \phi) &= \mathbb{E}^\mathbb{P} \left[\int_t^T e^{-\beta(s-t)} \left\{ \frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \left(dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \right. \right. \\
 & \quad \left. \left. \left. + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot|s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \right) \right. \right. \\
 & \quad \left. \left. + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \right\} \Big| X_t^{\pi^\phi} = x \right], \quad (t, x) \in [0, T] \times \mathbb{R}^d.
 \end{aligned} \tag{14}$$

Once again, all the terms inside the expectation above are all computable given samples (including action trajectories and the corresponding state trajectories) on $[t, T]$, together with an estimated value function J (obtained in the previous PE step). Note that the expectation (14) gives the gradient of the value function w.r.t. any policy, which is not 0 in general.

Observing (14) more closely, we can write $g(t, x; \phi) = g_1(t, x; \phi) + g_2(t, x; \phi)$ where

$$\begin{aligned}
 g_1(t, x; \phi) &= \mathbb{E}^\mathbb{P} \left[\int_t^T e^{-\beta(s-t)} \frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \left(dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \right. \\
 & \quad \left. \left. + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot|s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \right) \Big| X_t^{\pi^\phi} = x \right]
 \end{aligned}$$

and

$$g_2(t, x; \phi) = \mathbb{E}^{\mathbb{P}} \left[\int_t^T e^{-\beta(s-t)} \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \middle| X_t^{\pi^\phi} = x \right].$$

The integrand in the expression of g_1 is the discounted derivative of the log-likelihood (log-pdf) that determines the direction, multiplied by a scalar term. This scalar term is actually the TD error in the continuous setting (Jia and Zhou, 2022a) that also appears in the martingale orthogonality condition (9). Note that $g_1(t, x; \phi) \neq 0$ in general, because $\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi})$ depends on the realization of $a_s^{\pi^\phi}$, and hence is not $\mathcal{F}_s^{X^{\pi^\phi}}$ -measurable and does not qualify as a test function ξ in Theorem 4. On the other hand, g_2 comes entirely from the regularizer and vanishes should the latter be absent.

There are two equivalent forms of the representation (14), which can be used to add more flexibility in designing PG algorithms and to optimize their performance. The first one is to add a ‘‘baseline’’ action-independent function $B(t, x)$ to the integrand in (14). Precisely, it follows from $a_s^{\pi^\phi} \sim \pi^\phi(\cdot | s, X_s^{\pi^\phi})$ that

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}} \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) B(s, X_s^{\pi^\phi}) \middle| X_s^{\pi^\phi} \right] \\ &= B(s, X_s^{\pi^\phi}) \int_{\mathcal{A}} \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a | s, X_s^{\pi^\phi}) \right] \pi^\phi(a | s, X_s^{\pi^\phi}) da \\ &= B(s, X_s^{\pi^\phi}) \int_{\mathcal{A}} \frac{\partial \pi^\phi(a | s, X_s^{\pi^\phi})}{\partial \phi} da = B(s, X_s^{\pi^\phi}) \frac{\partial}{\partial \phi} \int_{\mathcal{A}} \pi^\phi(a | s, X_s^{\pi^\phi}) da = 0. \end{aligned}$$

Hence, an alternative representation of (14) is

$$\begin{aligned} & g(t, x; \phi) \\ &= \mathbb{E}^{\mathbb{P}} \left[\int_t^T e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] [dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \right. \\ & \quad + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi) - B(s, X_s^{\pi^\phi})] ds \\ & \quad \left. \left. + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \right\} \middle| X_t^{\pi^\phi} = x \right]. \end{aligned} \tag{15}$$

Including such a baseline function in the representation of PG goes back at least to Williams (1992). Sutton et al. (2000) and Zhao et al. (2011) find that adding an appropriate baseline function can reduce the variance of the learning process. In particular, a common choice of baseline function, though not theoretically optimal, is the current value function, which leads to the so-called advantage AC algorithms (Degris et al., 2012; Mnih et al., 2016). Interestingly, without including any exogenous baseline function, the PG algorithms out of (14) are exactly the continuous-time versions of the advantage AC algorithms. As such, we do not add other baseline functions for designing our algorithms below. More connections to the representation of policy gradient in discrete-time and detailed discussions of the baseline function can be found in Appendix A.

The second alternative form of (14) is to add an admissible test function to the derivative of the log-likelihood. Specifically, suppose $\zeta \in L^2_{\mathcal{F}^{X^{\pi^\phi}}}([0, T]; J(\cdot, X^{\pi^\phi}; \pi^\phi))$ is an \mathbb{R}^{L_ϕ} -valued

process. Then based on Theorem 4, the policy gradient can also be represented by

$$\begin{aligned}
 g(t, x; \phi) = & \mathbb{E}^{\mathbb{P}} \left[\int_t^T e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) + \zeta_s \right] \left(dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \right. \right. \\
 & + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \Big) \\
 & \left. \left. \left. + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \right\} \middle| X_t^{\pi^\phi} = x \right], \quad (t, x) \in [0, T] \times \mathbb{R}^d.
 \end{aligned} \tag{16}$$

As discussed before, we do not use (14) to approximate the function $g(\cdot, \cdot; \phi)$. Rather, at any current time–state (t, x) , (14) gives the gradient of $J(t, x; \phi)$ in ϕ so that we can update ϕ in the most promising direction (based on the gradient ascent algorithm) to improve the value of J . However, the right hand side of (14) involves only the *future* trajectories from t ; so Theorem 5 works only for the offline setting.

To treat the online case, assume that ϕ^* is the optimal point of $J(t, x; \pi^\phi)$ for any (t, x) and that the first-order condition holds (e.g., when ϕ^* is an interior point).⁵ Then $g(t, x; \phi^*) = 0$. It thus follows from (10) that

$$\begin{aligned}
 0 = & \int_{\mathcal{A}} \left\{ \left[\mathcal{L}^a J(t, x; \pi^{\phi^*}) + r(t, x, a) + \gamma p(t, x, a, \pi^{\phi^*}(\cdot | t, x)) - \beta J(t, x; \pi^{\phi^*}) \right] \frac{\partial \pi^{\phi^*}}{\partial \phi}(a | t, x) \right. \\
 & \left. + \gamma q(t, x, a, \phi^*) \pi^{\phi^*}(a | t, x) \right\} da \\
 = & \int_{\mathcal{A}} \check{r}(t, x, a; \phi^*) \pi^{\phi^*}(a | t, x) da.
 \end{aligned} \tag{17}$$

This is the same type of equation as (8) involved in the Feynman–Kac formula. In the same way as (8) leading to Theorem 4, we can prove the following conclusion.

Theorem 6 *If there exists an interior optimal point ϕ^* that maximizes $J(0, x; \pi^\phi)$ for any $x \in \mathbb{R}^d$, then*

$$\begin{aligned}
 0 = & \mathbb{E}^{\mathbb{P}} \left[\int_0^T \eta_s \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^{\phi^*}(a_s^{\pi^{\phi^*}} | s, X_s^{\pi^{\phi^*}}) + \zeta_s \right] [dJ(s, X_s^{\pi^{\phi^*}}; \pi^{\phi^*}) \right. \right. \\
 & + [r(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}) + \gamma p(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \pi^{\phi^*}(\cdot | s, X_s^{\pi^{\phi^*}})) - \beta J(s, X_s^{\pi^{\phi^*}}; \pi^{\phi^*})] ds \Big] \\
 & \left. \left. \left. + \gamma q(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \phi^*) ds \right\} \middle| X_0^{\pi^{\phi^*}} = x \right]
 \end{aligned} \tag{18}$$

for any $\eta, \zeta \in L^2_{\mathcal{F}X\pi^{\phi^*}}([0, T]; J(\cdot, X^{\pi^{\phi^*}}; \pi^{\phi^*}))$.

If we take $\eta_s = e^{-\beta s}$, then the right hand side of (18) coincides with $g(0, x, \phi^*)$. However, though only a necessary condition, (18) contains infinitely many equations with different

5. A theoretically optimal policy π^* indeed maximizes $J(t, x; \pi)$ for any (t, x) , based on the verification theorem; see Yong and Zhou (1999).

test functions η . More importantly, besides the flexibility of choosing different sets of test functions, (18) provides a way to derive a system of equations based on only *past* observations and, hence, enables online learning. For example, by taking $\eta_s = 0$ on $[t, T]$, (18) involves sample trajectories up to only the present time t . Thus, learning the optimal policy either offline or online boils down to solving a system of equations (with suitably chosen test functions) via stochastic approximation to find ϕ^* .

In sum, Theorems 5 and 6 foreshadow two different types of algorithms which we will develop in the next subsection.

3.3 Actor–Critic Algorithms

We now design actor–critic (AC) algorithms by combining the PE and the PG steps. For the former, Jia and Zhou (2022a) develop two methods, those of martingale loss function and martingale orthogonality conditions, to devise several online/offline PE algorithms for the continuous setting. As discussed in Subsection 3.1, one can adopt any of these algorithms that is suitable for the given learning context and computational resource to estimate the value function of any given policy. Here we focus on how to update the policy based on our previous theoretical analysis on PG.

First, in the offline setting where full state trajectories under any given policy can be repeatedly sampled and observed, the gradient of the value function w.r.t. the policy is given by (14), which can be estimated using future samples from any *current* time–state (t, x) . That is, $g(t, x; \phi)$ is the gradient direction that would maximally improve the total reward at (t, x) .

For online learning, as explained earlier, (14) is no longer implementable. Instead of computing gradients, we turn to (18) for directly solving the optimal policy. Specifically, at any current time t , we choose $\eta_s = 0$ for $s \in [t, T]$ so that the integral in (18) only utilizes past observations up to t , and hence is computable. Therefore, in the online setting one applies stochastic approximation to solve the optimal condition (18) in order to search for the optimal policy ϕ^* .

Recall that $J^\theta \equiv J^\theta(\cdot, \cdot)$, where $J^\theta(t, x) \in \mathbb{R}$, is a family of scalar functions on $(t, x) \in [0, T] \times \mathbb{R}^d$ parameterized by $\theta \in \Theta \subseteq \mathbb{R}^{L_\theta}$, and $\pi^\phi \equiv \pi^\phi(\cdot | \cdot, \cdot)$, where $\pi^\phi(\cdot | t, x) \in \mathcal{P}(\mathcal{A})$, is a family of pdf-valued policy functions on $(t, x) \in [0, T] \times \mathbb{R}^d$ parameterized by $\phi \in \Phi \subseteq \mathbb{R}^{L_\phi}$. The aim of an AC algorithm is to find the optimal (θ, ϕ) jointly, by updating the two parameters alternately. Note that, although our problem is continuous in time, the final algorithmic implementation requires discretizing time. For simplicity, we use equally spaced mesh grid $t_k = k\Delta t$, with $k = 0, \dots, K = \lfloor \frac{T}{\Delta t} \rfloor$.

We now present the following pseudo codes in Algorithms 1 and 2. Algorithm 1 is for offline-episodic learning, where full trajectories are sampled and observed repeatedly during different episodes and (θ, ϕ) are updated after one whole episode. Algorithm 2 is for online incremental learning, where only the past sample trajectory is available and (θ, ϕ) are updated in real-time incrementally.

Note that Algorithms 1 and 2 presented here are just for illustrative purpose; there is ample flexibility to devise their variants depending on the specific problems concerned. In particular, the choice of test functions dictates in which sense we approximate the value

Algorithm 1 Offline–Episodic Actor–Critic Algorithm

Inputs: initial state x_0 , horizon T , time step Δt , number of episodes N , number of mesh grids K , initial learning rates $\alpha_\theta, \alpha_\phi$ and a learning rate schedule function $l(\cdot)$ (a function of the number of episodes), functional form of the value function $J^\theta(\cdot, \cdot)$, functional form of the policy $\pi^\phi(\cdot|\cdot, \cdot)$, functional form of the regularizer $p(t, x, a, \pi(\cdot))$, functional forms of the test functions $\xi(t, x_{\wedge t}), \zeta(t, x_{\wedge t})$, and temperature parameter γ .

Required program: an environment simulator $(x', r) = \text{Environment}_{\Delta t}(t, x, a)$ that takes current time-state pair (t, x) and action a as inputs and generates state x' at time $t + \Delta t$ and the instantaneous reward r at time t .

Learning procedure:

 Initialize θ, ϕ .

for episode $j = 1$ **to** N **do**

 Initialize $k = 0$. Observe the initial state x_0 and store $x_{t_k} \leftarrow x_0$.

while $k < K$ **do**

 Compute and store the test function $\xi_{t_k} = \xi(t_k, x_{t_0}, \dots, x_{t_k}), \zeta_{t_k} = \zeta(t_k, x_{t_0}, \dots, x_{t_k})$.

 Generate action $a_{t_k} \sim \pi^\phi(\cdot|t_k, x_{t_k})$.

 Apply a_{t_k} to the environment simulator $(x, r) = \text{Environment}_{\Delta t}(t_k, x_{t_k}, a_{t_k})$, and observe the output new state x and reward r . Store $x_{t_{k+1}} \leftarrow x$ and $r_{t_k} \leftarrow r$.

 Update $k \leftarrow k + 1$.

end while

Compute

$$\Delta\theta = \sum_{i=0}^{K-1} \xi_{t_i} [J^\theta(t_{i+1}, x_{t_{i+1}}) - J^\theta(t_i, x_{t_i}) + r_{t_i} \Delta t + \gamma p(t_i, x_{t_i}, a_{t_i}, \pi^\phi(\cdot|t_i, x_{t_i})) \Delta t - \beta J^\theta(t_i, x_{t_i}) \Delta t],$$

$$\Delta\phi = \sum_{i=0}^{K-1} e^{-\beta t_i} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_{t_i}|t_i, x_{t_i}) + \zeta_{t_i} \right] [J^\theta(t_{i+1}, x_{t_{i+1}}) - J^\theta(t_i, x_{t_i}) + r_{t_i} \Delta t + \gamma p(t_i, x_{t_i}, a_{t_i}, \pi^\phi(\cdot|t_i, x_{t_i})) \Delta t - \beta J^\theta(t_i, x_{t_i}) \Delta t] + \gamma \frac{\partial p}{\partial \phi}(t_i, x_{t_i}, a_{t_i}, \pi^\phi(\cdot|t_i, x_{t_i})) \Delta t \right\}.$$

 Update θ (policy evaluation) by

$$\theta \leftarrow \theta + l(j) \alpha_\theta \Delta\theta.$$

 Update ϕ (policy gradient) by

$$\phi \leftarrow \phi + l(j) \alpha_\phi \Delta\phi.$$

end for

Algorithm 2 Online-Incremental Actor–Critic Algorithm

Inputs: initial state x_0 , horizon T , time step Δt , number of mesh grids K , initial learning rates $\alpha_\theta, \alpha_\phi$ and learning rate schedule function $l(\cdot)$ (a function of the number of episodes), functional form of the value function $J^\theta(\cdot, \cdot)$, functional form of the policy $\pi^\phi(\cdot|\cdot, \cdot)$, functional form of the regularizer $p(t, x, a, \pi(\cdot))$, functional forms of the test functions $\xi(t, x_{\cdot\wedge t}), \eta(t, x_{\cdot\wedge t}), \zeta(t, x_{\cdot\wedge t})$, and temperature parameter γ .

Required program: an environment simulator $(x', r) = \text{Environment}_{\Delta t}(t, x, a)$ that takes current time-state pair (t, x) and action a as inputs and generates state x' at time $t + \Delta t$ and the instantaneous reward r at time t .

Learning procedure:

Initialize θ, ϕ .

for episode $j = 1$ **to** ∞ **do**

Initialize $k = 0$. Observe the initial state x_0 and store $x_{t_k} \leftarrow x_0$.

while $k < K$ **do**

Compute test function $\xi_{t_k} = \xi(t_k, x_{t_0}, \dots, x_{t_k})$, $\eta_{t_k} = \eta(t_k, x_{t_0}, \dots, x_{t_k})$, and $\zeta_{t_k} = \zeta(t_k, x_{t_0}, \dots, x_{t_k})$.

Generate action $a_{t_k} \sim \pi^\phi(\cdot|t_k, x_{t_k})$.

Apply a_{t_k} to the environment simulator $(x, r) = \text{Environment}_{\Delta t}(t_k, x_{t_k}, a_{t_k})$, and observe the output new state x and reward r . Store $x_{t_{k+1}} \leftarrow x$ and $r_{t_k} \leftarrow r$.

Compute

$$\begin{aligned} \delta &= J^\theta(t_{k+1}, x_{t_{k+1}}) - J^\theta(t_k, x_{t_k}) + r_{t_k} \Delta t \\ &\quad + \gamma p(t_k, x_{t_k}, a_{t_k}, \pi^\phi(\cdot|t_k, x_{t_k})) \Delta t - \beta J^\theta(t_k, x_{t_k}) \Delta t, \\ \Delta\theta &= \xi_{t_k} \delta, \\ \Delta\phi &= \eta_{t_k} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_{t_k}|t_k, x_{t_k}) + \zeta_{t_k} \right] \delta + \gamma \frac{\partial p}{\partial \phi}(t_k, x_{t_k}, a_{t_k}, \pi^\phi(\cdot|t_k, x_{t_k})) \Delta t \right\}. \end{aligned}$$

Update θ (policy evaluation) by

$$\theta \leftarrow \theta + l(j) \alpha_\theta \Delta\theta.$$

Update ϕ (policy gradient) by

$$\phi \leftarrow \phi + l(j) \alpha_\phi \Delta\phi.$$

Update $k \leftarrow k + 1$

end while

end for

function and policy.⁶ For example, if we take the test functions $\xi_t = \frac{\partial J^\theta}{\partial \theta}(t, X_t)$, and $\eta_t = e^{-\beta t}$, then we have essentially TD(0) AC algorithms. If we take $\xi_t = \int_0^t \lambda^{t-s} \frac{\partial J^\theta}{\partial \theta}(s, X_s) ds$, $\zeta_t = \int_0^{t-\Delta t} \lambda^{t-s} \frac{\partial}{\partial \phi} \log \pi^\phi(a_s^\pi | s, X_s) ds$, then we end up with TD(λ) algorithms (Sutton and Barto, 2018). Moreover, in the PE part of the algorithms we can also use other methods (online or offline) as summarized in Subsection 3.1.

Finally, we reiterate that the main purpose of this paper is to provide a theoretical foundation to guide designing AC algorithms, instead of comparing which algorithm performs better. As such, we only present the TD-type algorithms for illustration, acknowledging that there are multiple ways to combine PE and the newly developed PG methods to design new learning algorithms.

4. Extension to Ergodic Tasks

In this section we extend our results and algorithms to ergodic (long-term average) tasks, which are also commonly studied in the RL literature. The ergodic objective is one possible formulation of continuing tasks, in which a learning algorithm is based on only one single trajectory.

Consider a regularized ergodic objective function

$$\begin{aligned} & \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} [r(\tilde{X}_s^\pi, a) + \gamma p(\tilde{X}_s^\pi, a, \pi(\cdot | \tilde{X}_s^\pi))] \pi(a | \tilde{X}_s^\pi) da ds \Big| \tilde{X}_t^\pi = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_t^T [r(X_s^\pi, a_s^\pi) + \gamma p(X_s^\pi, a_s^\pi, \pi(\cdot | X_s^\pi))] ds \Big| X_t^\pi = x \right], \end{aligned}$$

where p is the regularizer and $\gamma \geq 0$ is the temperature parameter. Note that now the running reward, the regularizer and the policy do not depend on time explicitly due to the stationary nature of ergodic tasks.

One way to study an ergodic task is to connect it to a discounted, infinite horizon problem:

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}^W} \left[\int_t^\infty e^{-\beta(s-t)} \int_{\mathcal{A}} [r(\tilde{X}_s^\pi, a) + \gamma p(\tilde{X}_s^\pi, a, \pi(\cdot | \tilde{X}_s^\pi))] \pi(a | \tilde{X}_s^\pi) da ds \Big| \tilde{X}_t^\pi = x \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\int_t^\infty e^{-\beta(s-t)} [r(X_s^\pi, a_s^\pi) + \gamma p(X_s^\pi, a_s^\pi, \pi(\cdot | X_s^\pi))] ds \Big| X_t^\pi = x \right]. \end{aligned}$$

It has been shown that, under suitable conditions, the optimal value function of the discounted infinite horizon problem converges to the optimal ergodic reward as the discount factor $\beta \rightarrow 0$; see, e.g., Borkar and Ghosh (1988, 1990); Bensoussan and Frehse (1992).

Here, we opt for a direct treatment of ergodic problems. According to Sutton and Barto (2018, page 249), ergodic tasks are actually better behaved than continuing tasks with discounting. For a systematic account of classical ergodic control theory in continuous time, see Arapostathis et al. (2012) and the references therein.

6. See Jia and Zhou (2022a) for detailed discussions on this point for the PE part. Also, to save computational and memory cost of algorithms, we usually choose test functions that can be computed incrementally. For example, in a TD(λ) algorithm, $\xi_{t_k} = \int_0^{t_k} \lambda^{t_k-s} \frac{\partial J^\theta}{\partial \theta}(s, X_s) ds \approx \lambda^{\Delta t} \xi_{t_{k-1}} + \frac{\partial J^\theta}{\partial \theta}(t_k, X_{t_k}) \Delta t$, and $\zeta_{t_k} \approx \lambda^{\Delta t} \zeta_{t_{k-1}} + \frac{\partial}{\partial \phi} \log \pi^\phi(a_{t_{k-1}}^\pi | t_{k-1}, X_{t_{k-1}}) \Delta t$, which can be calculated recursively.

We first present the ergodic version of the Feynman–Kac formula.

Lemma 7 *Let $\pi = \pi(\cdot|\cdot)$ be a given (time-invariant) policy. Suppose there is a function $J(\cdot; \pi) \in C^2(\mathbb{R}^d)$ and a scalar $V(\pi) \in \mathbb{R}$ satisfying*

$$\int_{\mathcal{A}} [\mathcal{L}^a J(x; \pi) + r(x, a) + \gamma p(x, a, \pi(\cdot|x))] \pi(a|x) da - V(\pi) = 0, \quad x \in \mathbb{R}^d. \quad (19)$$

Then for any $t \geq 0$,

$$\begin{aligned} V(\pi) &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} [r(\tilde{X}_s^\pi, a) + \gamma p(\tilde{X}_s^\pi, a, \pi(\cdot|\tilde{X}_s^\pi))] \pi(a|\tilde{X}_s^\pi) da ds \middle| \tilde{X}_t^\pi = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_t^T [r(X_s^\pi, a_s^\pi) + \gamma p(X_s^\pi, a_s^\pi, \pi(\cdot|X_s^\pi))] ds \middle| X_t^\pi = x \right]. \end{aligned} \quad (20)$$

Moreover, $J(X_t^\pi; \pi) + \int_0^t [r(X_s^\pi, a_s^\pi) + \gamma p(X_s^\pi, a_s^\pi, \pi(\cdot|X_s^\pi)) - V(\pi)] ds$ is an $(\mathcal{F}^{X^\pi}, \mathbb{P})$ -martingale.

We emphasize that the solution to (19) is a pair of (J, V) , where $J(\cdot; \pi)$ is a function of the state and $V(\pi) \in \mathbb{R}$ is a scalar. The long term average of the payoff does not depend on the initial state x nor the initial time t due to the ergodicity, and hence remains a constant as (20) implies. The function J , on the other hand, only represents the first-order approximation of long-run average and is not unique. Indeed, for any constant c , $(J + c, V)$ is also a solution to (19). We refer to V as the ‘‘value’’. Lastly, since the value does not depend on the initial time, we will fix the latter as 0 in the following discussions and applications of ergodic tasks.

For a given policy π , the PE problem is now to find a function $J(\cdot; \pi)$ and a value $V \in \mathbb{R}$, such that

$$J(X_t^\pi; \pi) + \int_0^t [r(X_s^\pi, a_s^\pi) + \gamma p(X_s^\pi, a_s^\pi, \pi(\cdot|X_s^\pi)) - V(\pi)] ds$$

is a martingale. Following Jia and Zhou (2022a), we can then design online PE algorithms based on the following martingale orthogonality conditions:

$$\mathbb{E}^{\mathbb{P}} \int_0^T \xi_t \left\{ dJ(X_t^\pi; \pi) + [r(X_t^\pi, a_t^\pi) + \gamma p(X_t^\pi, a_t^\pi, \pi(\cdot|X_t^\pi)) - V] dt \right\} = 0, \quad (21)$$

for any $T > 0$, any initial state x , and any test function $\xi \in L^2_{\mathcal{F}^{X^\pi}}([0, T]; J(X^\pi; \pi))$.

We now focus on PG. Suppose we parameterize the policy by π^ϕ , we aim to estimate $\frac{\partial V(\pi^\phi)}{\partial \phi}$. Taking the derivative in ϕ in (19), we obtain

$$\begin{aligned} \frac{\partial V(\pi^\phi)}{\partial \phi} &= \int_{\mathcal{A}} [\mathcal{L}^a J(x; \pi^\phi) + r(x, a) + \gamma p(x, a, \pi^\phi(\cdot|x))] \frac{\partial \pi^\phi(a|x)}{\partial \phi} da \\ &\quad + \int_{\mathcal{A}} \mathcal{L}^a \frac{\partial J(x; \pi^\phi)}{\partial \phi} \pi^\phi(a|x) da + \gamma \int_{\mathcal{A}} \frac{\partial p(x, a, \pi^\phi(\cdot|x))}{\partial \phi} \pi^\phi(a|x) da. \end{aligned}$$

Denote $q(x, a, \phi) := \frac{\partial}{\partial \phi} p(x, a, \boldsymbol{\pi}^\phi(\cdot|x))$,

$$\check{r}(x, a; \phi) := \left[\mathcal{L}^a J(x; \boldsymbol{\pi}^\phi) + r(x, a) + \gamma p(x, a, \boldsymbol{\pi}(\cdot|x)) \right] \frac{\frac{\partial \boldsymbol{\pi}^\phi(a|x)}{\partial \phi}}{\boldsymbol{\pi}^\phi(a|x)} + \gamma q(x, a, \phi),$$

and $g(x; \phi) := \frac{\partial}{\partial \phi} J(x; \boldsymbol{\pi}^\phi)$. Then

$$\int_{\mathcal{A}} [\mathcal{L}^a g(x; \phi) + \check{r}(x, a; \phi)] \boldsymbol{\pi}^\phi(a|x) da - \frac{\partial V(\boldsymbol{\pi}^\phi)}{\partial \phi} = 0.$$

Therefore, analogous to the case of episodic tasks, $\frac{\partial V(\boldsymbol{\pi}^\phi)}{\partial \phi}$ is the value corresponding to the long-term average of a different running reward, according to the ergodic Feynman–Kac formula (Lemma 7); that is

$$\begin{aligned} \frac{\partial V(\boldsymbol{\pi}^\phi)}{\partial \phi} &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_0^T \check{r}(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}; \phi) dt \middle| X_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \int_{\mathcal{A}} \check{r}(\tilde{X}_t^{\boldsymbol{\pi}^\phi}, a; \phi) \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) da dt \middle| \tilde{X}_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \int_{\mathcal{A}} \left\{ [\mathcal{L}^a J(\tilde{X}_t^{\boldsymbol{\pi}^\phi}; \boldsymbol{\pi}^\phi) + r(\tilde{X}_t^{\boldsymbol{\pi}^\phi}, a) + \gamma p(\tilde{X}_t^{\boldsymbol{\pi}^\phi}, a, \boldsymbol{\pi}^\phi(\cdot|\tilde{X}_t^{\boldsymbol{\pi}^\phi}))] \right. \right. \\ &\quad \left. \left. \times \frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) + \gamma q(\tilde{X}_t^{\boldsymbol{\pi}^\phi}, a, \phi) \right\} \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) da dt \middle| \tilde{X}_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_0^T \left\{ \frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a_t^{\boldsymbol{\pi}^\phi} | X_t^{\boldsymbol{\pi}^\phi}) [dJ(X_t^{\boldsymbol{\pi}^\phi}; \boldsymbol{\pi}^\phi) + r(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}) dt \right. \right. \\ &\quad \left. \left. + \gamma p(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \boldsymbol{\pi}^\phi(\cdot|X_t^{\boldsymbol{\pi}^\phi})) dt \right\} + \gamma q(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \phi) dt \right] \middle| X_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_0^T \left\{ \frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a_t^{\boldsymbol{\pi}^\phi} | X_t^{\boldsymbol{\pi}^\phi}) [dJ(X_t^{\boldsymbol{\pi}^\phi}; \boldsymbol{\pi}^\phi) + r(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}) dt \right. \right. \\ &\quad \left. \left. + \gamma p(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \boldsymbol{\pi}^\phi(\cdot|X_t^{\boldsymbol{\pi}^\phi})) dt - V dt \right\} + \gamma q(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \phi) dt \right] \middle| X_0^{\boldsymbol{\pi}^\phi} = x \right], \end{aligned} \tag{22}$$

where the last equality is due to

$$\begin{aligned} &\mathbb{E}^{\mathbb{P}} \left[\int_0^T V \frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a_t^{\boldsymbol{\pi}^\phi} | X_t^{\boldsymbol{\pi}^\phi}) dt \middle| X_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= V \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T dt \int_{\mathcal{A}} \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) \frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) da \middle| \tilde{X}_0^{\boldsymbol{\pi}^\phi} = x \right] \\ &= V \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T dt \frac{\partial}{\partial \phi} \int_{\mathcal{A}} \boldsymbol{\pi}^\phi(a|\tilde{X}_t^{\boldsymbol{\pi}^\phi}) da \middle| \tilde{X}_0^{\boldsymbol{\pi}^\phi} = x \right] = 0. \end{aligned}$$

An ergodic task is a continuing task so we are naturally interested in online algorithms only. We can design two algorithms based on the analysis above. The first one follows directly from the representation (22), in which the policy gradient is the expectation of a long-run average and hence can be estimated online incrementally by

$$\begin{aligned} &\frac{\partial}{\partial \phi} \log \boldsymbol{\pi}^\phi(a_t^{\boldsymbol{\pi}^\phi} | X_t^{\boldsymbol{\pi}^\phi}) \left[dJ(X_t^{\boldsymbol{\pi}^\phi}; \boldsymbol{\pi}^\phi) + [r(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}) \right. \\ &\quad \left. + \gamma p(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \boldsymbol{\pi}^\phi(\cdot|X_t^{\boldsymbol{\pi}^\phi})) - V] dt \right] + \gamma q(X_t^{\boldsymbol{\pi}^\phi}, a_t^{\boldsymbol{\pi}^\phi}, \phi) dt \end{aligned}$$

since it will converge to its stationary distribution as $t \rightarrow \infty$.⁷

Moreover, due to the martingale orthogonality condition (21), we can also add a test function ζ as we did in (16). Consequently, the algorithm updates ϕ by gradient ascent:

$$\begin{aligned} \phi \leftarrow \phi + \alpha_\phi \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_t^{\pi^\phi} | X_t^{\pi^\phi}) + \zeta_t \right] \left[dJ(X_t^{\pi^\phi}; \pi^\phi) \right. \right. \\ \left. \left. + [r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) + \gamma p(X_t^{\pi^\phi}, a_t^{\pi^\phi}, \pi^\phi(\cdot | X_t^a)) - V] dt \right] + \gamma q(X_t^{\pi^\phi}, a_t^{\pi^\phi}, \phi) dt \right\}. \end{aligned}$$

The second algorithm applies a test function η and stochastic approximation to solve the optimality condition as in Theorem 6, by updating

$$\begin{aligned} \phi \leftarrow \phi + \alpha_\phi \eta_t \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_t^{\pi^\phi} | X_t^{\pi^\phi}) + \zeta_t \right] \left[dJ(X_t^{\pi^\phi}; \pi^\phi) \right. \right. \\ \left. \left. + [r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) + \gamma p(X_t^{\pi^\phi}, a_t^{\pi^\phi}, \pi^\phi(\cdot | X_t^{\pi^\phi})) - V] dt \right] + \gamma q(X_t^{\pi^\phi}, a_t^{\pi^\phi}, \phi) dt \right\}. \end{aligned}$$

Observe the two algorithms above differ by only the presence of the test function η . To illustrate, we describe the second one in Algorithm 3.

5. Applications

In this section we report simulation experiments on our algorithms in two applications. The first one is mean–variance portfolio selection in a finite time horizon with multiple episodes of simulated stock price data. The second application is ergodic linear–quadratic control with a single sample trajectory.

5.1 Mean–Variance Portfolio Selection

We first review the formulation of the exploratory mean–variance portfolio selection problem proposed by Wang and Zhou (2020). The investment universe consists of one risky asset (e.g. a stock index) and one risk-free asset (e.g. a saving account) whose risk-free interest rate is r . The price of the risky asset is governed by a geometric Brownian motion with mean μ and volatility $\sigma > 0$ on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}^W; \{\mathcal{F}_t^W\}_{0 \leq t \leq T})$:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t. \tag{23}$$

Denote by $\rho = \frac{\mu-r}{\sigma}$ the Sharpe ratio of the risky asset.

An agent has a fixed investment horizon $0 < T < \infty$ and an initial endowment x_0 . A self-financing portfolio is represented by the real-valued adapted process $a = \{a_t, 0 \leq t \leq T\}$,

7. To be more specific, the reason why an infinitesimal increment of the (inner) integral can be used as an estimate for the gradient is due to the ergodicity of the state process. The expression of the gradient (22) is the long-time average of the integrand of the inner integral, which converges to its expectation with respect to the stationary measure. On the other hand, the distribution of the integrand itself also converges to its stationary measure. Therefore, the integrand itself becomes an asymptotically unbiased estimate for the gradient as time tends to infinity. For a brief summary of the ergodicity properties, see Sandrić (2017). More details can be found in Part III of Meyn and Tweedie (2012).

Algorithm 3 Actor–Critic Algorithm for Ergodic Tasks

Inputs: initial state x_0 , time step Δt , initial learning rates $\alpha_\theta, \alpha_\phi, \alpha_V$ and learning rate schedule function $l(\cdot)$ (a function of time), functional form of the value function $J^\theta(\cdot)$, functional form of the policy $\pi^\phi(\cdot|\cdot)$, functional form of the regularizer $p(x, a, \pi(\cdot))$, functional forms of test functions $\xi(x_{\cdot\wedge t}), \eta(x_{\cdot\wedge t}), \zeta(x_{\cdot\wedge t})$, and temperature parameter γ .

Required program: an environment simulator $(x', r) = \text{Environment}_{\Delta t}(x, a)$ that takes initial state x and action a as inputs and generates a new state x' (at Δt) and an instantaneous reward r .

Learning procedure:

Initialize θ, ϕ, V . Initialize $k = 0$. Observe the initial state x_0 and store $x_{t_k} \leftarrow x_0$.

loop

Compute test function $\xi_{t_k} = \xi(x_{t_0}, \dots, x_{t_k})$, $\eta_{t_k} = \eta(x_{t_0}, \dots, x_{t_k})$ and $\zeta_{t_k} = \zeta(x_{t_0}, \dots, x_{t_k})$.

Generate action $a \sim \pi^\phi(\cdot|x)$.

Apply a to the environment simulator $(x', r) = \text{Environment}_{\Delta t}(x, a)$, and observe the output new state x' and reward r . Store $x_{t_{k+1}} \leftarrow x'$.

Compute

$$\begin{aligned} \delta &= J^\theta(x') - J^\theta(x) + r\Delta t + \gamma p(x, a, \pi^\phi(\cdot|x))\Delta t - V\Delta t, \\ \Delta\theta &= \xi_{t_k} \delta, \\ \Delta V &= \delta, \\ \Delta\phi &= \eta_{t_k} \left\{ \left[\frac{\partial}{\partial\phi} \log \pi^\phi(a|x) + \zeta_{t_k} \right] \delta + \gamma \frac{\partial p}{\partial\phi}(x, a, \pi^\phi(\cdot|x))\Delta t \right\}. \end{aligned}$$

Update θ and V (policy evaluation) by

$$\theta \leftarrow \theta + l(k\Delta t)\alpha_\theta\Delta\theta,$$

$$V \leftarrow V + l(k\Delta t)\alpha_V\Delta V.$$

Update ϕ (policy gradient) by

$$\phi \leftarrow \phi + l(k\Delta t)\alpha_\phi\Delta\phi.$$

Update $x \leftarrow x'$ and $k \leftarrow k + 1$.

end loop

where a_t is the discounted dollar value invested in the risky asset at time t . Then the discounted value of this portfolio satisfies the wealth equation

$$dx_t^a = a_t[(\mu - r)dt + \sigma dW_t] = a_t \frac{d(e^{-rt}S_t)}{e^{-rt}S_t}, \quad x_0^a = x_0, \quad (24)$$

where $e^{-rt}S_t$ is the discounted stock price. We stress that the model on the stock price (23) is mainly for theoretical analysis and for generating samples in our simulation; we do not assume that the agent knows its parameters.

The agent has the mean–variance preference, namely, she aims to minimize the variance of the discounted value of the portfolio at T while achieving a given level of expected return:

$$\min_a \text{Var}(x_T^a), \quad \text{subject to } \mathbb{E}[x_T^a] = z, \quad (25)$$

where z is the target value, and the variance and expectation are w.r.t. the probability measure \mathbb{P}^W .

This problem is not a standard stochastic control problem and cannot be solved directly by the dynamic programming (DP) principle, or any DP-based reinforcement learning algorithms such as Q-learning. This is because the variance term causes *time-inconsistency* which violates the assumptions of DP. Strotz (1955) discusses three types of agents when facing time-inconsistency. Here, we consider one of them – the so-called *pre-committed* agent who solves the problem at time 0 and sticks to it afterwards.⁸ For this type of agent, to overcome the difficulty of DP not being directly applicable, Zhou and Li (2000) extend the embedding method, initially introduced by Li and Ng (2000) for the discrete-time mean–variance problem, to transform (25) into an equivalent, unconstrained, and expectation-only problem:

$$\min_a \mathbb{E}[(x_T^a)^2] - z^2 - 2w(\mathbb{E}[x_T^a] - z) = \min_a \mathbb{E}[(x_T^a - w)^2] - (w - z)^2,$$

where w is the Lagrange multiplier associated with the constraint $\mathbb{E}[x_T^a] = z$. This new problem is time-consistent and therefore can be solved by DP. Once the optimal a^* is derived, w can be obtained by the equation $\mathbb{E}[x_T^{a^*}] = z$.

In a reinforcement learning framework, Wang and Zhou (2020) allow randomized actions to incorporate exploration. A stochastic policy is denoted by $\pi = \pi(\cdot|t, x)$, namely, at any current time–wealth pair (t, x) , the total amount of discounted wealth invested in the stock is a random draw from the distribution with the density function $\pi(\cdot|t, x)$. Under such a policy, we denote by $\tilde{X}^\pi = \{\tilde{X}_s^\pi : t \leq s \leq T\}$ the solution to the following SDE

$$d\tilde{X}_s^\pi = (\mu - r) \int_{\mathbb{R}} a \pi(a|s, \tilde{X}_s^\pi) da ds + \sigma \sqrt{\int_{\mathbb{R}} a^2 \pi(a|s, \tilde{X}_s^\pi) da} dW_s; \quad \tilde{X}_t^\pi = x,$$

which is (5) specializing to the current case.

8. The other two types are the *naïve* one who re-optimizes at any given time and the *sophisticated* one who seeks subgame perfect Nash equilibria among her-selves at different times. The latter has been well studied in the continuous-time setting in recent years; see e.g. Ekeland and Lazrak (2006); Björk et al. (2014); Basak and Chabakauri (2010); Dai et al. (2021). The RL counterpart is studied in Dai et al. (2020).

Moreover, an entropy regularizer is added to incentivize exploration. Mathematically, the entropy-regularized mean–variance portfolio choice problem is to solve

$$V(t, x; w) = \min_{\pi} \mathbb{E} \left[(\tilde{X}_T^{\pi} - w)^2 - \gamma \int_t^T \mathcal{H}(\pi_s) ds \mid \tilde{X}_t^{\pi} = x \right] - (w - z)^2, \quad (26)$$

where z is the target expected terminal wealth, $\pi_s = \pi(\cdot | s, \tilde{X}_s^{\pi})$, $t \leq s \leq T$, \mathcal{H} is the differential entropy $\mathcal{H}(\pi) = - \int_{\mathcal{A}} \pi(a) \log \pi(a) da$, γ is the temperature parameter, and w is the Lagrange multiplier similar to that introduced earlier.

We follow Wang and Zhou (2020) to parameterize the value function by

$$J^{\theta}(t, x; w) = (x - w)^2 e^{-\theta_3(T-t)} + \theta_2(t^2 - T^2) + \theta_1(t - T) - (w - z)^2,$$

and parameterize the policy by

$$\pi^{\phi}(\cdot | t, x; w) = \mathcal{N}(\cdot | -\phi_1(x - w), e^{\phi_2 + \phi_3(T-t)}),$$

where $\mathcal{N}(\cdot | \alpha, \delta^2)$ is the pdf of the normal distribution with mean α and variance δ^2 . These function approximators are derived in Wang and Zhou (2020) by exploiting the special structure of the underlying problem; see also Appendix B1.

There is no running reward from the actions except the regularizer

$$\mathcal{H}(\pi^{\phi}(\cdot | t, x; w)) = -\frac{1}{2} \log(2\pi e) - \frac{1}{2}[\phi_2 + \phi_3(T - t)] =: \hat{p}(t, \phi).$$

Note that the regularizer turns out to be independent of the state x . Finally, the discount factor is $\beta = 0$.

From this point on, we depart from Wang and Zhou (2020) and instead apply the methods developed in this paper to solve the problem. We choose the test functions for PE as the following gradients, in accordance with the most popular $TD(0)$ algorithm:⁹

$$\frac{\partial J^{\theta}}{\partial \theta_1}(t, x; w) = t - T, \quad \frac{\partial J^{\theta}}{\partial \theta_2}(t, x; w) = t^2 - T^2, \quad \frac{\partial J^{\theta}}{\partial \theta_3}(t, x; w) = (x - w)^2 e^{-\theta_3(T-t)}(t - T).$$

The PE updating rule is

$$\theta \leftarrow \theta + \alpha_{\theta} \int_0^T \frac{\partial J^{\theta}}{\partial \theta}(t, X_t^{\pi^{\phi}}; w) \left[dJ^{\theta}(t, X_t^{\pi^{\phi}}; w) + \gamma \hat{p}(t, \phi) dt \right].$$

For the PG part, the gradients of log-likelihood are

$$\begin{aligned} \frac{\partial \log \pi^{\phi}(a | t, x; w)}{\partial \phi_1} &= -(a + \phi_1(x - w))(x - w) e^{-\phi_2 - \phi_3(T-t)}, \\ \frac{\partial \log \pi^{\phi}(a | t, x; w)}{\partial \phi_2} &= -\frac{1}{2} + \frac{(a + \phi_1(x - w))^2}{2} e^{-\phi_2 - \phi_3(T-t)}, \end{aligned}$$

9. Wang and Zhou (2020) employ a mean–square TD error (MSTDE) algorithm to do PE and a policy improvement theorem to update policies. However, it is shown in Jia and Zhou (2022a) that MSTDE only minimizes the quadratic variation of the martingale, which may not lead to the true solution of PE. As discussed earlier, other PE algorithms proposed in Jia and Zhou (2022a) can also be applied.

$$\frac{\partial \log \pi^\phi(a|t, x; w)}{\partial \phi_3} = -\frac{T-t}{2} + \frac{(a + \phi_1(x-w))^2}{2} e^{-\phi_2 - \phi_3(T-t)}(T-t),$$

and those of the regularizer are

$$\frac{\partial \hat{p}}{\partial \phi_1}(t, \phi) = 0, \quad \frac{\partial \hat{p}}{\partial \phi_2}(t, \phi) = -\frac{1}{2}, \quad \frac{\partial \hat{p}}{\partial \phi_3}(t, \phi) = -\frac{T-t}{2}.$$

Accordingly, the offline PG updating rule is

$$\phi \leftarrow \phi - \alpha_\phi \int_0^T \left\{ \frac{\partial \log \pi^\phi}{\partial \phi}(a_t|t, X_t^{\pi^\phi}; w) \left[dJ^\theta(t, X_t^{\pi^\phi}; w) + \gamma \hat{p}(t, X_t^{\pi^\phi}, \phi) dt \right] + \gamma \frac{\partial \hat{p}}{\partial \phi}(t, X_t^{\pi^\phi}, \phi) dt \right\}.$$

The online counterpart of this updating rule is to remove the integral “ \int_0^T ” in the above and use only the resulting increment to update the policy at every time step.

In addition, there is the Lagrange multiplier w we need to learn: we update w based on the same stochastic approximation scheme in Wang and Zhou (2020).

We present our offline and online algorithms as Algorithms 4 and 5 respectively. Then we replicate the simulation study of Wang and Zhou (2020) with the same basic setting: $x_0 = 1$, $z = 1.4$, $T = 1$, $\Delta t = \frac{1}{252}$. Choose temperature parameter $\gamma = 0.1$. The batch size $m = 10$ for updating the Lagrange multiplier. The learning rate parameters in Wang and Zhou (2020) are set to be $\alpha_w = 0.05$, and $\alpha_\theta = \alpha_\phi = 0.0005$ with decay rate $l(j) = j^{-0.51}$. In our experiment we adopt these learning rate values for the Wang and Zhou (2020) algorithm unless the algorithm does not converge, in which case we tune the initial learning rates to guarantee convergence. For our algorithm, we set $\alpha_w = 0.05$, and $\alpha_\theta = \alpha_\phi = 0.1$ with decay rate $l(j) = j^{-0.51}$ and tune the initial learning rate when necessary. The initialization of the parameters θ and ϕ is set to be all 0 for both algorithms (the initialization is not discussed in Wang and Zhou 2020). In particular, to mimic the real scenario, we choose a reasonable size of the training sample, with length of 20 years. In each iteration, we randomly sample 128 1-year trajectories to update the rest parameters, and we train the model for $N = 2 \times 10^4$ iterations. We calculate the performance metrics – the mean, variance and Sharpe ratio of the resulting terminal wealth – of the learned policies of both methods with the training set generated from the same distribution.¹⁰ We then repeat the experiment for 100 times and report the standard deviation of each metric.

Tables 1 and 2 present the test results of the algorithm in Wang and Zhou (2020) and the offline Algorithm 4 in this paper respectively, when stock price is generated from geometric Brownian motion under different specifications of the market parameters μ and σ . Our algorithm achieves significantly higher out-of-sample average Sharpe ratios in most scenarios. Underperformance of our strategy occurs mainly when the actual return of the stock is low ($\mu = 0, \pm 0.1$). In those cases, our learned policy yields larger volatility and less stable out-of-sample performance. However, although the average out-of-sample Sharpe ratios of the learned policy are lower than those of Wang and Zhou (2020), the standard

10. Wang and Zhou (2020) report in-sample performance of the last 2000 iterations in the training set but does not present out-of-sample test results.

Algorithm 4 Offline–Episodic Actor–Critic Mean–Variance Algorithm

Inputs: initial state x_0 , horizon T , time step Δt , number of episodes N , number of time grids K , initial learning rates $\alpha_\theta, \alpha_\phi, \alpha_w$ and learning rate schedule function $l(\cdot)$ (a function of the number of episodes), and temperature parameter γ .

Required program: a market simulator $x' = \text{Market}_{\Delta t}(t, x, a)$ that takes current time-state pair (t, x) and action a as inputs and generates state x' at time $t + \Delta t$.

Learning procedure:

Initialize θ, ϕ, w .

for episode $j = 1$ **to** N **do**

Initialize $k = 0$. Observe the initial state x and store $x_{t_k} \leftarrow x$.

while $k < K$ **do**

Compute and store the test function $\xi_{t_k} = \frac{\partial J^\theta}{\partial \theta}(t_k, x_{t_k}; w)$.

Generate action $a_{t_k} \sim \pi^\phi(\cdot | t_k, x_{t_k})$.

Apply a_{t_k} to the market simulator $x = \text{Market}_{\Delta t}(t_k, x_{t_k}, a_{t_k})$, and observe the output new state x . Store $x_{t_{k+1}} \leftarrow x$.

Update $k \leftarrow k + 1$.

end while

Store the terminal wealth $X_T^{(j)} \leftarrow x_{t_K}$.

Compute

$$\Delta\theta = \sum_{i=0}^{K-1} \xi_{t_i} [J^\theta(t_{i+1}, x_{t_{i+1}}; w) - J^\theta(t_i, x_{t_i}; w) + \gamma \hat{p}(t_i, x_{t_i}, \phi) \Delta t],$$

$$\begin{aligned} \Delta\phi = & \sum_{i=0}^{K-1} \frac{\partial}{\partial \phi} \log \pi^\phi(a_{t_i} | t_i, x_{t_i}) [J^\theta(t_{i+1}, x_{t_{i+1}}) - J^\theta(t_i, x_{t_i}) + \gamma \hat{p}(t_i, x_{t_i}, \phi) \Delta t] \\ & + \gamma \frac{\partial \hat{p}}{\partial \phi}(t_i, x_{t_i}, \phi) \Delta t. \end{aligned}$$

Update θ (policy evaluation) by

$$\theta \leftarrow \theta + l(j) \alpha_\theta \Delta\theta.$$

Update ϕ (policy gradient) by

$$\phi \leftarrow \phi - l(j) \alpha_\phi \Delta\phi.$$

Update w (Lagrange multiplier) every m episodes:

if $j \equiv 0 \pmod{m}$ **then**

$$w \leftarrow w - \alpha_w \frac{1}{m} \sum_{i=j-m+1}^j X_T^{(i)}.$$

end if

end for

Algorithm 5 Online–Episodic Actor–Critic Mean–Variance Algorithm

Inputs: initial state x_0 , horizon T , time step Δt , number of episodes N , number of time grids K , initial learning rates $\alpha_\theta, \alpha_\phi, \alpha_w$ and learning rate schedule function $l(\cdot)$ (a function of the number of episodes), and temperature parameter γ .

Required program: a market simulator $x' = \text{Market}_{\Delta t}(t, x, a)$ that takes current time-state pair (t, x) and action a as inputs and generates state x' at time $t + \Delta t$.

Learning procedure:

Initialize θ, ϕ, w .

for episode $j = 1$ **to** N **do**

Initialize $k = 0$. Observe the initial state x and store $x_{t_k} \leftarrow x$.

while $k < K$ **do**

Compute and store the test function $\xi_{t_k} = \frac{\partial J^\theta}{\partial \theta}(t_k, x_{t_k}; w)$.

Generate action $a_{t_k} \sim \pi^\phi(\cdot | t_k, x_{t_k})$.

Apply a_{t_k} to the market simulator $x = \text{Market}_{\Delta t}(t_k, x_{t_k}, a_{t_k})$, and observe the output new state x . Store $x_{t_{k+1}} \leftarrow x$.

Compute

$$\Delta\theta = \xi_{t_k} [J^\theta(t_{k+1}, x_{t_{k+1}}; w) - J^\theta(t_k, x_{t_k}; w) + \gamma \hat{p}(t_k, x_{t_k}, \phi) \Delta t],$$

$$\begin{aligned} \Delta\phi = & \frac{\partial}{\partial \phi} \log \pi^\phi(a_{t_k} | t_k, x_{t_k}) [J^\theta(t_{k+1}, x_{t_{k+1}}) - J^\theta(t_k, x_{t_k}) + \gamma \hat{p}(t, x_{t_k}, \phi) \Delta t] \\ & + \gamma \frac{\partial \hat{p}}{\partial \phi}(t_k, x_{t_k}, \phi) \Delta t. \end{aligned}$$

Update θ (policy evaluation) by

$$\theta \leftarrow \theta + l(j) \alpha_\theta \Delta\theta.$$

Update ϕ (policy gradient) by

$$\phi \leftarrow \phi - l(j) \alpha_\phi \Delta\phi.$$

Update $k \leftarrow k + 1$.

end while

Store the terminal wealth $X_T^{(j)} \leftarrow x_{t_K}$.

Update w (Lagrange multiplier) every m episodes:

if $j \equiv 0 \pmod{m}$ **then**

$$w \leftarrow w - \alpha_w \frac{1}{m} \sum_{i=j-m+1}^j X_T^{(i)}.$$

end if

end for

deviations of ours are large. Hence, even in those few scenarios it is still statistically inconclusive to determine which method is better.

We further carry out tests to compare Algorithm 4 with the online Algorithm 5. In implementing Algorithm 5, we update parameters at each time step and conduct learning for 20 years, with the same simulated stock prices as in offline learning. The batch size is set to be $m = 1$ for updating the Lagrange multiplier. We also repeat the experiment for 100 times to calculate the standard deviation of each metric. The results are presented in Table 3. Compared with Table 2, offline learning outperforms online one in terms of Sharpe ratio in most cases. Moreover, the former is always preferred when it comes to stably reaching the target return (set to be 40% annually in the experiments).

With a given training data set, it is not surprising that offline learning is typically preferred because it allows us to fully use the data set by bootstrapping multiple 1-year episodes. By contrast, online learning pretends the data to come sequentially without storing past data. For example, under our online setting, the 20-year training set only contains 20 complete episodes sequentially to adjust the final terminal wealth level, unlike in the offline setting where we bootstrap multiple 1-year episodes. Therefore, offline learning uses data thoroughly and efficiently. However, other important considerations motivate or even force us to use online learning. First and foremost, data distribution may not be stationary, so offline learning may suffer from overfitting. Second, for large-scale problems, online and incremental learning is more computationally efficient in reducing storage costs and computational time. Finally, there are also computational techniques to store a certain amount of past data to boost the efficiency of online learning, such as the experience replay with off-policy learning (Zhang and Sutton, 2017; Fedus et al., 2020).

5.2 Ergodic Linear–Quadratic Control

Consider the ergodic linear–quadratic (LQ) control problem where state responds to actions in a linear way

$$dX_t = (AX_t + Ba_t)dt + (CX_t + Da_t)dW_t, \quad X_0 = x_0, \quad (27)$$

and the goal is to maximize the long term average payoff

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T r(X_t, a_t) dt \mid X_0 = x_0 \right], \quad (28)$$

with $r(x, a) = -(\frac{M}{2}x^2 + Rxa + \frac{N}{2}a^2 + Px + Qa)$.

In the entropy-regularized RL formulation, the policy is denoted by $\boldsymbol{\pi}(\cdot|x)$ and actions are generated from this policy. The corresponding goal is to maximize

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^w} \left[\int_0^T \int_{\mathbb{R}} r(\tilde{X}_t^\pi, a) \boldsymbol{\pi}(a|\tilde{X}_t^\pi) da dt + \gamma \mathcal{H}(\boldsymbol{\pi}(\cdot|\tilde{X}_t^\pi)) dt \mid \tilde{X}_0^\pi = x_0 \right], \quad (29)$$

where \mathcal{H} is the differential entropy as before. Moreover, \tilde{X}^π satisfies

$$d\tilde{X}_t^\pi = \int_{\mathbb{R}} (A\tilde{X}_t^\pi + Ba) \boldsymbol{\pi}(a|\tilde{X}_t^\pi) da dt + \sqrt{\int_{\mathbb{R}} (C\tilde{X}_t^{\pi^2} + Da)^2 \boldsymbol{\pi}(a|\tilde{X}_t^\pi) da} dW_t. \quad (30)$$

Table 1: **Out-of-sample performance of algorithm proposed in Wang and Zhou (2020) for mean–variance problem when data are generated by geometric Brownian motion.** The columns “Mean” and “Variance” report the average mean and variance, respectively, of terminal wealth over 100 independent experiments. The column “Sharpe ratio” reports the corresponding average Sharpe ratio ($\frac{\text{Mean}-1}{\sqrt{\text{Variance}}}$). The numbers in the brackets are the standard deviations.

μ	σ	Mean	Variance	Sharpe ratio
-0.5	0.1	1.4 (0.015)	0 (0.00033)	6.69 (0.096)
-0.3	0.1	1.4 (0.027)	0.01 (0.002)	3.59 (0.064)
-0.1	0.1	1.4 (0.11)	0.11 (0.059)	1.25 (0.02)
0	0.1	1.04 (0.028)	0.06 (0.057)	0.2 (9.2e-05)
0.1	0.1	1.45 (0.29)	0.43 (0.55)	0.81 (0.0057)
0.3	0.1	1.41 (0.033)	0.02 (0.0032)	3.06 (0.046)
0.5	0.1	1.4 (0.017)	0 (0.00044)	6 (0.087)
-0.5	0.2	1.4 (0.03)	0.01 (0.0031)	3.38 (0.15)
-0.3	0.2	1.41 (0.067)	0.05 (0.022)	1.81 (0.074)
-0.1	0.2	1.44 (0.3)	0.74 (1.3)	0.61 (0.0041)
0	0.2	1.04 (0.032)	0.27 (0.27)	0.1 (6.3e-05)
0.1	0.2	1.25 (0.076)	0.43 (0.23)	0.4 (0.0014)
0.3	0.2	1.42 (0.089)	0.08 (0.04)	1.54 (0.055)
0.5	0.2	1.41 (0.034)	0.02 (0.004)	3.03 (0.13)
-0.5	0.3	1.4 (0.057)	0.03 (0.016)	2.25 (0.16)
-0.3	0.3	1.41 (0.14)	0.13 (0.12)	1.2 (0.057)
-0.1	0.3	1.32 (0.12)	0.71 (0.43)	0.41 (0.0023)
0	0.3	1.04 (0.031)	0.55 (0.58)	0.07 (3e-05)
0.1	0.3	1.19 (0.11)	0.67 (0.49)	0.27 (0.0011)
0.3	0.3	1.44 (0.14)	0.22 (0.21)	1 (0.018)
0.5	0.3	1.41 (0.055)	0.04 (0.016)	2.01 (0.13)
-0.5	0.4	1.41 (0.079)	0.07 (0.041)	1.67 (0.13)
-0.3	0.4	1.43 (0.15)	0.28 (0.23)	0.86 (0.011)
-0.1	0.4	1.28 (0.12)	1.04 (0.76)	0.3 (0.0016)
0	0.4	1.04 (0.028)	0.85 (0.89)	0.05 (2.3e-05)
0.1	0.4	1.17 (0.1)	0.93 (0.82)	0.2 (0.00069)
0.3	0.4	1.46 (0.17)	0.44 (0.43)	0.74 (0.012)
0.5	0.4	1.42 (0.082)	0.09 (0.046)	1.44 (0.058)

Table 2: **Out-of-sample performance of offline learning (Algorithm 4) for mean-variance problem when data are generated by geometric Brownian motion.** The columns “Mean” and “Variance” report the average mean and variance, respectively, of terminal wealth over 100 independent experiments. The column “Sharpe ratio” reports the corresponding average Sharpe ratio ($\frac{\text{Mean}-1}{\sqrt{\text{Variance}}}$). The numbers in the brackets are the standard deviations.

μ	σ	Mean	Variance	Sharpe ratio
-0.5	0.1	1.4 (0.012)	0 (0.00011)	8.15 (0.06)
-0.3	0.1	1.4 (0.023)	0.01 (0.00084)	4.37 (0.029)
-0.1	0.1	1.41 (0.08)	0.09 (0.037)	1.37 (0.0073)
0	0.1	1.13 (0.14)	0.91 (0.49)	0.12 (0.16)
0.1	0.1	1.51 (0.27)	0.47 (0.72)	0.84 (0.0023)
0.3	0.1	1.41 (0.028)	0.01 (0.0015)	3.71 (0.025)
0.5	0.1	1.4 (0.014)	0 (0.00016)	7.35 (0.055)
-0.5	0.2	1.4 (0.025)	0.01 (0.001)	3.98 (0.047)
-0.3	0.2	1.4 (0.049)	0.04 (0.0088)	2.1 (0.017)
-0.1	0.2	1.53 (0.27)	0.93 (1)	0.62 (0.0012)
0	0.2	1.06 (0.15)	2.51 (1.6)	0.04 (0.094)
0.1	0.2	1.44 (0.37)	2.02 (1.5)	0.35 (0.21)
0.3	0.2	1.42 (0.065)	0.06 (0.018)	1.78 (0.012)
0.5	0.2	1.41 (0.029)	0.01 (0.0015)	3.58 (0.041)
-0.5	0.3	1.4 (0.04)	0.03 (0.0047)	2.54 (0.026)
-0.3	0.3	1.41 (0.088)	0.1 (0.049)	1.32 (0.007)
-0.1	0.3	1.43 (0.33)	1.79 (1.9)	0.37 (0.18)
0	0.3	1.03 (0.12)	3.46 (2.3)	0.02 (0.064)
0.1	0.3	1.26 (0.36)	2.78 (2.3)	0.18 (0.2)
0.3	0.3	1.44 (0.13)	0.17 (0.14)	1.12 (0.012)
0.5	0.3	1.41 (0.048)	0.03 (0.0076)	2.28 (0.02)
-0.5	0.4	1.41 (0.061)	0.05 (0.017)	1.8 (0.01)
-0.3	0.4	1.43 (0.15)	0.24 (0.23)	0.93 (0.014)
-0.1	0.4	1.31 (0.44)	3.13 (4.8)	0.25 (0.17)
0	0.4	1.02 (0.096)	3.77 (2.7)	0.01 (0.049)
0.1	0.4	1.14 (0.36)	3.6 (2.8)	0.1 (0.18)
0.3	0.4	1.53 (0.34)	0.74 (1.3)	0.73 (0.0024)
0.5	0.4	1.42 (0.076)	0.07 (0.032)	1.6 (0.014)

Table 3: **Out-of-sample performance of online learning (Algorithm 5) for mean-variance problem when data are generated by geometric Brownian motion.** The columns “Mean” and “Variance” report the average mean and variance, respectively, of terminal wealth over 100 independent experiments. The column “Sharpe ratio” reports the corresponding average Sharpe ratio ($\frac{\text{Mean}-1}{\sqrt{\text{Variance}}}$). The numbers in the brackets are the standard deviations.

μ	σ	Mean	Variance	Sharpe ratio
-0.5	0.1	1.78 (0.0082)	0.01 (3e-04)	7.43 (0.04)
-0.3	0.1	1.55 (0.0077)	0.02 (0.00034)	3.84 (0.027)
-0.1	0.1	1.14 (0.022)	0.01 (0.0039)	1.24 (0.0073)
0	0.1	1.01 (0.0055)	0 (0.0016)	0.12 (0.16)
0.1	0.1	1.26 (0.052)	0.09 (0.032)	0.85 (0.012)
0.3	0.1	1.83 (0.021)	0.04 (0.0023)	4.31 (0.056)
0.5	0.1	1.92 (0.017)	0.01 (0.00032)	10.82 (0.15)
-0.5	0.2	1.77 (0.015)	0.04 (0.0023)	3.65 (0.038)
-0.3	0.2	1.54 (0.018)	0.08 (0.0036)	1.89 (0.025)
-0.1	0.2	1.14 (0.043)	0.05 (0.028)	0.62 (0.006)
0	0.2	1.01 (0.01)	0.01 (0.015)	0.04 (0.091)
0.1	0.2	1.12 (0.072)	0.12 (0.09)	0.36 (0.19)
0.3	0.2	1.79 (0.048)	0.17 (0.019)	1.94 (0.052)
0.5	0.2	1.92 (0.033)	0.04 (0.0032)	4.8 (0.11)
-0.5	0.3	1.76 (0.02)	0.1 (0.0073)	2.36 (0.035)
-0.3	0.3	1.52 (0.035)	0.18 (0.018)	1.23 (0.021)
-0.1	0.3	1.12 (0.061)	0.11 (0.079)	0.39 (0.14)
0	0.3	1.01 (0.013)	0.05 (0.051)	0.02 (0.063)
0.1	0.3	1.1 (0.12)	2.27 (20)	0.17 (0.21)
0.3	0.3	1.44 (0.049)	0.18 (0.032)	1.05 (0.018)
0.5	0.3	1.73 (0.018)	0.12 (0.0072)	2.1 (0.033)
-0.5	0.4	1.74 (0.028)	0.19 (0.016)	1.7 (0.033)
-0.3	0.4	1.48 (0.062)	0.29 (0.057)	0.9 (0.017)
-0.1	0.4	1.11 (0.075)	0.19 (0.15)	0.25 (0.17)
0	0.4	1.01 (0.016)	0.11 (0.12)	0.01 (0.048)
0.1	0.4	1.04 (0.06)	0.13 (0.13)	0.08 (0.18)
0.3	0.4	1.4 (0.08)	0.28 (0.087)	0.77 (0.015)
0.5	0.4	1.71 (0.027)	0.22 (0.015)	1.52 (0.03)

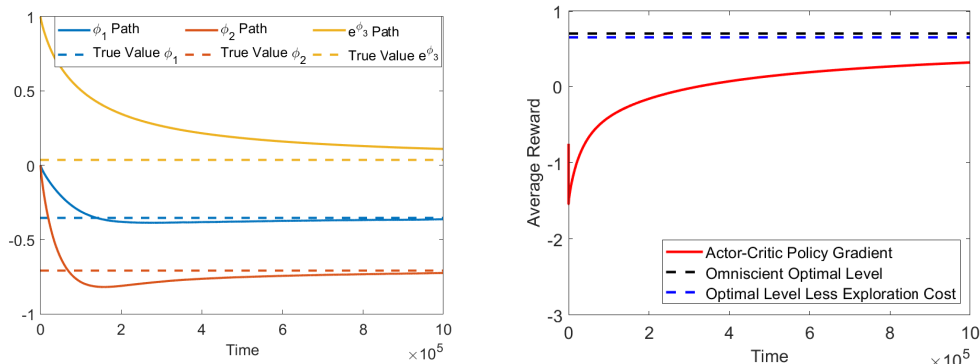
Following the same line of deductions as in Wang et al. (2020), we can show that the optimal policy is a normal distribution whose mean is linear in the state and whose variance is a constant. Therefore we parameterize the policy by $\pi^\phi(\cdot|x) = \mathcal{N}(\cdot|\phi_1x + \phi_2, e^{\phi_3})$. Moreover, the function J is parameterized as a quadratic function $J^\theta(x) = \frac{1}{2}\theta_0x^2 + \theta_1x$ (we ignore the constant term since J is unique up to a constant) and the optimal value V is an extra parameter.

This problem falls into the formulation of an ergodic task; so we directly implement Algorithm 3 in our simulation and then compare the learned parameters with the theoretically optimal ones. In addition, we compare the up-to-now average reward during the online learning process to two theoretical benchmarks. The first one is the omniscient optimal level, which is the maximum long term average reward that can be achieved by a hypothetical agent who knows completely about the environment (i.e. the correct model and model parameters) and acts optimally (the optimal policy is a deterministic one) without needing to explore (and hence there is no entropy regularization). The second benchmark is the omniscient optimal level less the exploration cost, which is the maximum long term average reward that can be achieved by the aforementioned hypothetical agent who is however forced to explore under entropy regularization.¹¹ Clearly, since exploration (rendering a stochastic policy) is inherent in the RL setting, our algorithm can at most achieve the second benchmark. In other words, after learning for a sufficiently long time, we can learn the correct optimal policy but can only expect the up-to-now average reward to approach the optimal level less the exploration cost.

To guarantee the stationarity of the controlled state process, we set $A = -1, B = C = 0$ and $D = 1$. Moreover, we set $x_0 = 0, M = N = Q = 2, R = P = 1$, and $\gamma = 0.1$. Learning rate is initialized as $\alpha_\theta = \alpha_\phi = 0.001$, and decays according to $l(t) = \frac{1}{\max\{1, \log t\}}$. All the parameters to be learned are initialized as 0 and time discretization is taken as $\Delta t = 0.01$. We repeat the experiment for 100 times.

We implement $TD(0)$ for both the PE and the PG parts of the AC algorithm, referred to as the Actor–Critic Policy Gradient algorithm in Figure 1. Namely, we choose test functions $\xi_t = \frac{\partial J^\theta}{\partial \theta}(X_t), \eta_t = 1, \zeta_t = 0$ in Algorithm 3. Figure 1 shows the convergence of the learned policy parameters along with that of the average reward along a single state sample trajectory. Observe that the average reward first decreases at the beginning of this particular trajectory. The reason may have been that during the initial iterations the underlying state process has not yet converged to the stationary distribution and the initial policies are still far away from the optimal one, and hence the average reward is dominated by a few “wrong trials”. After a sufficient amount of time, however, both the policies and the average reward start to converge to the theoretically optimal values. Between the two it takes a much longer time for the average reward to approach the optimal level as we wait for the contribution from the bad performance of the beginning period to diminish.

11. See Appendix B2 for precise definitions of these two benchmarks and detailed calculations of them.



(a) The learned parameters in the policy (b) The average reward along one sample trajectory.

Figure 1: Convergence of the learned policy and the average reward under the online learning algorithm. A single state trajectory is generated with length $T = 10^6$ under the online AC algorithm. The left panel illustrates the convergence of the policy parameters, where the dashed horizontal lines indicate the values of the respective parameters of the theoretically optimal policy to the entropy-regularized exploratory stochastic control problem. The right panel shows the convergence of the average reward, where the two dashed horizontal lines are respectively the omniscient optimal average reward without exploration when the model parameters are known, and the omniscient optimal average reward less the exploration cost. We repeat the experiment for 100 times to calculate the standard deviations of the predicted parameters, which are represented as the shaded areas. The width of each shaded area is twice the corresponding standard deviation, which is very small compared to the scale of the vertical axis.

6. Conclusion

This paper is the final installment of a “trilogy”, the first two being Wang et al. (2020) and Jia and Zhou (2022a), that endeavors to develop a systematic and unified theoretical foundation for RL in continuous time with continuous state space and possibly continuous action space. The previous two papers address exploration and PE, respectively, and this paper focuses on PG. A major finding of the current paper is that PG is intimately related to PE, and thus the martingale characterization of PE established in Jia and Zhou (2022a) can be applied to PG. Combining the theoretical results of the three papers, we propose online and offline actor–critic algorithms for general model-free RL tasks, where we learn value functions and stochastic policies simultaneously and alternately.

This series of papers are characterized by conducting all the theoretical analysis within the continuous setting and discretizing time only when implementing the algorithms. The advantages of this approach, versus discretizing time right at the start and then applying existing MDP results, are articulated in Doya (2000). Moreover, more analytical tools are at our disposal in the continuous setting, including calculus, stochastic calculus, stochastic control, and differential equations. The discrete-time versions of the various algorithms devised in the three papers are indeed well known in the discrete-time RL literature; hence

their convergence is well established. On the other hand, Jia and Zhou (2022a) prove that any convergent time-discretized PE algorithm also converges as the mesh size goes to zero. Because PG algorithms developed in the current paper are essentially derived from the martingality for PE, the same convergence also holds for them.

It is interesting to note that the derivation and representation of PG are not entirely analogous to that of MDPs. For example, the latter involves a state–action function (Q-function), whereas the former is essentially the expected integration of a term involving the value function.

The study on continuous-time RL is still in its infancy, and open questions abound. These include, to name but a few, regret bound of episodic RL problems in terms of the number of episodes, interpretation of Q-function and Q-learning in the continuous setting, and dependence of the performance of AC algorithms on the temperature parameter when there is an exploration regularizer.

Acknowledgments

We are indebted to the Action Editor and two anonymous referees for constructive comments which have led to an improved version of the paper. Zhou gratefully acknowledges financial support through a start-up grant and the Nie Center for Intelligent Asset Management at Columbia University.

Appendix A. Connections with Policy Gradient in Discrete Time

We review the classical policy gradient approach and results for discrete-time Markov decision processes (MDPs) here and compare them with their continuous-time counterparts developed in the main text.

For simplicity, we consider a time-homogeneous MDP $X = \{X_t, t = 0, 1, 2, \dots\}$ with a state space \mathcal{X} , an action space \mathcal{A} , and a transition matrix $\mathbb{P}(X_1 = x' | X_0 = x, a_0 = a) = p(x' | x, a)$. Both \mathcal{X} and \mathcal{A} are finite sets. The expected reward is $r(x, a)$ with a discount factor $\beta \in (0, 1)$. The agent's total expected reward is $\mathbb{E}[\sum_{t=0}^{\infty} \beta^t r(X_t, a_t)]$. A (stochastic) policy is denoted by $\pi^\phi(\cdot | x) \in \mathcal{P}(\mathcal{A})$, which is a probability density function on \mathcal{A} , with a suitable parameter vector $\phi \in \mathbb{R}^{L_\phi}$.

Define the value function associated with a given policy π^ϕ by

$$\begin{aligned} J(x; \pi^\phi) &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) \middle| X_0^{\pi^\phi} = x \right] \\ &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \mathbb{E}_{a_t \sim \pi^\phi(\cdot | X_t^{\pi^\phi})} [r(X_t^{\pi^\phi}, a_t)] \middle| X_0^{\pi^\phi} = x \right]. \end{aligned} \quad (31)$$

We are interested in the gradient of the value function with respect to the policy parameter ϕ , that is, $\frac{\partial J(x; \pi^\phi)}{\partial \phi}$. The classical policy gradient theorem (e.g., Sutton et al. 1999, Theorem 1) states that

$$\frac{\partial J(x; \pi^\phi)}{\partial \phi} = \sum_{x' \in \mathcal{X}} \mu^{\pi^\phi}(x') \sum_{a \in \mathcal{A}} \frac{\partial \pi^\phi}{\partial \phi}(a | x') Q(x', a; \pi^\phi), \quad (32)$$

where $Q(x, a; \pi^\phi) = r(x, a) + \mathbb{E}[\sum_{t=1}^{\infty} \beta^t r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) | X_0^{\pi^\phi} = x]$ is the Q-function, and $\mu^{\pi^\phi}(x') = \sum_{t=0}^{\infty} \beta^t \mathbb{P}(X_t^{\pi^\phi} = x' | X_0^{\pi^\phi} = x)$ is the (discounted) occupation time.

Define $\ell(x') = \sum_{a \in \mathcal{A}} \frac{\partial \pi^\phi}{\partial \phi}(a | x') Q(x', a; \pi^\phi)$, which is a deterministic function of x' . Since $\sum_{a \in \mathcal{A}} \frac{\partial \pi^\phi}{\partial \phi}(a | x') = \frac{\partial}{\partial \phi} \sum_{a \in \mathcal{A}} \pi^\phi(a | x') = 0$, $\ell(x')$ can be equivalently written as

$$\begin{aligned} \ell(x') &= \sum_{a \in \mathcal{A}} \frac{\partial \pi^\phi}{\partial \phi}(a | x') [Q(x', a; \pi^\phi) - B(x')] \\ &= \sum_{a \in \mathcal{A}} \frac{\partial \log \pi^\phi}{\partial \phi}(a | x') [Q(x', a; \pi^\phi) - B(x')] \pi^\phi(a | x') \\ &= \mathbb{E}_{a \sim \pi^\phi(\cdot | x')} \left[\frac{\partial \log \pi^\phi}{\partial \phi}(a | x') [Q(x', a; \pi^\phi) - B(x')] \right], \end{aligned}$$

for any function $B(\cdot)$, sometimes known as a baseline (Williams, 1992).

On the other hand,

$$\mu^{\pi^\phi}(x') = \sum_{t=0}^{\infty} \beta^t \mathbb{P}(X_t^{\pi^\phi} = x' | X_0^{\pi^\phi} = x) = \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \mathbb{1}_{\{X_t^{\pi^\phi} = x'\}} \middle| X_0^{\pi^\phi} = x \right].$$

Therefore, (32) is equivalent to

$$\begin{aligned}
 \frac{\partial J(x; \pi^\phi)}{\partial \phi} &= \sum_{x' \in \mathcal{X}} \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \mathbb{1}_{\{X_t^{\pi^\phi} = x'\}} \Big| X_0^{\pi^\phi} = x \right] \ell(x') \\
 &= \sum_{t=0}^{\infty} \beta^t \sum_{x' \in \mathcal{X}} \mathbb{E} \left[\mathbb{1}_{\{X_t^{\pi^\phi} = x'\}} \ell(x') \Big| X_0^{\pi^\phi} = x \right] \\
 &= \sum_{t=0}^{\infty} \beta^t \mathbb{E} \left[\sum_{x' \in \mathcal{X}} \mathbb{1}_{\{X_t^{\pi^\phi} = x'\}} \ell(X_t^{\pi^\phi}) \Big| X_0^{\pi^\phi} = x \right] \\
 &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \ell(X_t^{\pi^\phi}) \Big| X_0^{\pi^\phi} = x \right] \\
 &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \mathbb{E}_{a \sim \pi^\phi(\cdot | X_t^{\pi^\phi})} \left[\frac{\partial \log \pi^\phi}{\partial \phi}(a | X_t^{\pi^\phi}) \left[Q(X_t^{\pi^\phi}, a; \pi^\phi) - B(X_t^{\pi^\phi}) \right] \right] \Big| X_0^{\pi^\phi} = x \right] \\
 &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \frac{\partial \log \pi^\phi}{\partial \phi}(a_t^{\pi^\phi} | X_t^{\pi^\phi}) \left[Q(X_t^{\pi^\phi}, a_t^{\pi^\phi}; \pi^\phi) - B(X_t^{\pi^\phi}) \right] \Big| X_0^{\pi^\phi} = x \right] \\
 &= \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \frac{\partial \log \pi^\phi}{\partial \phi}(a_t^{\pi^\phi} | X_t^{\pi^\phi}) \left[Q(X_t^{\pi^\phi}, a_t^{\pi^\phi}; \pi^\phi) - J(X_t^{\pi^\phi}; \pi^\phi) \right] \Big| X_0^{\pi^\phi} = x \right].
 \end{aligned} \tag{33}$$

If we choose the baseline to be the value function $B(\cdot) = J(\cdot; \pi^\phi)$, then (33) gives the representation of policy gradient in the advantage actor–critic approach (Mnih et al., 2016).

If we are to extend the above derivation to the continuous-time setting, then an essential question is what the Q-function should be in continuous time. This question has been extensively studied in a recent paper Jia and Zhou (2022b) from which we realize that (using the notations in this paper with the discount factor $e^{-\beta t}$)

$$\begin{aligned}
 &Q(X_t^{\pi^\phi}, a_t^{\pi^\phi}; \pi^\phi) - J(X_t^{\pi^\phi}; \pi^\phi) \\
 &\approx \mathcal{L}^{a_t^{\pi^\phi}} J(X_t^{\pi^\phi}; \pi^\phi) \Delta t + r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) \Delta t - \beta J(X_t^{\pi^\phi}; \pi^\phi) \Delta t \\
 &\approx dJ(X_t^{\pi^\phi}; \pi^\phi) + r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) dt - \beta J(X_t^{\pi^\phi}; \pi^\phi) dt + \frac{\partial J}{\partial x}(X_t^{\pi^\phi}; \pi^\phi)^\top \sigma(X_t^{\pi^\phi}) dW_t.
 \end{aligned}$$

Therefore, (33) becomes

$$\begin{aligned}
 \frac{\partial J(x; \pi^\phi)}{\partial \phi} &= \mathbb{E} \left[\int_0^\infty e^{-\beta t} \frac{\partial \log \pi^\phi}{\partial \phi}(a_t^{\pi^\phi} | X_t^{\pi^\phi}) \left[dJ(X_t^{\pi^\phi}; \pi^\phi) + r(X_t^{\pi^\phi}, a_t^{\pi^\phi}) dt \right. \right. \\
 &\quad \left. \left. - \beta J(X_t^{\pi^\phi}; \pi^\phi) dt + \frac{\partial J}{\partial x}(X_t^{\pi^\phi}; \pi^\phi)^\top \sigma(X_t^{\pi^\phi}) dW_t \right] \Big| X_0^{\pi^\phi} = x \right].
 \end{aligned}$$

Note that there is no policy regularizer in the current discussion; so the above coincides with the expression of the policy gradient (14) with $\gamma = 0$.

Appendix B. Theoretical Results Employed in Simulation Experiments

For the reader's convenience, we summarize the theoretical results employed in the two simulation studies in Section 5. Their proofs are similar to those of the analogous results in Wang and Zhou (2020) and Wang et al. (2020) respectively.

Appendix B1. Mean-Variance Portfolio Selection

Let the true model be given as (23) and one aims to solve (25). An omniscient agent's optimal policy is a deterministic one, given by $a_t^* = -\frac{\mu-r}{\sigma^2}(x_t-w^*)$, where $w^* = \frac{z \exp\{\frac{(\mu-r)^2}{\sigma^2}T\}-x_0}{\exp\{\frac{(\mu-r)^2}{\sigma^2}T\}-1}$.

Given this policy, the discounted wealth process (24) becomes

$$dx_t^* = -\frac{(\mu-r)^2}{\sigma^2}(x_t^*-w^*)dt - \frac{\mu-r}{\sigma}(x_t^*-w^*)dW_t.$$

Hence $x_t^* - w$ is a geometric Brownian motion. We can compute $\mathbb{E}[x_T^*] = z$, and

$$\text{Var}(x_T^*) = (x_0 - w^*)^2 \exp\left\{-2\frac{(\mu-r)^2}{\sigma^2}T\right\} \left(\exp\left\{\frac{(\mu-r)^2}{\sigma^2}T\right\} - 1\right) = \frac{(x_0 - z)^2}{\left(\exp\left\{\frac{(\mu-r)^2}{\sigma^2}T\right\} - 1\right)^2}.$$

If the agent knows the true model but is forced to take a stochastic policy subject to the entropy regularizer, then the optimal policy is

$$\pi^*(\cdot|t, x) \sim \mathcal{N}\left(-\frac{\mu-r}{\sigma^2}(x-w^*), \frac{\gamma}{2\sigma^2} \exp\left\{\frac{(\mu-r)^2}{\sigma^2}(T-t)\right\}\right).$$

Under this policy, the dynamics of $\tilde{X}_t^{\pi^*}$ is

$$d\tilde{X}_t^{\pi^*} = -\frac{(\mu-r)^2}{\sigma^2}(\tilde{X}_t^{\pi^*}-w^*)dt + \sqrt{\frac{(\mu-r)^2}{\sigma^2}(\tilde{X}_t^{\pi^*}-w^*)^2 + \frac{\gamma}{2} \exp\left\{\frac{(\mu-r)^2}{\sigma^2}(T-t)\right\}}dW_t.$$

With the same value of w^* , one can show that $\mathbb{E}^{\mathbb{P}^W}[\tilde{X}_T^{\pi^*}] = z$, and

$$\text{Var}^{\mathbb{P}^W}(\tilde{X}_T^{\pi^*}) = \exp\left\{-\frac{(\mu-r)^2}{\sigma^2}T\right\}(x_0-w^*)^2 + \frac{\gamma}{2}T - (z-w^*)^2 = \frac{(x_0-z)^2}{\left(\exp\left\{\frac{(\mu-r)^2}{\sigma^2}T\right\}-1\right)^2} + \frac{\gamma}{2}T.$$

Appendix B2. Ergodic Linear-Quadratic Control

Two benchmarks, “omniscient optimal level” and “omniscient optimal level less exploration cost”, are used in Section 5.2 for comparison. We introduce their formal definitions here.

Definition 8 *The omniscient optimal level is the maximum value of (28) subject to (27) when all the model coefficients $(A, B, C, D, M, R, N, P, Q)$ are known to the agent. The omniscient optimal level less exploration cost is defined as*

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \int_{\mathbb{R}} r(\tilde{X}_t^{\pi^*}, a) \pi(a|\tilde{X}_t^{\pi^*}) da dt \Big| \tilde{X}_0^{\pi^*} = x_0 \right],$$

where π^* is the optimal policy to (29) with the entropy regularizer and subject to (30) when all the model coefficients are known to the agent.

We compute these two values using the Hamilton-Jacobi-Bellman (HJB) equation approach.

Let the true model be given by (27) and one aims to maximize the long-term average reward (28). Consider the associated HJB equation:

$$\begin{aligned} 0 &= \sup_a [\mathcal{L}^a \varphi(x) + r(x, a) - V] \\ &= \sup_a \left[(Ax + Ba)\varphi'(x) + \frac{1}{2}(Cx + Da)^2\varphi''(x) - \left(\frac{M}{2}x^2 + Rxa + \frac{N}{2}a^2 + Px + Qa \right) - V \right]. \end{aligned}$$

Conjecturing $\varphi(x) = \frac{1}{2}k_2x^2 + k_1x$ and plugging it to the HJB equation, we get the first-order condition $a^* = \frac{[k_2(B+CD)-R]x+k_1B-Q}{N-k_2D^2}$, assuming $N - k_2D^2 > 0$. The HJB equation now becomes

$$0 = \frac{1}{2}[k_2(2A + C^2) - M]x^2 + (k_1A - P)x - V + \frac{1}{2} \frac{\{[k_2(B + CD) - R]x + k_1B - Q\}^2}{N - k_2D^2}.$$

This leads to three algebraic equations by matching the coefficients of x^2 , x and the constant term:

$$\begin{cases} k_2(2A + C^2) - M + \frac{[k_2(B + CD) - R]^2}{N - k_2D^2} = 0, \\ k_1A - P + \frac{[k_2(B + CD) - R](k_1B - Q)}{N - k_2D^2} = 0, \\ V = \frac{(k_1B - Q)^2}{2(N - k_2D^2)}. \end{cases} \quad (34)$$

Note that (34) coincides with the system of equations in footnote 12 and Theorem 9 in Wang et al. (2020) when the discount factor is 0.

Solving these algebraic equation gives the omniscient optimal reward and the corresponding optimal policy $a^* = \frac{[k_2(B+CD)-R]x+k_1B-Q}{N-k_2D^2}$.

If the agent knows the true model but still adopts stochastic policies with a entropy regularizer, then the optimal policy is given by

$$\pi^*(\cdot|x) \sim \mathcal{N}\left(\frac{[k_2(B + CD) - R]x + k_1B - Q}{N - k_2D^2}, \frac{\gamma}{N - k_2D^2}\right),$$

where k_2, k_1 are determined by (34). This optimal solution is identical to that in Wang et al. (2020, Theorem 4) when the discount factor is 0.

Under this stochastic policy, the state dynamics become

$$\begin{aligned} d\tilde{X}_t^{\pi^*} &= \left(A\tilde{X}_t^{\pi^*} + B \frac{[k_2(B + CD) - R]\tilde{X}_t^{\pi^*} + k_1B - Q}{N - k_2D^2} \right) dt \\ &+ \left(C^2\tilde{X}_t^{\pi^{*2}} + CD\tilde{X}_t^{\pi^*} \frac{[k_2(B + CD) - R]\tilde{X}_t^{\pi^*} + k_1B - Q}{N - k_2D^2} \right. \\ &\left. + D^2 \left[\left(\frac{[k_2(B + CD) - R]\tilde{X}_t^{\pi^*} + k_1B - Q}{N - k_2D^2} \right)^2 + \frac{\gamma}{N - k_2D^2} \right] \right)^{1/2} dW_t. \end{aligned}$$

To calculate the long-term average value $\tilde{V} = \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \int_{\mathbb{R}} r(\tilde{X}_t^{\pi^*}, a) \pi^*(a | \tilde{X}_t^{\pi^*}) da dt \right]$, consider the corresponding HJB equation

$$\int_{\mathbb{R}} [(Ax + Ba)\tilde{\varphi}'(x) + \frac{1}{2}(Cx + Da)^2 \tilde{\varphi}''(x) - (\frac{M}{2}x^2 + Rxa + \frac{N}{2}a^2 + Px + Qa)] \pi^*(a|x) da - \tilde{V} = 0.$$

Starting with an ansatz $\tilde{\varphi}(x) = \frac{1}{2}\tilde{k}_2x^2 + \tilde{k}_1x$ and going through the same calculations as above we obtain three equations

$$\begin{aligned} & (A + B \frac{k_2(B + CD) - R}{N - k_2D^2})\tilde{k}_2 + \frac{1}{2} \left(C^2 + 2CD \frac{k_2(B + CD) - R}{N - k_2D^2} + D^2 \frac{[k_2(B + CD) - R]^2}{(N - k_2D^2)^2} \right) \tilde{k}_2 \\ & - \frac{M}{2} - R \frac{k_2(B + CD) - R}{N - k_2D^2} - \frac{N}{2} \frac{[k_2(B + CD) - R]^2}{(N - k_2D^2)^2} = 0, \\ & (A + B \frac{k_2(B + CD) - R}{N - k_2D^2})\tilde{k}_1 + B \frac{k_1B - Q}{N - k_2D^2} \tilde{k}_2 \\ & + \frac{1}{2} \left(2CD \frac{k_1B - Q}{N - k_2D^2} + 2D^2 \frac{[k_2(B + CD) - R](k_1B - Q)}{(N - k_2D^2)^2} \right) \tilde{k}_2 \\ & - R \frac{k_1B - Q}{N - k_2D^2} - N \frac{[k_2(B + CD) - R](k_1B - Q)}{(N - k_2D^2)^2} - P - \frac{Q[k_2(B + CD) - R]}{N - k_2D^2} = 0, \\ & \tilde{k}_1B \frac{k_1B - Q}{N - k_2D^2} + \frac{1}{2} \left(D^2 \frac{(k_1B - Q)^2}{(N - k_2D^2)^2} + D^2 \frac{\gamma}{N - k_2D^2} \right) \tilde{k}_2 \\ & - \frac{N}{2} \frac{(k_1B - Q)^2}{(N - k_2D^2)^2} - \frac{N}{2} \frac{\gamma}{N - k_2D^2} - Q \frac{k_1B - Q}{N - k_2D^2} = \tilde{V}. \end{aligned}$$

The solutions to the above equations are $\tilde{k}_2 = k_2$, $\tilde{k}_1 = k_1$. Hence

$$\begin{aligned} \tilde{V} &= k_1B \frac{k_1B - Q}{N - k_2D^2} + \frac{k_2D^2}{2} \left(\frac{(k_1B - Q)^2}{(N - k_2D^2)^2} + \frac{\gamma}{N - k_2D^2} \right) - \frac{N}{2} \left(\frac{(k_1B - Q)^2}{(N - k_2D^2)^2} + \frac{\gamma}{N - k_2D^2} \right) \\ & - Q \frac{k_1B - Q}{N - k_2D^2} \\ &= \frac{(k_1B - Q)^2}{2(N - k_2D^2)} - \frac{\gamma}{2} = V - \frac{\gamma}{2}. \end{aligned}$$

By definition, \tilde{V} is also the omniscient optimal level less exploration cost. The difference, $\frac{\gamma}{2}$, between \tilde{V} and V is hence the exploration cost due to randomization. Note that a parallel result when there is a discount factor is Theorem 10 in Wang et al. (2020).

Appendix C. Proofs of Statements

In all the proofs we use generic notations C_1, C_2, \dots to denote constants that are independent of other variables involved such as t, x, a . A same such notation may show up in different places but does not necessarily have the same value.

Proof of Lemma 2

We start by examining $\tilde{b}(t, x, \boldsymbol{\pi}(\cdot|t, x))$. Note that

$$\begin{aligned} & \tilde{b}(t, x, \boldsymbol{\pi}(\cdot|t, x)) - \tilde{b}(t, x', \boldsymbol{\pi}(\cdot|t, x')) \\ &= \int_{\mathcal{A}} [b(t, x, a) - b(t, 0, a)][\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')] da + \int_{\mathcal{A}} b(t, 0, a)[\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')] da \\ & \quad + \int_{\mathcal{A}} [b(t, x, a) - b(t, x', a)]\boldsymbol{\pi}(a|t, x') da. \end{aligned}$$

Hence

$$\begin{aligned} & \left| \tilde{b}(t, x, \boldsymbol{\pi}(\cdot|t, x)) - \tilde{b}(t, x', \boldsymbol{\pi}(\cdot|t, x')) \right| \\ & \leq \int_{\mathcal{A}} |b(t, x, a) - b(t, 0, a)| |\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')| da + \int_{\mathcal{A}} |b(t, 0, a)| |\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')| da \\ & \quad + \int_{\mathcal{A}} |b(t, x, a) - b(t, x', a)| \boldsymbol{\pi}(a|t, x') da \\ & \leq (C_1|x| + C_2) \int_{\mathcal{A}} |\boldsymbol{\pi}(a|t, x) - \boldsymbol{\pi}(a|t, x')| da + C_1|x - x'| \\ & \leq C_1|x - x'| + (C_1|x| + C_2)C_3|x - x'|. \end{aligned}$$

Moreover, note that

$$|\tilde{b}(t, x, \boldsymbol{\pi}(\cdot|t, x))| \leq \int_{\mathcal{A}} |b(t, x, a)| \boldsymbol{\pi}(a|t, x) da \leq \int_{\mathcal{A}} (C_1|x| + C_2) \boldsymbol{\pi}(a|t, x) da = C_1|x| + C_2.$$

Similarly, we can show that $\tilde{\sigma}(t, x, \boldsymbol{\pi}(\cdot|t, x))$ is locally Lipschitz continuous and has linear growth in x . The unique existence of the strong solution to (5) then follows from the standard SDE theory.

Next, the SDE (5) yields

$$\tilde{X}_s^\pi = x + \int_t^s \tilde{b}(\tau, \tilde{X}_\tau^\pi, \boldsymbol{\pi}(\cdot|\tau, \tilde{X}_\tau^\pi)) d\tau + \int_t^s \tilde{\sigma}(\tau, \tilde{X}_\tau^\pi, \boldsymbol{\pi}(\cdot|\tau, \tilde{X}_\tau^\pi)) dW_\tau.$$

Based on the proved growth condition on $\tilde{b}, \tilde{\sigma}$, Cauchy–Schwarz inequality, and Burkholder–Davis–Gundy inequalities, we obtain

$$\begin{aligned} & \mathbb{E} \left[\max_{t \leq s \leq T'} |\tilde{X}_s^\pi|^\mu \middle| \tilde{X}_t^\pi = x \right] \\ & \leq C_1 \mathbb{E} \left[|x|^\mu + \max_{t \leq s \leq T'} \left| \int_t^s \tilde{b}(\tau, \tilde{X}_\tau^\pi, \boldsymbol{\pi}(\cdot|\tau, \tilde{X}_\tau^\pi)) d\tau \right|^\mu \right. \\ & \quad \left. + \max_{t \leq s \leq T'} \left| \int_t^s \tilde{\sigma}(\tau, \tilde{X}_\tau^\pi, \boldsymbol{\pi}(\cdot|\tau, \tilde{X}_\tau^\pi)) dW_\tau \right|^\mu \middle| \tilde{X}_t^\pi = x \right] \\ & \leq C_1 \mathbb{E} \left[|x|^\mu + C_2 \int_t^{T'} (1 + \max_{t \leq s \leq \tau} |\tilde{X}_s^\pi|)^\mu d\tau + C_3 \left(\int_t^{T'} (1 + \max_{t \leq s \leq \tau} |\tilde{X}_s^\pi|)^2 d\tau \right)^{\mu/2} \middle| \tilde{X}_t^\pi = x \right] \\ & \leq C_4(1 + |x|^\mu) + C_5 \int_t^{T'} \mathbb{E} \left[\max_{t \leq s \leq \tau} |\tilde{X}_s^\pi|^\mu \middle| \tilde{X}_t^\pi = x \right] d\tau. \end{aligned}$$

Applying Gronwall's inequality to $\mathbb{E} \left[\max_{t \leq s \leq T'} |\tilde{X}_s^\pi|^\mu \middle| \tilde{X}_t^\pi = x \right]$ as a function of T' , we obtain the second desired result of the lemma. The final result is evident.

Proof of Lemma 3

Set $\tilde{v}(t, x) = e^{-\beta t} v(t, x)$. Then $\tilde{v}(T, x) = e^{-\beta T} h(x)$, and (8) implies

$$\int_{\mathcal{A}} [\mathcal{L}^a \tilde{v}(t, x) + e^{-\beta t} r(t, x, a) + \gamma e^{-\beta t} p(t, x, a, \boldsymbol{\pi}(\cdot|t, x))] \boldsymbol{\pi}(a|t, x) da = 0. \quad (35)$$

Similarly, consider $\tilde{J}(t, x; \boldsymbol{\pi}) = e^{-\beta t} J(t, x; \boldsymbol{\pi})$. Then (6) yields

$$\begin{aligned} \tilde{J}(t, x; \boldsymbol{\pi}) &= \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} [e^{-\beta s} r(s, \tilde{X}_s^\pi, a) + \gamma e^{-\beta s} p(s, \tilde{X}_s^\pi, a, \boldsymbol{\pi}(\cdot|s, \tilde{X}_s^\pi))] \boldsymbol{\pi}(a|s, \tilde{X}_s^\pi) da ds \right. \\ &\quad \left. + e^{-\beta T} h(\tilde{X}_T^\pi) \middle| \tilde{X}_t^\pi = x \right]. \end{aligned}$$

So it suffices to prove the (viscosity) solution to (35), \tilde{v} , coincides with $\tilde{J}(\cdot, \cdot; \boldsymbol{\pi})$. The proof now follows from applying Beck et al. (2021, Corollary 3.3) to the SDE (5): under Assumption 1 along with Definition 1, Lemma 2 verifies the sufficient conditions in Beck et al. (2021).

Proof of Theorem 4

Using the same discounting transformation as in the proof of Lemma 3, the first statement of Theorem 4 follows directly from Jia and Zhou (2022a, Proposition 1) along with the Markov property of the solution to the SDE (5).

For the second statement, according to Jia and Zhou (2022a, Proposition 4), we have the following martingale orthogonality condition for \tilde{X}^π :

$$\begin{aligned} \mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \xi_t [dJ(t, \tilde{X}_t^\pi; \boldsymbol{\pi}) - \beta J(t, \tilde{X}_t^\pi; \boldsymbol{\pi}) dt \right. \\ \left. + \int_{\mathcal{A}} [r(t, \tilde{X}_t^\pi, a) + \gamma p(t, \tilde{X}_t^\pi, a, \boldsymbol{\pi}(\cdot|t, \tilde{X}_t^\pi))] \boldsymbol{\pi}(a|t, \tilde{X}_t^\pi) da dt \middle| \tilde{X}_0^\pi = x \right] = 0 \end{aligned}$$

for all $\xi \in L^2_{\mathcal{F}\tilde{X}^\pi}([0, T]; J(\cdot, \tilde{X}^\pi; \boldsymbol{\pi}))$. Now, any $\xi \in L^2_{\mathcal{F}\tilde{X}^\pi}([0, T]; J(\cdot, \tilde{X}^\pi; \boldsymbol{\pi}))$ corresponds to a measurable functional $\boldsymbol{\xi} : [0, T] \times C([0, T]; \mathbb{R}^d) \mapsto \mathbb{R}$ such that $\xi_t = \boldsymbol{\xi}(t, \tilde{X}_{t\wedge\cdot}^\pi)$. However, \tilde{X}^π and X^π have the same distribution and $a_t^\pi \sim \boldsymbol{\pi}(\cdot|t, \tilde{X}_t^\pi)$; hence

$$\begin{aligned} &\mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \boldsymbol{\xi}(t, \tilde{X}_{t\wedge\cdot}^\pi) [dJ(t, \tilde{X}_t^\pi; \boldsymbol{\pi}) - \beta J(t, \tilde{X}_t^\pi; \boldsymbol{\pi}) dt] \middle| \tilde{X}_0^\pi = x \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\int_0^T \boldsymbol{\xi}(t, X_{t\wedge\cdot}^\pi) [dJ(t, X_t^\pi; \boldsymbol{\pi}) - \beta J(t, X_t^\pi; \boldsymbol{\pi}) dt] \middle| X_0^\pi = x \right], \end{aligned}$$

and

$$\begin{aligned} &\mathbb{E}^{\mathbb{P}^W} \left[\int_0^T \boldsymbol{\xi}(t, \tilde{X}_{t\wedge\cdot}^\pi) \int_{\mathcal{A}} [r(t, \tilde{X}_t^\pi, a) + \gamma p(t, \tilde{X}_t^\pi, a, \boldsymbol{\pi}(\cdot|t, \tilde{X}_t^\pi))] \boldsymbol{\pi}(a|t, \tilde{X}_t^\pi) da dt \middle| \tilde{X}_0^\pi = x \right] \\ &= \mathbb{E}^{\mathbb{P}} \left[\int_0^T \boldsymbol{\xi}(t, X_{t\wedge\cdot}^\pi) [r(t, X_t^\pi, a_t^\pi) + \gamma p(t, X_t^\pi, a_t^\pi, \boldsymbol{\pi}(\cdot|t, X_t^\pi))] dt \middle| X_0^\pi = x \right]. \end{aligned}$$

Combining the above two equations leads to (9).

Proof of Theorem 5

It suffices to prove (12) equals (14).

Fix t . Define a sequence of stopping times $\tau_n = \inf\{s \geq t : |X_s^{\pi^\phi}| \geq n\}$. Applying Itô's lemma to $J(s, X_s^{\pi^\phi})$, we obtain:

$$\begin{aligned}
 & \int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] \times [dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \\
 & \quad + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \\
 & \quad \left. + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \right\} \\
 = & \int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] \times \{ [\mathcal{L}^{a_s} J(s, X_s^{\pi^\phi}; \pi^\phi) + r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) \right. \\
 & \quad + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \\
 & \quad \left. + \frac{\partial J}{\partial x}(s, X_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) dW_s \right\} + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \Big\}.
 \end{aligned}$$

Taking expectation yields

$$\begin{aligned}
 & \mathbb{E}^\mathbb{P} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] \times [dJ(s, X_s^{\pi^\phi}; \pi^\phi) \right. \right. \\
 & \quad + [r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \\
 & \quad \left. \left. + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \right\} \Big| X_t^{\pi^\phi} = x \right] \\
 = & \mathbb{E}^\mathbb{P} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] \times \{ [\mathcal{L}^{a_s} J(s, X_s^{\pi^\phi}; \pi^\phi) + r(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) \right. \right. \\
 & \quad + \gamma p(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \pi^\phi(\cdot | s, X_s^{\pi^\phi})) - \beta J(s, X_s^{\pi^\phi}; \pi^\phi)] ds \} + \gamma q(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}, \phi) ds \Big\} \Big| X_t^{\pi^\phi} = x \right] \\
 & + \mathbb{E}^\mathbb{P} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \frac{\partial J}{\partial x}(s, X_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, X_s^{\pi^\phi}, a_s) dW_s \Big| X_t^{\pi^\phi} = x \right] \\
 = & \mathbb{E}^\mathbb{P} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \check{r}(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}; \phi) ds \Big| X_t^{\pi^\phi} = x \right].
 \end{aligned}$$

The second term above vanishes because when $t \leq s \leq T \wedge \tau_n$, it follows from Assumptions 1 that

$$\begin{aligned}
 & \left| \int_{\mathcal{A}} \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) \right] \times \frac{\partial J}{\partial x}(s, \tilde{X}_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, \tilde{X}_s^{\pi^\phi}, a) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da \right|^2 \\
 & \leq \int_{\mathcal{A}} \left| \frac{\partial}{\partial \phi} \log \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) \frac{\partial J}{\partial x}(s, \tilde{X}_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, \tilde{X}_s^{\pi^\phi}, a) \right|^2 \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da \\
 & \leq C \max_{|x| \leq n, 0 \leq t \leq T} \left| \frac{\partial}{\partial x} J(t, x; \pi^\phi) \right| \int_{\mathcal{A}} \left| \frac{\partial}{\partial \phi} \log \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) \right|^2 (1 + |\tilde{X}_s^{\pi^\phi}|)^2 \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da \\
 & \leq C(1+n)^2 \max_{|x| \leq n, 0 \leq t \leq T} \left| \frac{\partial}{\partial x} J(t, x; \pi^\phi) \right| \int_{\mathcal{A}} \left| \frac{\partial}{\partial \phi} \log \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) \right|^2 \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da,
 \end{aligned}$$

which is bounded by a function of n due to Assumption 3. Thus,

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a_s^{\pi^\phi} | s, X_s^{\pi^\phi}) \right] \frac{\partial J}{\partial x}(s, X_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}) dW_s \middle| X_t^{\pi^\phi} = x \right] \\
 & = \mathbb{E}^{\mathbb{P}^W} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \int_{\mathcal{A}} \left[\frac{\partial}{\partial \phi} \log \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) \right] \right. \\
 & \quad \times \left. \frac{\partial J}{\partial x}(s, \tilde{X}_s^{\pi^\phi}; \pi^\phi)^\top \sigma(s, \tilde{X}_s^{\pi^\phi}, a) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da dW_s \middle| \tilde{X}_t^{\pi^\phi} = x \right] = 0.
 \end{aligned}$$

Lastly, note

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \check{r}(s, X_s^{\pi^\phi}, a_s^{\pi^\phi}; \phi) ds \middle| X_t^{\pi^\phi} = x \right] \\
 & = \mathbb{E}^{\mathbb{P}^W} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \int_{\mathcal{A}} \check{r}(s, \tilde{X}_s^{\pi^\phi}, a; \phi) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da ds \middle| \tilde{X}_t^{\pi^\phi} = x \right].
 \end{aligned}$$

By Assumption 3 and Lemma 2, we get

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}^W} \left[\int_{\mathcal{A}} |\check{r}(s, \tilde{X}_s^{\pi^\phi}, a; \phi)| \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da \middle| \tilde{X}_t^{\pi^\phi} = x \right] \\
 & \leq C_1 \mathbb{E}^{\mathbb{P}^W} \left[1 + |\tilde{X}_t^{\pi^\phi}|^\mu \middle| \tilde{X}_t^{\pi^\phi} = x \right] \leq C_2(1 + |x|^\mu).
 \end{aligned}$$

Hence by the dominance convergence theorem, we conclude that as $n \rightarrow \infty$,

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}^W} \left[\int_t^{T \wedge \tau_n} e^{-\beta(s-t)} \int_{\mathcal{A}} \check{r}(s, \tilde{X}_s^{\pi^\phi}, a; \phi) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da ds \middle| \tilde{X}_t^{\pi^\phi} = x \right] \\
 & \rightarrow \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T e^{-\beta(s-t)} \int_{\mathcal{A}} \check{r}(s, \tilde{X}_s^{\pi^\phi}, a; \phi) \pi^\phi(a|s, \tilde{X}_s^{\pi^\phi}) da ds \middle| \tilde{X}_t^{\pi^\phi} = x \right].
 \end{aligned}$$

This proves the desired result.

Proof of Theorem 6

The proof is similar to the proof of Theorem 4 and that in Jia and Zhou (2022a, Proposition 4) by noticing π^* satisfies (17).

First, it suffices to consider the case when $\zeta = 0$ because of Theorem 4. For $\eta \in L^2_{\mathcal{F}X\pi^{\phi^*}}([0, T], J(\cdot, X_s^{\pi^{\phi^*}}; \pi^{\phi^*}))$, we write $\eta_s = \boldsymbol{\eta}(s, X_s^{\pi^{\phi^*}})$. Then the right hand side of (18) can be written as

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}} \left[\int_0^T \boldsymbol{\eta}(s, X_s^{\pi^{\phi^*}}) \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^{\phi^*}(a_s^{\pi^{\phi^*}} | s, X_s^{\pi^{\phi^*}}) \right] [dJ(s, X_s^{\pi}; \pi^{\phi^*}) \right. \right. \\
 & \quad \left. \left. + [r(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}) + \gamma p(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \pi^{\phi^*}(\cdot | s, X_s^{\pi^{\phi^*}})) - \beta J(s, X_s^{\pi^{\phi^*}}; \pi^{\phi^*})] ds \right. \right. \\
 & \quad \left. \left. + \gamma q(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \phi^*) ds \right\} \middle| X_0^{\pi^{\phi^*}} = x \right] \\
 &= \int_0^T \mathbb{E}^{\mathbb{P}} \left[\boldsymbol{\eta}(s, X_s^{\pi^{\phi^*}}) \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^{\phi^*}(a_s^{\pi^{\phi^*}} | s, X_s^{\pi^{\phi^*}}) \right] [\mathcal{L}^{a_s^{\pi^{\phi^*}}} J(s, X_s^{\pi}; \pi^{\phi^*}) \right. \right. \\
 & \quad \left. \left. + r(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}) + \gamma p(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \pi^{\phi^*}(\cdot | s, X_s^{\pi^{\phi^*}})) - \beta J(s, X_s^{\pi^{\phi^*}}; \pi^{\phi^*}) \right. \right. \\
 & \quad \left. \left. + \gamma q(s, X_s^{\pi^{\phi^*}}, a_s^{\pi^{\phi^*}}, \phi^*) \right\} \middle| X_0^{\pi^{\phi^*}} = x \right] ds \\
 &= \int_0^T \mathbb{E}^{\mathbb{P}^W} \left[\boldsymbol{\eta}(s, \tilde{X}_s^{\pi^{\phi^*}}) \int_{\mathcal{A}} \left\{ \left[\frac{\partial}{\partial \phi} \log \pi^{\phi^*}(a | s, \tilde{X}_s^{\pi^{\phi^*}}) \right] [\mathcal{L}^a J(s, \tilde{X}_s^{\pi}; \pi^{\phi^*}) \right. \right. \\
 & \quad \left. \left. + r(s, \tilde{X}_s^{\pi^{\phi^*}}, a) + \gamma p(s, \tilde{X}_s^{\pi^{\phi^*}}, a, \pi^{\phi^*}(\cdot | s, \tilde{X}_s^{\pi^{\phi^*}})) - \beta J(s, \tilde{X}_s^{\pi^{\phi^*}}; \pi^{\phi^*}) \right. \right. \\
 & \quad \left. \left. + \gamma q(s, \tilde{X}_s^{\pi^{\phi^*}}, a, \phi^*) \right\} \pi^{\phi^*}(a | s, \tilde{X}_s^{\pi^{\phi^*}}) da \middle| \tilde{X}_0^{\pi^{\phi^*}} = x \right] ds = 0,
 \end{aligned}$$

where the last equality follows from (17).

Proof of Lemma 7

Apply Itô's lemma to $J(\tilde{X}_s^{\pi}; \pi)$ on $s \in [t, T]$ to obtain

$$\begin{aligned}
 & \mathbb{E}^{\mathbb{P}^W} \left[J(\tilde{X}_T^{\pi}; \pi) \middle| \tilde{X}_t^{\pi} = x \right] - J(x; \pi) \\
 &= \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} \mathcal{L}^a J(\tilde{X}_s^{\pi}; \pi) \pi(a | \tilde{X}_s^{\pi}) da ds \middle| \tilde{X}_t^{\pi} = x \right] \\
 &= \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} -r(\tilde{X}_s^{\pi}, a) - \gamma p(\tilde{X}_s^{\pi}, a, \pi(\cdot | \tilde{X}_s^{\pi})) \pi(a | \tilde{X}_s^{\pi}) da ds \middle| \tilde{X}_t^{\pi} = x \right] + V(\pi)(T - t).
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 & \frac{1}{T} \mathbb{E}^{\mathbb{P}} \left[\int_t^T [r(X_s^{\pi}, a_s^{\pi}) + \gamma p(X_s^{\pi}, a_s^{\pi}, \pi(\cdot | X_s^{\pi}))] ds \middle| X_t^{\pi} = x \right] \\
 &= \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[\int_t^T \int_{\mathcal{A}} [r(\tilde{X}_s^{\pi}, a) + \gamma p(\tilde{X}_s^{\pi}, a, \pi(\cdot | \tilde{X}_s^{\pi}))] \pi(a | \tilde{X}_s^{\pi}) da ds \middle| \tilde{X}_t^{\pi} = x \right] \\
 &= V(\pi) \frac{T-t}{T} + \frac{1}{T} J(x; \pi) - \frac{1}{T} \mathbb{E}^{\mathbb{P}^W} \left[J(\tilde{X}_T^{\pi}; \pi) \middle| \tilde{X}_t^{\pi} = x \right].
 \end{aligned}$$

By a similar localization argument as in the proof of Theorem 5, we can show that $\limsup_{T \rightarrow \infty} \mathbb{E}^{\mathbb{P}^W} \left[J(\tilde{X}_T^\pi; \pi) \mid \tilde{X}_t^\pi = x \right]$ is finite and independent of x . Taking limit $T \rightarrow \infty$ on both sides of the above yields (20).

The above analysis also implies

$$J(\tilde{X}_t^\pi; \pi) + \int_0^t \int_{\mathcal{A}} [r(\tilde{X}_s^\pi, a) + \gamma p(\tilde{X}_s^\pi, a, \pi(\cdot | \tilde{X}_s^\pi)) - V(\pi)] \pi(a | \tilde{X}_s^\pi) da ds$$

is an $(\mathcal{F}^{\tilde{X}^\pi}, \mathbb{P}^W)$ -martingale. For the same reason as in the proof of Theorem 4, we arrive at the second desired conclusion of the lemma.

References

- V. Aleksandrov, V. Sysoev, and V. Shemenева. Stochastic optimization. *Engineering Cybernetics*, 5:11–16, 1968.
- R. Aragon-Gómez and J. B. Clempner. Traffic-signal control reinforcement learning approach for continuous-time Markov games. *Engineering Applications of Artificial Intelligence*, 89:103415, 2020.
- A. Arapostathis, V. S. Borkar, and M. K. Ghosh. *Ergodic control of diffusion processes*, volume 143. Cambridge University Press, 2012.
- L. C. Baird. Advantage updating. Technical report, Write Lab Wright-Patterson Air Force Base, OH 45433-7301, USA, 1993.
- A. G. Barto, R. S. Sutton, and C. W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, (5):834–846, 1983.
- S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *The Review of Financial Studies*, 23(8):2970–3016, 2010.
- M. Basei, X. Guo, A. Hu, and Y. Zhang. Logarithmic regret for episodic continuous-time linear-quadratic reinforcement learning over a finite-time horizon. *arXiv preprint arXiv:2006.15316*, 2020.
- C. Beck, M. Hutzenthaler, and A. Jentzen. On nonlinear Feynman–Kac formulas for viscosity solutions of semilinear parabolic partial differential equations. *Stochastics and Dynamics*, page 2150048, 2021.
- A. Bensoussan and J. Frehse. On Bellman equations of ergodic control in R^n . In *Applied Stochastic Analysis*, pages 21–29. Springer, 1992.
- S. Bhatnagar, R. S. Sutton, M. Ghavamzadeh, and M. Lee. Natural actor–critic algorithms. *Automatica*, 45(11):2471–2482, 2009.
- T. Björk, A. Murgoci, and X. Y. Zhou. Mean–variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance*, 24(1):1–24, 2014.

- V. S. Borkar and M. K. Ghosh. Ergodic control of multidimensional diffusions I: The existence results. *SIAM Journal on Control and Optimization*, 26(1):112–126, 1988.
- V. S. Borkar and M. K. Ghosh. Ergodic control of multidimensional diffusions II: Adaptive control. *Applied Mathematics and Optimization*, 21(1):191–220, 1990.
- S. J. Bradtke and A. G. Barto. Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22(1):33–57, 1996.
- M. Dai, Y. Dong, and Y. Jia. Learning equilibrium mean-variance strategy. *SSRN preprint SSRN:3770818*, 2020.
- M. Dai, H. Jin, S. Kou, and Y. Xu. A dynamic mean-variance analysis for log returns. *Management Science*, 67(2):1093–1108, 2021.
- T. Degris, P. M. Pilarski, and R. S. Sutton. Model-free reinforcement learning with continuous action in practice. In *2012 American Control Conference (ACC)*, pages 2177–2182. IEEE, 2012.
- K. Doya. Reinforcement learning in continuous time and space. *Neural Computation*, 12(1):219–245, 2000.
- Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338. PMLR, 2016.
- I. Ekeland and A. Lazrak. Being serious about non-commitment: subgame perfect equilibrium in continuous time. *arXiv preprint math/0604264*, 2006.
- W. Fedus, P. Ramachandran, R. Agarwal, Y. Bengio, H. Larochelle, M. Rowland, and W. Dabney. Revisiting fundamentals of experience replay. In *International Conference on Machine Learning*, pages 3061–3071. PMLR, 2020.
- W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*, volume 25. Springer Science & Business Media, 2006.
- N. Frémaux, H. Sprekeler, and W. Gerstner. Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Computational Biology*, 9(4): e1003024, 2013.
- X. Gao, Z. Q. Xu, and X. Y. Zhou. State-dependent temperature control for Langevin diffusions. *SIAM Journal on Control and Optimization*, 60(3):1250–1268, 2022.
- P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.
- X. Guo, R. Xu, and T. Zariwopoulou. Entropy regularization for mean field games with learning. *Mathematics of Operations Research*, 2022.

- T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Y. Jia and X. Y. Zhou. Policy evaluation and temporal-difference learning in continuous time and space: A martingale approach. *Journal of Machine Learning Research*, 23(154): 1–55, 2022a.
- Y. Jia and X. Y. Zhou. q-Learning in continuous time. *arXiv preprint; <http://arxiv.org/abs/2207.00713>*, 2022b.
- I. Karatzas and S. Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer, 2014.
- J. Kim, J. Shin, and I. Yang. Hamilton-Jacobi deep Q-Learning for deterministic continuous-time systems with Lipschitz continuous controls. *Journal of Machine Learning Research*, 22:206–1, 2021.
- B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- J. Lee and R. S. Sutton. Policy iterations for reinforcement learning problems in continuous time and space—Fundamental theory and methods. *Automatica*, 126:109421, 2021.
- D. Li and W.-L. Ng. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10(3):387–406, 2000.
- T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- H. R. Maei, C. Szepesvari, S. Bhatnagar, D. Precup, D. Silver, and R. S. Sutton. Convergent temporal-difference learning with arbitrary smooth function approximation. In *NIPS*, pages 1204–1212, 2009.
- P. Marbach and J. N. Tsitsiklis. Simulation-based optimization of Markov reward processes. *IEEE Transactions on Automatic Control*, 46(2):191–209, 2001.
- S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer Science & Business Media, 2012.
- V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937. PMLR, 2016.
- R. Munos. Policy gradient in continuous time. *Journal of Machine Learning Research*, 7: 771–791, 2006.
- R. Munos and P. Bourgin. Reinforcement learning for continuous stochastic control problems. *Advances in Neural Information Processing Systems*, 10, 1997.

- N. Sandrić. A note on the Birkhoff ergodic theorem. *Results in Mathematics*, 72(1):715–730, 2017.
- J. Schulman, X. Chen, and P. Abbeel. Equivalence between policy gradients and soft Q-learning. *arXiv preprint arXiv:1704.06440*, 2017a.
- J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017b.
- D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, pages 387–395. PMLR, 2014.
- D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- R. H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955.
- R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, 1988.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 2018.
- R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 12, 1999.
- R. S. Sutton, S. Singh, and D. McAllester. Comparing policy-gradient algorithms. *IEEE Transactions on Systems, Man, and Cybernetics*, 2000.
- R. S. Sutton, C. Szepesvári, and H. R. Maei. A convergent $o(n)$ temporal-difference algorithm for off-policy learning with linear function approximation. In *NIPS*, 2008.
- R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 993–1000, 2009.
- L. Szpruch, T. Treetanthiploet, and Y. Zhang. Exploration-exploitation trade-off for continuous-time episodic reinforcement learning with linear-convex models. *arXiv preprint arXiv:2112.10264*, 2021.
- C. Tallec, L. Blier, and Y. Ollivier. Making deep Q-learning methods robust to time discretization. In *International Conference on Machine Learning*, pages 6096–6104. PMLR, 2019.
- W. Tang, P. Y. Zhang, and X. Y. Zhou. Exploratory HJB equations and their convergence. *arXiv preprint arXiv:2109.10269*, 2021.

- K. G. Vamvoudakis and F. L. Lewis. Online actor–critic algorithm to solve the continuous-time infinite horizon optimal control problem. *Automatica*, 46(5):878–888, 2010.
- H. Wang and X. Y. Zhou. Continuous-time mean–variance portfolio selection: A reinforcement learning framework. *Mathematical Finance*, 30(4):1273–1308, 2020.
- H. Wang, T. Zariphopoulou, and X. Y. Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 21(198):1–34, 2020.
- W. Wang, J. Han, Z. Yang, and Z. Wang. Global convergence of policy gradient for linear-quadratic mean-field control/game in continuous time. In *International Conference on Machine Learning*, pages 10772–10782. PMLR, 2021.
- P. Wawrzynski. Learning to control a 6-degree-of-freedom walking robot. In *EUROCON 2007-The International Conference on "Computer as a Tool"*, pages 698–705. IEEE, 2007.
- R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.
- C. Yildiz, M. Heinonen, and H. Lähdesmäki. Continuous-time model-based reinforcement learning. In *International Conference on Machine Learning*, pages 12009–12018. PMLR, 2021.
- J. Yong and X. Y. Zhou. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. New York, NY: Springer, 1999.
- D. Zambrano, P. R. Roelfsema, and S. M. Bohte. Continuous-time on-policy neural reinforcement learning of working memory tasks. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2015.
- S. Zhang and R. S. Sutton. A deeper look at experience replay. *arXiv preprint arXiv:1712.01275*, 2017.
- T. Zhao, H. Hachiya, G. Niu, and M. Sugiyama. Analysis and improvement of policy gradient estimation. In *NIPS*, pages 262–270. Citeseer, 2011.
- X. Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42(1):19–33, 2000.