

Learning Stochastic Optimal Policies via Gradient Descent

Stefano Massaroli^{1,*}, Michael Poli^{2,*}, Stefano Peluchetti³
Jinkyoo Park², Atsushi Yamashita¹ and Hajime Asama¹

Abstract—We systematically develop a learning-based treatment of stochastic optimal control (SOC), relying on direct optimization of parametric control policies. We propose a derivation of adjoint sensitivity results for stochastic differential equations through direct application of variational calculus. Then, given an objective function for a predetermined task specifying the desiderata for the controller, we optimize their parameters via iterative gradient descent methods. In doing so, we extend the range of applicability of classical SOC techniques, often requiring strict assumptions on the functional form of system and control. We verify the performance of the proposed approach on a continuous-time, finite horizon portfolio optimization with proportional transaction costs.

I. INTRODUCTION

In this work we consider the following class of controlled stochastic dynamical systems:

$$\dot{x}_t = f(t, x_t, u_t) + g(t, x_t, u_t)\xi(t) \quad (1)$$

with state $x_t \in \mathbb{X} \subset \mathbb{R}^{n_x}$ and control policy $u_t \in \mathbb{U} \subset \mathbb{R}^{n_u}$. $\xi(t) \in \mathbb{R}^{n_\xi}$ is a stationary δ -correlated Gaussian noise, i.e. $\forall t > 0 \quad \mathbb{E}[\xi(t)] = 0$ and $\forall s, t$ such that $0 < s < t$ it holds $\mathbb{E}[\xi(s)\xi(t)] = \delta(s-t)$. The RHS of (1) comprises a *drift* term $f : \mathbb{R} \times \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}^{n_x}$ and a *diffusion* term $g : \mathbb{R} \times \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}^{n_x \times n_\xi}$. This paper develops a novel, systematic approach to learning optimal control policies for systems in the form (1), with respect to smooth scalar objective functions.

The link between *stochastic optimal control* (SOC) and learning has been explored in the discrete-time case [1] with policy iteration and value function approximation methods [2], [3] seeing widespread utilization in *reinforcement learning* [4]. Adaptive stochastic control has obtained explicit solutions through strict assumptions on the class of systems and objectives [5], preventing its applicability to the general case. Forward-backward SDEs (FBSDEs) have been proposed to solve classic SOC problems [6], even employing neural approximators for value function dynamics [7], [8]. A further connection between SOC and machine learning has also been discussed within the *continuous-depth* paradigm of neural networks [9], [10], [11], [12]. [13] e.g. showed that fully connected residual networks converge, in the infinite depth and width limit, to diffusion processes.

Here, we explore a different direction, motivated by the affinity between neural network training and optimal control which involve the reliance on carefully crafted objective functions encoding task-dependent desiderata.

The *raison d'être* of synthesizing optimal stochastic controllers through gradient-based techniques is to enrich the class of objectives for which an optimal controller can be found, as well scale the tractability to high-dimensional regimes. This is in line with the empirical results of modern machine learning research where large deep neural networks are often optimized on high-dimensional non-convex problems with outstanding performance [14], [15]. Gradient-based methods are also being explored in classic control settings to obtain a rigorous characterization of optimal control problems in the *linear-quadratic* regime [16], [17].

Notation: Let (Ω, \mathcal{F}, P) be a probability space. If a property (event) $A \in \mathcal{F}$ holds with $P(A) = 1$, we say that such property holds *almost surely*. A family of \mathbb{X} -valued random variables defined on a compact time domain $\mathbb{T} \subset \mathbb{R}$ $\{x_t\}_{t \in \mathbb{T}}$ is called *stochastic process* and is *measurable* if $x_t(A)$ is measurable with respect to the σ -algebra $\mathcal{B}(\mathbb{T}) \times \mathcal{F}$ being $\mathcal{B}(\mathbb{T})$ the Borel-algebra of \mathbb{T} . As a convention, we use $\int_s^t = -\int_t^s$ if $s < t$ and we denote with δ the Krocnecker delta function.

II. STOCHASTIC DIFFERENTIAL EQUATIONS

Although (1) “looks like a differential equation, it is really a meaningless string of symbols” [18]. This relation should be hence treated only as a *pre-equation* and cannot be studied in this form. Such ill-posedness arises from the fact that, being $\xi(t)$ a δ -autocorrelated process, the noise fluctuates an infinite number of times with infinite variance¹ in any time interval. Therefore, a rigorous treatment of the model requires a different *calculus* to interpret integration of the RHS of (1). The resulting well-defined version of (1) is known as a *Stochastic Differential Equation* (SDE) [19].

A. Itô–Stratonovich Dilemma

According to Van Kampen [18], $\xi(t)$ might be thought as a random sequence of δ functions causing, at each time t , a sudden jump of x_t . The controversies on the interpretation of (1) arises over the fact that it does not specify which value of x should be used to compute g when the δ functions are applied. There exist two main interpretations of the issue, namely Itô’s [20] and Stratonovich’s [21]. Itô prescribes that g should be computed with the value of x *before* the jump while Stratonovich uses the *mean value* of x before and after the jump. This choice leads to two different (yet, both admissible and equivalent) types of integration. Formally, consider

¹ The University of Tokyo, massaroli@robot.t.u-tokyo.ac.jp

² KAIST, polim@kaist.ac.kr

³ Cogent Labs, speluchetti@cogent.co.jp

* equal contribution authors

¹From a control theoretical perspective, if $\xi(t) \in \mathbb{R}$ is, for instance, a white noise signal, its energy $\int_{-\infty}^{\infty} |F(\xi)(\omega)| d\omega$ would not be finite ($F(\cdot)$ denotes the Fourier transform)

a compact time horizon $\mathbb{T} = [0, T]$, $T > 0$ and let $\{B_t\}_{t \in \mathbb{T}}$ be the standard n_ξ -dimensional Wiener process defined on a filtered probability space $(\Omega, \mathcal{F}, P; \{\mathcal{F}_t\}_{t \in \mathbb{T}})$ which is such that $B_0 = 0$, B_t is almost surely continuous in t , nowhere differentiable, has independent Gaussian increments, namely $\forall s, t \in \mathbb{T} \ s < t \Rightarrow B_t - B_s \sim \mathcal{N}(0, t - s)$ and, for all $t \in \mathbb{T}$, we have $\xi(t)dt = dB_t$. Moreover, let $\phi_t := g(t, x_t, u_t)$. Itô and Stratonovich integral calculi are then defined as $\int_0^T \phi_t dB_t = \lim_{|D| \rightarrow 0} \sum_{k=1}^K \phi_{t_{k-1}} (B_{t_k} - B_{t_{k-1}})$ and $\int_0^T \phi_t \circ dB_t = \lim_{|D| \rightarrow 0} \frac{1}{2} \sum_{k=1}^K (\phi_{t_k} - \phi_{t_{k-1}}) (B_{t_k} - B_{t_{k-1}})$, respectively, where D is a given partition of \mathbb{T} , $D := \{t_k : 0 = t_0 < t_1 < \dots < t_K = T\}$, $|D| = \max_k (t_k - t_{k-1})$ and the limit is intended in the mean-square sense, (see [22]) if exists. Note that the symbol “ \circ ” in $\circ dB_t$ is used only to indicate that the integral is interpreted in the Stratonovich sense and does not stand for function composition. In the Stratonovich convention, we may use the standard rules of calculus² while this is not the case for Itô’s. This is because Stratonovich calculus corresponds to the limit case of a smooth process with a small finite auto-correlation approaching B_t [23]. Therefore, there are two different interpretations of (1): $dx_t = f(t, x_t, u_t)dt + g(t, x_t, u_t)dB_t$ and $dx_t = f(t, x_t, u_t)dt + g(t, x_t, u_t) \circ dB_t$. Despite their name, SDEs are formally defined as integral equations due to the non-differentiability of the Brownian paths B_t .

From a control systems perspective, on one hand Stratonovich interpretation seems more appealing for noisy physical systems as it captures the limiting case of multiplicative noise processes with a very small finite auto-correlation, often a reasonable assumption in real applications. However, if we need to treat diffusion terms as intrinsic components of the system’s dynamics or impose martingale assumptions on the solutions of the SDE, the Itô convention should be adopted. In particular, this is always the case of financial applications, where stochastic optimal control is widely applied (see e.g. [24]).

Within the scope of this paper, we will mainly adopt the Stratonovich convention for consistency of notation. Nonetheless, all results presented in this manuscript can be equivalently derived for the Itô case. Indeed an Itô SDE can be transformed into a Stratonovich one (and viceversa) by the equivalence relation between the two calculi: for all $t \in \mathbb{T}$ $\int_0^t \phi_s \circ dB_s = \int_0^t \phi_s dB_s + \frac{1}{2} \int_0^t \sum_{i=1}^{n_\xi} \frac{\partial \phi_s^{(i)}}{\partial x} \phi_s^{(i)} ds$.

B. Stratonovich SDE

We will formally introduce Stratonovich SDEs following the treatment of Kunita [25], [22]. For $0 \leq s \leq t \leq T$ we denote by $\mathcal{F}_{s,t}$ the smallest σ -algebra containing all null sets of \mathcal{F} with respect to which, for all $v, w \in \mathbb{T}$ such that $s \leq v \leq w \leq t$, $B_w - B_v$ is measurable. $\{\mathcal{F}_{s,t}\}_{s \leq t; s, t \in \mathbb{T}}$ is called a *two-sided filtration* generated by the Wiener process and we set write $\mathcal{F}_t = \mathcal{F}_{0,t}$ for compactness. Then, $\{\mathcal{F}_t\}_{t \in \mathbb{T}}$ is itself a filtration with respect to which B_t is \mathcal{F}_t -adapted. Further, we let f, g to be bounded in \mathbb{X} , infinitely

differentiable in x , continuously differentiable in t , uniformly continuous in u and we assume the controller $\{u_t\}_{t \in \mathbb{T}}$ to be a \mathcal{F}_t -adapted process. Given an initial condition $x_0 \in \mathbb{X}$ assumed to be a \mathcal{F}_0 measurable random variable, we suppose that there exists a \mathbb{X} -valued continuous \mathcal{F}_t -adapted semi-martingale $\{x_t\}_{t \in \mathbb{T}}$ such that

$$x_T = x_0 + \int_0^T f(t, x_t, u_t)dt + \int_0^T g(t, x_t, u_t) \circ dB_t, \quad (2)$$

almost surely. Path-wise existence and uniqueness of solutions, i.e. if, for all $t \in \mathbb{T}$, two solutions x_t, x'_t such that $x_0 = x'_0$ satisfy $\forall t \in \mathbb{T} \ x_t = x'_t$ almost surely, is guaranteed under our class assumptions on f, g and the process $\{u_t\}_{t \in \mathbb{T}}$ ³.

If, as assumed here, f, g are functions of class $\mathcal{C}^{1,\infty}$ in (x, t) uniformly continuous in u with bounded derivatives w.r.t x and t , and $\{u_t\}_{t \in \mathbb{T}}$ belongs to some admissible control set \mathbb{A} of $\mathbb{T} \rightarrow \mathbb{U}$ functions, then, given a realization of the Wiener process, there exists a \mathcal{C}^∞ mapping Φ called *stochastic flow* from $\mathbb{X} \times \mathbb{A}$ to the space of absolutely continuous functions $[s, t] \rightarrow \mathbb{X}$ such that

$$x_t = \Phi_s(x_s, \{u_s\}_{s \leq t; s, t \in \mathbb{T}})(t) \quad s \leq t; s, t \in \mathbb{T}; x_s \in \mathbb{X} \quad (3)$$

almost surely. For the sake of compactness we denote the RHS of (3) with $\Phi_{s,t}(x_s)$. It is worth to be noticed that the collection $\{\Phi_{s,t}\}_{s \leq t; s, t \in \mathbb{T}}$ satisfies the *flow property* (see [22, Sec. 3.1]) $\forall s, t, v \in \mathbb{T} : s < v < t \Rightarrow \Phi_{v,t}(\Phi_{s,v}(x_s)) = \Phi_{s,t}(x_s)$ and that it is also a diffeomorphism [22, Theorem 3.7.1], i.e. there exists an *inverse flow* $\Psi_{s,t} := \Phi_{s,t}^{-1}$ which satisfies the *backward SDE*

$$\begin{aligned} \Psi_{s,t}(x_s) = x_s - & \int_s^t f(v, \Psi_{s,v}(x_s), u_v)dv \\ & - \int_s^t g(v, \Psi_{s,v}(x_s), u_v) \circ d\check{B}_v. \end{aligned}$$

being $\{\check{B}_t\}_{t \in \mathbb{T}}$ the realization of the *backward* Wiener process defined as $\check{B}_t := B_t - B_T$ for all $t \in \mathbb{T}$. The diffeomorphism property of $\Phi_{s,t}$ thus yields: $\Psi_{s,t}(\Phi_{s,t}(x_s)) = x_s$.

Therefore, it is possible to *reverse* the solutions of Stratonovich SDEs in a similar fashion to *ordinary* differential equations (ODEs) by storing/generating identical realizations of the noise process. From a practical point of view, under mild assumptions on the chosen numerical SDE scheme, approximated solutions $\hat{\Phi}_{s,t}(x_s)$ satisfy $\forall x_s \in \mathbb{X} \ \hat{\Psi}_{s,t}(\hat{\Phi}_{s,t}(x_s)) \rightarrow x_s$ in probability as the discretization time-step $\epsilon \rightarrow 0$ [26, Theorem 3.3].

III. DIRECT STOCHASTIC OPTIMAL CONTROL

We consider the optimal control problem for systems of class (1) interpreted in the Stratonovich sense. In particular, we aim at determining a control process $\{u_t\}_{t \in \mathbb{T}}$ within an admissible set of functions \mathbb{A} to minimize or maximize in expectation some scalar criterion J_u of the form

$$J_u(x_0) = \int_{\mathbb{T}} \gamma(t, \Phi_{0,t}(x_0), u_t)dt + \Gamma(\Phi_{0,T}(x_0), u_T), \quad (4)$$

²see e.g. Itô’s Lemma in Stratonovich form [22, Theorem 2.4.1], where the *chain rule* of classic Riemann–Stieltjes calculus is shown to be preserved under the sign of Stratonovich integrals.

³Note that less strict sufficient conditions only required uniform Lipschitz continuity of $f(t, x_t, u_t) + \frac{1}{2} \sum_{i=1}^{n_\xi} g^{(i)}(t, x_t, u_t)[\partial g^{(i)}(t, x_t, u_t)/\partial x]$ and $g(t, x_t, u_t)$ w.r.t. x and uniform continuity w.r.t. t and u .

measuring the performance of the system. Within the scope of the paper, we characterize \mathbb{A} as set of all control processes $\{u_r\}_{t \in \mathbb{T}}$ which are \mathcal{F}_t -adapted, have values in \mathbb{U} and such that $\mathbb{E}[\int_{\mathbb{T}} \|u_t\|_2^2 dt] \leq \infty$. To emphasize the (implicit) dependence of $J_u(x_0)$ on the realization of the Wiener process, we often explicitly write $J_u(x_0, \{B_t\}_{t \in \mathbb{T}})$.

A. Stochastic Gradient Descent

In this work we propose to directly seek the optimal controller via mini-batch *stochastic gradient descent* (SGD) (or *ascent*) optimization [27], [28]. The algorithm works as follows: given a scalar criterion $J_\alpha(x_0, \{B_t\}_{t \in \mathbb{T}})$ dependent on the variable α and the Brownian path, it attempts at computing $\alpha^* = \arg \min_\alpha \mathbb{E}[J_\alpha(x_0, \{B_t\}_{t \in \mathbb{T}})]$ or $\alpha^* = \arg \max_\alpha \mathbb{E}[J_\alpha(x_0, \{B_t\}_{t \in \mathbb{T}})]$ starting from an initial guess α_0 by updating a candidate solution α_k recursively as $\alpha_{k+1} = \alpha_k \pm \frac{\eta_k}{N} \sum_{i=1}^N \frac{d}{d\alpha} J_{\alpha_k}(x_0, \{B_t\}_{t \in \mathbb{T}}^i)$, where N is a predetermined number of independent and identically distributed samples $\{B_t\}_{t \in \mathbb{T}}^i$ of the Wiener process, η_k is a positive scalar *learning rate* and the sign of the update depends on the minimization or maximization nature of the problem. If η_k is suitably chosen and J_α is convex, α_k converges in expectation to the minimizer of J_α as $k \rightarrow \infty$. Although global convergence is no longer guaranteed in the non-convex case, SGD-based techniques have been employed across application areas due to their scaling and unique convergence characteristics. These methods have further been refined over time to show remarkable results even in non-convex settings [29], [30].

B. Finite Dim. Optimization via Neural Approximators

Consider the case in which the criterion J_u has to be minimized. Locally optimal control processes $u = \{u_t\}_{t \in \mathbb{T}} \in \mathbb{A}$ could be obtained, in principle, by iterating the SGD algorithm for the criterion J_u in the set \mathbb{A} of admissible control processes. Since \mathbb{A} is at least in $L_2(\mathbb{T} \rightarrow \mathbb{U})$ we could preserve local convergence of SGD in the function-space [31]. An idealized function-space SGD algorithm would compute iterates of the form

$$u^{k+1} = u^k - \frac{\eta_k}{N} \sum_{i=1}^N \frac{\delta}{\delta u} J_u(x_s, \{B_t\}_{t \in \mathbb{T}}^i)$$

where u^k is the solution at the k th SGD iteration, $\delta/\delta(\cdot)$ is the variational or functional derivative and J_u satisfies $\delta \mathbb{E}[J_u(x_s, \{B_t\}_{t \in \mathbb{T}})]/\delta u = \mathbb{E}[\delta J_u(x_s, \{B_t\}_{t \in \mathbb{T}})/\delta u]$.

At each *training step* of the controller, this approach would thus perform N independent path simulations, compute the criterion J_u and apply the gradient descent update. While local convergence to optimal solutions is still ensured under mild regularity and smoothness assumptions, obtaining derivatives in function space turns out to be computationally intractable. Any infinite-dimensional algorithm needs to be discretized in order to be practically implementable. We therefore seek to reduce the problem to finite dimension by approximating u_t^k with a *neural network*⁴.

⁴Here by *neural network* we mean any *learnable* parametric function u_θ with some approximation power for specific classes of functions.

Let $u_{\theta,t} : \theta, t \mapsto u_{\theta,t}$ be a neural network with parameters $\theta \in \mathbb{R}^{n_\theta}$ which we use as functional approximators for candidate optimal controllers, i.e. $u_{\theta,t} \approx u_t^* \forall t \in \mathbb{T}$. Further, we denote $f_\theta(t, x_t) := f(t, x_t, u_{\theta,t})$ and $g_\theta(t, x_t) := g(t, x_t, u_{\theta,t})$ and the optimization criterion

$$J_\theta(x_0) = \int_{\mathbb{T}} \gamma_\theta(t, \Phi_{s,t}(x_0)) dt + \Gamma_\theta(T, \Phi_{s,T}(x_0)) \quad (5)$$

with $\gamma_\theta(\cdot, \cdot) = \gamma(\cdot, \cdot, u_{\theta,\cdot})$ and $\Gamma_\theta(\cdot, \cdot) = \Gamma(\cdot, \cdot, u_{\theta,T})$. Then, the optimization problem turns into finding via gradient descent the optimal parameters θ , by iterating

$$\theta^{k+1} = \theta^k - \frac{\eta_k}{N} \sum_{i=1}^N \frac{d}{d\theta} J_{\theta^k}(x_0, \{B_t\}_{t \in \mathbb{T}}^i). \quad (6)$$

If strong constraints over the set of admissible controllers \mathbb{A} are imposed, the approximation problem can be rewritten onto a complete orthogonal basis of \mathbb{A} and u_θ is parameterized by a truncation of the eigenfunction expansion, rather than a neural network.

Remark 1 (Heuristics): As common practice in machine learning, the proposed approach relies heavily on the following empirical assumptions

- i. The numerical estimate of the mean gradients will be accurate enough to track the direction of the true $dJ_\theta/d\theta$. Here, estimation errors of the gradient will be introduced by the numerical discretization of the SDE and the finiteness of independent path simulations.
- ii. The local optima controller reached by gradient descent/ascent will be *good enough* to satisfy the performance requirements of the control problem.

In order to perform the above gradient descent we therefore need to compute the gradient (i.e. the sensitivity) of J_u with respect to the parameters θ in a computationally efficient manner. In the following we detail different approaches to differentiate through SDE solutions.

IV. COST GRADIENTS AND ADJOINT DYNAMICS

The most straightforward approach for computing path-wise gradients (i.e. independently for each realization of $\{B_t\}_{t \in \mathbb{T}}$) is by directional differentiation⁵ i.e.

$$\frac{dJ_\theta}{d\theta} = \int_{\mathbb{T}} \left[\frac{\partial \gamma_\theta}{\partial x_t} \frac{dx_t}{d\theta} + \frac{\partial \gamma_\theta}{\partial \theta} \right] dt + \frac{\partial \Gamma_\theta}{\partial x_T} \frac{dx_T}{d\theta} + \frac{\partial \Gamma_\theta}{\partial \theta}$$

where the quantities $dx_t/d\theta$ and $dx_T/d\theta$ can be obtained with the following result.

Proposition 1 (Path-wise Forward Sensitivity): Let $S_t = dx_t/d\theta$. Then, S_t is a $\{\mathcal{F}_t\}$ -adapted process satisfying

$$dS_t = \left[\frac{\partial f_\theta}{\partial x_t} S_t + \frac{\partial f_\theta}{\partial \theta} \right] dt + \sum_{i=1}^{n_\xi} \left[\frac{\partial g_\theta^{(i)}}{\partial x_t} S_t + \frac{\partial g_\theta^{(i)}}{\partial \theta} \right] \circ dB_t^{(i)}$$

Proof: The proof is an extension of the forward sensitivity analysis for ODEs (see [32, Sec. 3.4]) to the stochastic case. Given the SDE of interest

$$x_t = x_0 + \int_{\mathbb{T}} f_\theta(t, x_t) dt + \int_{\mathbb{T}} g_\theta(t, x_t) \circ dB_t$$

⁵Differentiability under the integral sign follows by our smoothness assumptions.

differentiating under the integral sign w.r.t. θ gives

$$\frac{dx_t}{d\theta} = \int_{\mathbb{T}} \left[\frac{\partial f_\theta}{\partial x_t} \frac{dx_t}{d\theta} + \frac{\partial f_\theta}{\partial \theta} \right] dt + \sum_{i=1}^{n_\xi} \int_{\mathbb{T}} \left[\frac{\partial g_\theta^{(i)}}{\partial x_t} \frac{dx_t}{d\theta} + \frac{\partial g_\theta^{(i)}}{\partial \theta} \right] \circ dB_t^{(i)}$$

and result follows setting $S_t = dx_t/d\theta$. That differentiating under the integral sign is allowed follows by our assumptions on f_θ and g_θ and by an application of [22, Lemma 3.7.1] to the augmented SDE (x_t, θ) . ■

The main issue with the forward sensitivity approach is its curse of dimensionality with respect to the number of parameters in θ and state dimensions n_x as it requires us to solve an $n_x \times n_\theta$ matrix-valued SDE for the whole time horizon \mathbb{T} . At each integration step the full Jacobians $\partial f_\theta/\partial \theta$, $\partial g_\theta/\partial \theta$ are required, causing memory and computational overheads in software implementations. Such issue can be overcome by introducing *adjoint* backward gradients.

Theorem 1 (Path-wise Backward Adjoint Gradients):

Consider the cost function (5) and let $\lambda \in C^1(\mathbb{T} \rightarrow \mathbb{X})$ be a Lagrange multiplier. Then, $dJ_\theta/d\theta$ is given by

$$\frac{dJ_\theta}{d\theta} = \frac{\partial \Gamma_\theta}{\partial \theta} + \int_{\mathbb{T}} \left[\lambda_t^\top \frac{\partial f_\theta}{\partial \theta} + \frac{\partial \gamma_\theta}{\partial \theta} \right] dt + \sum_{i=1}^{n_\xi} \int_{\mathbb{T}} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial \theta} \circ dB_t^{(i)}$$

where the Lagrange multiplier λ_t satisfies the following backward Stratonovich SDE:

$$d\lambda_t^\top = - \left[\lambda_t^\top \frac{\partial f_\theta}{\partial x_t} + \frac{\partial \gamma_\theta}{\partial x_t} \right] dt - \sum_{i=1}^{n_\xi} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial x_t} \circ d\check{B}_t^{(i)} \\ \lambda_T^\top = \frac{\partial \Gamma_\theta}{\partial x_T}$$

Lemma 1 (Line integration by parts): Let \mathbb{T} be a compact subset of \mathbb{R} and $r : \mathbb{T} \rightarrow \mathbb{R}^n$, $v : \mathbb{T} \rightarrow \mathbb{R}^n$ such that $\forall t \in \mathbb{T}$, $t \mapsto r_t := r(t) \in C^1$ and $t \mapsto v_t := v(t) \in C^1$. We have $\int_{\mathbb{T}} \langle r_t, dv_t \rangle = \langle r_T, v_T \rangle - \langle r_0, v_0 \rangle - \int_{\mathbb{T}} \langle dr_t, v_t \rangle$.

Proof: The proof follows immediately from integration by parts using $dr_t = \dot{r}_t dt$, $dv_t = \dot{v}_t dt$. ■

Proof: (Theorem 1) Let \mathcal{L} be the Lagrangian of the optimization problem and the process $\lambda_t \in C^1(\mathbb{T} \rightarrow \mathbb{X})$ an \mathcal{F}_t -adapted Lagrange multiplier. We consider \mathcal{L} of the form

$$\mathcal{L} = J_\theta(x_0) - \int_{\mathbb{T}} \langle \lambda_t, dx_t - f_\theta(t, x_t)dt - g_\theta(t, x_t) \circ dB_t \rangle \\ = \Gamma_\theta(x_T) + \int_{\mathbb{T}} \gamma_\theta(t, x_t)dt \\ - \int_{\mathbb{T}} \langle \lambda_t, dx_t - f_\theta(t, x_t)dt - g_\theta(t, x_t) \circ dB_t \rangle$$

Since $dx_t - f_\theta(t, x_t)dt - g_\theta(t, x_t) \circ dB_t = 0$ by construction, the integral term is always null, $d\mathcal{L}/d\theta = dJ_\theta(x_0)/d\theta$ and the Lagrange multiplier process can be freely assigned. Thus,

$$\frac{dJ_\theta}{d\theta} = \frac{d\mathcal{L}}{d\theta} = \frac{\partial \Gamma_\theta}{\partial \theta} + \frac{\partial \Gamma_\theta}{\partial x_T} \frac{dx_T}{d\theta} + \frac{d}{d\theta} \int_{\mathbb{T}} \gamma_\theta(t, x_t)dt \\ - \frac{d}{d\theta} \int_{\mathbb{T}} \langle \lambda_t, dx_t - f_\theta(t, x_t)dt - g_\theta(t, x_t) \circ dB_t \rangle$$

Note that, by Lemma 1, we have $\int_{\mathbb{T}} \langle \lambda_t, dx_t \rangle = \langle \lambda_t, x_t \rangle|_0^T - \int_{\mathbb{T}} \langle d\lambda_t, x_t \rangle$. For compactness, we will omit the argument of

all functions unless special cases. We have

$$\frac{dJ_\theta}{d\theta} = \frac{\partial \Gamma_\theta}{\partial \theta} + \frac{\partial \Gamma_\theta}{\partial x_T} \frac{dx_T}{d\theta} - \lambda_T^\top \frac{dx_T}{d\theta} + \lambda_0^\top \frac{dx_0}{d\theta} \\ + \int_{\mathbb{T}} \left[\left(\frac{\partial \gamma_\theta}{\partial \theta} + \frac{\partial \gamma_\theta}{\partial x_t} \frac{dx_t}{d\theta} \right) dt + d\lambda_t^\top \frac{dx_t}{d\theta} \right] \\ + \int_{\mathbb{T}} \lambda_t^\top \left[\frac{\partial f_\theta}{\partial \theta} + \frac{\partial f_\theta}{\partial x_t} \frac{dx_t}{d\theta} \right] dt \\ + \sum_{i=1}^{n_\xi} \int_{\mathbb{T}} \lambda_t^\top \left[\frac{\partial g_\theta^{(i)}}{\partial \theta} + \frac{\partial g_\theta^{(i)}}{\partial x_t} \frac{dx_t}{d\theta} \right] \circ dB_t^{(i)} \quad (7)$$

which, by reorganizing the terms, leads to

$$\frac{dJ_\theta}{d\theta} = \frac{\partial \Gamma_\theta}{\partial \theta} + \int_{\mathbb{T}} \left[\frac{\partial \gamma_\theta}{\partial \theta} + \lambda_t^\top \frac{\partial f_\theta}{\partial \theta} \right] dt \\ + \sum_{i=1}^{n_\xi} \int_{\mathbb{T}} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial \theta} \circ dB_t^{(i)} + \left[\frac{\partial \Gamma_\theta}{\partial x_T} - \lambda_T^\top \right] \frac{dx_T}{d\theta} \\ + \int_{\mathbb{T}} \left[d\lambda_t^\top + \left(\lambda_t^\top \frac{\partial f_\theta}{\partial x_t} + \frac{\partial \gamma_\theta}{\partial x_t} \right) dt + \sum_{i=1}^{n_\xi} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial \theta} \circ dB_t^{(i)} \right] \frac{dx_t}{d\theta}$$

Finally, if the process λ_t satisfies the backward SDE $d\lambda_t^\top = - \left[\lambda_t^\top \frac{\partial f_\theta}{\partial x_t} + \frac{\partial \gamma_\theta}{\partial x_t} \right] dt - \sum_{i=1}^{n_\xi} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial x_t} \circ d\check{B}_t^{(i)}$, $\lambda_T^\top = \frac{\partial \Gamma_\theta}{\partial x_T}$. The criterion function gradient is then simply obtained by

$$\frac{dJ_\theta}{d\theta} = \frac{\partial \Gamma_\theta}{\partial \theta} + \int_{\mathbb{T}} \left[\lambda_t^\top \frac{\partial f_\theta}{\partial \theta} + \frac{\partial \gamma_\theta}{\partial \theta} \right] dt + \sum_{i=1}^{n_\xi} \int_{\mathbb{T}} \lambda_t^\top \frac{\partial g_\theta^{(i)}}{\partial \theta} \circ dB_t^{(i)}$$

Note that if the integral term of the objective function is defined point-wise at a finite number of time instants t_k , i.e. $\int_s^T \sum_{k=1}^K \gamma_\theta(t_k, \Phi_{s,t_k}^{\theta,b}(x_s)) \delta(t - t_k) dt$, then the RHS of the adjoint state SDE becomes piece-wise continuous in $(t_k, t_{k+1}]$, yielding the *hybrid* stochastic dynamics

$$d\lambda_t = - \frac{\partial f_\theta}{\partial x_t} \lambda_t dt - \sum_{i=1}^{n_\xi} \frac{\partial g_\theta^{(i)}}{\partial x_t} \lambda_t \circ dB_t^{(i)} \quad t \in (t_k, t_{k+1}] \\ \lambda_t^- = \lambda_t + \frac{\partial \gamma_\theta}{\partial x_t} \quad t = t_k$$

where λ_t^- indicates the value of λ_t after a discrete *jump*, or, formally, $\lambda_t^- = \lim_{s \nearrow t} \lambda_s$.

V. EXPERIMENTAL EVALUATION

We evaluate the proposed approach on a classical problem in financial mathematics, continuous-time portfolio optimization [24]. We consider the challenging finite-horizon, transaction cost case [33].

A. Optimal Portfolio Allocation

Consider a two-asset market with proportional transaction costs, all securities are perfectly divisible. Suppose that in such a market all assets are traded continuously and that one of the them is *riskless*, i.e. it evolves according to the ODE

$$dV_t = r_t V_t dt, \quad V_s = v \in \mathbb{R}$$

where the *interest rate* r_t is any \mathcal{F}_t -adapted non-negative scalar-valued process. The other asset is a *stock* $S_t \in \mathbb{R}$ whose dynamics satisfy the Itô SDEs

$$dS_t = \mu_t S_t + \sigma_t S_t dB_t, \quad S_s = s \in \mathbb{R}$$

with *expected rate of return* $\mu : \mathbb{T} \rightarrow \mathbb{R}$ and *instantaneous volatility* $\sigma : \mathbb{T} \rightarrow \mathbb{R}_+$ \mathcal{F}_t -adapted processes.

Now, let us consider an agent who invests in such market and is interested in rebalancing a two-asset portfolio through buy and sell actions. We obtain

$$\begin{aligned} dV_t &= r_t V_t dt - d\pi_t^i + (1 - \alpha)d\pi_t^d \\ dS_t &= \mu_t S_t + \sigma_t S_t dB_t + d\pi_t^i - d\pi_t^d \end{aligned} \quad (8)$$

where π_t^i, π_t^d are nondecreasing, \mathcal{F}_t -adapted processes indicating the cumulative amount of sales and purchases of the risky asset, and $\alpha > 0$ is the proportional transaction cost. Here we let $d\pi_t^i = u_t^i dt$, $d\pi_t^d = u_t^d dt$ and $u_t = (u_t^i, u_t^d)$. An optimal strategy u_t for such a portfolio requires specification of a utility function, encoding risk aversion and final objective of the investor. Explicit results for the finite horizon, transaction cost case exist only on specific classes of utilities, such as isoelastic [33]. As a further complication, the horizon length itself is known to greatly affect the optimal strategy. As an illustrative example, we consider here a portfolio optimality criteria where high levels of stock value are penalized, perhaps due to hedging considerations related to other already existing portfolios in possession of the hypothetical trader:

$$J_u = \int_0^T (r_t V_t + \mu_t S_t - \nu S_t^2 \sigma_t) dt + r_T V_T + \mu_T S_T$$

to be maximized in expectation for some constant $\nu > 0$. The same methodology can be directly extended to cover generic objectives, including e.g different risk measures, complex stock and interest rate dynamics, and more elaborate transaction cost accounting rules. We then obtain a parameterization $u_{\theta,t}^i, u_{\theta,t}^d$ of the policies u_t^i, u_t^d as follows. In particular we define feed-forward neural networks $N_\theta : \mathbb{R}^2 \rightarrow \mathbb{R}^2$; $(S_t, V_t) \mapsto (u_{\theta,t}^i, u_{\theta,t}^d)$ taking as input the instantaneous values of the assets, $u_{\theta,t} = N_\theta(V_t, S_t) = \ell_L \circ \varphi \circ \ell_{L-1} \circ \dots \circ \varphi \circ \ell_0(V_t, S_t)$, where ℓ_k are linear affine finite operators, $\ell_k x = A_k x + b_k$, $A_k \in \mathbb{R}^{n_{k+1} \times n_k}$, $b_k \in \mathbb{R}^{n_{k+1}}$ and $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is any nonlinear *activation* function thought to be acting element-wise. The vector of parameters θ is thus the flattened collection of all *weights* A_k and *biases* b_k .

B. Numerical Results

In particular, we choose three-layer feed-forward neural networks with thirty-two neurons per layer, capped with a `softplus` activation $\varphi(x) = \log(1 + \exp(x))$ to ensure $\pi_\theta^i, \pi_\theta^d$ be nondecreasing processes.

The optimization procedure for strategy parameters θ has been carried out through 100 iterations of the Adam [34] gradient-based optimizer with step size 0.03. The asset dynamics and cost parameters are set as $\alpha = 0.05$, $r = 0.04$, $\mu = 0.23$, $\sigma = 0.18$. We produce 50 path realizations during each training iteration and average the gradients. The asset dynamics are interpreted in the Itô sense and have been solved forward with the Itô-Milstein method [35], which has been converted to Stratonovich-Milstein for the backward adjoint SDE. Figure 1 shows different realization of portfolio trajectories in the (S_t, V_t) state-space for four different values of ν , $[0, 0.25, 0.5, 1]$. The trajectories agree with plausible strategies for each investment style; risk-seeking investors ($\nu = 0$) sell V_t to maximize the amount of stock owned at final time T , and hence potential portfolio

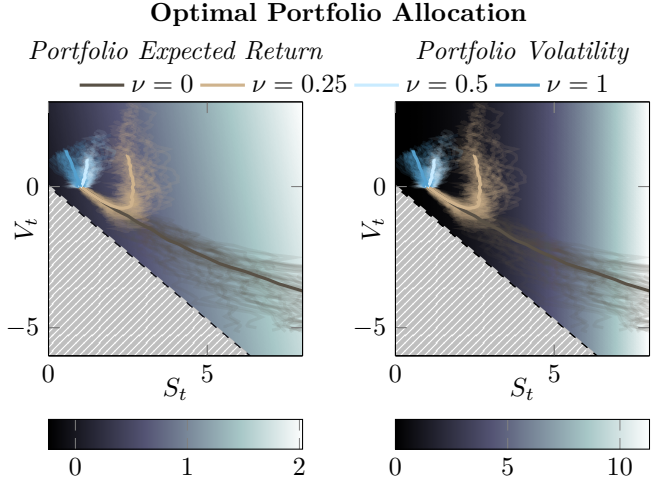


Fig. 1. Realization of (S_t, V_t) trajectories under learned strategies for four investors with different degrees of risk aversion, encoded in ν , with a portfolio initialized at $(1, 0)$ at time 0. Different path realizations generated by the strategies are shown, along with their mean. [Left] Portfolio evolution against $r_t V_t + \mu_t S_t$ [Right] Portfolio evolution against $S_t^2 \sigma$. Each strategy suits the corresponding risk aversion level.

growth. On the other hand, moderately risk-averse investors ($\nu = 0.25$) learn hybrid strategies where an initial phase of stock purchasing is followed by a movement towards a more conservative portfolio with a balanced amount of S_t and V_t . The policies u_θ^i, u_θ^d are visualized in Figure 2, confirming the highly non-linear nature of the learned controller. It should be noted that for a strategy to be allowed, the portfolio has to remain above the solvency line $V_t + (1 - \alpha) S_t \geq 0$ [33], indicated in the plot. We empirically observe automatic satisfaction of this specific constraint for all but the investors with lowest risk aversion. For the case $\nu = 0$, we introduce a simple logarithmic barrier function as a soft constraint [36].

VI. RELATED WORK

Computing path-wise sensitivities through stochastic differential equations has been extensively explored in the literature. In particular, *forward* sensitivities with respect to initial conditions of solutions analogous to Proposition 1 were first introduced in [25] and extended for Itô case to parameters sensitivities in [37]. These approaches rely on integrating forward matrix-valued SDEs whose drift and diffusion require full Jacobians of f_θ and g_θ at each integration step. Thus, such methods poorly scales in computational cost with large number of parameters and high-dimensional-state regimes. This issue is averted by using *backward* adjoint sensitivity [38] where a vector-valued SDE is integrated backward and only requires vector-Jacobian products to be evaluated. In this direction, [26] proposed to solve the system's backward SDE alongside the adjoint SDE to recover the value of the state x_t and thus improve the algorithm memory footprint. Further, the approach is derived as a special case of [22, Theorem 2.4.1] and only considers criteria J_θ depending on the final state x_T of the system.

Our stochastic adjoint process extends the results of [26] to *integral* criteria potentially exhibiting explicit parameter

Optimal Portfolio Allocation

Learned Allocation Policies

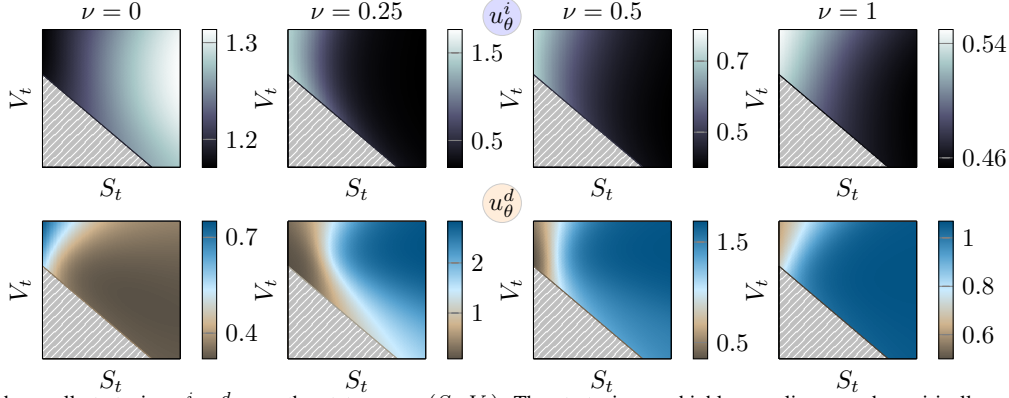


Fig. 2. Learned buy-sell strategies u_θ^i, u_θ^d over the state-space (S_t, V_t) . The strategies are highly non-linear, and empirically agree with investment strategies at different risk aversion degrees. In example, the risk-seeking investor $\nu = 0$ balances u_θ^i, u_θ^d to purchase the maximum amount of stock possible while remaining above the solvency line $V_t + (1 - \alpha)S_t \geq 0$.

dependence. Further, we proposed a novel proving strategy based on classic variational analysis.

REFERENCES

- [1] D. P. Bertsekas and S. Shreve, *Stochastic optimal control: the discrete-time case*, 2004.
- [2] H. Yu and D. P. Bertsekas, “Q-learning and policy iteration algorithms for stochastic shortest path problems,” *Annals of Operations Research*, vol. 208, no. 1, pp. 95–132, 2013.
- [3] Y. Wu and T. Shen, “Policy iteration algorithm for optimal control of stochastic logical dynamical systems,” *IEEE transactions on neural networks and learning systems*, vol. 29, no. 5, pp. 2031–2036, 2017.
- [4] S. Levine, “Reinforcement learning and control as probabilistic inference: Tutorial and review,” *arXiv preprint arXiv:1805.00909*, 2018.
- [5] E. Tse and M. Athans, “Adaptive stochastic control for a class of linear systems,” *IEEE Transactions on Automatic Control*, vol. 17, no. 1, pp. 38–52, 1972.
- [6] S. Peng and Z. Wu, “Fully coupled forward-backward stochastic differential equations and applications to optimal control,” *SIAM Journal on Control and Optimization*, vol. 37, no. 3, pp. 825–843, 1999.
- [7] M. A. Pereira and Z. Wang, “Learning deep stochastic optimal control policies using forward-backward sdes,” in *Robotics: science and systems*, 2019.
- [8] Z. Wang, K. Lee, M. A. Pereira, I. Exarchos, and E. A. Theodorou, “Deep forward-backward sdes for min-max control,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 6807–6814.
- [9] E. Weinan, “A proposal on machine learning via dynamical systems,” *Communications in Mathematics and Statistics*, vol. 5, no. 1, pp. 1–11, 2017.
- [10] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, “Neural ordinary differential equations,” *arXiv preprint arXiv:1806.07366*, 2018.
- [11] S. Massaroli, M. Poli, J. Park, A. Yamashita, and H. Asama, “Dissecting neural odes,” *arXiv preprint arXiv:2002.08071*, 2020.
- [12] S. Massaroli, M. Poli, M. Bin, J. Park, A. Yamashita, and H. Asama, “Stable neural flows,” *arXiv preprint arXiv:2003.08063*, 2020.
- [13] S. Peluchetti and S. Favaro, “Infinitely deep neural networks as diffusion processes,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1126–1136.
- [14] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural networks*, vol. 61, pp. 85–117, 2015.
- [15] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. MIT press Cambridge, 2016, vol. 1, no. 2.
- [16] L. Frieri, Y. Zheng, and M. Kamgarpour, “Learning the globally optimal distributed lq regulator,” in *Learning for Dynamics and Control*. PMLR, 2020, pp. 287–297.
- [17] J. Bu, A. Mesbahi, M. Fazel, and M. Mesbahi, “Lqr through the lens of first order methods: Discrete-time case,” *arXiv preprint arXiv:1907.08921*, 2019.
- [18] N. G. Van Kampen, “Itô versus stratonovich,” *Journal of Statistical Physics*, vol. 24, no. 1, pp. 175–187, 1981.
- [19] B. Øksendal, “Stochastic differential equations,” in *Stochastic differential equations*. Springer, 2003, pp. 65–84.
- [20] K. Itô, “Differential equations determining a markoff process,” *J. Pan-Japan Math. Coll.*, vol. 1077, pp. 1352–1400, 1942.
- [21] R. Stratonovich, “A new representation for stochastic integrals and equations,” *SIAM Journal on Control*, vol. 4, no. 2, pp. 362–371, 1966.
- [22] H. Kunita, *Stochastic Flows and Jump-Diffusions*. Springer.
- [23] E. Wong and M. Zakai, “On the convergence of ordinary integrals to stochastic integrals,” *The Annals of Mathematical Statistics*, vol. 36, no. 5, pp. 1560–1564, 1965.
- [24] R. C. Merton, “Lifetime portfolio selection under uncertainty: The continuous-time case,” *The review of Economics and Statistics*, pp. 247–257, 1969.
- [25] H. Kunita, *Stochastic flows and stochastic differential equations*. Cambridge university press, 1997, vol. 24.
- [26] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud, “Scalable gradients for stochastic differential equations,” *arXiv preprint arXiv:2001.01328*, 2020.
- [27] H. Robbins and S. Monro, “A stochastic approximation method,” *The annals of mathematical statistics*, pp. 400–407, 1951.
- [28] S. Ruder, “An overview of gradient descent optimization algorithms,” *arXiv preprint arXiv:1609.04747*, 2016.
- [29] C. De Sa, C. Re, and K. Olukotun, “Global convergence of stochastic gradient descent for some non-convex matrix problems,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2332–2341.
- [30] Z. Allen-Zhu and E. Hazan, “Variance reduction for faster non-convex optimization,” in *International conference on machine learning*. PMLR, 2016, pp. 699–707.
- [31] G. Smyrlis and V. Zisis, “Local convergence of the steepest descent method in hilbert spaces,” *Journal of mathematical analysis and applications*, vol. 300, no. 2, pp. 436–453, 2004.
- [32] H. K. Khalil, *Nonlinear systems*, vol. 3.
- [33] H. Liu and M. Loewenstein, “Optimal portfolio selection with transaction costs and finite horizons,” *The Review of Financial Studies*, vol. 15, no. 3, pp. 805–835, 2002.
- [34] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [35] G. Mil’shtein, “Approximate integration of stochastic differential equations,” *Theory of Probability & Its Applications*, vol. 19, no. 3, pp. 557–562, 1975.
- [36] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.
- [37] E. Gobet and R. Munos, “Sensitivity analysis using itô-malliavin calculus and martingales, and application to stochastic optimal control,” *SIAM Journal on control and optimization*, vol. 43, no. 5, pp. 1676–1713, 2005.
- [38] M. Giles and P. Glasserman, “Smoking adjoints: fast evaluation of greeks in monte carlo calculations,” 2005.