

000 001 002 003 004 005 A SCHRÖDINGER EIGENFUNCTION METHOD FOR 006 LONG-HORIZON STOCHASTIC OPTIMAL CONTROL 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028

ABSTRACT

029
030 High-dimensional stochastic optimal control (SOC) becomes harder with longer
031 planning horizons: existing methods scale linearly in the horizon T , with per-
032 formance often deteriorating exponentially. We overcome these limitations for
033 a subclass of linearly-solvable SOC problems—those whose uncontrolled drift
034 is the gradient of a potential. In this setting, the Hamilton-Jacobi-Bellman equa-
035 tion reduces to a linear PDE governed by an operator \mathcal{L} . We prove that, under
036 the gradient drift assumption, \mathcal{L} is unitarily equivalent to a Schrödinger operator
037 $\mathcal{S} = -\Delta + \mathcal{V}$ with purely discrete spectrum, allowing the long-horizon control to
038 be efficiently described via the eigensystem of \mathcal{L} . This connection provides two
039 key results: first, for a symmetric linear-quadratic regulator (LQR), \mathcal{S} matches the
040 Hamiltonian of a quantum harmonic oscillator, whose closed-form eigensystem
041 yields an analytic solution to the symmetric LQR with *arbitrary* terminal cost.
042 Second, in a more general setting, we learn the eigensystem of \mathcal{L} using neural
043 networks. We identify implicit reweighting issues with existing eigenfunction
044 learning losses that degrade performance in control tasks, and propose a novel
045 loss function to mitigate this. We evaluate our method on several long-horizon
046 benchmarks, achieving an order-of-magnitude improvement in control accuracy
047 compared to state-of-the-art methods, while reducing memory usage and runtime
048 complexity from $\mathcal{O}(Td)$ to $\mathcal{O}(d)$.
049
050

1 INTRODUCTION

051 Stochastic optimal control (SOC) concerns the problem of directing a stochastic system, typically
052 modeled by a stochastic differential equation (SDE), to minimize an expected total cost. SOC has
053 found applications in various domains, e.g. stochastic filtering (Mitter, 1996), rare event simulation
054 in molecular dynamics (Hartmann & Schütte, 2012; Hartmann et al., 2014), robotics (Gorodetsky
055 et al., 2018) and finance (Pham, 2009).

056 Built on the principle of dynamic programming, the global optimality condition of SOC can be
057 expressed by the Hamilton-Jacobi-Bellman (HJB) equation. In this paper, we focus on the *affine*
058 *control* setting commonly considered in the literature (Fleming & Rishel, 1975; Kappen, 2005b;
059 Fleming & Soner, 2006; Yong & Zhou, 1999; Domingo-Enrich et al., 2024b; Nüsken & Richter, 2021;
060 Carius et al., 2022; Holdijk et al., 2023), where the control affects the state of the system linearly. This
061 setting is of interest since the optimal control will exactly match the gradient of the value function of
062 SOC problem and hence the corresponding HJB equation can be drastically simplified.

063 A canonical special case of this affine-control framework is the linear–quadratic regulator (LQR), in
064 which the uncontrolled dynamics follow an Ornstein–Uhlenbeck linear SDE and both the running cost
065 and the terminal cost are quadratic. One can show that in the LQR setting the value function retains a
066 quadratic form—indeed, it is quadratic at terminal time because the terminal cost is quadratic—and
067 that the optimal feedback control is linear w.r.t. the state. Consequently, the associated SOC problem
068 admits an explicit solution via the finite-dimensional matrix Riccati differential equation (van Handel,
069 2007)

070 To obtain the optimal control in more general settings requires numerical procedures. For low-
071 dimensional problems, classical grid-based PDE solvers may be used, but these suffer from the curse
072 of dimensionality. This has led to several works proposing the use of neural networks (NN) to solve
073 the HJB equation in more complex high-dimensional settings, either through a forward-backward

stochastic differential equation (FBSDE) approach (Han et al., 2018; Ji et al., 2022; Andersson et al., 2023; Beck et al., 2019) or so-called iterative diffusion optimization (IDO) methods, which sample controlled trajectories through simulation and update the NN parameter using stochastic gradients from automatic differentiation (Nüsken & Richter, 2021; Domingo-Enrich et al., 2024b;a). A more comprehensive review on existing methods for short-horizon SOC can be found in Appendix C.

While these methods have proven successful, their performance suffers as the time horizon T grows. Both the memory requirement and per-iteration runtime increase at least linearly in T . Additionally, it holds that error estimates for the deep FBSDE method worsen as T increases (Han & Long, 2020, Theorem 4), and for IDO methods using importance sampling the weight variance may increase exponentially in T (Liu et al., 2018). These limitations were observed empirically in Assabumrungrat et al. (2024); Domingo-Enrich et al. (2024b), and were reproduced in our experiments (see Figure 1).

Linearly-solvable HJB. The HJB equation is in general nonlinear. However, in the special case where the system’s diffusion coefficient matches the affine-control mapping, the Cole–Hopf transform can eliminate the nonlinearity (Evans, 2022). Specifically, let $V(x, t)$ denote the value function of the SOC problem and define a new function $\psi := \exp(-V)$.

Under this transformation, the HJB is equivalently rewritten as the following *linear PDE*

$$\partial_t \psi(x, t) = \mathcal{L} \psi(x, t), \quad \psi(x, T) = \psi_T(x) \quad (1)$$

for some *linear* operator \mathcal{L} . Moreover, the optimal control u , which exactly matches the gradient of the value function $-\partial_x V$, can be obtained as $u^* = \partial_x \log \psi$ (Kappen, 2005b).

Working with a linear PDE brings several clear benefits over a nonlinear one, such as a simplified analysis — questions of well-posedness and solution regularity of the PDE become much more tractable — but more importantly algorithmic insight: One can borrow ideas from the finite-dimensional linear ODE $\dot{u}(t) + A u(t) = 0$, where A is a real symmetric matrix. This ODE has a closed-form solution $u(t) = e^{-At} u(0)$, and because the matrix exponential e^{-At} acts simply on A ’s eigenvalues, expanding $u(0)$ in the corresponding eigenvectors yields an efficient numerical scheme.

While the domain of the operator \mathcal{L} is infinite-dimensional (acting on functions rather than finite vectors), the same exponential-integrator principle applies (Theorem 1), and we can write $\psi_t = e^{(t-T)\mathcal{L}} \psi_T$, where $(e^{s\mathcal{L}})_{s \geq 0}$ is a semi-group (Renardy & Rogers, 2006, Chapter 12). Of course, carrying it out in an infinite-dimensional setting introduces additional technical challenges that must be carefully addressed. In particular, expanding ψ_T as a series of eigenfunctions requires \mathcal{L} to possess a *discrete spectrum*, i.e. one can find an orthonormal basis of eigenfunctions $(\phi_i)_{i \in \mathbb{N}}$ and corresponding eigenvalues $\lambda_0 < \lambda_1 \leq \dots$ such that $\mathcal{L}\phi_i = \lambda_i \phi_i$ for all $i \in \mathbb{N}$. Under this assumption, the optimal control can be informally written as

$$u^*(x, t) = \partial_x \log \phi_0(x) + \mathcal{O}\left(e^{-(\lambda_1 - \lambda_0)(T-t)}\right) \quad \text{for any fixed } t \text{ as } T \rightarrow \infty, \quad (2)$$

where \mathcal{O} hides eigenfunctions ϕ_i with $i \geq 1$. A precise statement can be found in Theorem 3. To turn the above formula into a practical, long-horizon SOC algorithm, we need (1) **Spectral verification**. Prove that \mathcal{L} indeed has a discrete spectrum; (2) **Eigenfunction identification**. Compute the spatial derivative of the principal eigenfunction ϕ_0 .

Our approach: Reduction to Schrödinger operator. In this paper, to guarantee the spectral verification, we assume that the drift of the dynamics is *the gradient of a potential*. Such problems

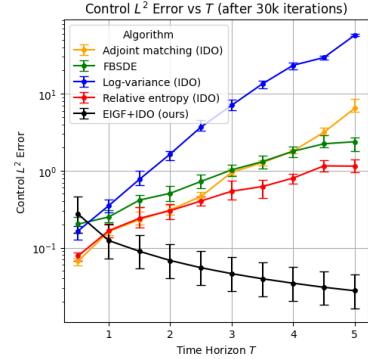


Figure 1: Performance degradation as time horizon T increases for different methods (see Appendix E for details).

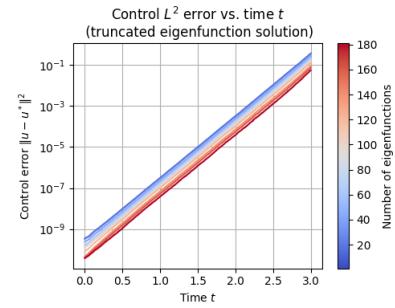


Figure 2: Diminishing returns from increasing the number of eigenfunctions for an LQR in $d = 20$ dimensions.

108 of controlling a diffusion process with gradient drift show up in overdamped molecular dynamics
 109 (Schütte et al., 2012), mean-field games (Bakshi et al., 2019; Grover et al., 2018), the control of
 110 particles with interaction potentials (Carrillo et al., 2020; Totzeck & Pinna, 2020), as well as social
 111 models for opinion formation (Castellano et al., 2009; Albi et al., 2017).

112 In this setting, the operator \mathcal{L} in (1) is unitarily equivalent to the Schrödinger operator $\tilde{\mathcal{L}} = -\Delta + \mathcal{V}$,
 113 where Δ is the Laplacian and \mathcal{V} is an effective potential determined by the original drift and running
 114 cost. Because $\tilde{\mathcal{L}}$ is known to have a purely discrete spectrum on L^2 and unitary equivalence preserves
 115 the spectral properties, the operator \mathcal{L} in (1) likewise enjoys a discrete spectrum.
 116

117 This result forms the basis of our new framework for long-horizon SOC, where the problem is
 118 reformulated as learning the eigensystem of a Schrödinger operator. Indeed, (2) shows that the top
 119 eigenfunction ϕ_0 determines the long-term control, with corrections decaying exponentially with rate
 120 $\lambda_1 - \lambda_0$. We address the problem of eigenfunction identification in the following two scenarios:

121 • **Closed-form solution for LQR with non-quadratic terminal cost.** When the drift is linear with a
 122 *symmetric* coefficient matrix and the running cost is quadratic in the state, the resulting Schrödinger
 123 operator $\tilde{\mathcal{L}}$ coincides with the Hamiltonian of the quantum harmonic oscillator. Its eigenvalues and
 124 eigenfunctions are known explicitly (see Lemma 4 or Griffiths & Schroeter (2018)), i.e. $\{\lambda_i, \phi_i\}_{i \in \mathbb{N}}$
 125 are available in closed form. Consequently, our framework yields a fully explicit expression for the
 126 corresponding SOC. This removes the requirement of quadratic terminal cost in the the classical
 127 LQR solution.

128 • **Neural network-based approach for general gradient drift.** For general gradient-drift dynamics,
 129 we introduce a hybrid neural-network method to approximate the optimal control efficiently: Rather
 130 than attempting to learn the full spectrum of \mathcal{L} —which is prohibitively expensive and yields rapidly
 131 diminishing returns (fig. 2)—we exploit the exponential decay of the higher modes w.r.t. $T - t$ in
 132 eq. (2). Concretely, whenever $T - t$ exceeds a modest threshold (in our experiments $T - t \geq 1$), it
 133 suffices to approximate the control using only the top eigenfunction ϕ_0 . For the remaining period
 134 (t very close to T), we switch to established FBSDE/IDO solvers to handle the short-horizon SOC.

135 We propose a novel deep learning strategy for the task of learning the eigenfunction, tailored to
 136 SOC. While such a task has been extensively studied in the literature, previous approaches either
 137 optimize a variational Ritz objective (E & Yu, 2018; Zhang et al., 2022; Cabannes & Bach, 2024)
 138 or minimize the residual norm $\|\mathcal{L}\psi - \lambda\psi\|^2$ (Jin et al., 2022; Zhang et al., 2024; Nusken & Richter,
 139 2023). However, these losses implicitly reweight spatial regions, causing them to fail to learn
 140 the control in regions where the value function V is large—the most crucial areas. To eliminate
 141 this bias, we introduce a *relative* eigenfunction loss, $\|\mathcal{L}\psi/\psi - \lambda\|^2$, which removes the undesired
 142 weighting and robustly recovers the dominant eigenpair needed for control synthesis.

143 Our contributions are summarized as follows:

144 • We provide a new perspective on finite-horizon gradient-drift SOC problems, linking their solution
 145 to the eigensystem of a Schrödinger operator (Theorem 3). This yields a previously unreported
 146 closed-form solution to the symmetric LQR with arbitrary terminal cost (Theorem 4).
 147 • With this framework, we introduce a new procedure for solving SOC problems over long horizons
 148 by learning the operator’s eigenfunctions with neural networks. We show that existing eigenfunction
 149 solvers can be ill-suited for this task due to an implicit reweighting in the used loss, and propose a
 150 new loss function to remedy this.
 151 • We perform experiments in different settings to evaluate the proposed method against state-of-
 152 the-art SOC solvers, showing roughly an order of magnitude improvement in control L^2 error on
 153 several high-dimensional ($d = 20$) long-horizon problems.

155 2 PRELIMINARIES

157 2.1 STOCHASTIC OPTIMAL CONTROL

159 Fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$, and denote $(W_t)_{t \geq 0}$ a Brownian motion in this
 160 space. Let $(X_t^u)_{t \geq 0}$ denote the random variable taking values in \mathbb{R}^d defined through the SDE
 161

$$162 \quad dX_t^u = (b(X_t^u) + \sigma u(X_t^u, t)) dt + \sqrt{\beta^{-1}} \sigma dW_t, \quad X_0^u \sim p_0 \quad (3)$$

162 where $u : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is the control, $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the base drift, $\sigma \in \mathbb{R}^{d \times d}$ is the diffusion
 163 coefficient (assumed invertible) and $\beta \in \mathbb{R}_0^+$ is an inverse temperature characterizing the noise level.
 164 Note that we assume the drift and noise to be time-independent, in contrast to Nüsken & Richter
 165 (2021); Domingo-Enrich et al. (2024b). Under some regularity conditions on the coefficients and
 166 control u described in Appendix A, the SDE (3) has a unique strong solution. In stochastic optimal
 167 control, we view the dynamics (b, σ, β) as given and consider the problem of finding a control u
 168 which minimizes the cost functional

$$169 \quad J(u; x, t) = \mathbb{E} \left[\int_t^T \left(\frac{1}{2} \|u(X_t^u, t)\|^2 + f(X_t^u) \right) dt + g(X_T^u) \middle| X_t = x \right] \quad (4)$$

172 where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the running cost and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the terminal cost. We denote
 173 this optimal control as $u^*(x, t) = \arg \min_{u \in \mathcal{U}} J(u; x, t)$, with \mathcal{U} the set of admissible controls. To
 174 analyze this problem, one defines the *value function* V , which is defined as the minimum achievable
 175 cost when starting from x at time t ,

$$176 \quad V(x, t) := \inf_{u \in \mathcal{U}} J(u; x, t). \quad (5)$$

178 In this case, the optimal control u^* that minimizes the objective (4) is obtained from the value function
 179 through the relation $u^* = -\sigma^T \nabla V$, as described in (Nüsken & Richter, 2021, Theorem 2.2).

181 **Hamilton-Jacobi-Bellman equation.** A well-known fundamental result is that when the value
 182 function V is sufficiently regular, it satisfies the following partial differential equation, called the
 183 Hamilton-Jacobi-Bellman (HJB) equation (Fleming & Rishel, 1975):

$$184 \quad \partial_t V + \mathcal{K}V = 0 \quad \text{in } \mathbb{R}^d \times [0, T], \quad V(\cdot, T) = g \quad \text{on } \mathbb{R}^d, \quad (6)$$

$$186 \quad \text{where } \mathcal{K}V = \frac{1}{2\beta} \text{Tr}(\sigma \sigma^T \nabla^2 V) + b^T \nabla V - \frac{1}{2} \|\sigma^T \nabla V\|^2 + f. \quad (7)$$

188 The so-called *verification theorem* states (in some sense) the converse: if a function V satisfying the
 189 above PDE is sufficiently regular, it coincides with the value function (5) corresponding to (3)-(4),
 190 see (Fleming & Rishel, 1975, Section VI.4) and (Pavliotis, 2014, Sec. 2.3).

192 **A linear PDE reformulation** Although the HJB equation (6) is nonlinear in general, it was shown in
 193 Kappen (2005b) that for a specific class of optimal control problems (which includes the formulation
 194 (3)-(4)), a suitable transformation allows for a linear reformulation of (6). More specifically, when
 195 parametrizing $V(x, t) = -\beta^{-1} \log \psi(x, \frac{1}{2\beta}(T-t))$, the nonlinear terms cancel, and (6) becomes

$$197 \quad \begin{cases} \partial_\tau \psi + \mathcal{L}\psi = 0, \\ \psi(\cdot, 0) = \psi_0, \end{cases} \quad \text{where } \mathcal{L}\psi = -\text{Tr}(\sigma \sigma^T \nabla^2 \psi) - 2\beta b^T \nabla \psi + 2\beta^2 f \cdot \psi, \quad \psi_0 = \exp(-\beta g), \quad (8)$$

199 and we have introduced the variable $\tau = (2\beta)^{-1}(T-t)$. This is precisely the abstract form (1),
 200 but with a time reversal. For more details on this result, we refer to Appendix B. To simplify the
 201 presentation, we will often assume w.l.o.g. that $\sigma = I$ (see Appendix A), so that $\text{Tr}(\sigma \sigma^T \nabla^2) = \Delta$.

203 2.2 EIGENFUNCTION SOLUTIONS

205 Consider a Hilbert space \mathcal{H} with inner product $\langle \cdot, \cdot \rangle$, and a linear operator $\mathcal{L} : D(\mathcal{L}) \rightarrow \mathcal{H}$ defined on
 206 a dense subspace $D(\mathcal{L}) \subset \mathcal{H}$. Then we have the following standard definition:

207 **Definition 1** An element $\phi \in \mathcal{H}$ with $\phi \neq 0$ is an eigenfunction of \mathcal{L} if there exists a $\lambda \in \mathbb{C}$ such that
 208 $\mathcal{L}\phi = \lambda\phi$. The value λ is called an eigenvalue of \mathcal{L} (corresponding to ϕ), and the dimension of the
 209 null space of $\mathcal{L} - \lambda I$ is called the multiplicity of λ .

211 In a finite-dimensional setting, the study of a linear operator A is drastically simplified when we have
 212 access to an orthonormal basis of eigenvectors. Similarly, some operators admit a countable set of
 213 eigenfunctions $(\phi_i)_{i \in \mathbb{N}}$ which forms an orthonormal basis of the Hilbert space \mathcal{H} . When such an
 214 eigensystem exists, the following theorem proven in Appendix B gives a solution to the PDE (8) in
 215 terms of the eigensystem. A rigorous connection between this solution to (8) and solutions to (6) is
 explored in Appendix B.

216 **Theorem 1** Let \mathcal{L} be an essentially self-adjoint¹, densely defined operator on \mathcal{H} which admits an
 217 orthonormal basis of eigenfunctions $(\phi_i, \lambda_i)_{i \in \mathbb{N}}$. Assume further that the λ_i are bounded from below
 218 (write $\lambda_0 \leq \lambda_1 \leq \dots$) and do not have a finite accumulation point. Then a solution to the abstract
 219 evolution problem in (8) is given by

$$220 \quad 221 \quad 222 \quad \psi(\tau) = \sum_{i \in \mathbb{N}} e^{-\lambda_i \tau} \langle \phi_i, \psi_0 \rangle \phi_i. \quad (9)$$

223 3 OUR FRAMEWORK

225 3.1 SPECTRAL PROPERTIES OF THE SCHRÖDINGER OPERATOR

227 In order to apply Theorem 1, we must establish conditions under which the operator \mathcal{L} in (8) satisfies
 228 the required properties. In order for \mathcal{L} to be symmetric, we assume

229 **(A1)** The drift b in (3) is described by a gradient field: $b(x) = -\nabla E(x)$.

231 Define the measure μ on \mathbb{R}^d with density $\mu(x) = \exp(-2\beta E(x))$, and consider the weighted
 232 Lebesgue space $L^2(\mu)$. Note that we do *not* require μ to be a finite measure. Under the assumption
 233 **(A1)**, the operator appearing in (8) becomes the following operator on $L^2(\mu)$:

$$234 \quad \mathcal{L} : D(\mathcal{L}) \subset L^2(\mu) \rightarrow L^2(\mu) : \psi \mapsto \mathcal{L}\psi = -\Delta\psi + 2\beta \langle \nabla E, \nabla\psi \rangle + 2\beta^2 f\psi. \quad (10)$$

235 Under mild regularity conditions, we can further show that \mathcal{L} is essentially self-adjoint (Appendix B).
 236 Furthermore, it can be shown (Appendix B) that

$$237 \quad 238 \quad U\mathcal{L}U^{-1} = -\Delta + \beta^2 \|\nabla E\|^2 - \beta\Delta E + 2\beta^2 f \quad (11)$$

239 where $U : L^2(\mu) \rightarrow L^2(\mathbb{R}^d) : \psi \mapsto e^{-\beta E}\psi$ is a unitary operator, so that \mathcal{L} is unitarily equivalent
 240 to the well-known Schrödinger operator $\mathcal{S} = -\Delta + \mathcal{V}$ on $L^2(\mathbb{R}^d)$, which forms the cornerstone of
 241 the mathematical formulation of quantum physics. Its properties have been studied to great extent
 242 (Reed & Simon, 1978), allowing us to invoke well-known results on the properties of the Schrödinger
 243 operator to study the behavior of \mathcal{L} . In particular, the following assumption on E and f is enough to
 244 guarantee that \mathcal{L} satisfies all the desired properties (see Appendix B).

245 **(A2)** For the energy E and running cost f , $\mathcal{V} := \beta \|\nabla E\|^2 - \Delta E + 2\beta f$ satisfies $\mathcal{V} \in L^2_{loc}(\mathbb{R}^d)$,
 246 $\exists C \in \mathbb{R}, \forall x \in \mathbb{R}^d : C \leq \mathcal{V}(x)$, and $\mathcal{V}(x) \rightarrow \infty$ as $\|x\| \rightarrow \infty$.

247 **Theorem 2** Suppose **(A2)** is satisfied. Then the operator \mathcal{L} in (10) is densely defined and essentially
 248 self-adjoint. Moreover, it admits a countable, orthonormal basis of eigenfunctions. In addition,
 249 the eigenvalues are bounded from below and do not have a finite accumulation point, the lowest
 250 eigenvalue λ_0 has multiplicity one, and the associated eigenfunction (called the ground state, or in
 251 our context the top eigenfunction) can be chosen to be strictly positive.

253 **Remark 1** Previous studies have linked the Schrödinger operator to optimal control in contexts
 254 distinct from ours: for example, Schiitte et al. (2012) and Bakshi et al. (2020) analyze the stationary
 255 HJB equation $\mathcal{K}V_\infty = \lambda$ (see Remark 2), whereas Kalise et al. (2025) explores distribution-level
 256 control of the Fokker–Planck equation.

258 3.2 EIGENFUNCTION CONTROL

260 From the previous discussion, we obtain the following result, which links the eigenfunctions of \mathcal{L}
 261 with the optimal control problem.

262 **Theorem 3** Suppose **(A1)**–**(A2)** are satisfied. Then \mathcal{L} satisfies all assumptions in Theorem 1, hence
 263 the solution of the optimal control problem (3)–(4) is given by

$$264 \quad 265 \quad 266 \quad u^*(x, t) = \beta^{-1} \left(\nabla \log \phi_0(x) + \nabla \log \left(1 + \sum_{i>0} \frac{\langle e^{-\beta g}, \phi_i \rangle_\mu}{\langle e^{-\beta g}, \phi_0 \rangle_\mu} e^{-\frac{1}{2\beta}(\lambda_i - \lambda_0)(T-t)} \frac{\phi_i(x)}{\phi_0(x)} \right) \right). \quad (12)$$

267 where $(\phi_i, \lambda_i)_{i \in \mathbb{N}}$ is the orthonormal eigensystem of \mathcal{L} defined in (10), and $\lambda_0 < \lambda_1 \leq \dots$

268 ¹An operator is called essentially self-adjoint if its closure is self-adjoint. See Reed & Simon (1980) and
 269 Reed & Simon (1975) for more details.

270 Thus, the long-term optimal control ($t \ll T$) is given by $\beta^{-1} \nabla \log \phi_0$, and the corrections decay
 271 exponentially, motivating truncation of the series in (12).
 272

273 3.3 CLOSED-FORM SOLUTION FOR THE SYMMETRIC LQR
 274

275 When E and f are quadratic, the Schrödinger operator associated with the optimal control system is
 276 the Hamiltonian for the harmonic oscillator, which has an exact solution (Appendix B):
 277

278 **Theorem 4** Suppose $b(x) = -Ax$ for some symmetric matrix $A \in \mathbb{R}^{d \times d}$, and $f(x) = x^T Px$ for
 279 some matrix $P \in \mathbb{R}^{d \times d}$ such that $A^T A + 2P$ is positive definite, and has diagonalization $U^T \Lambda U$.
 280 Then the orthonormal eigensystem of the operator \mathcal{L} given in (10) is described through
 281

$$\phi_\alpha(x) = \frac{\exp\left(-\frac{\beta}{2}x^T(-A + U^T \Lambda^{1/2} U)x\right)}{(\lambda\pi)^{d/4}} \prod_{i=1}^d \frac{\Lambda_{ii}^{1/8}}{\sqrt{2^{\alpha_i}(\alpha_i!)}} H_{\alpha_i}\left(\sqrt{\beta}(\Lambda^{1/4} U x)_i\right), \quad (13)$$

$$\lambda_\alpha = \beta \left(-\text{Tr}(A) + \sum_{i=1}^d \Lambda_{ii}^{1/2} (2\alpha_i + 1) \right). \quad (14)$$

282 where $\alpha \in \mathbb{N}^d$ and H_i denotes the i th physicist's Hermite polynomial. We can bijectively map $\alpha \in \mathbb{N}^d$
 283 to $i \in \mathbb{N}$ by ordering the eigenvalues (14), yielding the same representation as before.
 284

285 Combined with Theorem 3, this yields a closed-form solution for the optimal control problem with
 286 symmetric linear drift, quadratic running cost and arbitrary terminal reward.
 287

292 4 NUMERICAL METHODS
 293

294 We propose a hybrid method with two components: Far from the terminal time T , we learn the top
 295 eigenfunction ϕ_0 and simply use $\partial_x \log \phi_0$ as the control (c.f. eq. (2)); Close to the terminal time, e.g.
 296 $t \geq T - 1$, we use an existing solver to learn an additive short-horizon correction to the control.
 297

298 4.1 LEARNING EIGENFUNCTIONS
 299

300 In absence of a closed-form solution, a wide range of numerical techniques exist for solving the
 301 eigenvalue problem for the operator \mathcal{L} in (10). Classically, the eigenfunction equation is projected
 302 onto a finite-dimensional subspace to yield a Galerkin/finite element method, see Chaitin-Chatelin
 303 (1983). In high dimensions, these methods often perform poorly, motivating deep learning approaches
 304 which differ from each other mainly in the loss function used. We will only discuss methods for
 305 learning a single eigenfunction, referring to Appendix C for extensions to multiple eigenfunctions.
 306 An overview of the deep learning algorithm for learning eigenfunctions is given in Algorithm 1 in
 307 Appendix E.
 308

PINN loss Based on the success of physics-informed neural networks (PINNs) (Raissi et al., 2019),
 309 one idea is to design a loss function that attempts to enforce the equation $\mathcal{L}\phi = \lambda\phi$ via an L^2 loss, as
 310 done in Jin et al. (2022). The idea is to consider some density ρ on \mathbb{R}^d , and define the loss function
 311

$$\mathcal{R}_{\text{PINN}}^\rho(\phi) = \|\mathcal{L}[\phi] - \lambda\phi\|_\rho^2 + \alpha\mathcal{R}_{\text{reg}}^\rho(\phi), \quad \mathcal{R}_{\text{reg}}^\rho(\phi) = (\|\phi\|_\rho^2 - 1)^2 \quad (15)$$

312 where $\alpha > 0$ is a regularizer to avoid the trivial solution $\phi = 0$. The eigenvalue λ is typically also
 313 modeled as a trainable parameter of the model, or obtained through other estimation procedures.
 314

315 **Variational loss** A second class of loss functions is based on the variational characterization of the
 316 eigensystem of \mathcal{L} . Since \mathcal{L} is essentially self-adjoint with orthogonal eigenbasis in a subset of $L^2(\mu)$
 317 (μ is defined below (A1)), it holds that (see (Reed & Simon, 1978, Theorem XIII.1))
 318

$$\lambda_0 = \inf_{\psi \in L^2(\mu)} \frac{\langle \psi, \mathcal{L}\psi \rangle_\mu}{\langle \psi, \psi \rangle_\mu} \quad (16)$$

319 where the infimum is obtained when $\mathcal{L}\psi = \lambda_0\psi$. This motivates the loss functions
 320

$$\mathcal{R}_{\text{Var},1}(\phi) = \frac{\langle \phi, \mathcal{L}\phi \rangle_\mu}{\langle \phi, \phi \rangle_\mu} + \alpha\mathcal{R}_{\text{reg}}^\mu(\phi), \quad \mathcal{R}_{\text{Var},2}(\phi) = \langle \phi, \mathcal{L}\phi \rangle_\mu + \alpha\mathcal{R}_{\text{reg}}^\mu(\phi). \quad (17)$$

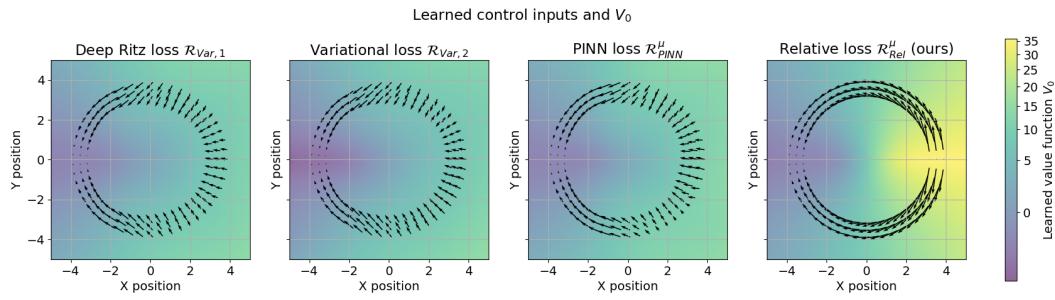


Figure 3: Learned controls (arrows) and V_0 for different eigenfunction losses. Existing methods fail to learn the correct control in regions where V_0 is large due to implicit reweighting.

The first of these, sometimes called the *deep Ritz loss*, was introduced in E & Yu (2018), and the second was described in Cabannes & Bach (2024); Zhang et al. (2022). These loss functions do not require prior knowledge of the eigenvalue λ_0 .

4.1.1 IMPLICIT REWEIGHTING IN PREVIOUS APPROACHES

The exponential decay of the correction term in eq. (2) suggests that the optimal control u^* can be approximated by the spatial derivative of *logarithm* of the top eigenfunction. We therefore parameterize in our implementation $\phi = \exp(-\beta V_0)$, where V_0 is some neural network and β is the temperature constant. This choice also enforces the *strict positivity* of ϕ , which matches the same property of ϕ_0 established in Theorem 2.

Adopting such a parameterization, for the PINN and variational losses, it holds that

$$\mathcal{R}_{\text{PINN}}^{\rho}(e^{-\beta V_0}) = 4\beta^4 \left\| e^{-\beta V_0} \left(\mathcal{K}V_0 - \frac{\lambda_0}{2\beta^2} \right) \right\|_{\rho}^2 + \alpha \mathcal{R}_{\text{reg}}^{\rho}(e^{-\beta V_0}) \quad (18)$$

$$\mathcal{R}_{\text{Var},2}(e^{-\beta V_0}) = 2\beta^2 \int e^{-2\beta V_0} \mathcal{K}V_0 \, d\mu + \alpha \mathcal{R}_{\text{reg}}^{\mu}(e^{-\beta V_0}) \quad (19)$$

where \mathcal{K} is the HJB operator (7). Because both (18) and (19) incorporate an exponential factor that vanishes where V_0 is large, these losses become effectively blind to errors in high- V_0 regions and are only able to learn where V_0 is small. This pathology is illustrated in Figure 3: Consider a 2D RING energy landscape E , whose minimizers lie on a circle, and a cost f that grows linearly with the x -coordinate (see the full setup in Section E). The true optimal control remains tangential to the circle. In contrast, the controls obtained via the PINN and variational eigenfunction losses collapse in regions of large V_0 , deviating sharply from the expected direction.

4.1.2 OUR APPROACH: REMOVING IMPLICIT REWEIGHTING VIA RELATIVE LOSS

Based on this observation, we propose to modify (15) to

$$\mathcal{R}_{\text{Rel}}^{\rho}(\phi) = \left\| \frac{\mathcal{L}\phi}{\phi} - \lambda \right\|_{\rho}^2 + \alpha \mathcal{R}_{\text{reg}}^{\rho}(\phi). \quad (20)$$

This loss function, which we call the *relative loss*, eliminates the implicit reweighting of the stationary HJB equation. Indeed, the same computation as before yields

$$\mathcal{R}_{\text{Rel}}^{\rho}(e^{-\beta V_0}) = 4\beta^4 \left\| \mathcal{K}V_0 - \frac{\lambda_0}{2\beta^2} \right\|_{\rho}^2 + \alpha \mathcal{R}_{\text{reg}}^{\rho}(e^{-\beta V_0}). \quad (21)$$

As a result, the relative loss remains sensitive even in regions where $\phi = e^{-\beta V_0}$ becomes small. This can also be empirically observed in the RING task, as illustrated in Figure 3. Instead of learning V_0 and λ_0 jointly, a natural idea is to combine the benefits of the above loss functions by first training with a variational loss (17) to obtain an estimate for the eigenvalue λ_0 and a good initialization of V_0 , and then ‘fine-tune’ using (20). In practice, we also observed that this initialization is necessary for the relative loss (20) to converge.

378 **Remark 2** We note that there is an alternative interpretation of (21) based on a separate class of
 379 control problems in which there is no terminal cost g , and an infinite-horizon cost is minimized,
 380 yielding a stationary or ergodic optimal control (Kushner, 1978). In this setting, u is related to a
 381 time-independent value function V_∞ which satisfies a stationary HJB equation of the form $\mathcal{K}V_\infty = \lambda$.
 382 In low dimensions, these problems are solved through classical techniques such as basis expansions
 383 or grid-based methods, with no involvement of neural networks (Todorov, 2009).

385 4.2 OUR HYBRID METHOD: COMBINING EIGENFUNCTIONS AND SHORT-HORIZON SOLVERS

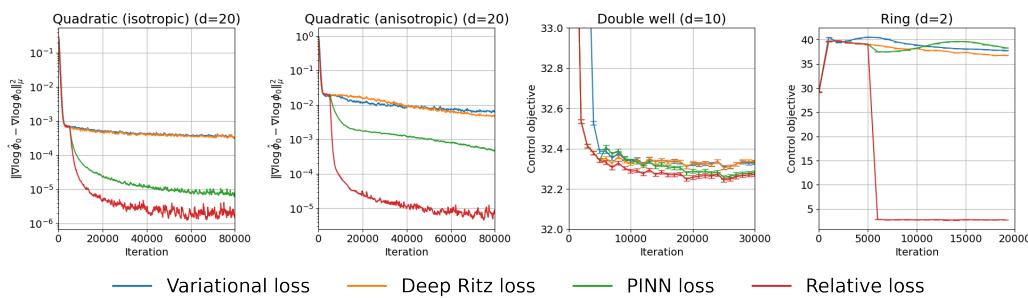
387 For both IDO and FBSDE methods, every iteration requires the numerical simulation of an SDE,
 388 yielding a linear increase in computation cost with the time horizon T . We propose to leverage the
 389 eigenfunction solution given in Theorem 3 in order to scale these methods to longer time horizons as
 390 follows: first, parametrize the top eigenfunction as $\phi_0^{\theta_0} = \exp(-\beta V_0^{\theta_0})$ for a neural network $V_0^{\theta_0}$,
 391 and learn the parameters θ_0 using the relative loss, as well as the first two eigenvalues λ_0, λ_1 (see
 392 Appendix E). Next, choose some cutoff time $T_{cut} < T$ and parametrize the control as

$$393 \quad u_\theta(x, t) = \begin{cases} \beta^{-1} \nabla \log \phi_0^{\theta_0} & 0 \leq t \leq T_{cut}, \\ 394 \quad \beta^{-1} \left(\nabla \log \phi_0^{\theta_0}(x) + e^{-\frac{1}{2\beta}(\lambda_1 - \lambda_0)(T-t)} v^{\theta_1}(x, t) \right) & T_{cut} < t \leq T. \end{cases} \quad (22)$$

396 This control can then be used in an IDO/FBSDE algorithm to optimize the parameters θ_1 of the
 397 additive correction v^{θ_1} , a second neural network, near the terminal time. Crucially, this only requires
 398 simulation of the system in the interval $[T_{cut}, T]$, significantly reducing the overall computational
 399 burden and reducing the time complexity of the algorithm from $\mathcal{O}(Td)$ to $\mathcal{O}(d)$.

401 5 EXPERIMENTS

403 To evaluate the benefits of the proposed method, we consider four different settings, QUADRATIC
 404 (ISOTROPIC), QUADRATIC (ANISOTROPIC), DOUBLE WELL, and RING. An additional setting
 405 QUADRATIC (REPULSIVE) with nonconfining energy is discussed in Appendix E. The first three are
 406 high-dimensional benchmark problems adapted from Nüsken & Richter (2021), modified to be long-
 407 horizon problems, where a ground truth can be obtained. Detailed information on the experimental
 408 setups, including computational costs, is given in Appendix E.



419 420 Figure 4: Comparison of the different eigenfunction losses (EMA).

422 Figure 4 shows the results of the various eigenfunction losses. For the QUADRATIC settings, we can
 423 compute $\nabla \log \phi_0$ exactly, and see that the relative loss significantly improves upon existing loss
 424 functions for approximating this quantity (with the error measured in $L^2(\mu)$). For the other settings,
 425 the resulting control $\nabla \log \phi_0$ yields the lowest value of the control objective for the relative loss.

426 In Figure 5, we show the result of using the learned eigenfunctions in the IDO algorithm using the
 427 combined algorithm described in the previous section, and compare it with the standard IDO/FBSDE
 428 methods. In each setting, we obtain a lower L^2 error using the combined method, typically by an
 429 order of magnitude. The bottom row of Figure 5 shows how the error behaves as a function of
 430 $t \in [0, T]$: the pure eigenfunction method achieves superior performance for $t \rightarrow 0$, but performs
 431 worse closer to the terminal time T . The IDO method has constant performance in $[0, T]$, and the
 combined method combines the merits of both to provide the lowest overall L^2 error.

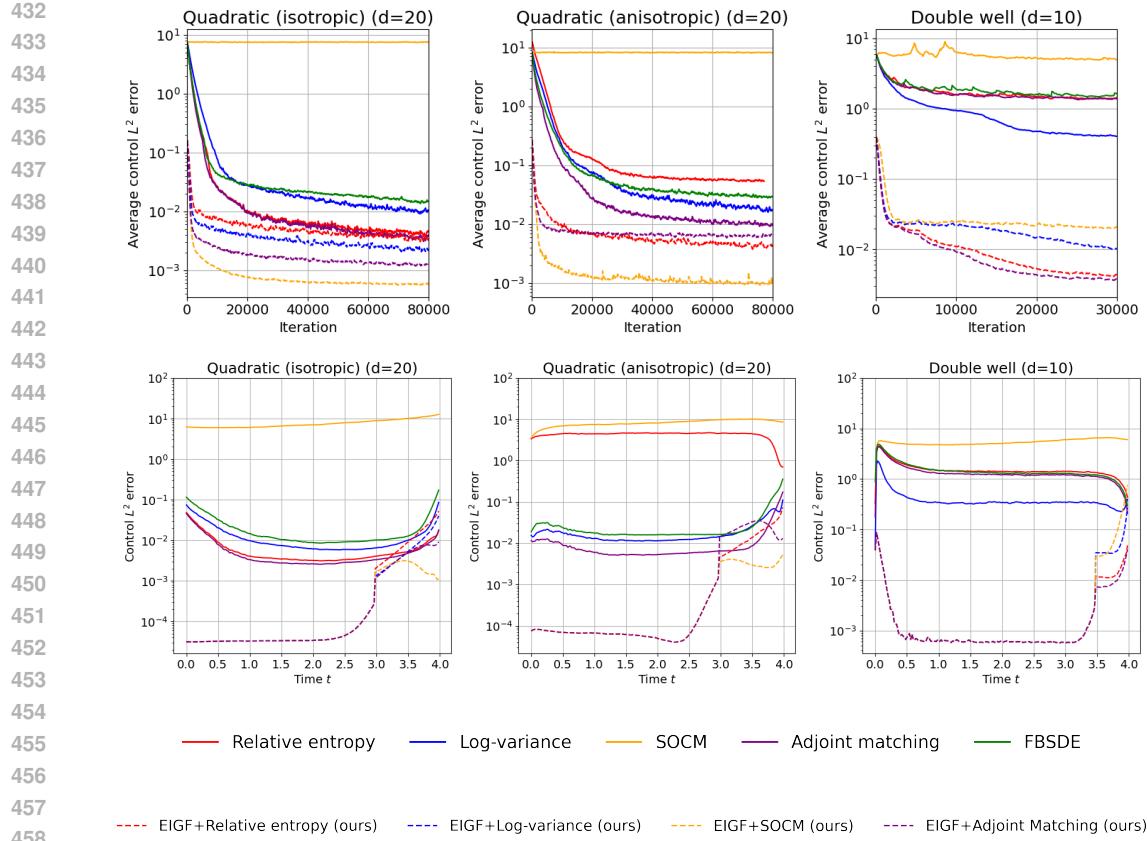


Figure 5: Average L^2 control error (EMA) as a function of iteration (top row) and L^2 error as a function of $t \in [0, T]$ (bottom row).

6 CONCLUSION

In this work, we have introduced a new perspective on a class of stochastic optimal control problems with gradient drift, showing that the optimal control can be obtained from the eigensystem of a Schrödinger operator. We have investigated the use of deep learning methods to learn the eigensystem, introducing a new loss function for this task. We have shown how this approach can be combined with existing IDO methods, yielding an improvement in L^2 error of roughly an order of magnitude over state-of-the-art methods in several long-horizon experiments, and overcoming the increase in computation cost typically associated with longer time horizons.

Limitations The main drawback of the proposed approach is that it is currently limited to problems with gradient drift. When the operator \mathcal{L} is not even symmetric, it may no longer have real eigenvalues. Nonetheless, the top eigenfunction may still be real and nondegenerate with real eigenvalue, so that the long-term behaviour of the control is still described by an eigenfunction (Evans, 2022, Theorem 6.3). A second limitation is that there is no a priori method for determining an appropriate cutoff time T_{cut} , this is a hyperparameter that should be decided based on the application and the spectral gap $\lambda_1 - \lambda_0$.

486 REFERENCES
487

488 Giacomo Albi, Young Pil Choi, Massimo Fornasier, and Dante Kalise. Mean Field Control Hierarchy.
489 *Applied Mathematics and Optimization*, 76(1):93–135, August 2017. ISSN 0095-4616. doi:
490 10.1007/s00245-017-9429-x.

491 Kristoffer Andersson, Adam Andersson, and C. W. Oosterlee. Convergence of a Robust Deep
492 FBSDE Method for Stochastic Control. *SIAM Journal on Scientific Computing*, 45(1):A226–A255,
493 February 2023. ISSN 1064-8275. doi: 10.1137/22M1478057.

494 Rawin Assabumrungrat, Kentaro Minami, and Masanori Hirano. Error Analysis of Option Pricing
495 via Deep PDE Solvers: Empirical Study. In *2024 16th IIAI International Congress on Advanced
496 Applied Informatics (IIAI-AAI)*, pp. 329–336. IEEE Computer Society, July 2024. ISBN 979-8-
497 3503-7790-3. doi: 10.1109/IIAI-AAI63651.2024.00068.

498 Kaivalya Bakshi, Piyush Grover, and Evangelos A Theodorou. On mean field games for agents with
499 langevin dynamics. *IEEE Transactions on Control of Network Systems*, 6(4):1451–1460, 2019.

500 Kaivalya Bakshi, David D Fan, and Evangelos A Theodorou. Schrödinger approach to optimal
501 control of large-size populations. *IEEE Transactions on Automatic Control*, 66(5):2372–2378,
502 2020.

503 Miłosz Baraniewicz. Ground state decay for schrödinger operators with confining potentials. *arXiv
504 preprint arXiv:2407.09267*, 2024.

505 Christian Beck, Weinan E, and Arnulf Jentzen. Machine learning approximation algorithms for high-
506 dimensional fully nonlinear partial differential equations and second-order backward stochastic
507 differential equations. *Journal of Nonlinear Science*, 29(4):1563–1619, August 2019. ISSN
508 0938-8974, 1432-1467. doi: 10.1007/s00332-018-9525-3.

509 Dimitri Bertsekas and Steven E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*.
510 Athena Scientific, December 1996. ISBN 978-1-886529-03-8.

511 Samuel Biton. Nonlinear monotone semigroups and viscosity solutions. In *Annales de l’Institut
512 Henri Poincaré C, Analyse non linéaire*, volume 18, pp. 383–402. Elsevier, 2001.

513 J. Frédéric Bonnans, Élisabeth Ottenwaelter, and Housnaa Zidani. A fast algorithm for the two
514 dimensional HJB equation of stochastic control. *ESAIM: Modélisation mathématique et analyse
515 numérique*, 38(4):723–735, 2004. ISSN 1290-3841. doi: 10.1051/m2an:2004034.

516 Susanne C. Brenner and L. Ridgway Scott. *The Mathematical Theory of Finite Element Methods*,
517 volume 15 of *Texts in Applied Mathematics*. Springer, New York, NY, 2008. ISBN 978-0-387-
518 75933-3 978-0-387-75934-0. doi: 10.1007/978-0-387-75934-0.

519 Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng (eds.). *Handbook of Markov Chain
520 Monte Carlo*. Chapman and Hall/CRC, New York, May 2011. ISBN 978-0-429-13850-8. doi:
521 10.1201/b10905.

522 Vivien Cabannes, Bobak Kiani, Randall Balestrieri, Yann LeCun, and Alberto Bietti. The ssl
523 interplay: Augmentations, inductive bias, and generalization. In *International Conference on
524 Machine Learning*, pp. 3252–3298. PMLR, 2023.

525 Vivien A. Cabannes and Francis Bach. The Galerkin method beats Graph-Based Approaches for
526 Spectral Algorithms. In *Proceedings of The 27th International Conference on Artificial Intelligence
527 and Statistics*, pp. 451–459. PMLR, April 2024.

528 Elisa Calzola, Elisabetta Carlini, Xavier Dupuis, and Francisco Silva. A semi-Lagrangian scheme for
529 Hamilton–Jacobi–Bellman equations with oblique derivatives boundary conditions. *Numerische
530 Mathematik*, 153(1):49–84, January 2023. doi: 10.1007/s00211-022-01336-6.

531 Jan Carius, René Ranftl, Farbod Farshidian, and Marco Hutter. Constrained stochastic optimal control
532 with learned importance sampling: A path integral approach. *The International Journal of Robotics
533 Research*, 41(2):189–209, 2022.

540 E. Carlini, A. Festa, and N. Forcadel. A Semi-Lagrangian Scheme for Hamilton–Jacobi–Bellman
 541 Equations on Networks. *SIAM Journal on Numerical Analysis*, 58(6):3165–3196, January 2020.
 542 ISSN 0036-1429. doi: 10.1137/19M1260931.

543
 544 José A. Carrillo, Edgard A. Pimentel, and Vardan K. Voskanyan. On a mean field optimal control
 545 problem. *Nonlinear Analysis*, 199:112039, October 2020. ISSN 0362-546X. doi: 10.1016/j.na.
 546 2020.112039.

547 Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics.
 548 *Reviews of Modern Physics*, 81(2):591–646, May 2009. doi: 10.1103/RevModPhys.81.591.

549
 550 Françoise Chaitin-Chatelin. *Spectral Approximation of Linear Operators*. Academic Press, 1983.
 551 ISBN 978-0-12-170620-3.

552 Michael G Crandall, Hitoshi Ishii, and Pierre-Louis Lions. User’s guide to viscosity solutions of
 553 second order partial differential equations. *Bulletin of the American mathematical society*, 27(1):
 554 1–67, 1992.

555
 556 Francesca Da Lio and Olivier Ley. Uniqueness results for second-order bellman–isaacs equations
 557 under quadratic growth assumptions and applications. *SIAM journal on control and optimization*,
 558 45(1):74–106, 2006.

559
 560 Francesca Da Lio and Olivier Ley. Convex hamilton-jacobi equations under superlinear growth
 561 conditions on data. *Applied Mathematics & Optimization*, 63(3):309–339, 2011.

562
 563 Edward Brian Davies and Barry Simon. Ultracontractivity and the heat kernel for schrödinger
 564 operators and dirichlet laplacians. *Journal of Functional Analysis*, 59(2):335–395, 1984.

565
 566 Carles Domingo-Enrich. A Taxonomy of Loss Functions for Stochastic Optimal Control, October
 567 2024.

568
 569 Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint Matching:
 570 Fine-tuning Flow and Diffusion Generative Models with Memoryless Stochastic Optimal Control.
 571 In *The Thirteenth International Conference on Learning Representations*, October 2024a.

572
 573 Carles Domingo-Enrich, Jiequn Han, Brandon Amos, Joan Bruna, and Ricky T. Chen. Stochastic
 574 Optimal Control Matching. *Advances in Neural Information Processing Systems*, 37:112459–
 575 112504, December 2024b.

576
 577 Weinan E and Bing Yu. The Deep Ritz Method: A Deep Learning-Based Numerical Algorithm for
 578 Solving Variational Problems. *Communications in Mathematics and Statistics*, 6(1):1–12, March
 579 2018. ISSN 2194-671X. doi: 10.1007/s40304-018-0127-z.

580
 581 Weinan E, Jiequn Han, and Arnulf Jentzen. Deep Learning-Based Numerical Methods for High-
 582 Dimensional Parabolic Partial Differential Equations and Backward Stochastic Differential Equa-
 583 tions. *Communications in Mathematics and Statistics*, 5(4):349–380, December 2017. ISSN
 584 2194-671X. doi: 10.1007/s40304-017-0117-6.

585
 586 Weinan E, Martin Hutzenthaler, Arnulf Jentzen, and Thomas Kruse. Multilevel Picard iterations for
 587 solving smooth semilinear parabolic heat equations. *Partial Differential Equations and Applica-
 588 tions*, 2(6):80, November 2021. ISSN 2662-2971. doi: 10.1007/s42985-021-00089-5.

589
 590 Alexandre Ern and Jean-Luc Guermond. *Theory and Practice of Finite Elements*, volume 159
 591 of *Applied Mathematical Sciences*. Springer, New York, NY, 2004. ISBN 978-1-4419-1918-2
 592 978-1-4757-4355-5. doi: 10.1007/978-1-4757-4355-5.

593
 594 Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2022.

595
 596 Wendell Fleming and Raymond Rishel. *Deterministic and Stochastic Optimal Control*. Springer, New
 597 York, NY, 1975. ISBN 978-1-4612-6382-1 978-1-4612-6380-7. doi: 10.1007/978-1-4612-6380-7.

598
 599 Wendell H Fleming and H Mete Soner. *Controlled Markov processes and viscosity solutions*. Springer,
 600 2006.

594 Emmanuel Gobet, Jean-Philippe Lemor, and Xavier Warin. A regression-based Monte Carlo method
 595 to solve backward stochastic differential equations. *The Annals of Applied Probability*, 15(3):
 596 2172–2202, August 2005. ISSN 1050-5164, 2168-8737. doi: 10.1214/105051605000000412.
 597

598 Alex Gorodetsky, Sertac Karaman, and Youssef Marzouk. High-dimensional stochastic optimal
 599 control using continuous tensor decompositions. *The International Journal of Robotics Research*,
 600 37(2-3):340–377, 2018.

601 David J. Griffiths and Darrell F. Schroeter. *Introduction to Quantum Mechanics*. Cambridge University
 602 Press, 3 edition, 2018.

603 Piyush Grover, Kaivalya Bakshi, and Evangelos A. Theodorou. A mean-field game model for
 604 homogeneous flocking. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(6):061103,
 605 June 2018. ISSN 1054-1500. doi: 10.1063/1.5036663.

606 Jiequn Han and Jihao Long. Convergence of the deep BSDE method for coupled FBSDEs.
 607 *Probability, Uncertainty and Quantitative Risk*, 5(1):5, July 2020. ISSN 2367-0126. doi:
 608 10.1186/s41546-020-00047-w.

609 Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations
 610 using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510,
 611 August 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1718942115.

612 Carsten Hartmann and Christof Schütte. Efficient rare event simulation by optimal nonequilibrium
 613 forcing. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(11):P11004, November
 614 2012. ISSN 1742-5468. doi: 10.1088/1742-5468/2012/11/P11004.

615 Carsten Hartmann, Ralf Banisch, Marco Sarich, Tomasz Badowski, and Christof Schütte. Charac-
 616 terization of Rare Events in Molecular Dynamics. *Entropy*, 16(1):350–376, January 2014. ISSN
 617 1099-4300. doi: 10.3390/e16010350.

618 Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs), June 2023.

619 Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochastic
 620 Optimal Control for Collective Variable Free Sampling of Molecular Transition Paths. *Advances
 621 in Neural Information Processing Systems*, 36:79540–79556, December 2023.

622 Max Jensen and Iain Smears. On the Convergence of Finite Element Methods for Hamilton–Jacobi–
 623 Bellman Equations. *SIAM Journal on Numerical Analysis*, 51(1):137–162, January 2013. ISSN
 624 0036-1429. doi: 10.1137/110856198.

625 Shaolin Ji, Shige Peng, Ying Peng, and Xichuan Zhang. Solving Stochastic Optimal Control Problem
 626 via Stochastic Maximum Principle with Deep Learning Method. *J. Sci. Comput.*, 93(1), October
 627 2022. ISSN 0885-7474. doi: 10.1007/s10915-022-01979-5.

628 Henry Jin, Marios Mattheakis, and Pavlos Protopapas. Physics-informed neural networks for quantum
 629 eigenvalue problems. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pp.
 630 1–8. IEEE, 2022.

631 Kamil Kaleta, Mateusz Kwaśnicki, and József Lőrinczi. Contractivity and ground state domination
 632 properties for non-local schrödinger operators. *Journal of Spectral Theory*, 8(1):165–189, 2018.

633 Dante Kalise, Lucas M. Moschen, Grigorios A. Pavliotis, and Urbain Vaes. A Spectral Approach to
 634 Optimal Control of the Fokker-Planck Equation, March 2025.

635 H J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical
 636 Mechanics: Theory and Experiment*, 2005(11):P11011, November 2005a. ISSN 1742-5468. doi:
 637 10.1088/1742-5468/2005/11/P11011.

638 Hilbert J. Kappen. Linear Theory for Control of Nonlinear Stochastic Systems. *Physical Review
 639 Letters*, 95(20):200201, November 2005b. doi: 10.1103/PhysRevLett.95.200201.

640 H. J. Kushner. Optimality conditions for the average cost per unit time problem with a diffusion
 641 model. *SIAM Journal on Control and Optimization*, 16(2):330–346, 1978. doi: 10.1137/0316021.

648 Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics.
 649 Springer, New York, NY, 2004. ISBN 978-0-387-76369-9 978-0-387-76371-2. doi: 10.1007/
 650 978-0-387-76371-2.

651 Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the Curse of Horizon: Infinite-
 652 Horizon Off-Policy Estimation. In *Advances in Neural Information Processing Systems*, volume 31.
 653 Curran Associates, Inc., 2018.

654 Francis A. Longstaff and Eduardo S. Schwartz. Valuing American Options by Simulation: A Simple
 655 Least-Squares Approach. May 2001.

656 S.K. Mitter. Filtering and stochastic control: A historical perspective. *IEEE Control Systems
 657 Magazine*, 16(3):67–76, June 1996. ISSN 1941-000X. doi: 10.1109/37.506400.

658 Nikolas Nüsken and Lorenz Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs
 659 using neural networks: Perspectives from the theory of controlled diffusions and measures on path
 660 space. *Partial Differential Equations and Applications*, 2(4):48, June 2021. ISSN 2662-2971. doi:
 661 10.1007/s42985-021-00102-x.

662 Nikolas Nüsken and Lorenz Richter. Interpolating Between BSDEs and PINNs: Deep Learning
 663 for Elliptic and Parabolic Boundary Value Problems. *Journal of Machine Learning*, 2(1):31–64,
 664 March 2023. ISSN 2790-203X.

665 Grigorios A. Pavliotis. *Stochastic Processes and Applications: Diffusion Processes, the Fokker-
 666 Planck and Langevin Equations*, volume 60 of *Texts in Applied Mathematics*. Springer, New York,
 667 NY, 2014. ISBN 978-1-4939-1322-0 978-1-4939-1323-7. doi: 10.1007/978-1-4939-1323-7.

668 Huyêñ Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*,
 669 volume 61 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, Heidelberg, 2009.
 670 ISBN 978-3-540-89499-5 978-3-540-89500-8. doi: 10.1007/978-3-540-89500-8.

671 M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning
 672 framework for solving forward and inverse problems involving nonlinear partial differential
 673 equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2018.10.045>. URL <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.

674 M. Reed and B. Simon. *IV: Analysis of Operators*. Methods of Modern Mathematical Physics.
 675 Elsevier Science, 1978. ISBN 9780125850049.

676 Michael Reed and Barry Simon. *II: Fourier analysis, self-adjointness*, volume 2. Elsevier, 1975.

677 Michael Reed and Barry Simon. *Methods of modern mathematical physics: Functional analysis*,
 678 volume 1. Gulf Professional Publishing, 1980.

679 Michael Renardy and Robert C Rogers. *An introduction to partial differential equations*, volume 13.
 Springer Science & Business Media, 2006.

680 Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal Scaling of Discrete Approximations to Langevin
 681 Diffusions. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 60(1):
 682 255–268, 1998. ISSN 1369-7412.

683 Christof Schütte, Stefanie Winkelmann, and Carsten Hartmann. Optimal control of molecular
 684 dynamics using Markov state models. *Mathematical Programming*, 134(1):259–282, August 2012.
 685 ISSN 1436-4646. doi: 10.1007/s10107-012-0547-6.

686 Emanuel Todorov. Eigenfunction approximation methods for linearly-solvable optimal control
 687 problems. In *2009 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement
 688 Learning*, pp. 161–168, March 2009. doi: 10.1109/ADPRL.2009.4927540.

689 Surya T. Tokdar and Robert E. Kass. Importance sampling: A review. *WIREs Computational
 690 Statistics*, 2(1):54–60, 2010. ISSN 1939-0068. doi: 10.1002/wics.56.

702 Claudia Totzeck and René Pinnau. Space mapping-based receding horizon control for stochastic
 703 interacting particle systems: Dogs herding sheep. *Journal of Mathematics in Industry*, 10(1):11,
 704 April 2020. ISSN 2190-5983. doi: 10.1186/s13362-020-00077-1.

705

706 A Tychonoff. Théorèmes d'unicité pour l'équation de la chaleur. *Sb. Math.*, 42(2):199–216, 1935.

707

708 Ramon van Handel. *Stochastic Calculus, Filtering and Stochastic Control*. Lecture Notes for ACM
 709 217: Advanced Topics in Stochastic Analysis, Princeton. 2007.

710

711 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau,
 712 Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt,
 713 Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric
 714 Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas,
 715 Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris,
 716 Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0
 717 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature
 Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

718

719 Jiongmin Yong and Xun Yu Zhou. *Stochastic Controls*. Springer, New York, NY, 1999. ISBN
 978-1-4612-7154-3 978-1-4612-1466-3. doi: 10.1007/978-1-4612-1466-3.

720

721 Sen Zhang, Jian Zu, and Jingqi Zhang. Deep learning method for finding eigenpairs in Sturm-
 722 Liouville eigenvalue problems. *Electronic Journal of Differential Equations*, 2024(01-83):53–17,
 723 September 2024. ISSN 1072-6691. doi: 10.58997/ejde.2024.53.

724

725 Wei Zhang, Tiejun Li, and Christof Schütte. Solving eigenvalue PDEs of metastable diffusion
 726 processes using artificial neural networks. *Journal of Computational Physics*, 465, 2022. doi:
 10.1016/j.jcp.2022.111377.

727

728 Xun Yu Zhou, Jiongmin Yong, and Xunjing Li. Stochastic verification theorems within the framework
 729 of viscosity solutions. *SIAM Journal on Control and Optimization*, 35(1):243–253, 1997.

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756	CONTENTS	
757		
758	1 Introduction	1
759		
760	2 Preliminaries	3
761	2.1 Stochastic optimal control	3
762	2.2 Eigenfunction solutions	4
763		
764	3 Our framework	5
765	3.1 Spectral properties of the Schrödinger operator	5
766	3.2 Eigenfunction control	5
767	3.3 Closed-form solution for the symmetric LQR	6
768		
769	4 Numerical methods	6
770	4.1 Learning eigenfunctions	6
771	4.1.1 Implicit reweighting in previous approaches	7
772	4.1.2 Our approach: Removing implicit reweighting via relative loss	7
773	4.2 Our hybrid method: combining eigenfunctions and short-horizon solvers	8
774		
775	5 Experiments	8
776		
777	6 Conclusion	9
778		
779	A Technical details/assumptions	16
780	A.1 Regularity conditions	16
781	A.2 Simplifying assumption: $\sigma\sigma^T = I$	16
782		
783	B Proofs/derivations	16
784	B.1 Cole-Hopf transformation (8)	16
785	B.2 Proof of Theorem 1	16
786	B.3 Unitary equivalence (11)	18
787	B.4 Essential self-adjointness of \mathcal{L}	18
788	B.5 Proof of Theorem 2	18
789	B.6 Proof of Theorem 4	19
790	B.7 Semigroup and Viscosity solutions	20
791		
792	C Loss functions	23
793	C.1 Existing methods for short-horizon SOC	23
794	C.2 Extending to multiple eigenfunctions	23
795	C.3 Estimating λ_i	24
796	C.4 Empirical loss & sampling	24
797	C.5 Logarithmic regularization	25
798	C.6 FBSDE Loss	25
799	C.7 IDO loss functions	26
800		
801	D Algorithms	28
802		
803	E Experimental Details	29
804	E.1 Setting configurations	29
805	E.2 Model architecture and training	30
806	E.3 Computational cost	31
807	E.4 Further details on the RING setting	31
808	E.5 Details on Figure 1	31
809	E.6 Additional experiment: repulsive potential	32
	E.7 Control objectives	32

810 A TECHNICAL DETAILS/ASSUMPTIONS
811812 A.1 REGULARITY CONDITIONS
813814 Following Fleming & Soner (2006), we make the following assumptions, which guarantee that the
815 SDE (3) has a unique strong solution.816 1. The coefficient b is Lipschitz in and satisfies the *linear growth* condition
817

818
$$\exists C > 0 : \|b(x) - b(y)\| \leq C\|x - y\|. \quad (23)$$

819

820
$$\exists C' > 0 : \|b(x)\| \leq C'(1 + \|x\|). \quad (24)$$

821

822 A.2 SIMPLIFYING ASSUMPTION: $\sigma\sigma^T = I$
823824 Since $\sigma \in \mathbb{R}^{d \times d}$ was assumed invertible, the matrix $\sigma\sigma^T$ is positive definite, and hence there exists a
825 diagonalization $\sigma\sigma^T = U\Lambda U^T$ where $UU^T = I$ and $\Lambda = \text{diag}((\lambda_i)_{i=1}^d)$ for $\lambda_i > 0$. Consider now
826 the change of variables $y = \Lambda^{-1/2}U^T x$. Then it holds that

827
$$\nabla_x \psi = U\Lambda^{-1/2} \nabla_x \psi, \quad \nabla_x^2 \psi = U\Lambda^{-1/2} \nabla_y^2 \psi \Lambda^{-1/2} U^T \quad (25)$$

828

829 and in particular $\text{Tr}(\sigma\sigma^T \nabla_x^2 \psi) = \text{Tr}(U\Lambda U^T \nabla_x^2 \psi) = \text{Tr}(\nabla_y^2 \psi) = \Delta_y \psi$. Thus the PDE (8) can be
830 written in terms of y as

831
$$\partial_t \psi + \left(-\Delta_y - 2\beta b^T U\Lambda^{-1/2} \nabla_y + 2\beta^2 f \right) \psi = 0 \quad (26)$$

832

833 B PROOFS/DERIVATIONS
834835 B.1 COLE-HOPF TRANSFORMATION (8)
836837 Let V denote the value function defined in (5), satisfying the HJB equation (6). The so-called *Cole-Hopf transformation* consists of setting $V = -\beta^{-1} \log \tilde{\psi}$, leading to a linear PDE for $\tilde{\psi}$, which is
838 called the *desirability function*. To obtain the form (8), we set $V(x, t) = -\beta^{-1} \log \psi \left(x, \frac{1}{2\beta}(T-t) \right)$.
839 The derivatives of V and ψ are then related through
840

841
$$\partial_t V = \frac{1}{2\beta^2} \frac{\partial_t \psi}{\psi}, \quad \nabla V = -\beta^{-1} \frac{\nabla \psi}{\psi}, \quad (27)$$

842

843
$$\frac{\partial^2 V}{\partial x_i \partial x_j} = -\beta^{-1} \frac{1}{\psi^2} \left(\psi \frac{\partial^2 \psi}{\partial x_i \partial x_j} - \left(\frac{\partial \psi}{\partial x_i} \frac{\partial \psi}{\partial x_j} \right) \right), \quad (28)$$

844

845
$$\implies \text{Tr}(\sigma\sigma^T \nabla^2 V) = \sum_{i,j} \sigma_{ij}^2 \frac{\partial^2 V}{\partial x_i \partial x_j} = -\beta^{-1} \frac{1}{\psi} \text{Tr}(\sigma\sigma^T \nabla^2 \psi) + \beta^{-1} \frac{1}{\psi^2} \|\sigma^T \nabla \psi\|^2. \quad (29)$$

846

847 Plugging these expressions into (6) gives
848

849
$$\frac{1}{\psi} \left(\frac{1}{2\beta^2} \partial_t \psi - \frac{1}{2\beta^2} \text{Tr}(\sigma\sigma^T \nabla^2 \psi) - \beta^{-1} b^T \nabla \psi \right) + f = 0. \quad (30)$$

850

851 Multiplying with $2\beta^2 \psi$ and combining with the terminal condition $V(x, T) = g(x)$ shows that the
852 linear PDE for ψ given in (8) is equivalent to the HJB equation (6).
853854 B.2 PROOF OF THEOREM 1
855856 We begin with an informal argument. Recall that we are interested in solving the abstract evolution
857 equation
858

859
$$\text{Find } \psi : [0, T] \rightarrow \mathcal{D}(\mathcal{L}) \text{ such that } \begin{cases} \partial_t \psi(t) + \mathcal{L}\psi(t) &= 0, \\ \psi(0) &= \psi_0. \end{cases} \quad (31)$$

860

Now, since $(\phi_i)_{i \in \mathbb{N}}$ forms an orthonormal basis, it holds that $\psi(t) = \sum_{i \in \mathbb{N}} a_i(t) \phi_i$, where $a_i(t) = \langle \psi_i(t), \phi \rangle$. Hence the PDE in (31) becomes

$$\sum_{i \in \mathbb{N}} \left(\frac{da_i(t)}{dt} + \lambda_i a_i(t) \right) \phi_i = 0 \quad (32)$$

after formally interchanging the series expansion and derivatives. Since the ϕ_i are orthogonal, this equation is satisfied if and only if $a_i(t) = a_i(0)e^{-\lambda_i t}$ for each i .

Formally establishing (9) can be achieved through the theory of semigroups. Essentially, we want to formally define the semigroup $(e^{-t\mathcal{L}})_{t \geq 0}$ and show that it forms the solution operator to (31). We begin by recalling the following definition.

Definition 2 *The set of bounded linear operators on \mathcal{H} is given by*

$$\mathcal{B}(\mathcal{H}) = \left\{ \mathcal{L} : \mathcal{H} \rightarrow \mathcal{H} : \mathcal{L} \text{ is linear and } \|\mathcal{L}\|_{op} = \sup_{f \in \mathcal{H}, \|f\|=1} \|\mathcal{L}(f)\| < \infty \right\} \quad (33)$$

The following result, sometimes called the *functional calculus form* of the spectral theorem, allows us to define $h(\mathcal{L})$ for bounded Borel functions h .

Lemma 1 (Reed & Simon, 1980, Theorem VIII.5) Suppose \mathcal{L} is a self-adjoint operator on \mathcal{H} . Then there exists a unique map $\hat{\phi}$ from the bounded Borel functions on \mathbb{R} into $\mathcal{B}(\mathcal{H})$ which satisfies

- If $\mathcal{L}\psi = \lambda\psi$, then $\hat{\phi}(h)\psi = h(\lambda)\psi$.

In particular, when \mathcal{L} admits a countable orthonormal basis of eigenfunctions (ϕ_i) with eigenvalues λ_i , this operator is given by

$$\hat{\phi}(h)\psi = h(\mathcal{L})\psi := \sum_{i \in \mathbb{N}} h(\lambda_i) \langle \phi_i, \psi \rangle \phi_i. \quad (34)$$

The next result makes use of the defined semigroup and leads to the desired representation. Note that we indeed consider a bounded function of the operator, as we consider it only for $t \geq 0$ and assume that the operator is bounded from below.

Theorem 5 (Reed & Simon, 1980, Theorem VIII.7) Suppose \mathcal{L} is self-adjoint and bounded from below, and define $T(t) = e^{-t\mathcal{L}}$ for $t \geq 0$. Then

- $T(0) = I$.
- For every $\phi \in D(\mathcal{L})$, it holds that

$$\left(\frac{d}{dt} T(t)\psi \right) \Big|_{t=0} = \lim_{h \rightarrow 0} \frac{T(0+h)\psi - T(0)\psi}{h} = -\mathcal{L}T(0)\psi = -\mathcal{L}\psi \quad (35)$$

Proof of Theorem 1 Combining (a) and (b) of the Theorem 5 shows that $\psi(t) := T(t)\psi_0$ satisfies

$$\frac{d}{dt} \psi(t) = \left(\frac{d}{dh} T(t+h)\psi_0 \right) \Big|_{h=0} = -\mathcal{L}T(t)\psi_0 = -\mathcal{L}\psi(t), \quad \psi(0) = T(0)\psi_0 = \psi_0, \quad (36)$$

which is exactly the claim of Theorem 1. The last remark we make is that in Theorem 1 we only assume essential self-adjointness of \mathcal{L} , while above results operate with self-adjoint operators. Thus, the above theorems hold true for the closure $\bar{\mathcal{L}}$ of the operator \mathcal{L} . However, note that \mathcal{L} is densely defined and has eigenfunctions which make up an orthonormal basis of both its domain and $L^2(\mu)$. Thus, the same representation (9) as for $\bar{\mathcal{L}}$ holds for \mathcal{L} . \square

918 B.3 UNITARY EQUIVALENCE (11)
919920 First, notice that $U : L^2(\mu) \rightarrow L^2(\mathbb{R}^d) : \psi \mapsto e^{-\beta E} \psi$ is indeed a unitary transformation, since
921

922
$$\forall \psi, \varphi \in L^2(\mu) : \langle U\psi, U\varphi \rangle_{L^2(\mathbb{R}^d)} = \int e^{-2\beta E} \psi \varphi \, dx = \langle \psi, \varphi \rangle_\mu. \quad (37)$$

923

924 To establish equivalence, we compute
925

926
$$\nabla(e^{\beta E} \psi) = e^{\beta E} (\nabla \psi + \beta \nabla E \psi), \quad (38)$$

927

928
$$\Delta(e^{\beta E} \psi) = e^{\beta E} (\Delta \psi + 2\beta \langle \nabla E, \nabla \psi \rangle + (\beta \Delta E + \beta^2 \|\nabla E\|^2) \psi), \quad (39)$$

929

930
$$2\beta \langle \nabla E, \nabla(e^{\beta E} \psi) \rangle = e^{\beta E} (2\beta \langle \nabla E, \nabla \psi \rangle + 2\beta^2 \|\nabla E\|^2 \psi). \quad (40)$$

931 Putting this together gives
932

933
$$\mathcal{L}(U^{-1}\psi) = \mathcal{L}(e^{\beta E} \psi) = e^{\beta E} (-\Delta \psi + \beta^2 \|\nabla E\|^2 \psi - \beta \Delta E \psi) + 2\beta^2 f e^{\beta E} \psi, \quad (41)$$

934 from which the result (11) follows.
935

B.4 ESSENTIAL SELF-ADJOINTNESS OF \mathcal{L} 937 We will first show the following relation
938

939
$$\langle \varphi, \mathcal{L}\psi \rangle_\mu = \langle \nabla \varphi, \nabla \psi \rangle_\mu + 2\beta^2 \langle \varphi, f\psi \rangle_\mu \quad (42)$$

940 from which it is clear that \mathcal{L} is symmetric on $C_0^\infty(\mathbb{R}^d)$. Indeed, using the divergence theorem, for
941 $\psi, \phi \in C_0^\infty(\mathbb{R}^d)$ one obtains that
942

943
$$\langle \psi, -\Delta \varphi \rangle_\mu = - \int \psi \Delta \varphi e^{-2\beta E} \, dx \quad (43)$$

944

945
$$= \int (\langle \nabla \psi, \nabla \varphi \rangle - 2\beta \psi \langle \nabla \varphi, \nabla E \rangle) e^{-2\beta E} \, dx \quad (44)$$

946

947
$$= \langle \nabla \psi, \nabla \varphi \rangle_\mu - 2\beta \langle \psi, \langle \nabla E, \nabla \varphi \rangle \rangle_\mu, \quad (45)$$

948

949 which immediately shows the result.
950951 While for matrices the notions of symmetry and self-adjointness are equivalent, the situation becomes
952 more delicate for general (possibly unbounded) linear operators. In our case we can use the following
953 result on the essential self-adjointness of the Schrödinger operator.
954955 **Lemma 2** (Reed & Simon, 1975, Theorem X.28) *Let $\mathcal{V} \in L^2_{loc}(\mathbb{R}^d)$ with $\mathcal{V} \geq 0$ pointwise. Then
956 $\mathcal{S} = -\Delta + \mathcal{V}$ is essentially self-adjoint on $C_0^\infty(\mathbb{R}^d)$.*
957958 Note, that $C_0^\infty(\mathbb{R}^d)$ is dense in $L^2(\mathbb{R}^d)$. Thus, operator \mathcal{L} is also densely defined as its domain
959 contains $U^{-1}(C_0^\infty(\mathbb{R}^d))$.
960961 The essential self-adjointness of \mathcal{L} follows from the unitary equivalence with the Schrödinger operator
962 as unitary transformation preserves essential self-adjointness of the transformed operator.
963

B.5 PROOF OF THEOREM 2

964 As we have shown, the operator \mathcal{L} is unitarily equivalent to the Schrödinger operator defined through
965

966
$$\mathcal{S} : D(\mathcal{S}) \subset L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d) : h \mapsto -\Delta h + \mathcal{V}h \quad (46)$$

967

968 with $\mathcal{V} = \beta^2 \|\nabla E\|^2 - \beta \Delta E + 2\beta^2 f$. All the properties mentioned in the statement of Theorem 2
969 remain unchanged under unitary equivalence (both the assumptions on the operator and the conclusions
970 regarding the spectral properties). Hence it is sufficient to show the result for the Schrödinger
971 operator \mathcal{S} associated with \mathcal{L} .972 This is precisely the content of the following theorem, proven in Reed & Simon (1978), combined
973 with observations in section B.4.

Theorem 6 (Reed & Simon, 1978, Theorem XIII.67, XIII.64, XIII.47) Suppose $\mathcal{V} \in L^1_{loc}(\mathbb{R}^d)$ is bounded from below and $\lim_{|x| \rightarrow \infty} \mathcal{V}(x) = \infty$. Then the Schrödinger operator $\mathcal{S} = -\Delta + \mathcal{V}$ has compact resolvent, and in particular a purely discrete spectrum and an orthonormal basis of eigenfunctions in $L^2(\mathbb{R}^d)$. The spectrum $\sigma(\mathcal{S})$ is bounded from below and the eigenvalues do not have a finite accumulation point. If in addition $\mathcal{V} \in L^2_{loc}(\mathbb{R}^d)$, the smallest eigenvalue has multiplicity one, and the corresponding eigenfunction can be taken to be strictly positive.

B.6 PROOF OF THEOREM 4

The result relies on the fact that the Schrödinger operator associated with \mathcal{L} has a quadratic potential, for which the eigenfunctions can be computed exactly. Indeed, the following is a standard result that can be found in many textbooks on quantum mechanics (Griffiths & Schroeter, 2018):

Lemma 3 Consider the Schrödinger operator in $d = 1$ with quadratic potential,

$$\mathcal{S} : D(\mathcal{S}) \subset L^2(\mathbb{R}) \rightarrow L^2(\mathbb{R}) : \psi \mapsto -\frac{d^2\psi}{dx^2} + x^2\psi. \quad (47)$$

The eigenvalues of \mathcal{S} are given by $\lambda_n = 2n + 1$, and the normalized eigenfunctions are given by

$$\phi_n(x) = \frac{1}{\pi^{1/4}} \frac{1}{\sqrt{2^n n!}} H_n(x) e^{-x^2/2} \quad (48)$$

where H_n denotes the n -th physicist's Hermite polynomial defined through

$$H_0(x) = 1, \quad H_1(x) = 2x, \quad H_{n+1}(x) = 2xH_n(x) - 2nH_{n-1}(x). \quad (49)$$

These satisfy the generating function relation

$$e^{2xt-t^2} = \sum_{n=0}^{\infty} H_n(x) \frac{t^n}{n!}. \quad (50)$$

This result can be easily extended to higher dimensions:

Lemma 4 Consider now the d -dimensional Harmonic oscillator:

$$\mathcal{S} : D(\mathcal{S}) \subset L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d) : \psi \mapsto -\Delta\psi + x^T A x \psi. \quad (51)$$

where $A \in \mathbb{R}^{d \times d}$ is symmetric and positive definite. Denote $A = U^T \Lambda U$ the diagonalization of A . Then the eigenvalues are indexed by the multi-index $\alpha \in \mathbb{N}^d$, and are given by

$$\lambda_\alpha = \sum_{i=1}^d \Lambda_i^{1/2} (2\alpha_i + 1). \quad (52)$$

The associated normalized eigenfunctions are

$$\phi_\alpha(x) = \frac{1}{\pi^{d/4}} \exp\left(-\frac{1}{2}x^T U^T \Lambda^{1/2} U x\right) \prod_{i=1}^d \frac{\Lambda_i^{1/8}}{\sqrt{2_i^\alpha(\alpha_i)!}} H_{\alpha_i}((\Lambda^{1/4} U x)_i) \quad (53)$$

Proof. We begin by introducing the variable $y = \Lambda^{1/4} U x$. This gives

$$x^T A x = y^T \Lambda^{1/2} y \quad (54)$$

$$\Delta_x \psi = \text{Tr}(\Lambda^{1/2} \nabla_y^2 \psi) \quad (55)$$

so that we can write

$$-\Delta_x \psi + x^T A x \psi = \sum_{i=1}^d \Lambda_i^{1/2} \left(-\frac{\partial^2 \psi}{\partial y^2} + y_i^2 \psi \right) \quad (56)$$

Thus, the d -dimensional harmonic oscillator decouples into d rescaled one-dimensional oscillators. Since $L^2(\mathbb{R}^d) \cong \bigotimes_{i=1}^d L^2(\mathbb{R})$ (after completion), this means that the eigenfunctions are given by the

product of the one-dimensional eigenfunctions. Hence the spectrum is indexed by $\alpha \in \mathbb{N}^d$ and the eigenfunctions are, expressed in the variable y ,

$$\phi_\alpha(y) = \frac{1}{\pi^{d/4}} \prod_{i=1}^d \frac{1}{\sqrt{2^{\alpha_i}(\alpha_i!)}} H_{\alpha_i}(y_i) e^{-y_i^2/2} \quad (57)$$

To get the eigenfunctions in the x variable, we transform back and note that the above expression is normalized as $\int \phi_\alpha(y)^2 dy = 1$, while we need $\int \phi_\alpha(x)^2 dx = 1$. Using the fact that $dy = \det(\Lambda^{1/4})dx$, this yields the final normalized eigenfunctions as

$$\phi_\alpha(x) = \frac{1}{\pi^{d/4}} \exp\left(-\frac{1}{2}x^T U^T \Lambda^{1/2} U x\right) \prod_{i=1}^d \frac{\Lambda_i^{1/8}}{\sqrt{2^{\alpha_i}(\alpha_i!)}} H_{\alpha_i}((\Lambda^{1/4} U x)_i) \quad (58)$$

and the eigenvalues as

$$\lambda_\alpha = \sum_{i=1}^d \Lambda_i^{1/2} (2\alpha_i + 1). \quad (59)$$

□

Now consider the operator \mathcal{L} as defined in (10), with $b(x) = -Ax$ and $f(x) = x^T Px$. Since A is symmetric, it holds that $b = -\nabla E$ with $E(x) = \frac{1}{2}x^T Ax$. It follows that \mathcal{L} is unitarily equivalent with the Schrödinger operator with potential

$$\mathcal{V} = -\beta \Delta E + \beta^2 \|\nabla E\|^2 + 2\beta^2 f = -\beta \text{Tr}(A) + \beta^2 x^T (A^T A + 2P)x. \quad (60)$$

The first term gives a constant shift to the eigenvalues, but otherwise we are precisely dealing with the d -dimensional harmonic oscillator, whose eigensystem is described in Lemma 4. Hence the eigenvalues are precisely (14), and the eigenfunctions of the original operator \mathcal{L} are then obtained by multiplying with $e^{\frac{\beta}{2}x^T Ax}$, which yields (13).

B.7 SEMIGROUP AND VISCOSITY SOLUTIONS

Plan of the section. We work with two linked equations on $\mathbb{R}^d \times [0, T]$: the HJB equation (6) and the linear parabolic PDE (8) obtained from (6) via the (monotone) Cole–Hopf transform. Our goal is to produce the optimal control u^* by first solving the linear PDE, then applying the inverse Cole–Hopf transform to obtain a value function candidate V' . This program raises two issues which have to be addressed:

1. *Nonuniqueness in unbounded domains.* On \mathbb{R}^d (even with the same terminal/boundary data), second-order finite-horizon HJB and linear parabolic equations may admit multiple solutions unless one restricts to an appropriate growth class; see Tychonoff (1935) for a classical example of this phenomenon.
2. *Verification.* To identify u^* from a solution of (6) one uses a verification theorem. These theorems require the candidate V' to be a (sufficiently regular) viscosity solution (Crandall et al., 1992) lying in a class where comparison (hence uniqueness) holds, see (Fleming & Soner, 2006, Section V.9) for an example of such result.

In our approach we fix a specific *semigroup solution* of the linear PDE and use it to define V' via the inverse Cole–Hopf map. Specifically, we have $V'(x, t) = -\beta^{-1} \log \psi(x, \frac{1}{2\beta}(T-t))$, where $\psi(x, t) = (e^{-t\mathcal{L}}\psi_0)(x)$. The central task is therefore to (a) place this V' in a uniqueness class for (6) and (b) confirm the hypotheses of a verification theorem. In the following text we establish sufficient conditions under which the task is solved.

Verification theorem and Comparison (uniqueness) in \mathbb{R}^d under growth constraints. There are several ways to establish that the value function V (5) is a unique solution of HJB equation (6) in some growth class for the given terminal conditions. In general it is obtained with a use of Dynamic Programming Principle. However, this approach is not so simple in the case of unbounded controls and non-smooth viscosity solutions. We refer to (Zhou et al., 1997; Fleming & Soner, 2006; Da Lio & Ley, 2006; 2011) for an overview of ways to establish connection between value function and

1080 viscosity solutions, as well as results regarding existence and uniqueness of viscosity solutions in
 1081 appropriate growth classes. Following Da Lio & Ley (2011) we introduce class of functions $\tilde{\mathcal{C}}_p$. We
 1082 say that a locally bounded function $u : \mathbb{R}^d \times [0; T] \rightarrow \mathbb{R}$ is in the class $\tilde{\mathcal{C}}_p$ if for some $C > 0$ we have
 1083

$$1084 |u(x, t)| \leq C(1 + \|x\|^p), \forall (x, t) \in \mathbb{R}^d \times [0; T]. \quad (61)$$

1085 Now, we can formulate the following

1087 **Theorem 7** (Da Lio & Ley, 2006, Thm. 3.1)+(Da Lio & Ley, 2011, Thm. 3.2) Suppose that in (6)
 1088 we have

1090 • b satisfies (23) and (24), and $\sigma = I$
 1091 • Running cost f satisfies: $\exists C_1 > 0 : \forall x \in \mathbb{R}^d : |f(x)| \leq C_1(1 + \|x\|^p)$
 1093 • Terminal condition satisfies $g \in \tilde{\mathcal{C}}_p$

1095 Then there exists $0 \leq \tau < T$ such that (6) has a unique continuous viscosity solution in $\mathbb{R}^d \times [\tau; T]$
 1096 in the class $\tilde{\mathcal{C}}_p$. Moreover, this unique solution is the value function V (5).

1098 **Remark 3** See discussion in (Da Lio & Ley, 2011, Remark 3.1) why we can't hope to obtain existence
 1099 in Theorem 7 on a whole interval $[0; T]$.

1101 **Growth of semigroup solution** From previous paragraphs it is clear that one way to provide a link
 1102 between proposed V' and solutions of (6) is to establish growth rates on V' . Specifically, following
 1103 Theorem 7 we want to show that $V' \in \tilde{\mathcal{C}}_p$.

1104 There are other ways to establish connection between solutions, see Biton (2001), (Fleming & Soner,
 1105 2006, Section VI). However, we choose to pursue a path related to growth conditions of spectral
 1106 elements of Schrödinger operator (Davies & Simon, 1984; Baraniewicz, 2024).

1108 **Theorem 8** (Davies & Simon, 1984, Thm. 6.1, Thm. 6.3) Let $\mathcal{S} = -\Delta + \mathcal{V}$ on \mathbb{R}^d with ground state
 1109 $\phi_0 > 0$ (normalized in $L^2(\mathbb{R}^d)$) and suppose that there exist $C_1, C_3 > 0$ and $C_2, C_4 \in \mathbb{R}$ such that,
 1110 for all x with $\|x\|$ large,

$$1111 C_3\|x\|^b + C_4 \leq \mathcal{V}(x) \leq C_1\|x\|^a + C_2, \quad \text{where } \frac{a}{2} + 1 < b \leq a.$$

1113 Then \mathcal{S} is *intrinsically ultracontractive* (IUC). Moreover, there exist constants $C_5, C_7 > 0$ and
 1114 $C_6, C_8 \in \mathbb{R}$ such that, as $\|x\| \rightarrow \infty$,

$$1116 C_5\|x\|^{\frac{b}{2}+1} + C_6 \leq -\log \phi_0(x) \leq C_7\|x\|^{\frac{a}{2}+1} + C_8. \quad (62)$$

1118 *Proof.* Intrinsic ultracontractivity under the stated growth is exactly (Davies & Simon, 1984,
 1119 Thm. 6.3). The upper growth bound in (62) is the estimate $-\log \phi_0(x) \leq C\|x\|^{\frac{a}{2}+1}$ stated explicitly
 1120 in (Davies & Simon, 1984, Eq. (6.4)). For the lower bound in (62), take a comparator $W(x) = \tilde{c}\|x\|^b$
 1121 with $0 < \tilde{c} < C_3$; then $W \rightarrow \infty$ and $\mathcal{V} - W \rightarrow \infty$, so the ground states satisfy $\phi_0^{(V)} \leq C\phi_0^{(W)}$
 1122 by the comparison Lemma (Davies & Simon, 1984, Lem. 6.2). (Davies & Simon, 1984, App. B)
 1123 constructs WKB-type barriers for $-\Delta + W$ and gives pointwise upper bounds of the form $\phi_0^{(W)}(x) \leq$
 1124 $C'\|x\|^{-\beta} \exp\{-\kappa\|x\|^{1+\frac{b}{2}}\}$ for large $\|x\|$ (see the JWKB ansatz and subharmonic comparison in
 1125 (Davies & Simon, 1984, App. B, esp. Lem. B.1–B.3)); combining yields the stated lower bound on
 1126 $-\log \phi_0$. \square

1127 **Remark 4** Note that we use ultracontractivity properties of \mathcal{S} which require growth of \mathcal{V} to be at
 1128 least as $\|x\|^{2+\varepsilon}, \varepsilon > 0$. The characterization of operator contractivity properties is fully given in
 1129 (Davies & Simon, 1984, Thm. 6.1) and does not allow for slower orders of growth. This restriction
 1130 on growth of \mathcal{V} is encoded through condition $\frac{a}{2} + 1 < b \leq a$.

1132 **Theorem 9** Let

$$1133 \mathcal{V} = \beta\|\nabla E\|^2 - \Delta E + 2\beta f, \quad \mathcal{S} = U\mathcal{L}U^{-1} = -\Delta + \mathcal{V}.$$

1134 where $U^{-1}f = e^{\beta E}f$. Let $\psi(x, t) = (e^{-t\mathcal{L}}\psi_0)(x)$ and define the candidate value function
 1135

$$1136 \quad V'(x, t) = -\beta^{-1} \log \psi\left(x, \frac{1}{2\beta}(T-t)\right).$$

1138 Assume the hypotheses of Theorem 8 for \mathcal{V} . Finally, let us suppose that there exist $0 < m \leq M$ so
 1139 that

$$1140 \quad m \leq \frac{\psi_0}{U^{-1}\phi_0} \leq M, \quad \forall x \in \mathbb{R}^d, \quad (63)$$

1142 where ϕ_0 is a ground state of Shrödinger operator \mathcal{S} . Then, uniformly for all $t \in [0, T]$, there exist
 1143 constants $C_1, C_3 > 0$ and $C_2, C_4 \in \mathbb{R}$ such that, as $\|x\| \rightarrow \infty$,

$$1144 \quad C_1 \|x\|^{\frac{b}{2}+1} + C_2 \leq V'(x, t) + E(x) \leq C_3 \|x\|^{\frac{a}{2}+1} + C_4. \quad (64)$$

1146 *Proof.* Recall that the ground state of \mathcal{L} is an eigenfunction $\varphi_0 = U^{-1}\phi_0$, where ϕ_0 is a ground state
 1147 of Shrödinger operator \mathcal{S} . Let us introduce the ground state transformed semigroup:
 1148

$$1149 \quad P_t^{\varphi_0}g(x) := \frac{e^{\lambda_0 t}}{\varphi_0(x)}(e^{-t\mathcal{L}}(\varphi_0 g))(x)$$

1151 Thus, we have a representation
 1152

$$1153 \quad \psi(x, t) = (e^{-t\mathcal{L}}\psi_0)(x) = e^{-\lambda_0 t}\varphi_0(x)P_t^{\varphi_0}\left(\frac{\psi_0}{\varphi_0}\right)(x)$$

1155 It is known, that semigroup $P_t^{\varphi_0}g(x)$ is Markov, see Kaleta et al. (2018). Using the fact that $P_t^{\varphi_0}g(x)$
 1156 is Markov and (63) we have the two-sided bound for all $t \in [0; T]$

$$1158 \quad e^{-\lambda_0 t}\varphi_0(x)m \leq (e^{-t\mathcal{L}}\psi_0)(x) \leq e^{-\lambda_0 t}\varphi_0(x)M$$

1159 which leads to a bound
 1160

$$1161 \quad \lambda_0 \frac{T-t}{2\beta^2} - \beta^{-1} \log M - \beta^{-1} \log \phi_0 - E(x) \leq V'(x, t) \\ 1162 \quad \leq \lambda_0 \frac{T-t}{2\beta^2} - \beta^{-1} \log m - \beta^{-1} \log \phi_0 - E(x) \quad (65)$$

1166 Applying growth bounds on ground state of Shrödinger operator from Theorem 8 leads to the desired
 1167 bound. \square

1168 **Remark 5** Note that definition of $\tilde{\mathcal{C}}_p$ (61) requires a growth bound on coordinate x uniformly for
 1169 all $t \in [0; T]$. Unfortunately, analysis based purely on spectral properties of operator leads to
 1170 non-uniform growth bounds depending on t , which blow up when $t \rightarrow 0$ for $e^{-t\mathcal{L}}$ (for example,
 1171 note that constants in (Davies & Simon, 1984, Thm. 3.2) depend on time). For this reason and for
 1172 simplicity we introduced assumption (63) on boundary conditions.

1174 **Relationship between semigroup-based and viscosity solution** Finalising everything in this
 1175 section, we formulate the following result which addresses the questions raised above.

1177 **Theorem 10** Consider HJB equation (6). Let $\mathcal{V} = \beta\|\nabla E\|^2 - \Delta E + 2\beta f$. Assume the following:

- 1179 • $b(x) = -\nabla E(x)$ satisfies (23) and (24), and $\sigma = I$.
- 1180 • \mathcal{V} satisfies Assumption (A2). Moreover, there exist $C_1, C_3 > 0$ and $C_2, C_4 \in \mathbb{R}$ such that,
 1181 for all x with $\|x\|$ large,

$$1183 \quad C_3\|x\|^b + C_4 \leq \mathcal{V}(x) \leq C_1\|x\|^a + C_2, \quad \text{where } \frac{a}{2} + 1 < b \leq a. \quad (66)$$

- 1185 • Let ϕ_0 be the ground state of operator $\mathcal{S} = -\Delta + \mathcal{V}$. Assume $g \in \tilde{\mathcal{C}}_a$ and there exist
 1186 $0 < m \leq M$ so that

$$1187 \quad m \leq \frac{\exp(-\beta g)}{\exp(\beta E(x))\phi_0} \leq M, \quad \forall x \in \mathbb{R}^d, \quad (67)$$

1188 Then there exists $0 \leq \tau < T$ such that function $V'(x, t) = -\beta^{-1} \log \psi(x, \frac{1}{2\beta}(T-t))$, where
 1189 $\psi(x, t) = (e^{-t\mathcal{L}}\psi_0)(x)$ and \mathcal{L} and ψ_0 are from (8), is a unique continuous viscosity solution of (6)
 1190 on $\mathbb{R}^d \times [\tau; T]$ in the class $\tilde{\mathcal{C}}_{a/2+1}$. Moreover, it coincides with value function V (5) on $\mathbb{R}^d \times [\tau; T]$.
 1191

1192 *Proof.* Under Assumption **(A2)** Theorem 6 holds. Theorem 6 ensures that operators \mathcal{S} and \mathcal{L} have
 1193 some good properties and ground state ϕ_0 is positive.
 1194

1195 Under the stated assumptions Theorem 7 and Theorem 9 hold. Under growth assumptions (24) and
 1196 (66) Theorem 9 gives us that $V' \in \tilde{\mathcal{C}}_{a/2+1} \subseteq \tilde{\mathcal{C}}_a$ which has to be the unique viscosity solution in $\tilde{\mathcal{C}}_a$
 1197 given by Theorem 7 on $\mathbb{R}^d \times [\tau; T]$, which coincides with value function V (5) on $\mathbb{R}^d \times [\tau; T]$. \square
 1198

1199 **Remark 6** Note that assumption (67) is realistic under rapid growth of \mathcal{V} , taking in account growth
 1200 bounds on the ground state (62) and $b(x) = \nabla E(x)$ in (24), and does not contradict $g \in \tilde{\mathcal{C}}_a$.
 1201

C LOSS FUNCTIONS

C.1 EXISTING METHODS FOR SHORT-HORIZON SOC

1205 **Grid-based solvers** In low dimensions ($d \leq 3$), classical techniques for numerically solving PDEs
 1206 can be used. These include finite difference (Bonnans et al., 2004) and finite element methods (Jensen
 1207 & Smeears, 2013; Ern & Guermond, 2004; Brenner & Scott, 2008), as well as semi-Langrangian
 1208 schemes (Calzola et al., 2023; Carlini et al., 2020) and multi-level Picard iteration (E et al., 2021).
 1209

1210 **FBSDE solvers** In another line of work, the SOC problem is transformed into a pair of forward-
 1211 backward SDEs (FBSDEs). These are solved through dynamic programming (Gobet et al., 2005;
 1212 Longstaff & Schwartz, 2001) or deep learning methods which parametrize the solution to the FBSDE
 1213 using a neural network (Han et al., 2018; E et al., 2017; Andersson et al., 2023).
 1214

1215 **IDO methods** In recent years, many methods have been proposed which parametrize the control u_θ
 1216 directly, and optimize it by rolling out simulations of the system (3) under the current control. Authors
 1217 of Nüsken & Richter (2021) coined the term *iterative diffusion optimization (IDO)*, arguing that many
 1218 of these methods can be viewed from a common perspective given in Algorithm 2 (Appendix E). This
 1219 class of algorithms contains state-of-the art methods such as SOCM and adjoint matching (Holdijk
 1220 et al., 2023; Domingo-Enrich et al., 2024a). We describe the most commonly used loss functions in
 1221 Appendix C.
 1222

C.2 EXTENDING TO MULTIPLE EIGENFUNCTIONS

1224 **PINN loss** The most common way to extend the PINN loss (15) to multiple eigenfunctions is to
 1225 define

$$1226 \mathcal{R}_{\text{PINN}}^k(\phi) = \sum_{i=0}^k (\|\mathcal{L}[\phi_i] - \lambda_i \phi_i\|_\rho^2 + \alpha_{\text{norm}}(\|\phi_i\|_\rho^2 - 1)^2) + \alpha_{\text{orth}} \sum_{j \neq i} \langle \phi_i, \phi_j \rangle_\mu^2 \quad (68)$$

1227 for $\alpha_{\text{norm}}, \alpha_{\text{orth}} > 0$. Here we have denoted $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^{k+1}$. The main difference is the addition
 1228 of the orthogonal regularization term, which both ensures that the same eigenfunction is not learned
 1229 twice, and attempts to speed up learning by enforcing the known property that eigenfunctions with
 1230 different eigenvalues are orthogonal w.r.t. the inner product $\langle \cdot, \cdot \rangle_\mu$.
 1231

1232 **Variational loss** A similar idea is used to generalize the variational loss (17). The following result
 1233 shows that the variational principle (16) can be extended to multiple eigenfunctions by imposing
 1234 orthogonality.
 1235

1236 **Theorem 11** (Zhang et al., 2022, Theorem 1) Let $k \in \mathbb{N}$, and let \mathcal{L} be a self-adjoint operator with
 1237 discrete spectrum which admits an orthonormal basis of eigenfunctions, and whose eigenvalues are
 1238 bounded below. Furthermore, let $\omega_0 \geq \dots \geq \omega_k > 0$ be real numbers. Then it holds that
 1239

$$1240 \sum_{i=0}^k \omega_i \lambda_i = \inf_{f_0, \dots, f_k \in \mathcal{H}} \sum_{i=0}^k \omega_i \langle f_i, \mathcal{L} f_i \rangle \quad (69)$$

1242 where the infimum is taken over all $(f_i)_{i=0}^k \subset \mathcal{H}$ such that
 1243
 1244

$$\forall i, j \in \{0, \dots, k\} : \langle f_i, f_j \rangle = \delta_{ij} \quad (70)$$

1245 *Proof.* The proof of this result is given in Zhang et al. (2022). \square
 1246
 1247 Based on this result, the following generalization of the variational loss is proposed in Zhang et al.
 1248 (2022):

$$1249 \quad \mathcal{R}_{\text{Var}}^k(\phi) = \sum_{i=0}^k \langle \phi_i, \mathcal{L}\phi_i \rangle + \alpha \left\| \mathbb{E}_\mu [\phi\phi^T] - I \right\|_F^2, \quad (71)$$

1252 where we have written $\|\cdot\|_F$ for the Frobenius norm and $\alpha > 0$. This loss was also studied in
 1253 Cabannes et al. (2023), where it was noted that the minimizers of the variational loss (17)-(71) are
 1254 not obtained at the eigenfunctions. Instead, the following characterization of the minimizers of (71)
 1255 was obtained.

1256 **Lemma 5** (Cabannes et al., 2023, Lemma 2) Suppose that $\mathcal{H} = L^2(\mu)$. Then it holds that
 1257

$$1258 \quad \arg \min_{\psi_0, \dots, \psi_k \in \mathcal{H}} \mathcal{R}_{\text{Var}}^k(\psi) = \left\{ U\tilde{\phi} \mid UU^T = I, \tilde{\phi}_i = \sqrt{\left(1 - \frac{\lambda_i}{2\alpha}\right)_+} \phi_i \right\} \quad (72)$$

1260 where (ϕ_i, λ_i) denotes the orthonormal eigensystem of \mathcal{L} .
 1261

1262 C.3 ESTIMATING λ_i

1264 The main advantage of the variational loss (17), (71) is that it does not require the eigenvalues of \mathcal{L}
 1265 to be available. However, Lemma 5 shows that the minimizers of $\mathcal{R}_{\text{Var}}^k$ do not coincide exactly with
 1266 the eigenfunctions of \mathcal{L} , so we cannot naively compute $\langle \phi, \mathcal{L}\phi \rangle_\mu$ to obtain the eigenvalues. Instead,
 1267 the following lemma shows how to obtain the eigenvalues and eigenfunctions from an element in the
 1268 minimizing set (72).

1269 **Lemma 6** Suppose it holds that $\psi = U\tilde{\phi}$, where $UU^T = I$ and $\tilde{\phi}_i = \sqrt{\left(1 - \frac{\lambda_i}{2\alpha}\right)} \phi_i$ for each i , and
 1270 $\alpha > \frac{\lambda_k}{2}$. Then the first $k+1$ eigenfunctions and eigenvalues are given by
 1271

$$1272 \quad \phi = D^{-1/2}U^T\psi, \quad \lambda_i = \frac{2}{\beta}(1 - D_{ii}) \quad (73)$$

1274 where $UDU^T = \mathbb{E}_\mu[\psi\psi^T]$ is the diagonalization of the second moment matrix of ψ .
 1275

1276 *Proof.* By definition of ψ , we have

$$1277 \quad \mathbb{E}_\mu[\psi\psi^T] = U\mathbb{E}_\mu[\tilde{\phi}\tilde{\phi}^T]U^T \quad (74)$$

$$1279 \quad = U \left(I - \frac{1}{2\alpha} \Lambda \right) U^T \quad (75)$$

1281 where $\Lambda = \text{diag}(\lambda_i)$ and we have used the orthonormality property of the eigenfunctions. From this,
 1282 we obtain that

$$1283 \quad D = I - \frac{\beta}{2} \Lambda, \quad \psi = UD^{1/2}\phi \quad (76)$$

1285 which concludes the proof. \square

1286 In light of this result, we can estimate the second moment matrix $\mathbb{E}_\mu[\psi\psi^T]$, apply a diagonalization
 1287 algorithm, and obtain the eigenfunctions and eigenvalues using (73).

1289 C.4 EMPIRICAL LOSS & SAMPLING

1290 **Rewriting variational loss** Recalling the equation (42),
 1291

$$1292 \quad \langle \varphi, \mathcal{L}\psi \rangle_\mu = \langle \nabla\varphi, \nabla\psi \rangle_\mu + 2\beta^2 \langle \varphi, f\psi \rangle_\mu, \quad (77)$$

1293 we see that it is possible to evaluate inner products of the form $\langle \varphi, \mathcal{L}\psi \rangle_\mu$ by only evaluating φ, ψ and
 1294 its derivatives. This avoids the expensive computation of the second order derivatives of the neural
 1295 network, making the variational losses less memory-intensive than the other loss functions, which
 requires explicit computation of $\mathcal{L}\phi$.

1296 **Estimation of inner products** All of the loss functions discussed for learning eigenfunctions
 1297 contain inner products of the form $\langle \psi, \phi \rangle_\mu$. To obtain an empirical loss, these quantities must be
 1298 approximated. When μ is a density, this is done using a Monte Carlo estimate
 1299

$$1300 \quad \langle \varphi, \psi \rangle_\mu = \int \varphi \psi \, d\mu = \mathbb{E}_\mu[\varphi \psi] \approx \frac{1}{m} \sum_{i=1}^m \varphi(X_i) \psi(X_i), \quad X_i \sim \mu. \quad (78)$$

1302 Since $\mu(x) \propto \exp(-2\beta E(x))$, we can employ Markov Chain Monte Carlo (MCMC) techniques in
 1303 order to obtain the samples (X_i) (Brooks et al., 2011). In particular, when training the eigenfunctions,
 1304 we can store m samples in memory and apply some number of MCMC steps in each iteration to
 1305 update these samples. Alternatively, one can pre-sample a large dataset of samples from μ .
 1306

1307 **Non-confining energy.** In general, μ may not be a finite measure, for instance when we have a
 1308 repulsive LQR ($E = -\frac{1}{2}\|x\|^2$). In this case, the inner products can no longer be estimated directly
 1309 from samples, but may still be computed by using importance sampling techniques (Liu, 2004; Tokdar
 1310 & Kass, 2010). In the case where the measure defined through $\mu(x) = \exp(-2\beta E(x))$ is not finite
 1311 but $\bar{\mu}(x) = \exp(2\beta E(x))$ is, we can apply importance sampling with $\bar{\mu}$, so that the inner products
 1312 are obtained as
 1313

$$\langle \varphi, \psi \rangle_\mu = \int \varphi \psi e^{-2\beta E(x)} \, dx = \mathbb{E}_{\bar{\mu}} [\bar{\varphi} \bar{\psi}] \quad (79)$$

1314 where we have defined $\bar{\varphi} = e^{-2\beta E} \varphi$ and $\bar{\psi} = e^{-2\beta E} \psi$. For stable training, it is then advisable to
 1315 parametrize $\bar{\phi} := \phi e^{-2\beta E}$ instead of ϕ .
 1316

1317 C.5 LOGARITHMIC REGULARIZATION

1319 As discussed in the main text, the PINN and relative eigenfunction losses (15)-(20) have the form
 1320

$$\mathcal{R}(\phi) = \mathcal{R}_{\text{main}}(\phi) + \alpha \mathcal{R}_{\text{reg}}(\phi), \quad \mathcal{R}_{\text{reg}}(\phi) = (\|\phi\|_\rho^2 - 1)^2. \quad (80)$$

1322 for some $\alpha > 0$. We observe in our experiments that this regularizer is sometimes not strong enough,
 1323 and the network may still converge to $\phi = 0$. For this reason, we instead use a logarithmic regularizer
 1324

$$\mathcal{R}_{\text{reg}}(\phi) = (\log \|\phi\|_\rho)^2 \quad (81)$$

1325 The behaviour is exactly the same as before for $\|\phi\|_\rho \approx 1$, since $\log(1 + x) = x + \mathcal{O}(x^2)$ as $x \rightarrow 0$,
 1326 but this regularizer avoids the convergence to $\phi = 0$.
 1327

1328 C.6 FBSDE LOSS

1330 We briefly discuss the Robust FBSDE loss for stochastic optimal control introduced in Andersson
 1331 et al. (2023). The main idea is that the solution to the HJB equation (6) can be written down as a pair
 1332 of SDEs, as the following lemma illustrates.

1333 **Lemma 7** Suppose $\sigma = I$. Then it holds that the solution to the pair of FBSDEs

$$1335 \quad dX_t = (b(X_t, t) - Z_t) dt + \sqrt{\lambda} dW_t, \quad X_0 \sim p_0, \quad (82)$$

$$1336 \quad dY_t = \left(-f(X_t, t) - \frac{1}{2} \|Z_t\|^2 \right) dt + \sqrt{\lambda} \langle Z_t, dW_t \rangle, \quad Y_T = g(X_T) \quad (83)$$

1338 is given by $Z_t = \partial_x V(X_t, t)$ and $Y_t = V(X_t, t)$.
 1339

1340 This pair of SDEs can be transformed in a variational formulation which is amenable to deep learning,
 1341 as described in detail in Andersson et al. (2023). Using the Markov property of the FBSDE, one can
 1342 show that $Z_t = \zeta(t, X_t)$ for some function $\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, and that the FBSDE problem can
 1343 be reformulated as the following variational problem:

$$1344 \quad \begin{cases} \underset{\zeta}{\text{minimize}} \quad \Psi_\alpha(\zeta) = \mathbb{E}[Y_0^\zeta] + \alpha \mathbb{E} \left[|Y_T^\zeta - g(X_T^\zeta)| \right], \quad \text{where} \\ 1345 \quad \mathcal{Y}_0^\zeta = g(X_T^\zeta) + \int_0^T \left(f(X_s^\zeta, t) + \frac{1}{2} \|Z_s^\zeta\|^2 \right) dt - \int_0^T \langle Z_s^\zeta, dW_t \rangle \\ 1346 \quad X_t^\zeta = X_0 + \int_0^t (b(s, X_s^\zeta) - Z_s^\zeta) dt + \int_0^t \lambda dW_s \\ 1347 \quad Y_t^\zeta = \mathbb{E}[Y_0^\zeta] - \int_0^t (f(X_s^\zeta, s) + \frac{1}{2} \|Z_s^\zeta\|^2) ds + \int_0^t \langle Z_s^\zeta, dW_s \rangle, \quad Z_t^\zeta = \zeta(t, X_t^\zeta), \quad t \in [0, T] \end{cases} \quad (84)$$

1350 Parametrizing ζ as a neural network, we can then simulate the stochastic processes in (84) and define
 1351 the loss function as

$$1353 \quad \mathcal{R}_{FBSDE}(u; X^u) = \mathbb{E}[\mathcal{Y}_0^\zeta] + \alpha \mathbb{E} \left[|Y_T^\zeta - g(X_T^u)| \right], \quad \text{with } \zeta = -u \quad (85)$$

1354 For more details we refer to the relevant work Andersson et al. (2023).
 1355

1356 C.7 IDO LOSS FUNCTIONS

1358 The IDO algorithm described in Algorithm 2 (Appendix D) is rather general, and a large number of
 1359 algorithms can be obtained by specifying different loss functions. For a detailed discussion of the
 1360 various loss functions and the relations between the resulting algorithms, we refer to Domingo-Enrich
 1361 (2024); Domingo-Enrich et al. (2024b). Here, we go over the loss functions that were used for the
 1362 experiments in this paper.

1363 **Adjoint loss** The most straightforward choice of loss is to simply use the objective of the control,
 1364 and define

$$1366 \quad \mathcal{R}(u; X^u) = \int_0^T \left(\frac{1}{2} \|u(X_t^u, t)\|^2 + f(X_t^u) \right) dt + g(X_T^u). \quad (86)$$

1368 This loss is also called the *relative entropy loss*. When converting this to an empirical loss, there
 1369 are two options. The *discrete adjoint method* consists of first discretizing the objective and then
 1370 differentiating w.r.t. the parameters. However, this requires storing the numerical solver in memory
 1371 and can hence be quite memory-intensive.

1372 The *continuous adjoint method* instead analytically computes derivatives w.r.t. the state space. To
 1373 this end, define the adjoint state as

$$1374 \quad a(t, X^u; u) := \nabla_{X_t} \left(\int_t^T \left(\frac{1}{2} \|u(X_t^u, t)\|^2 + f(X_t^u) \right) dt' + g(X_T^u) \right), \quad (87)$$

1377 where X^u solves (3). Then the dynamics of a can be computed as

$$1378 \quad da_t = \left(\nabla_{X_t^u} (b(X_t^u, t) + \sigma(t)u(X_t^u, t)) \right)^T a(t; X^u, u) dt \\ 1379 \quad + \left(\nabla_{X_t^u} (f(X_t^u, t) + \frac{1}{2} \|u(X_t^u, t)\|^2) \right) dt, \quad (88)$$

$$1382 \quad a(T; X^u, u) = \nabla g(X_T^u). \quad (89)$$

1383 The path of a is obtained by solving the above equations backwards in time given a trajectory (X_t^u)
 1384 and control u . Finally, the derivative of the relative entropy loss (86) is then computed as

$$1386 \quad \partial_\theta \mathcal{R} = \frac{1}{2} \int_0^T \frac{\partial}{\partial \theta} \|u(X_t^{\bar{u}}, t)\|^2 dt + \int_0^T \left(\frac{\partial u(X_t^{\bar{u}}, t)}{\partial \theta} \right)^T \sigma(t)a(t; X_t^{\bar{u}}, \bar{u}) dt, \quad (90)$$

1388 where the notation \bar{u} means that we do stop gradients w.r.t. θ from flowing through these values.
 1389 Both the discrete and continuous adjoint method have been shown to work well in practice (Nüsken
 1390 & Richter, 2021; Bertsekas & Shreve, 1996; Domingo-Enrich, 2024). However, their training can be
 1391 unstable due to the non-convexity of the problem.

1392 **Variance and log-variance** For the second class of loss functions, let (X_t^v) denote the solution to
 1394 (3), with u replaced by v . Then we can define

$$1395 \quad \tilde{Y}_T^{u,v} = - \int_0^T (u \cdot v)(X_s^v, s) ds - \int_0^T f(X_s^v, s) ds - \int_0^T u(X_s^v, s) \cdot dW_s + \int_0^T \|u(X_s^v, s)\|^2 ds. \quad (91)$$

1398 The *variance* and *log-variance* loss are then defined as

$$1399 \quad \mathcal{R}_{\text{Var}}(u) = \text{Var}(e^{\tilde{Y}_T^{u,v} - g(X_T^v)}), \quad (92)$$

$$1400 \quad \mathcal{R}_{\text{log-var}}(u) = \text{Var}(\tilde{Y}_T^{u,v} - g(X_T^v)). \quad (93)$$

1402 It can be shown that these losses are minimized when $u = u^*$, irrespective of the choice of v . In
 1403 Nüsken & Richter (2021), it was shown that these loss functions are closely connected to the FBSDE
 formulation of the SOC problem.

SOCM The Stochastic Optimal Control Matching (SOCM) loss, introduced in Domingo-Enrich et al. (2024b), is one of the most recently introduced IDO losses. It is described in the following theorem, where we adapt the notation $\lambda = \beta^{-1}$:

Theorem 12 (Domingo-Enrich et al., 2024b, Theorem 1) For each $t \in [0, T]$, let $M_t : [t, T] \rightarrow \mathbb{R}^{d \times d}$ be an arbitrary matrix-valued differentiable function such that $M_t(t) = \text{Id}$. Let $v \in \mathcal{U}$ be an arbitrary control. Let $\mathcal{R}_{\text{SOCM}} : L^2(\mathbb{R}^d \times [0, T]; \mathbb{R}^d) \times L^2([0, T]^2; \mathbb{R}^{d \times d}) \rightarrow \mathbb{R}$ be the loss function defined as

$$\mathcal{R}_{\text{SOCM}}(u, M) := \mathbb{E} \left[\frac{1}{T} \int_0^T \|u(X_t^v, t) - w(t, v, X^v, W, M_t)\|^2 dt \times \alpha(v, X^v, W) \right], \quad (94)$$

where X^v is the process controlled by v (i.e., $dX_t^v = (b(X_t^v, t) + \sigma(t)v(X_t^v, t)) dt + \sqrt{\lambda}\sigma(t)dW_t$ and $X_0^v \sim p_0$), and

$$\begin{aligned} w(t, v, X^v, W, M_t) &= \sigma(t)^\top \left(- \int_t^T M_t(s) \nabla_x f(X_s^v, s) ds - M_t(T) \nabla g(X_T^v) \right. \\ &\quad + \int_t^T (M_t(s) \nabla_x b(X_s^v, s) - \partial_s M_t(s)) (\sigma^{-1})^\top(s) v(X_s^v, s) ds \\ &\quad \left. + \lambda^{1/2} \int_t^T (M_t(s) \nabla_x b(X_s^v, s) - \partial_s M_t(s)) (\sigma^{-1})^\top(s) dW_s \right), \\ \alpha(v, X^v, B) &= \exp \left(- \lambda^{-1} \int_0^T f(X_t^v, t) dt - \lambda^{-1} g(X_T^v) \right. \\ &\quad \left. - \lambda^{-1/2} \int_0^T \langle v(X_t^v, t), dW_t \rangle - \frac{\lambda^{-1}}{2} \int_0^T \|v(X_t^v, t)\|^2 dt \right). \end{aligned} \quad (95)$$

$\mathcal{L}_{\text{SOCM}}$ has a unique optimum (u^*, M^*) , where u^* is the optimal control.

The proof hinges on the path integral representation described in Kappen (2005a) and using a reparametrization trick to compute its gradients. We refer the reader to the relevant work Domingo-Enrich et al. (2024b) for more details. The reason why this method outperforms other IDO losses is as follows: the loss (94) can be seen as minimizing the discrepancy between u and a target vector field w . The addition of the parametrized matrix M allows the variance of this weight to be reduced, making it easier to learn. The downside of this method is that the variance of the importance weight α can blow up in more complex settings or with poor initialization, with the method failing to converge as a result.

Adjoint Matching The most recently introduced IDO loss, *adjoint matching*, was proposed in Domingo-Enrich et al. (2024a), and was proposed in the context of finetuning diffusion models. It is based on two observations. Firstly, one can write down a regression objective that does not have an importance weighting α by using the adjoint state a defined earlier.

Lemma 8 (Domingo-Enrich et al., 2024a, Proposition 2) Define the basic adjoint matching objective as

$$\mathcal{R}_{\text{Basic-Adj-Match}}(u; X^u) = \frac{1}{2} \int_0^T \|u(X_t, t) + \sigma(t)^T a(t; X^u, \bar{u})\|^2 dt, \quad \bar{u} = \text{stopgrad}(u). \quad (96)$$

where $\bar{u} = \text{stopgrad}(u)$ means that the gradients of \bar{u} w.r.t. the parameters θ of the control u are artificially set to zero. Then the gradient of this loss w.r.t. θ is equal to (90), the gradient of the loss in the continuous adjoint method. Consequently, the only critical point of $\mathbb{E}_{X^u \sim \mathbb{P}^u} [\mathcal{R}_{\text{Basic-Adj-Match}}]$ is the optimal control u^* .

Secondly, it is observed that some terms in the SDE for the adjoint state (88) have expectation zero under the trajectories of the optimal control. Indeed, it holds by definition of the value function and adjoint state that

$$\nabla V(x, t) = \mathbb{E}[a(t, X^{u^*}, u^*) \mid X_t = x] \quad (97)$$

1458 and hence, since $u^* = -\sigma \nabla V$, we get
 1459

$$\mathbb{E}_{X \sim \mathbb{P}^{u^*}} [u^*(x, t)^T \nabla_x u^*(x, t) + a(t, X, u^*)^T \sigma(t) u^*(x, t) \mid X_t = x] = 0 \quad (98)$$

1460 This motivates dropping the terms with expectation zero in (88), yielding the loss function
 1461

$$\mathcal{R}_{Adj-Match}(u; X) = \frac{1}{2} \int_0^T \|u(X_t, t) + \sigma(t) \tilde{a}(t; X, \bar{u})\|^2 dt, \quad \bar{u} = \text{stopgrad}(u), \quad (99)$$

$$\text{where } \frac{d}{dt} \tilde{a}(t; X) = -(\tilde{a}(t; X)^T \nabla_x b(X_t, t) + \nabla_x f(X_t, t)), \quad (100)$$

$$\tilde{a}(T, X) = \nabla_x g(X_T). \quad (101)$$

1462 The value \tilde{a} is called the *lean adjoint state*. The claimed benefit of this method is that the resulting
 1463 loss is a simple least-squares regression objective with no importance weighting, allowing it to avoid
 1464 the problem of high variance importance weights. We refer to the original work Domingo-Enrich
 1465 et al. (2024a) for a more in-depth discussion and proofs related to the adjoint matching loss.
 1466

D ALGORITHMS

1467 We give an overview of the algorithm used for eigenfunction learning in Algorithm 1, and for the
 1468 IDO method in Algorithm 2.

Algorithm 1 Deep learning for eigenfunctions

- 1: Parametrize the eigenfunctions $(\phi_i^{\theta_i})$, choose loss functions (\mathcal{R}_i) .
- 2: Fix a batch size m , number of iterations N , regularization $\alpha > 0$, learning rate $\eta > 0$.
- 3: Generate m samples $(X_i)_{i=1}^m$ from μ , for instance using an MCMC scheme.
- 4: **for** $n = 0, \dots, N - 1$ **do**
- 5: (optional) Update the samples $(X_i)_{i=1}^m$ using a sampling algorithm.
- 6: Compute an m -sample Monte-Carlo estimate $\widehat{\mathcal{R}}_i(\phi_i^{\theta_i})$ of the loss $\mathcal{R}_i(\phi_i^{\theta_i})$.
- 7: Compute the gradients $\nabla_{\theta_i} \widehat{\mathcal{R}}_i(\phi_i^{\theta_i})$ for the current parameters θ_i , and update θ_i using Adam.
- 8: **end for**
- 9: Estimate the eigenvalues λ_i from the learned functions $\phi_i^{\theta_i}$
- 10: **return** eigenfunction estimates $(\phi_i^{\theta_i})$, eigenvalue estimates (λ_i) .

Algorithm 2 Iterative Diffusion Optimization (IDO)

- 1: Parametrize the control $u_\theta \in \mathcal{U}, \theta \in \mathbb{R}^p$, and choose a loss function \mathcal{R} .
- 2: Fix a batch size m , number of iterations N , number of timesteps K , learning rate $\eta > 0$.
- 3: **for** $i = 0, \dots, N - 1$ **do**
- 4: Simulate m trajectories of (3) with control $u = u_\theta$ using a K -step discretization scheme.
- 5: Compute an m -sample Monte-Carlo estimate $\widehat{\mathcal{R}}(u_\theta)$ of the loss $\mathcal{R}(u_\theta)$.
- 6: Compute the gradients $\nabla_\theta \widehat{\mathcal{R}}(u_\theta)$ for the current parameters θ and update using Adam.
- 7: **end for**
- 8: **return the learned control** $u_\theta \approx u^*$.

1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

1512 **E EXPERIMENTAL DETAILS**
15131514 **E.1 SETTING CONFIGURATIONS**
15151516 To evaluate the learned controls, we will use three different metrics:
15171518 1. The *control objective* is the value we are trying to minimize,
1519

1520
$$\mathbb{E} \left[\int_0^T \left(\frac{1}{2} \|u(X_t^u, t)\|^2 + f(X_t^u) \right) dt + g(X_T^u) \right]. \quad (102)$$

1521

1522 It can be estimated by simulating trajectories of (3). Unless mentioned otherwise, we
1523 estimate it using 65536 trajectories and report the standard deviation of the estimate.
15241525 2. The *control L^2 error at time t* is given by
1526

1527
$$\mathbb{E}_{x \sim \mathbb{P}_t^{u^*}} [\|u(x, t) - u^*(x, t)\|^2], \quad (103)$$

1528

1529 where $\mathbb{P}_t^{u^*}$ denotes the density over \mathbb{R}^d induced by the trajectory (3) under the optimal
1530 control at time t .
15311532 3. The *average control L^2 error* is the above quantity averaged over the entire trajectory,
1533

1534
$$\mathbb{E}_{t \sim [0, T]} \mathbb{E}_{x \sim \mathbb{P}_t^{u^*}} [\|u(x, t) - u^*(x, t)\|^2]. \quad (104)$$

1535

1536 These are also the metrics reported in previous works Domingo-Enrich et al. (2024b); Domingo-
1537 Enrich (2024); Nüsken & Richter (2021). The L^2 error is introduced because, after the control u is
1538 sufficiently close to the optimal control u^* , the expectation (102) requires increasingly more Monte
1539 Carlo samples to distinguish u from the optimal control. In this case the L^2 error is a more precise
1540 metric for determining how close a given control is to the optimal control.
15411542 For all experiments, we trained the neural networks involved using Adam with learning rate $\eta = 10^{-4}$.
1543 For the IDO methods, we use a batch size of $m = 64$. For the eigenfunction learning, we sample
1544 from μ using the Metropolis-Adjusted Langevin Algorithm (Roberts & Rosenthal, 1998). Unless
1545 mentioned otherwise, we sample $m = 65536$ samples, updating these in every iteration with 100
1546 MCMC steps with a timestep size of $\Delta t = 0.01$ after a warm-up phase of 1000 steps. For more
1547 details, we refer to the code in the supplementary material, which contains a description of all
1548 hyperparameters used.
15491550 **QUADRATIC** The first setting we consider has
1551

1552
$$E(x) = \frac{1}{2} x^T A x, \quad f(x) = x^T P x, \quad g(x) = x^T Q x, \quad (105)$$

1553

1554 where $A \in \mathbb{R}^{d \times d}$ is symmetric, $Q \in \mathbb{R}^{d \times d}$ is positive definite and $P \in \mathbb{R}^{d \times d}$. This type of control
1555 problem is more widely known as the linear quadratic regulator. The optimal control is given by
1556 $u^*(x, t) = -2F_t x$, where F_t solves the Riccati equation (see (van Handel, 2007, Theorem 6.5.1))
1557

1558
$$\frac{dF_t}{dt} - A^T F_t - F_t A - 2F_t F_t^T + P = 0, \quad F_T = Q. \quad (106)$$

1559

1560 We consider three different configurations:
15611562

- (ISOTROPIC) We set $d = 20$, $A = I$, $P = I$, $Q = 0.5I$, $\beta = 1$, $T = 4$, $x_0 \sim \mathcal{N}(0, 0.5I)$,
1563 taking $K = 200$ time discretization steps for the simulation of the SDE.
- (REPULSIVE) Exactly the same as isotropic, but with $A = -I$.
- (ANISOTROPIC) We set $d = 20$, $A = \text{diag}(e^{a_i})$, $P = U \text{diag}(e^{p_i}) U^T$, $Q = 0.5I$, $\beta = 1$,
1564 $T = 4$, $x_0 \sim \mathcal{N}(0, 0.5I)$, taking $K = 200$ time discretization steps for the simulation of
1565 the SDE. The values a_i, p_i are sampled i.i.d. from $\mathcal{N}(0, 1)$, and the matrix U is a random
1566 orthogonal matrix sampled using `scipy.stats.ortho_group` (Virtanen et al., 2020)
1567 at the start of the simulation.

1566 DOUBLE WELL The second setting we consider is the d -dimensional double well defined through
 1567

$$1568 \quad E(x) = \sum_{i=1}^d \kappa_i (x_i^2 - 1)^2, \quad f(x) = \sum_{i=1}^d \nu_i (x_i^2 - 1)^2, \quad g(x) = 0. \quad (107)$$

$$1569$$

$$1570$$

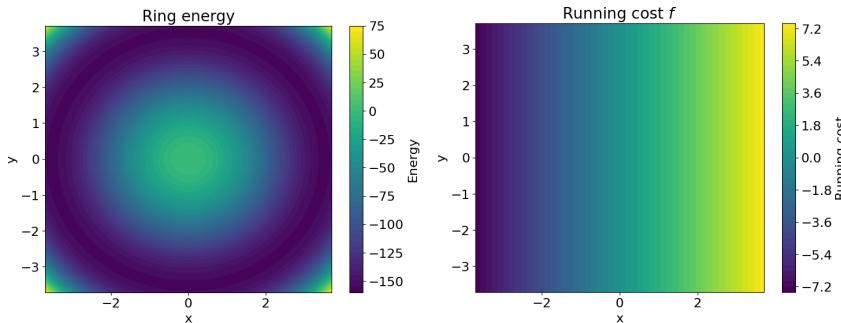
1571 where $d = 10$, and $\kappa_i = 5, \nu_i = 3$ for $i = 1, 2, 3$ and $\kappa_i = \nu_i = 1$ for $i \geq 4$. In addition, we again set
 1572 $T = 4, \beta = 1$ and use $K = 400$ discretization steps. This problem is similar to the one considered in
 1573 Nüsken & Richter (2021) and Domingo-Enrich et al. (2024b), the difference being that we consider
 1574 a longer horizon ($T = 4$ instead of $T = 1$) and consider a nonzero running cost in order to have
 1575 nontrivial long-term behaviour. This problem is considered a highly nontrivial benchmark problem,
 1576 since the double well in each dimension creates a total of $2^d = 1024$ local minima, making this
 1577 setting highly multimodal. The ground truth is *not* available in closed form, but can be approximated
 1578 efficiently by noticing that the energy and running cost are a sum of one-dimensional terms, and
 1579 hence we can compute the solution by solving d one-dimensional problems using a classical solver.
 1580

RING The final setting considers the setup

$$1581 \quad E(x) = \alpha (\|x\|^2 - 2R^2) \|x\|^2, \quad f(x) = 2x_1, \quad g(x) = 0, \quad (108)$$

$$1582$$

1583 where $\alpha = 1, R = 5/\sqrt{2}$. This energy is nonconvex and has its minimizers lying on the hypersphere
 1584 with radius R . This setting serves to highlight the difference between the relative eigenfunction
 1585 loss (20) and the other eigenfunction losses, and for visualization purposes will be done in $d = 2$.
 1586 In addition, we set $T = 5, \beta = 1$ and $x_0 = (R, 0)$. The energy landscape and running cost are
 1587 visualized in Figure 6. The goal of the controller is essentially to guide the system to the left hand
 1588 side of the xy -plane while being constrained by the potential to stay close to the circle of radius R .
 1589 This setting requires a smaller timestep for stable simulation, we take $K = 500$.
 1590



1601 Figure 6: Energy function (left) and running cost (right) for the RING setting in $d = 2$.
 1602

E.2 MODEL ARCHITECTURE AND TRAINING

1606 **Architecture** For the IDO methods, we use the exact same architecture as in Domingo-Enrich et al.
 1607 (2024b). They argue that the control can be viewed as the analog of a score function in diffusion
 1608 models, and hence they use a simplified U-Net architecture, where each of the up-/downsampling
 1609 steps is a fully connected layer with ReLU activations. As in their work, we use three downsampling
 1610 and upsampling steps, with widths 256, 128 and 64.

1611 For the eigenfunction models, we use the same architecture, but replace the ReLU activation with
 1612 the GELU activation function $x \mapsto x\Phi(x)$, where Φ is the cdf of a standard normal distribution
 1613 (Hendrycks & Gimpel, 2023). This is done because the eigenfunction losses require evaluating the
 1614 derivatives of the network w.r.t. the inputs, hence requiring a smoother activation function.
 1615

1616 **Training** For the eigenfunction method, we train using the following procedure:
 1617

1. Start training the top eigenfunction using the deep Ritz loss (17). Every 100 iterations,
 we estimate the eigenvalue λ_0 , and continue training until the variance of these estimates
 (computed with EMA 0.5) is below 10^{-4} and we have reached at least 5000 iterations.

Table 1: Iteration times by method and loss

Method	Experiment Loss	Iteration time (s)
COMBINED	Adjoint Matching	0.332
	Log variance	0.328
	Relative entropy	0.419
	SOCM	0.432
EIGF	Deep ritz loss	0.227
	PINN	0.662
	Relative loss	0.662
	Variational loss	0.228
FBSDE	FBSDE	0.443
IDO	Adjoint Matching	0.230
	Log variance	0.212
	Relative entropy	0.413
	SOCM	0.799

2. Fix λ_0 , and continue training the top eigenfunction using a loss function of choice among (15),(17), (20). Start training the excited state using the variational loss \mathcal{R}_{Var} in (71) with $k = 1$ and regularization parameter $\alpha = |\lambda_0|$.

For the combined methods, we first train the eigenfunction and eigenvalues using the method above with \mathcal{R}_{Rel} for 80 000 iterations, and then start training with an IDO loss using the control parametrization (22). We resample the trajectories in $[0, T_{\text{cut}}]$ (which only use the eigenfunction control) every $L = 100$ iterations in order to have a diverse set of starting positions at $T = T_{\text{cut}}$.

E.3 COMPUTATIONAL COST

Table 1 shows the computation cost per iteration for the different algorithms for the QUADRATIC (REPULSIVE) setting, measured in seconds/iteration when ran in isolation on a single GPU. All experiments were carried out on an NVIDIA H100 NV GPU.

E.4 FURTHER DETAILS ON THE RING SETTING

The RING setting serves as an illustrative example the difference between the 'absolute' eigenfunction losses and our relative eigenfunction loss. As shown in Figure 4, the relative loss obtains a drastically lower control objective than the absolute losses. To further understand this, consider the shape of the learned eigenfunctions for the different losses, shown in Figure 7. From the top row, it is clear that the learned eigenfunctions for the different methods are all very close in terms of the distance induced by $\|\cdot\|_{\mu}$. However, while the learned eigenfunctions look similar in $L^2(\mu)$, the logarithm of the learned eigenfunctions varies drastically, and hence the resulting control $\nabla \log \phi$ (shown in Figure 3) differs greatly - the control learned by the relative loss correctly guides the system along the circle in the negative x direction, while the other controls are not learned correctly for $x \geq 0$. The name 'relative loss' comes from the analogy with *absolute* and *relative* errors,

$$|x - x^*| < \epsilon \quad \text{vs.} \quad \left| \frac{x}{x^*} - 1 \right| < \epsilon \iff |\log x - \log x^*| < \epsilon + \mathcal{O}(\epsilon^2). \quad (109)$$

In essence, since the resulting control is given by $\nabla \log \phi$, the quantity of interest is the *relative* error of the learned eigenfunction, while existing methods are designed to minimize the *absolute* error.

E.5 DETAILS ON FIGURE 1

We give here some details on the motivating plot shown in Figure 1. The value shown is the control L^2 error of each of the different methods in the QUADRATIC (REPULSIVE) setting. Each algorithm was run for 30k iterations, and the values reported are the mean and 5%-95% quantiles over the last

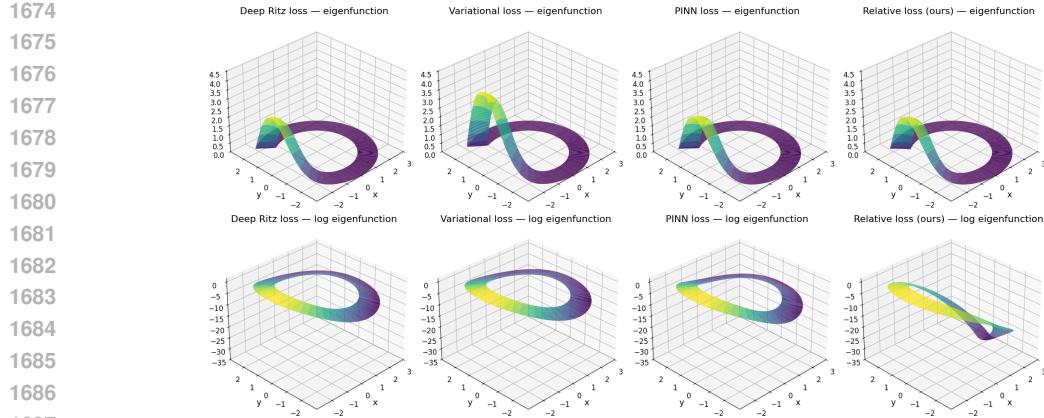


Figure 7: Learned eigenfunctions (top row) and their logarithms (bottom row) for the RING setting with different loss functions.

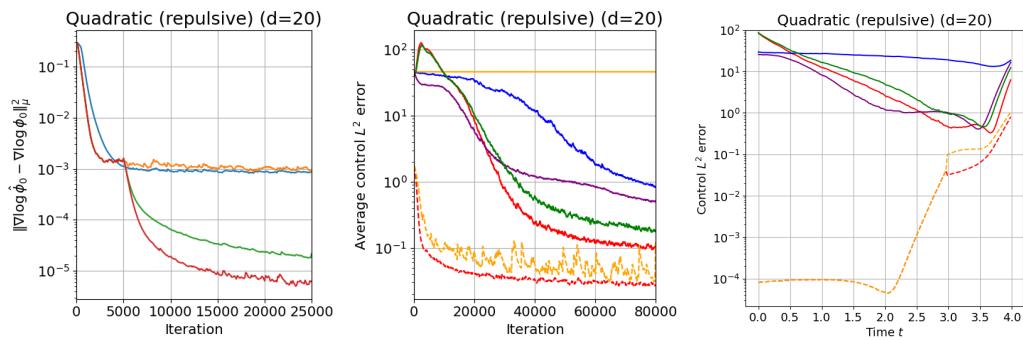


Figure 8: L^2 error of $\nabla \log \phi$ for different eigenfunctions (left), average control L^2 error for different methods (middle) and control L^2 error over time (right). Legend is the same as in Figure 4 and Figure 5.

1000 iterations. The EIGF+IDO (OURS) label refers to the combined method, where we first train the eigenfunctions until convergence and then train using the relative entropy loss.

E.6 ADDITIONAL EXPERIMENT: REPULSIVE POTENTIAL

Figure 8 shows the same plots discussed in the main text for the QUADRATIC (REPULSIVE) setting, where we obtain similar results. The reported eigenfunction error is measured in $L^2(\bar{\mu})$ instead of $L^2(\mu)$.

Additionally, for this experiment setting we have tested the performance of “ergodic” control estimator. It is based only on the first eigenfunction, i.e. $\beta^{-1} \nabla \log \phi_0^{\theta_0}$ is used for the whole time range. The performance of such an estimator is presented in Figure 9 by a curve labeled EIGF. It can be seen, that this approach reduces the error down with growth of time horizon T . However, using our proposed approach with control as in (22), which corresponds to EIGF+IDO curve, leads to a significant improvement.

E.7 CONTROL OBJECTIVES

As mentioned before, the control objective is the final metric for evaluating the performance of the different algorithms, but due to variance of the Monte Carlo estimator the difference between methods can be quite small. We report the control objectives for all experiments at convergence in Table 2 and Table 3. The value reported is the mean value of the control objective over $N = 65536$ simulations, and the error is the standard deviation of these estimates, divided by \sqrt{N} .

1728

1729

1730

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

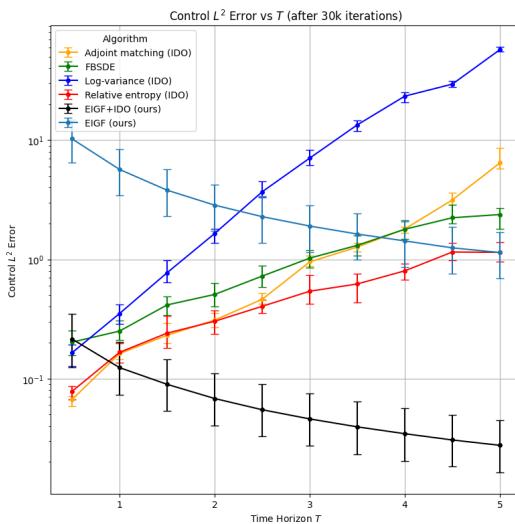


Figure 9: In this experiment EIFG method uses an ergodic estimator based only on the first eigenfunction. EIFG+IDO curve corresponds to the application of the proposed controller (22) with Relative Entropy loss. The figure shows L^2 control error for different methods after 30000 iterations.

Table 2: Control objective for the different methods in the QUADRATIC (ISOTROPIC) and QUADRATIC (REPULSIVE) settings. The SOCM method did not converge, and hence the dynamics diverge.

Method	Loss	QUADRATIC (ISOTROPIC)	QUADRATIC (REPULSIVE)
IDO	Relative entropy	32.7870 ± 0.014	112.5172 ± 0.051
	Log variance	32.7850 ± 0.014	114.0552 ± 0.053
	SOCM	73.1062 ± 0.062	nan \pm nan
	Adjoint matching	32.7754 ± 0.014	113.4554 ± 0.052
COMBINED (ours)	Relative entropy	32.7763 ± 0.014	112.3444 ± 0.050
	Log variance	32.7740 ± 0.014	150.0721 ± 0.12
	SOCM	32.7717 ± 0.014	112.3960 ± 0.050
	Adjoint matching	32.7725 ± 0.014	114.7020 ± 0.055
FBSDE	FBSDE	32.7979 ± 0.014	112.6393 ± 0.051

Table 3: Control objective for the different methods in the QUADRATIC (ANISOTROPIC) and DOUBLE WELL settings.

Method	Loss	QUADRATIC (ANISOTROPIC)	DOUBLE WELL
IDO	Relative entropy	38.9967 ± 0.022	35.2688 ± 0.010
	Log variance	31.3664 ± 0.016	32.8645 ± 0.0094
	SOCM	112.2728 ± 0.18	41.7215 ± 0.013
	Adjoint matching	31.3584 ± 0.016	34.8713 ± 0.010
COMBINED (ours)	Relative entropy	31.3476 ± 0.016	32.6130 ± 0.0088
	Log variance	32.5115 ± 0.047	32.9080 ± 0.0088
	SOCM	31.3483 ± 0.016	32.4421 ± 0.0088
	Adjoint matching	31.3497 ± 0.016	32.5638 ± 0.0088
FBSDE	FBSDE	31.3854 ± 0.016	35.1890 ± 0.011