

Drift Control of High-Dimensional RBM: A Computational Method Based on Neural Networks

Baris Ata

Michael Harrison

Nian Si

Sunday 17th September, 2023

Abstract

Motivated by applications in queueing theory, we consider a stochastic control problem whose state space is the d -dimensional positive orthant. The controlled process Z evolves as a reflected Brownian motion whose covariance matrix is exogenously specified, as are its directions of reflection from the orthant's boundary surfaces. A system manager chooses a drift vector $\theta(t)$ at each time t based on the history of Z , and the cost rate at time t depends on both $Z(t)$ and $\theta(t)$. In our initial problem formulation, the objective is to minimize expected discounted cost over an infinite planning horizon, after which we treat the corresponding ergodic control problem. Extending earlier work by Han et al. (Proceedings of the National Academy of Sciences, 2018, 8505-8510), we develop and illustrate a simulation-based computational method that relies heavily on deep neural network technology. For test problems studied thus far, our method is accurate to within a fraction of one percent, and is computationally feasible in dimensions up to at least $d = 30$.

1 Introduction

Beginning with the seminal work of Iglehart and Whitt [26, 27], there has developed over the last 50+ years a large literature that justifies the use of reflected Brownian motions as approximate models of queueing systems under “heavy traffic” conditions. In particular, a limit theorem proved by Reiman [39] justifies the use of d -dimensional reflected Brownian motion (RBM) as an approximate model of a d -station queueing network. Reiman's theory is restricted to networks of the generalized Jackson type, also called single-class networks, or networks with homogeneous customer populations, but it has been extended to more complex multi-class networks under certain restrictions, most notably by Peterson [37] and Williams [45]. The survey papers by Williams [44] and by Harrison and Nguyen [20] provide an overview of heavy traffic limit theory through its first 25 years.

Many authors have commented on the compactness and simplicity of RBM as a mathematical model, at least in comparison with the conventional discrete-flow models that it replaces. For example, in the preface to Kushner [32]'s book on heavy traffic analysis one finds the following passage:

“These approximating [Brownian] models have the basic structure of the original problem, but are significantly simpler. Much inessential detail is eliminated ... They greatly

simplify analysis, design, and optimization, [yielding] good approximations to problems that would otherwise be intractable ...”

Of course, having adopted RBM as a system model, one still confronts the question of how to do performance analysis, and in that regard there has been an important recent advance: Blanchet et al. [10] have developed a simulation-based method to estimate steady-state performance measures for RBM in dimensions up to 200, and those estimates come with performance guarantees.

Descriptive performance analysis versus optimal control. Early work on heavy traffic approximations, including the papers cited above, focused on descriptive performance analysis under fixed operating policies. Harrison [18, 19] expanded the framework to include consideration of dynamic control, using informal arguments to justify Brownian approximations for queueing network models where a system manager can make sequencing, routing and/or input control decisions. Early papers in that vein by Harrison and Wein [22, 23] and by Wein [43] dealt with Brownian models simple enough that their associated control problems could be solved analytically. But for larger systems and/or more complex decisions, the Brownian control problem that approximates an original queueing control problem may only be solvable numerically. Such stochastic control problems may be of several different types, depending on context.

At one end of the spectrum are drift control problems, in which the controlling agent can effect changes in system state only at bounded finite rates. At the other end of the spectrum are impulse control problems, in which the controlling agent can effect instantaneous jumps in system state, usually with an associated fixed cost. In between are singular control problems, in which the agent can effect instantaneous state changes of any desired size, usually at a cost proportional to the size of the displacement; see for example, Karatzas [28]. In this paper we develop a computational method for the first of those three problem classes, and then illustrate its use on selected test problems. Our method is a variant of the one developed by Han et al. [17] for solution of semi-linear partial differential equations, and in its implementation we have re-used substantial amounts of the code provided by Han et al. [17] and Zhou et al. [49].

Literature Review. Two of the most relevant streams of literature are *i*) drift rate control problems, and *ii*) solving PDEs using deep learning. Ata et al. [5] considers a one-dimensional drift rate control problem on a bounded interval under a general cost of control but no state costs. The authors characterize the optimal policy in closed form; and they discuss the application of their model to a power control problem in wireless communication. Ormeci Matoglu and Vande Vate [36] consider a drift rate control problem where a system controller incurs a fixed cost to change the drift rate. The authors prove that a deterministic, non-overlapping control band policy is optimal; also see Vande Vate [42]. Ghosh and Weerasinghe [15, 16] extend Ata et al. [5] by incorporating state costs, abandonments and optimally choosing the interval where the process lives.

Drift control problems arise in a broad range of applications in practice. Rubino and Ata [40] studies a dynamic scheduling problem for a make-to-order manufacturing system. The authors model order cancellations as abandonments from their queueing system. This model feature gives rise to a drift rate control problem in the heavy traffic limit. Ata et al. [6] uses a drift control model to

study a dynamic staffing problem in order to determine the number of volunteer gleaners, who sign up to help but may not show up, for harvesting leftover crops donated by farmers for the purpose of feeding food-insecure individuals. Bar-Ilan et al. [7] use a drift control model to study international reserves.

All of the papers mentioned above study one-dimensional drift-rate control problems. To the best of our knowledge, there have not been any papers studying such problems in high dimensions. One exception to this is the recent working paper Ata and Kasikaralar [4] that studies dynamic scheduling of a multiclass queue motivated by call center industry. Focusing on the Halfin-Whitt asymptotic regime, the authors derive a (limiting) drift rate control problem whose state space is \mathbb{R}^d , where d is the number of buffers in their queueing model. Similar to us, the authors build on Han et al. [17] to solve their (high-dimensional) drift rate control problem. However, our work differs from theirs significantly, because their control problem has no state space constraints.

As mentioned earlier, our work builds on the seminal paper Han et al. [17]. In the last five years, there have been many papers written on solving PDEs using deep neural networks. We refer the reader to the recent survey Beck et al. [8]; also see E et al. [14].

The remainder of this paper. Section 2 recapitulates essential background knowledge from RBM theory, after which Section 3 states in precise mathematical terms the discounted control and ergodic control problems that are the object of our study. In each case, the problem statement is expressed in probabilistic terms initially, and then re-expressed analytically in the form of an equivalent Hamilton-Jacobi-Bellman equation (hereafter abbreviated to HJB equation). Section 4 derives key identities, that significantly contribute to the subsequent development of our computational method. Section 5 describes our computational method in detail.

Section 6 specifies three families of drift control test problems, each of which has members of dimensions $d = 1, 2, \dots$. The first two families arise as heavy traffic limits of certain queueing network control problems, and we explain that motivation in some detail. Drift control problems in the third family have a separable structure that allows them to be solved exactly by analytical means, which is of obvious value for assessing the accuracy of our computational method. Section 7 presents numerical results obtained with our method for all three families of test problems. In that admittedly limited context, our computed solutions are accurate to within a fraction of one percent, and our method remains computationally feasible up to at least dimension $d = 30$, and in some cases up to dimension 100 or more. In Section 8 we describe variations and generalizations of the problems formulated in Section 3 that are of interest for various purposes, and which we expect to be addressed in future work. Finally, there are a number of appendices that contain proofs or other technical elaboration for arguments or procedures that have only been sketched in the body of the paper.

2 RBM preliminaries

We consider here a reflected Brownian motion $Z = \{Z(t), t \geq 0\}$ with state space \mathbb{R}_+^d , where $d \geq 1$. The data of Z are a (negative) drift vector $\mu \in \mathbb{R}^d$, a $d \times d$ positive-definite covariance matrix $A = (a_{ij})$, and a $d \times d$ reflection matrix R of the form

$$R = I - Q, \text{ where } Q \text{ has non-negative entries and spectral radius } \rho(Q) < 1. \quad (1)$$

The restriction to reflection matrices of the form (1) is not essential for our purposes, but it simplifies the technical development and is consistent with usage in the related earlier paper by Blanchet et al. [10]. Denoting by $W = \{W(t), t \geq 0\}$ a d -dimensional Brownian motion with zero drift, covariance matrix A , and $W(0) = 0$, we then have the representation

$$Z(t) = Z(0) + W(t) - \mu t + RY(t), \quad t \geq 0, \text{ where} \quad (2)$$

$$Y_i(\cdot) \text{ is continuous and non-decreasing with } Y_i(0) = 0 \quad (i = 1, 2, \dots, d), \text{ and} \quad (3)$$

$$Y_i(\cdot) \text{ only increases at those times } t \text{ when } Z_i(t) = 0 \quad (i = 1, 2, \dots, d). \quad (4)$$

Harrison and Reiman [21] showed that the relationships (1) to (4) determine Y and Z as pathwise functionals of W , and that the mapping $W \rightarrow (Y, Z)$ is continuous in the topology of uniform convergence. We interpret the i^{th} column of R as the direction of reflection on the boundary surface $\{z \in \mathbb{R}_+^d : z_i = 0\}$, and call $Y_i = \{Y_i(t), t \geq 0\}$ the “pushing process” on that boundary surface.

In preparation for future developments, let f be an arbitrary C^2 (that is, twice continuously differentiable) function $\mathbb{R}^d \rightarrow \mathbb{R}$, and let ∇f denote its gradient vector as usual. Also, we define a second-order differential operator \mathcal{L} via

$$\mathcal{L}f = \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d a_{ij} \frac{\partial^2}{\partial z_i \partial z_j} f, \quad (5)$$

and a first-order differential operator $\mathcal{D} = (\mathcal{D}_1, \dots, \mathcal{D}_d)^\top$ via

$$\mathcal{D}f = R^\top \nabla f, \quad (6)$$

where \top in (6) denotes transpose. Thus $\mathcal{D}_i f(\cdot)$ is the directional derivative of f in the direction of reflection on the boundary surface $\{z_i = 0\}$. With these definitions, an application of Ito’s formula now gives the following identity, cf. Harrison and Reiman [21], Section 3:

$$df(X(t)) = \nabla f(Z(t)) \cdot dW(t) + (\mathcal{L}f - \mu \cdot \nabla f)(Z(t)) dt + \mathcal{D}f(Z(t)) \cdot dY(t), \quad t \geq 0. \quad (7)$$

In the obvious way, the first inner product on the right side of (7) is shorthand for a sum of d Ito differentials, while the last one is shorthand for a sum of d Riemann-Stieltjes differentials.

3 Problem statements and HJB equations

Let us now consider a stochastic control problem whose state space is \mathbb{R}_+^d ($d \geq 1$). The controlled process Z has the form

$$Z(t) = Z(0) + W(t) - \int_0^t \theta(s) ds + RY(t), \quad t \geq 0, \quad (8)$$

where (i) $W = \{W(t), t \geq 0\}$ is a d -dimensional Brownian motion with zero drift, covariance matrix A , and $W(0) = 0$ as in Section 2, (ii) $\theta = \{\theta(t), t \geq 0\}$ is a non-anticipating control, or non-anticipating drift process, chosen by a system manager and taking values in a bounded set $\Theta \subset \mathbb{R}^d$, and (iii) $Y = \{Y(t), t \geq 0\}$ is a d -dimensional pushing process with components Y_i that satisfy (3) and (4). Note that our sign convention on the drift in the basic system equation (8) is *not* standard. That is, we denote by $\theta(t)$ the *negative* drift vector at time t .

The control θ is chosen to optimize an economic objective (see below), and attention will be restricted to *stationary* Markov controls, or stationary control policies, by which we mean that

$$\theta(t) = u(Z(t)), \quad t \geq 0 \text{ for some measurable policy function } u : \mathbb{R}_+^d \rightarrow \Theta. \quad (9)$$

Hereafter the set Θ of drift vectors available to the system manager will be referred to as the *action space* for our control problem, a function $u : \mathbb{R}_+^d \rightarrow \Theta$ will simply be called a *policy*, and we denote by Z^u the controlled RBM defined via (8) and (9). With regard to the system manager's objective, we take as given a continuous cost function $c : \mathbb{R}_+^d \times \Theta \rightarrow \mathbb{R}$ with polynomial growth (see below for the precise meaning of that phrase), and assume that the cumulative cost incurred over the time interval $[0, t]$ under policy u is

$$C^u(t) \equiv \int_0^t c(Z^u(s), u(Z^u(s))) ds, \quad t \geq 0. \quad (10)$$

To be more specific, for $m, n \geq 1$, a function $g : D \subset \mathbb{R}^m \rightarrow \mathbb{R}^n$ is said to have polynomial growth if there exist constants $\alpha_1, \beta_1 > 0$ such that

$$|g(z)| \leq \alpha_1 \left(1 + |z|^{\beta_1}\right), \quad z \in D.$$

Because the action space Θ is bounded, the polynomial growth assumption on c reduces to the following:

$$|c(z, \theta)| \leq \alpha_2 \left(1 + |z|^{\beta_2}\right) \text{ for all } z \in \mathbb{R}_+^d \text{ and } \theta \in \Theta, \quad (11)$$

where α_2, β_2 are positive constants.

Because our action space Θ is bounded by assumption, the controlled RBM Z^u has bounded drift under any policy u , from which one can prove the following mild but useful property; see Appendix A for its proof.

Proposition 1. *Under any policy u and for any integer $n = 1, 2, \dots$ the function*

$$g_n(z, t) = \mathbb{E}_z \{|Z^u(t)|^n\}, \quad t \geq 0,$$

has polynomial growth in t for each fixed $z \in \mathbb{R}_+^d$.

3.1 Discounted control

In our first problem formulation, an interest rate $r > 0$ is taken as given, and we adopt the following discounted cost objective: choose a policy u to minimize

$$V^u(z) \equiv \mathbb{E}_z \left[\int_0^\infty e^{-rt} dC^u(t) \right] = \mathbb{E}_z \left[\int_0^\infty e^{-rt} c(Z^u(t), u(Z^u(t))) dt \right], \quad (12)$$

where $\mathbb{E}_z(\cdot)$ denotes a conditional expectation given that $Z(0) = z$. Given the polynomial growth condition (11), it follows from Proposition 1 that the moments of $Z(t)$ are polynomially bounded as functions of t for each fixed initial state z . Given the assumed positivity of the interest rate r , the expectation in (12) is therefore well defined and finite for each $z \in \mathbb{R}_+^d$.

Hereafter we refer to $V^u(\cdot)$ as the *value function* under policy u , and define the *optimal value function*

$$V(z) = \min_{u \in \mathcal{U}} V^u(z) \text{ for each } z \in \mathbb{R}_+^d, \quad (13)$$

where \mathcal{U} is the set of stationary Markov control policies. To solve for the value function $V^u(\cdot)$ under an arbitrary policy u , a standard argument gives the following PDE with boundary conditions, where \mathcal{L} and \mathcal{D}_i are the differential operators defined via (5) and (6), respectively:

$$\mathcal{L}V^u(z) - u(z) \cdot \nabla V^u(z) + c(z, u(z)) = rV^u(z), \quad z \in \mathbb{R}_+^d, \quad (14)$$

with boundary conditions

$$\mathcal{D}_i V^u(z) = 0 \text{ if } z_i = 0 \text{ (} i = 1, 2, \dots, d \text{)}. \quad (15)$$

The corresponding HJB equation, to be solved for the *optimal* value function $V(\cdot)$, is

$$\mathcal{L}V(z) - \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\} = rV(z), \quad z \in \mathbb{R}_+^d, \quad (16)$$

with boundary conditions

$$\mathcal{D}_i V(z) = 0 \text{ if } z_i = 0 \text{ (} i = 1, 2, \dots, d \text{)}. \quad (17)$$

Moreover, the policy

$$u^*(z) = \arg \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\} \quad (18)$$

is optimal, meaning that $V^{u^*}(z) = V(z)$ for $z \in \mathbb{R}_+^d$.

There will be no attempt here to prove existence of C^2 solutions, but our computational method

proceeds as if that were the case, striving to compute a C^2 function V that satisfies (16)-(17) as closely as possible in a certain sense. In Appendix B.1 we use (7) to verify that a sufficiently regular solution of the PDE (14)-(15) does in fact satisfy (12) as intended, and similarly, that a sufficiently regular solution of (16)-(17) does in fact satisfy (13).

3.2 Ergodic control

For our second problem formulation, it is assumed that

$$c(z, \theta) \geq 0 \text{ for all } (z, \theta) \in \mathbb{R}_+^d \times \Theta. \quad (19)$$

Readers will see that our analysis can be extended to cost functions that take on negative values in at least some states, but to do so one must deal with certain irritating technicalities. To be specific, the issue is whether the expected values involved in our formulation are well defined.

In preparation for future developments, let us recall that a square matrix R of the form (1), called a *Minkowski matrix* in linear algebra (or just *M-matrix* for brevity), is non-singular, and its inverse is given by the Neumann expansion

$$R^{-1} = I + Q + Q^2 + \dots \geq 0.$$

Hereafter, we assume that

$$\text{there exists at least one } \theta \in \Theta \text{ such that } R^{-1}\theta > 0. \quad (20)$$

It is known that an RBM with a non-singular covariance matrix, reflection matrix R , and negative drift vector θ has a stationary distribution if and only if the inequality in (20) holds, cf. Section 6 of Harrison and Williams [24]. Of course, our statement of this "stability condition" reflects the non-standard sign convention used in this paper. That is, θ denotes the *negative* drift vector of the RBM under discussion.

For our ergodic control problem, a policy function $u : \mathbb{R}_+^d \rightarrow \Theta$ is said to be *admissible* if, first, the corresponding controlled RBM Z^u has a unique stationary distribution π^u , and if, moreover,

$$\int_{\mathbb{R}_+^d} |f(z)| \pi^u(dz) < \infty \quad (21)$$

for any function $f : \mathbb{R}_+^d \rightarrow \mathbb{R}$ with polynomial growth. Our assumption (20) ensures the existence of at least one admissible policy u , as follows. Let $\theta \in \Theta$ be a negative drift vector satisfying (20), and consider the constant policy $u(\cdot) \equiv \theta$. The corresponding controlled process Z^u is then an RBM having a unique stationary distribution π^u , as noted above. It has been shown in Budhiraja and Lee [11] that the moment generating function of π^u is finite in a neighborhood of the origin, from which it follows that π^u has finite moments of all orders. Thus π^u satisfies (21) for any function f with polynomial growth, so u is admissible.

Because our cost function $c(z, \theta)$ has polynomial growth and our action space Θ is bounded, the steady-state average cost

$$\xi^u \equiv \int_{\mathbb{R}_+^d} c(z, u(z)) \pi^u(dz) \quad (22)$$

is well defined and finite under any admissible policy u . The objective in our ergodic control problem is to find an admissible policy u for which ξ^u is minimal.

To solve for the steady-state average cost ξ^u and the corresponding *relative value function*, denoted by $v^u(\cdot)$, under an admissible policy u , a standard argument gives the following PDE:

$$\mathcal{L}v^u(z) - u(z) \cdot \nabla v^u(z) + c(z, \theta) = \xi^u \text{ for each } z \in \mathbb{R}_+^d, \quad (23)$$

with boundary conditions

$$\mathcal{D}_i v^u(z) = 0 \text{ if } z_i = 0 \text{ (} i = 1, 2, \dots, d \text{)}. \quad (24)$$

The HJB equation for ergodic control is again of a standard form, involving a constant ξ (interpreted as the minimum achievable steady-state average cost) and a relative value function $v : \mathbb{R}_+^d \rightarrow \mathbb{R}$. To be specific, the HJB equation is

$$\mathcal{L}v(z) - \max_{\theta \in \Theta} \{\theta \cdot \nabla v(z) - c(z, \theta)\} = \xi \text{ for each } z \in \mathbb{R}_+^d, \quad (25)$$

with boundary conditions

$$\mathcal{D}_i v(z) = 0 \text{ if } z_i = 0 \text{ (} i = 1, 2, \dots, d \text{)}. \quad (26)$$

Paralleling the previous development for discounted control, we show the following in Appendix B.2: if a C^2 function v and a constant ξ jointly satisfy (25)-(26), then

$$\xi = \inf_{u \in \mathcal{U}} \xi^u, \quad (27)$$

where \mathcal{U} denotes the set of admissible controls for the ergodic cost formulation. Moreover, the policy

$$u^*(z) = \arg \max_{\theta \in \Theta} \{\theta \cdot \nabla v(z) - c(z, \theta)\}, \quad z \in \mathbb{R}_+^d, \quad (28)$$

is optimal, meaning that $\xi^{u^*} = \xi$. Again paralleling the previous development for discounted control, there is no attempt to prove that such a solution for (25)-(26) exists. In Appendix B.2 we use (7) to verify that a sufficiently regular solution of the PDE (23)-(24) does in fact satisfy (22) as intended, and similarly, that a sufficiently regular solution of (25)-(26) does in fact satisfy (27).

4 Equivalent SDEs

In this section we prove two key identities, Equations (31) and (40) below, that are closely patterned after results used by Han et al. [17] to justify their "deep BSDE method" for solution of certain

non-linear PDEs. That earlier work provided both inspiration and detailed guidance for our study, but we include these derivations to make the current account as nearly self-contained as possible. Sections 4.1 and 4.2 treat the discounted and ergodic cases, respectively.

Our method begins by specifying what we call a *reference policy*. This is a nominal or default policy, specified at the outset but possibly revised in light of computational experience, that we use to generate sample paths of the controlled RBM Z . Roughly speaking, one wants to choose the reference policy so that its paths tend to occupy parts of the state space thought to be most frequently visited by an optimal policy.

4.1 Discounted control

Our reference policy for the discounted case chooses a constant action $u(z) = \tilde{\theta} > 0$ in every state $z \in \mathbb{R}_+^d$. (Again we stress that, given the non-standard sign convention embodied in (8) and (19), this means that \tilde{Z} has a constant drift vector $-\tilde{\theta}$, with all components negative.) Thus the corresponding *reference process* \tilde{Z} is a d -dimensional RBM which, in combination with its d -dimensional pushing process \tilde{Y} and the d -dimensional Brownian motion W defined in Section 2, satisfies

$$\tilde{Z}(t) = \tilde{Z}(0) + W(t) - \tilde{\theta}t + R\tilde{Y}(t), \quad t \geq 0, \quad (29)$$

plus the obvious analogs of Equations (3) and (4). For the key identity (31) below, let

$$F(z, x) = \tilde{\theta} \cdot x - \max_{\theta \in \Theta} \{\theta \cdot x - c(z, \theta)\} \quad \text{for } z \in \mathbb{R}_+^d \text{ and } x \in \mathbb{R}^d. \quad (30)$$

Proposition 2. *If $V(\cdot)$ satisfies the HJB equation (16) - (17), then it also satisfies the following identity almost surely for any $T > 0$:*

$$e^{-rT}V(\tilde{Z}(T)) - V(\tilde{Z}(0)) = \int_0^T e^{-rt}\nabla V(\tilde{Z}(t)) \cdot dW(t) - \int_0^T e^{-rt}F(\tilde{Z}(t), \nabla V(\tilde{Z}(t)))dt. \quad (31)$$

Proof. Applying Ito's formula to $e^{-rt}V(\tilde{Z}(t))$ and using Equation (7) yield

$$\begin{aligned} e^{-rT}V(\tilde{Z}(T)) - V(\tilde{Z}(0)) &= \int_0^T e^{-rt}\nabla V(\tilde{Z}(t)) \cdot dW(t) + \int_0^T e^{-rt}\mathcal{D}V(\tilde{Z}(t)) \cdot d\tilde{Y}(t) \\ &\quad + \int_0^T e^{-rt} \left(\mathcal{L}V(\tilde{Z}(t)) - \tilde{\theta} \cdot \nabla V(\tilde{Z}(t)) - rV(\tilde{Z}(t)) \right) dt. \end{aligned} \quad (32)$$

Using the HJB boundary condition (17), plus the complementarity condition (4) for \tilde{Y} and \tilde{Z} , one has

$$\int_0^T e^{-rt}\mathcal{D}V(\tilde{Z}(t)) \cdot d\tilde{Y}(t) = 0.$$

Furthermore, substituting $z = \tilde{Z}(t)$ in the HJB equation (16), multiplying both sides by e^{-rt} , rear-

ranging the terms, and integrating over $[0, T]$ yields

$$\int_0^T e^{-rt} \left(\mathcal{L}V(\tilde{Z}(t)) - rV(\tilde{Z}(t)) \right) dt = \int_0^T e^{-rt} \max_{\theta \in \Theta} \left(\theta \cdot \nabla V(\tilde{Z}(t)) - c(\tilde{Z}(t), \theta) \right) dt. \quad (33)$$

Substituting Equation (33) into Equation (32) gives Equation (31). \square

Proposition 2 provides the motivation for the loss function that we strive to minimize in our computational method (see Section 5). Before developing that approach, we prove the following, which can be viewed as a converse of Proposition 2.

Proposition 3. *Suppose that $V : \mathbb{R}_+^d \rightarrow \mathbb{R}$ is a C^2 function, $G : \mathbb{R}_+^d \rightarrow \mathbb{R}^d$ is continuous, and $V, \nabla V$, and G all have polynomial growth. Also assume that the following identity holds almost surely for some fixed $T > 0$ and every $Z(0) = z \in \mathbb{R}_+^d$:*

$$e^{-rT}V(\tilde{Z}(T)) - V(\tilde{Z}(0)) = \int_0^T e^{-rt}G(\tilde{Z}(t)) \cdot dW(t) - \int_0^T e^{-rt}F(\tilde{Z}(t), G(\tilde{Z}(t))) dt. \quad (34)$$

Then $G(\cdot) = \nabla V(\cdot)$ and V satisfies the HJB equation (16) - (17).

Remark 1. The surprising conclusion that (34) implies $\nabla V(\cdot) = G(\cdot)$, without any *a priori* relationship between G and ∇V being assumed, motivates the “double parametrization” method in Section 5.

Proof of Proposition 3. Because \tilde{Z} is a time-homogeneous Markov process, we can express (34) equivalently as follows for any $k = 0, 1, \dots$:

$$\begin{aligned} e^{-rT}V(\tilde{Z}((k+1)T)) - V(\tilde{Z}(kT)) &= \int_{kT}^{(k+1)T} e^{-r(t-kT)}G(\tilde{Z}(t)) \cdot dW(t) \\ &\quad - \int_{kT}^{(k+1)T} e^{-r(t-kT)}F(\tilde{Z}(t), G(\tilde{Z}(t))) dt. \end{aligned} \quad (35)$$

Now multiply both sides of (35) by e^{-rkT} , then add the resulting relationships for $k = 0, 1, \dots, n-1$ to arrive at the following:

$$e^{-rnT}V(\tilde{Z}(nT)) = V(\tilde{Z}(0)) + \int_0^{nT} e^{-rt}G(\tilde{Z}(t)) \cdot dW(t) - \int_0^{nT} e^{-rt}F(\tilde{Z}(t), G(\tilde{Z}(t)))dt. \quad (36)$$

Because G has polynomial growth, one can show that

$$\mathbb{E}_z \left[\int_0^{nT} e^{-2rt}G(\tilde{Z}(t))^2 dz \right] < \infty$$

for all $n \geq 1$. Thus, when we take \mathbb{E}_z of both sides of (36), the stochastic integral (that is, the second

term) on the right side vanishes, and then rearranging terms gives the following:

$$V(z) = e^{-rnT} \mathbb{E}_z \left[V(\tilde{Z}(nT)) \right] + \mathbb{E}_z \left[\int_0^{nT} e^{-rt} F(\tilde{Z}(t), G(\tilde{Z}(t))) dt \right],$$

for arbitrary positive integer n .

By Proposition 1 and the polynomial growth condition of V , we have $e^{-rnT} \mathbb{E}_z[V(\tilde{Z}(nT))] \rightarrow 0$ as $n \rightarrow \infty$. Therefore,

$$V(z) = \lim_{n \rightarrow \infty} \mathbb{E}_z \left[\int_0^{nT} e^{-rt} F(\tilde{Z}(t), G(\tilde{Z}(t))) dt \right] \text{ for } z \in \mathbb{R}_+^d.$$

Similarly, since F and G have polynomial growth, we conclude that

$$\begin{aligned} \mathbb{E}_z \left[\int_0^\infty e^{-rt} \left| F(\tilde{Z}(t), G(\tilde{Z}(t))) \right| dt \right] &< +\infty \text{ for } z \in \mathbb{R}_+^d, \text{ and} \\ \int_0^{nT} e^{-rt} F(\tilde{Z}(t), G(\tilde{Z}(t))) dt &\leq \int_0^\infty e^{-rt} \left| F(\tilde{Z}(t), G(\tilde{Z}(t))) \right| dt < +\infty \text{ for } z \in \mathbb{R}_+^d. \end{aligned}$$

Thus, by dominated convergence, we have

$$V(z) = \mathbb{E}_z \left[\int_0^\infty e^{-rt} F(\tilde{Z}(t), G(\tilde{Z}(t))) dt \right] \text{ for } z \in \mathbb{R}_+^d. \quad (37)$$

In other words, $V(z)$ can be viewed as the expected discounted cost associated with the RBM under the reference policy starting in state $\tilde{Z}(0) = z$, where $F(\cdot, G(\cdot))$ is the state-cost function. Therefore, it follows from Equations (14) - (15) with $u(z) = \tilde{\theta}$ and $c(z, u(z)) = c(z, \tilde{\theta}) = F(z, G(z))$ for $z \in \mathbb{R}_+^d$ that V satisfies the following PDE:

$$\mathcal{L}V(z) - \tilde{\theta} \cdot \nabla V(z) + F(z, G(z)) = rV(z), \quad z \in \mathbb{R}_+^d, \quad (38)$$

with boundary conditions (15), and that it has polynomial growth.

Suppose that $G(\cdot) = \nabla V(\cdot)$ (which we will prove later). Substituting this into Equation (38) and using the definition of F , it follows that

$$\mathcal{L}V(z) - \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\} = rV(z), \quad z \in \mathbb{R}_+^d,$$

which along with the boundary condition (15) gives the desired result.

To complete the proof, it remains to show that $G(\cdot) = \nabla V(\cdot)$. By applying Ito's formula to $e^{-rt}V(\tilde{Z}(t))$ and using Equations (3)-(4) and (17), we conclude that

$$\begin{aligned} e^{-rT}V(\tilde{Z}(T)) - V(\tilde{Z}(0)) &= \int_0^T e^{-rt} \left(\mathcal{L}V(\tilde{Z}(t)) - \tilde{\theta} \cdot \nabla V(\tilde{Z}(t)) - rV(\tilde{Z}(t)) \right) dt \\ &\quad + \int_0^T e^{-rt} \nabla V(\tilde{Z}(t)) \cdot dW(t). \end{aligned}$$

Then, using Equation (38), we rewrite the preceding equation as follows:

$$e^{-rT}V(\tilde{Z}(T)) - V(\tilde{Z}(0)) = \int_0^T e^{-rt} \nabla V(\tilde{Z}(t)) \cdot dW(t) - \int_0^T e^{-rt} F(\tilde{Z}(t), G(\tilde{Z}(t))) dt.$$

Comparing this with Equation (34) yields

$$\int_0^T e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right) \cdot dW(t) = 0,$$

which yields the following:

$$\mathbb{E}_z \left[\left(\int_0^T e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right) \cdot dW(t) \right)^2 \right] = 0. \quad (39)$$

Thus, provided that $e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right)$ is square integrable, Ito's isometry [48, Lemma D.1] yields the following:

$$\mathbb{E}_z \left[\left(\int_0^T e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right) \cdot dW(t) \right)^2 \right] = \mathbb{E}_z \left[\int_0^T \left\| e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right) \right\|_A^2 dt \right] = 0,$$

where $\|x\|_A := x^\top Ax$. The square integrability of $e^{-rt} \left(G(\tilde{Z}(t)) - \nabla V(\tilde{Z}(t)) \right)$ follows because G and ∇V have polynomial growth and the action space Θ is bounded. Because of A is a positive definite matrix, we then have $\nabla V(\tilde{Z}(t)) = G(\tilde{Z}(t))$ almost surely. By the continuity of $\nabla V(\cdot)$ and $G(\cdot)$, we conclude that $\nabla V(\cdot) = G(\cdot)$. \square

4.2 Ergodic control

Again we use a reference policy with constant (negative) drift vector $\tilde{\theta}$, and now we assume that $R^{-1}\tilde{\theta} > 0$, which ensures that the reference policy is admissible for our ergodic control formulation. For the following analogs of Propositions 2 and 3, let

$$f(z, x) = \tilde{\theta} \cdot x - \max_{\theta \in \Theta} \{ \theta \cdot x - c(z, \theta) \} \text{ for } x \in \mathbb{R}^d, z \in \mathbb{R}_+^d.$$

Proposition 4. *If $v(\cdot)$ and ξ solve the HJB equation (25) - (26), then we have*

$$v(\tilde{Z}(T)) - v(\tilde{Z}(0)) = \int_0^T \nabla v(\tilde{Z}(t)) \cdot dW(t) + T\xi - \int_0^T f(\tilde{Z}(t), \nabla v(\tilde{Z}(t))) dt. \quad (40)$$

Proof. Applying Ito's formula to $v(z)$ yields

$$\begin{aligned} v(\tilde{Z}(T)) - v(\tilde{Z}(0)) & \\ &= \int_0^T \nabla v(\tilde{Z}(t)) \cdot dW(t) + \int_0^T \mathcal{D}v(\tilde{Z}(t)) \cdot d\tilde{Y}(t) + \int_0^T \left(\mathcal{L}v(\tilde{Z}(t)) - \tilde{\theta} \cdot \nabla v(\tilde{Z}(t)) \right) dt. \end{aligned} \quad (41)$$

Recall the boundary condition of the HJB equation is $\mathcal{D}_j v(z) = 0$ if $z_j = 0$. Thus Equations (3)-(4) jointly imply

$$\int_0^T \mathcal{D}v(\tilde{Z}(t)) \cdot d\tilde{Y}(t) = 0.$$

Then, substituting the HJB equation (25) into Equation (41) gives (40). \square

Proposition 5. *Suppose that $v : \mathbb{R}_+^d \rightarrow \mathbb{R}$ is a C^2 function, $g : \mathbb{R}_+^d \rightarrow \mathbb{R}^d$ is continuous, and $v, \nabla v, g$ all have polynomial growth. Also assume that the following identity holds almost surely for some fixed $T > 0$, a scalar ξ and every $Z(0) = z \in \mathbb{R}_+^d$:*

$$v(\tilde{Z}(T)) - v(\tilde{Z}(0)) = \int_0^T g(\tilde{Z}(t)) \cdot dW(t) + T\xi - \int_0^T f(\tilde{Z}(t), g(\tilde{Z}(t))) dt. \quad (42)$$

Then, $g(\cdot) = \nabla v(\cdot)$ and (v, ξ) satisfies the HJB equation (16) - (17).

Proof. Let $\tilde{\pi}$ be the stationary distribution of the RBM \tilde{Z} under the reference policy and $\tilde{Z}(\infty)$ be a random variable with the distribution $\tilde{\pi}$. Then, assuming the initial distribution of the RBM under the reference policy is $\tilde{\pi}$, i.e. $\tilde{Z}(0) \sim \tilde{\pi}$, its marginal distribution at time t is also $\tilde{\pi}$, i.e. $\tilde{Z}(t) \sim \tilde{\pi}$ for every $t \geq 0$.

Because g has polynomial growth, one can show that the expectation of the stochastic integral (that is, the first term) on the right side of (42) vanishes. Then, by taking the expectation over $\tilde{Z}(0) \sim \tilde{\pi}$, Equation (42) implies

$$\mathbb{E}_{\tilde{\pi}} [v(\tilde{Z}(0))] = \mathbb{E}_{\tilde{\pi}} [v(\tilde{Z}(T))] + \mathbb{E}_{\tilde{\pi}} \left[\int_0^T f(\tilde{Z}(t), g(\tilde{Z}(t))) dt \right] - T\xi. \quad (43)$$

By observing that $\mathbb{E}_{\tilde{\pi}}[v(\tilde{Z}(0))] = \mathbb{E}_{\tilde{\pi}}[v(\tilde{Z}(T))]$ and

$$\mathbb{E}_{\tilde{\pi}} [f(\tilde{Z}(t), g(\tilde{Z}(t)))] = \mathbb{E} [f(\tilde{Z}(\infty), g(\tilde{Z}(\infty)))] \text{ for } t \geq 0,$$

we conclude that $\xi = \mathbb{E}[f(\tilde{Z}(\infty), g(\tilde{Z}(\infty)))]$. In other words, ξ can be viewed as the expected steady-state cost associated with the RBM under the reference policy, where $f(\cdot, g(\cdot))$ is the state-cost function. Therefore, it follows (by assumption) from Equations (23) - (24) with $u(z) = \tilde{\theta}$ and $c(z, u(z)) = c(z, \theta) = f(z, g(z))$ for $z \in \mathbb{R}_+^d$ that there exists a C^2 relative value function \tilde{v} with polynomial growth that satisfies the following PDE:

$$\mathcal{L}\tilde{v}(z) - \tilde{\theta} \cdot \nabla \tilde{v}(z) + f(z, g(z)) = \xi, \quad z \in \mathbb{R}_+^d, \quad (44)$$

with boundary conditions $\mathcal{D}_i \tilde{v}(z) = 0$ if $z_i = 0$ ($i = 1, \dots, d$). Furthermore, applying Ito's formula to $\tilde{v}(\tilde{Z}(t))$ yields

$$\tilde{v}(\tilde{Z}(T)) - \tilde{v}(\tilde{Z}(0)) \quad (45)$$

$$= \int_0^T \nabla \tilde{v}(\tilde{Z}(t)) \cdot dW(t) + \int_0^T \mathcal{D}\tilde{v}(\tilde{Z}(t)) \cdot d\tilde{Y}(t) + \int_0^T \left(\mathcal{L}\tilde{v}(\tilde{Z}(t)) - \tilde{\theta} \cdot \nabla \tilde{v}(\tilde{Z}(t)) \right) dt.$$

Since $\tilde{v}(z)$ also satisfies the boundary conditions $\mathcal{D}_i \tilde{v}(z) = 0$ if $z_i = 0$ ($i = 1, \dots, d$), it follows from Equations (3)-(4) that

$$\int_0^T \mathcal{D}\tilde{v}(\tilde{Z}(t)) \cdot d\tilde{Y}(t) = 0.$$

Then, substituting Equation (44) into Equation (45) gives

$$\tilde{v}(\tilde{Z}(T)) - \tilde{v}(\tilde{Z}(0)) = \int_0^T \nabla \tilde{v}(\tilde{Z}(t)) \cdot dW(t) + T\xi - \int_0^T f(\tilde{Z}(t), g(\tilde{Z}(t))) dt. \quad (46)$$

In the proof of Proposition 3, we first showed that, because \tilde{Z} is a time-homogeneous Markov process, the assumed stochastic relationship (34) can be extended to the more general form (36) with n an arbitrary positive integer. In the current context one can argue in exactly the same way to establish the following. First, the assumed stochastic relationship (42) actually holds in the more general form where T is replaced by nT , with n an arbitrary positive integer. As a consequence, the derived stochastic relationship (46) also holds with nT in place of T . And finally, after taking expectations on both sides of those generalized versions of (42) and (46) we arrive at the following:

$$v(z) = \mathbb{E}_z \left[v(\tilde{Z}(nT)) \right] - nT\xi + \mathbb{E}_z \left[\int_0^{nT} f(\tilde{Z}(t), g(\tilde{Z}(t))) dt \right], \text{ and} \quad (47)$$

$$\tilde{v}(z) = \mathbb{E}_z \left[\tilde{v}(\tilde{Z}(nT)) \right] - nT\xi + \mathbb{E}_z \left[\int_0^{nT} f(\tilde{Z}(t), g(\tilde{Z}(t))) dt \right], \quad (48)$$

for $z \in \mathbb{R}_+^d$ and an arbitrary positive integer n . Note that the expectation of the stochastic integral vanishes because $\nabla \tilde{v}$ has polynomial growth. Subtracting (48) from (47) further yields

$$v(z) - \tilde{v}(z) = \mathbb{E}_z [v(Z(nT))] - \mathbb{E}_z [\tilde{v}(Z(nT))].$$

Without loss of generality, we assume $\mathbb{E} \left[\tilde{v}(\tilde{Z}(\infty)) \right] = \mathbb{E} \left[v(\tilde{Z}(\infty)) \right]$. Since $v(\cdot)$ and $\tilde{v}(\cdot)$ have polynomial growth and the reference policy is admissible, we have that $\sup_{n>0} \mathbb{E}_z \left[(v(\tilde{Z}(nT)))^2 \right] < \infty$ and $\sup_{n>0} \mathbb{E}_z \left[(\tilde{v}(\tilde{Z}(nT)))^2 \right] < \infty$ by Equation (21). Then, by the Vitali convergence theorem, we have

$$\begin{aligned} \lim_{n \rightarrow +\infty} \mathbb{E}_z [v(Z(nT))] &= \mathbb{E} \left[v(\tilde{Z}(\infty)) \right] \text{ and} \\ \lim_{n \rightarrow +\infty} \mathbb{E}_z [\tilde{v}(Z(nT))] &= \mathbb{E} \left[\tilde{v}(\tilde{Z}(\infty)) \right]. \end{aligned}$$

Therefore, we have

$$v(z) - \tilde{v}(z) = \lim_{n \rightarrow +\infty} (\mathbb{E}_z [v(Z(nT))] - \mathbb{E}_z [\tilde{v}(Z(nT))]) = 0 \text{ for } z \in \mathbb{R}_+^d,$$

which means $v(\cdot)$ also satisfies the PDE (44) and the associated boundary conditions. That is,

$$\mathcal{L}v(z) - \tilde{\theta} \cdot \nabla v(z) + f(z, g(z)) = \xi, \text{ for } z \in \mathbb{R}_+^d. \quad (49)$$

$$\mathcal{D}_i v(z) = 0 \text{ if } z_i = 0 \quad (i = 1, \dots, d). \quad (50)$$

Suppose that $g(\cdot) = \nabla v(\cdot)$ (which we will prove later). Substituting this into Equation (49) and using the definition of f , it follows that

$$\mathcal{L}v(z) - \max_{\theta \in \Theta} \{\theta \cdot \nabla v(z) - c(z, \theta)\} = \xi, \quad z \in \mathbb{R}_+^d,$$

which along with the boundary condition (50) gives the desired result.

To complete the proof, it remains to show that $g(\cdot) = \nabla v(\cdot)$. By applying Ito's formula to $v(\tilde{Z}(t))$ and using Equations (3)-(4) and (50), we conclude that

$$v(\tilde{Z}(T)) - v(\tilde{Z}(0)) = \int_0^T \left(\mathcal{L}v(\tilde{Z}(t)) - \tilde{\theta} \cdot \nabla v(\tilde{Z}(t)) \right) dt + \int_0^T \nabla v(\tilde{Z}(t)) \cdot dW(t).$$

Then, using Equation (49), we rewrite the preceding equation as follows:

$$v(\tilde{Z}(T)) - v(\tilde{Z}(0)) = T\xi - \int_0^T f(\tilde{Z}(t), g(\tilde{Z}(t))) dt + \int_0^T \nabla v(\tilde{Z}(t)) \cdot dW(t).$$

Comparing this with Equation (42) yields

$$\int_0^T \left(g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t)) \right) \cdot dW(t) = 0,$$

which yields the following:

$$\mathbb{E}_z \left[\left(\int_0^T \left(g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t)) \right) \cdot dW(t) \right)^2 \right] = 0. \quad (51)$$

Thus, provided $g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t))$ is square integrable, Ito's isometry [48, Lemma D.1] gives the following:

$$\mathbb{E}_z \left[\left(\int_0^T \left(g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t)) \right) \cdot dW(t) \right)^2 \right] = \mathbb{E}_z \left[\int_0^T \left\| g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t)) \right\|_A^2 dt \right] = 0.$$

The square integrability of $g(\tilde{Z}(t)) - \nabla v(\tilde{Z}(t))$ follows because g and ∇v have polynomial growth, and $\mathbb{E}_z \left(|\tilde{Z}(nT)|^k \right)$ is finite for all k because our action space Θ is bounded. Then, since A is positive definite, $\nabla v(\tilde{Z}(t)) = g(\tilde{Z}(t))$ almost surely. By the continuity of $\nabla v(\cdot)$ and $g(\cdot)$, we conclude that $\nabla v(\cdot) = g(\cdot)$. \square

5 Computational method

We follow in the footsteps of Han et al. [17], who developed a computational method to solve semilinear parabolic partial differential equations (PDEs). Those authors focused on a backward stochastic differential equation (BSDE) associated with their PDE, and in similar fashion, we focus on the stochastic differential equations (34) and (42) that are associated with our two stochastic control formulations (see Section 4). Our method differs from that of Han et al. [17], because they consider PDEs on a finite-time interval with an unbounded state space and a specified terminal condition, whereas our stochastic control problem has an infinite time horizon and state space constraints. As such, it leads to a PDE on a polyhedral domain with oblique derivative boundary conditions. We modify the approach of [17] to incorporate those additional features, treating the discounted and ergodic formulations in Sections 5.1 and 5.2, respectively.

5.1 Discounted control

We approximate the value function $V(\cdot)$ and its gradient $\nabla V(\cdot)$ by deep neural networks $V_{w_1}(\cdot)$ and $G_{w_2}(\cdot)$, respectively, with associated parameter vectors w_1 and w_2 . Seeking an approximate solution of the stochastic equation (34), we define the loss function

$$\begin{aligned} \ell(w_1, w_2) = \mathbb{E} & \left[\left(e^{-rT} V_{w_1}(\tilde{Z}(T)) - V_{w_1}(\tilde{Z}(0)) \right. \right. \\ & \left. \left. - \int_0^T e^{-rt} G_{w_2}(\tilde{Z}(t)) \cdot dW(t) + \int_0^T e^{-rt} F(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t))) dt \right)^2 \right], \end{aligned} \quad (52)$$

Here the expectation is taken with respect to the sample path distribution of our reference process \tilde{Z} , which will be specified in Algorithm 3 below. Our definition (52) of the loss function does not explicitly enforce the consistency requirement $\nabla V_{w_1}(\cdot) = G_{w_2}(\cdot)$, but Proposition 3 provides the justification for this separate parametrization. This type of double parametrization has also been implemented by Zhou et al. [49].

Our computational method seeks a neural network parameter combination (w_1, w_2) that minimizes an approximation of the loss defined via (52). Specifically, we first simulate multiple discretized paths of the reference RBM \tilde{Z} , restricted to a fixed and finite time domain $[0, T]$. To do that, we sample discretized paths of the underlying Brownian motion W , and then solve a discretized Skorohod problem for each path of W (this is the purpose of Subroutine 2) to obtain the corresponding path of \tilde{Z} . Thereafter, our method computes a discretized version of the loss (52), summing over sampled paths to approximate the expectation and over discrete time steps to approximate the integral over $[0, T]$, and minimizes it using stochastic gradient descent; see Algorithm 3. In Subroutine 2, given the index set B , $R_{B,B}$ is the submatrix derived by deleting the rows and columns of R with indices in $\{1, \dots, d\} \setminus B$. Similarly, $R_{\cdot,B}$ is the matrix that one arrives at by deleting the columns of R whose indices are in the set $\{1, \dots, d\} \setminus B$.

Subroutine 1 Euler discretization scheme

Input: The drift vector $-\tilde{\theta}$, the covariance matrix A , the reflection matrix R , the time horizon T , a step-size h (for simplicity, we assume $N \triangleq T/h$ is an integer), and a starting point $\tilde{Z}(0) = z$.

Output: A discretized reflected Brownian motion and the Brownian increments at times $h, 2h, \dots, Nh$.

```
1: function DISCRETIZE( $T, h, z$ )
2:   For time interval  $[0, T]$  and  $N = T/h$ , construct the partition  $0 = t_0 < t_1 < \dots < t_N = T$ ,
   where  $\Delta t_n = t_{n+1} - t_n = h$  for  $n = 0, 1, \dots, N - 1$ .
3:   Generate  $N$  i.i.d.  $d$ -dimensional Gaussian random variables with mean zero and covariance
   matrix  $hA$ , denoted by  $\delta_0, \dots, \delta_{N-1}$ .
4:   for  $k \leftarrow 0$  to  $N - 1$  do
5:      $x \leftarrow \tilde{Z}(kh) + \delta_k - \tilde{\theta}h$ 
6:      $\tilde{Z}((k + 1)h) \leftarrow \text{SKOROKHOD}(x)$ 
7:   end for
8:   return  $\tilde{Z}(h), \tilde{Z}(2h), \dots, \tilde{Z}(Nh)$  and  $\delta_0, \dots, \delta_{N-1}$ .
9: end function
```

Subroutine 2 Solve the Skorokhod problem (linear complementarity problem)

Input: A vector $x \in \mathbb{R}^d$ and the reflection matrix R .

Output: A solution to the Skorokhod problem $y \in \mathbb{R}_+^d$

```
1: Set  $\epsilon = 10^{-8}$ ;
2: function SKOROKHOD( $x$ )
3:    $y = x$ ;
4:   while Exists  $y_i < -\epsilon$  do
5:     Compute the set  $B = \{i : y_i < \epsilon\}$ ;
6:     Compute  $L_B = -R_{B,B}^{-1}x_B$ ;
7:     Compute  $y = x + R_{\cdot,B} \times L_B$ ;
8:   end while
9:   return  $y$ .
10: end function
```

After the parameter values w_1 and w_2 have been determined, our proposed policy is as follows:

$$\theta_{w_2}(z) = \arg \max_{\theta \in \Theta} (\theta \cdot G_{w_2}(z) - c(z, \theta)), \quad z \in \mathbb{R}_+^d. \quad (54)$$

Remark 2. One can also consider the policy using $\nabla V_{w_1}(\cdot)$ instead of $G_{w_2}(\cdot)$. That is,

$$\arg \max_{\theta \in \Theta} (\theta \cdot \nabla V_{w_1}(z) - c(z, \theta)), \quad z \in \mathbb{R}_+^d. \quad (55)$$

However, our numerical experiments suggest that this policy is inferior to (54).

Algorithm 3 Method for the discounted control case

Input: The number of iteration steps M , a batch size B , a learning rate α , a time horizon T , a discretization step-size h (for simplicity, we assume $N \triangleq T/h$ is an integer), a starting point z , and an optimization solver (SGD, ADAM, RMSProp, etc).

Output: A neural network approximation of the value function V_{w_1} and the gradient function G_{w_2} .

- 1: Initialize the neural networks V_{w_1} and G_{w_2} ; set $z_0^{(i)} = z$ for $i = 1, 2, \dots, B$.
- 2: **for** $k \leftarrow 0$ to $M - 1$ **do**
- 3: Simulate B discretized RBM paths and the Brownian increments $\{\tilde{Z}^{(i)}, \delta^{(i)}\}$ with a time horizon T and a discretization step-size h starting from $\tilde{Z}^{(i)}(0) = z_k^{(i)}$ by invoking DISCRETIZE($T, h, z_k^{(i)}$) for $i = 1, 2, \dots, B$.
- 4: Compute the empirical loss

$$\hat{\ell}(w_1, w_2) = \frac{1}{B} \sum_{i=1}^B \left(e^{-rT} V_{w_1}(\tilde{Z}^{(i)}(T)) - V_{w_1}(\tilde{Z}^{(i)}(0)) - \sum_{j=0}^{N-1} e^{-rjh} G_{w_2}(\tilde{Z}^{(i)}(hj)) \cdot \delta_j^{(i)} + \sum_{j=0}^{N-1} e^{-rjh} F(\tilde{Z}^{(i)}(hj), G_{w_2}(\tilde{Z}^{(i)}(hj))) h \right)^2. \quad (53)$$

- 5: Compute the gradient $\partial \hat{\ell}(w_1, w_2) / \partial w_1, \partial \hat{\ell}(w_1, w_2) / \partial w_2$ and update w_1, w_2 using the chosen optimization solver.
 - 6: Update $z_{k+1}^{(i)}$ as the end point of the path $\tilde{Z}^{(i)}$: $z_{k+1}^{(i)} \leftarrow \tilde{Z}^{(i)}(T)$.
 - 7: **end for**
 - 8: **return** Functions $V_{w_1}(\cdot)$ and $G_{w_2}(\cdot)$.
-

5.2 Ergodic control

We parametrize $v(\cdot)$ and $\nabla v(\cdot)$ using deep neural networks $v_{w_1}(\cdot)$ and $g_{w_2}(\cdot)$ with parameters w_1 and w_2 , respectively, and then use Equation (42) to define an auxiliary loss function

$$\tilde{\ell}(w_1, w_2, \xi) = \mathbb{E} \left[\left(v_{w_1}(\tilde{Z}(T)) - v_{w_1}(\tilde{Z}(0)) - \int_0^T g_{w_2}(\tilde{Z}(t)) dW(t) - T\xi + \int_0^T f(\tilde{Z}(t), g_{w_2}(\tilde{Z}(t))) dt \right)^2 \right]. \quad (56)$$

Then, defining the loss function $\ell(w_1, w_2) = \min_{\xi} \ell(w_1, w_2, \xi)$ and noting that

$$\text{Var}(X) = \min_{\xi} \mathbb{E}[(X - \xi)^2],$$

we arrive at the following expression for the loss function

$$\ell(w_1, w_2) = \text{Var} \left(v_{w_1}(\tilde{Z}(T)) - v_{w_1}(\tilde{Z}(0)) - \int_0^T g_{w_2}(\tilde{Z}(t)) dW(t) + \int_0^T f(\tilde{Z}(t), g_{w_2}(\tilde{Z}(t))) dt \right). \quad (57)$$

We present our method for the ergodic control case formally in Algorithm 4.

Algorithm 4 Method for the ergodic control case

Input: The number of iteration steps M , a batch size B , a learning rate α , a time horizon T , a discretization step-size h (for simplicity, we assume $N \triangleq T/h$ is an integer), a starting point z , and an optimization solver (SGD, ADAM, RMSProp, etc).

Output: A neural network approximation of the value function v_{w_1} and the gradient function g_{w_2} .

- 1: Initialize the neural networks v_{w_1} and g_{w_2} ; set $z_0^{(i)} = z$ for $i = 1, 2, \dots, B$.
- 2: **for** $k \leftarrow 0$ to $M - 1$ **do**
- 3: Simulate B discretized RBM paths and the Brownian increments $\{\tilde{Z}^{(i)}, \delta^{(i)}\}$ with a time horizon T and a discretization step-size h starting from $\tilde{Z}^{(i)}(0) = z_k^{(i)}$ by invoking DISCRETIZE($T, h, z_k^{(i)}$), for $i = 1, 2, \dots, B$.
- 4: Compute the empirical loss

$$\begin{aligned} \hat{\ell}(w_1, w_2) = & \widehat{\text{Var}} \left(v_{w_1}(\tilde{Z}^{(i)}(T)) - v_{w_1}(\tilde{Z}^{(i)}(0)) - \sum_{j=0}^{N-1} g_{w_2}(\tilde{Z}^{(i)}(hj)) \cdot \delta_j^{(i)} \right. \\ & \left. + \sum_{j=0}^{N-1} f \left(\tilde{Z}^{(i)}(hj), g_{w_2} \left(\tilde{Z}^{(i)}(hj) \right) \right) h \right). \end{aligned} \quad (58)$$

- 5: Compute the gradient $\partial \hat{\ell}(w_1, w_2) / \partial w_1, \partial \hat{\ell}(w_1, w_2) / \partial w_2$ and update w_1, w_2 using the chosen optimization solver.
 - 6: Update $z_{k+1}^{(i)}$ as the end point of the path $\tilde{Z}^{(i)}$: $z_{k+1}^{(i)} \leftarrow \tilde{Z}^{(i)}(T)$.
 - 7: **end for**
 - 8: **return** Functions $v_{w_1}(\cdot)$ and $g_{w_2}(\cdot)$.
-

After the parameters values w_1 and w_2 have been determined, our proposed policy is the following:

$$\bar{\theta}_{w_2}(z) = \arg \max_{\theta \in \Theta} (\theta \cdot g_{w_2}(z) - c(z, \theta)), \quad z \in \mathbb{R}_+^d.$$

6 Three families of test problems

Here we specify three families of test problems for which numerical results will be presented later (see Section 7). Each family consists of RBM drift control problems indexed by $d = 1, 2, \dots$, where d is the dimension of the orthant that serves as the problem's state space. The first of the three problem families, specified in Section 6.1, is characterized by a feed-forward network structure and linear cost of control. Recapitulating earlier work by Ata [2], Section 6.2 explains the interpretation of such problems as "heavy traffic" limits of input control problems for certain feed-forward queueing networks.

Our second family of test problems is identical to the first one except that now the cost of control is quadratic rather than linear. The exact meaning of that phrase will be spelled out in Section 6.3, where we also explain the interpretation of such problems as heavy traffic limits of dynamic pricing problems for queueing networks. In Section 6.4, we describe two parametric families of policies with special structure that will be used later for comparison purposes in our numerical study. Finally,

Section 6.5 specifies our third family of test problems, which have a separable structure that allows them to be solved exactly by analytical means. Such problems are of obvious value for evaluating the accuracy of our computational method.

6.1 Main example with linear cost of control

We consider a family of test problems with parameters $K = 0, 1, \dots$, attaching to each such problem the index d (mnemonic for *dimension*) $= K + 1$. Problem d has state space \mathbb{R}_+^d and the $d \times d$ reflection matrix

$$R = \begin{bmatrix} 1 & & & \\ -p_1 & 1 & & \\ \vdots & & \ddots & \\ -p_K & & & 1 \end{bmatrix}, \quad (59)$$

where $p_1, \dots, p_K > 0$ and $p_1 + \dots + p_K = 1$. Also, the set of drift vectors available in each state is

$$\Theta = \prod_{k=0}^K [\underline{\theta}_k, \bar{\theta}_k]. \quad (60)$$

where the lower limit $\underline{\theta}_k$ and upper limit $\bar{\theta}_k$ are as specified in Section 6.2 below. Similarly, the $d \times d$ covariance matrix A for problem d is as specified in Section 6.2. Finally, the cost function for problem d has the linear form

$$c(z, \theta) = h^\top z + c^\top \theta \quad \text{where } h, c \in \mathbb{R}_+^d. \quad (61)$$

That is, the cost rate $c(Z(t), u(Z(t)))$ that the system manager incurs under policy u at time t is linear in both the state vector $Z(t)$ and the chosen drift rate $u(Z(t))$.

In either the discounted control setting or the ergodic control setting, inspection of the HJB equation displayed earlier in Section 3 shows that, given this linear cost structure, there exists an optimal policy $u^*(\cdot)$ such that

$$\text{either } u_k^*(z) = \underline{\theta}_k \quad \text{or} \quad u_k^*(z) = \bar{\theta}_k \quad (62)$$

for each state $z \in \mathbb{R}_+^{K+1}$ and each component $k = 0, 1, \dots, K$.

In the next section we explain how drift control problems of the form specified here arise as heavy traffic limits in queueing theory. Strictly speaking, however, that interpretation of the test problems is inessential to the main subject of this paper: the computational results presented in Section 7 can be read without reference to the queueing theoretic interpretations of our test problems.

6.2 Interpretation as heavy traffic limits of queueing network control problems

Let us consider the feed-forward queueing network model of a make-to-order production system portrayed in Figure 1. There are $d = K + 1$ buffers, represented by the open-ended rectangles, indexed by $k = 0, 1, \dots, K$. Each buffer has a dedicated server, represented by the circles in Figure 1. Arriving jobs wait in their designated buffer if the server is busy. There are two types of jobs arriving to the system: regular versus thin streams. Thin stream jobs have the same service time distributions as the regular jobs, but they differ from the regular jobs in two important ways: First, thin stream jobs can be turned away upon arrival. That is, a system manager can exercise admission control in this manner, but in contrast, she must admit all regular jobs arriving to the system. Second, the volume of thin stream jobs is smaller than that of the regular jobs; see Assumption 1.

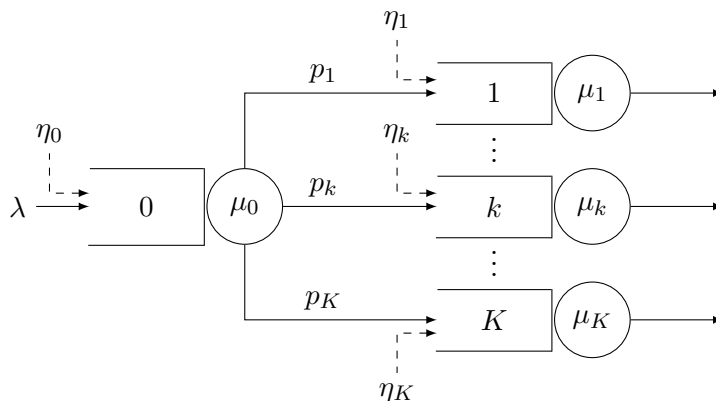


Figure 1: A feedforward queueing network with thin arrival streams.

Regular jobs enter the systems only through buffer zero, as shown by the solid arrow pointing to buffer zero in Figure 1. A renewal process $E = \{E(t) : t \geq 0\}$ models the cumulative number of regular jobs arriving to the system over time. We let λ denote the arrival rate and a^2 denote the squared coefficient of variation of the interarrival times for the regular jobs. The thin stream jobs arrive to buffer k (as shown by the dashed arrows in Figure 1) according to the renewal process $A_k = \{A_k(t) : t \geq 0\}$ for $k = 0, 1, \dots, K$. We let η_k denote the arrival rate and b_k^2 denote the squared coefficient of variation of the interarrival times for renewal process A_k .

Jobs in buffer k have i.i.d. general service time distributions with mean m_k and squared coefficient of variation $s_k^2 \geq 0$, $k = 0, 1, \dots, K$; $\mu_k = 1/m_k$ is the corresponding service rate. We let $S_k = \{S_k(t) : t \geq 0\}$ denote the renewal process associated with the service completions by server k for $k = 1, \dots, K$. To be specific, $S_k(t)$ denotes the number of jobs server k processes by time t if it incurs no idleness during $[0, t]$. The jobs in each buffer are served on a first-come-first-served (FCFS) basis, and servers work continuously unless their buffer is empty. After receiving service, jobs in buffer zero join buffer k with probability p_k , $k = 1, 2, \dots, K$, independently of other events. This probabilistic routing structure is captured by a vector-valued process $\Phi(\cdot)$ where $\Phi_k(\ell)$ denotes the total number of jobs routed to buffer k among the first ℓ jobs served by server zero for $k = 1, \dots, K$ and $\ell \geq 1$. We let $p = (p_k)$ denote the K -dimensional vector of routing probabilities. Jobs in buffers $1, \dots, K$

leave the system upon receiving service.

As stated earlier, the system manager makes admission control decisions for thin stream jobs. Turning away a thin stream job arriving to buffer k (externally) results in a penalty of c_k . For mathematical convenience, we model admission control decisions as if the system manager can simply “turn off” each of the thin stream arrival processes as desired. In particular, we let $\Delta_k(t)$ denote the cumulative amount of time that the (external) thin stream input to buffer k is turned off during the interval $[0, t]$. Thus, the vector-valued process $\Delta = (\Delta_k)$ represents the admission control policy. Similarly, we let $T_k(t)$ denote the cumulative amount of time server k is busy during the time interval $[0, t]$, and $I_k(t) = t - T_k(t)$ denotes the cumulative amount of idleness that server K incurs during $[0, t]$.

Letting $Q_k(t)$ denote the number of jobs in buffer k at time t , the vector-valued process $Q = (Q_k)$ will be called the queue-length process. Given a control $\Delta = (\Delta_k)$, assuming $Q(0) = 0$, it follows that

$$Q_0(t) = E(t) + A_0(t - \Delta_0(t)) - S_0(T_0(t)) \geq 0, \quad t \geq 0, \quad (63)$$

$$Q_k(t) = A_k(t - \Delta_k(t)) + \Phi_k(S_0(T_0(t))) - S_k(T_k(t)) \geq 0, \quad t \geq 0, \quad k = 1, \dots, K. \quad (64)$$

Moreover, the following must hold:

$$I(\cdot) \text{ is continuous and nondecreasing with } I(0) = 0, \quad (65)$$

$$I_k(\cdot) \text{ only increases at those times } t \text{ when } Q_k(t) = 0, \quad k = 0, 1, \dots, K, \quad (66)$$

$$\Delta_k(t) - \Delta_k(s) \leq t - s, \quad 0 \leq s \leq t < \infty, \quad k = 0, 1, \dots, K. \quad (67)$$

$$I, \Delta \text{ are non-anticipating.} \quad (68)$$

The system manager also incurs a holding cost at rate h_k per job in buffer k per unit of time. We use the processes $\xi = \{\xi(t), t \geq 0\}$ as a proxy for the cumulative cost under a given admission control policy $\Delta(\cdot)$, where

$$\xi(t) = \sum_{k=0}^K c_k \eta_k \Delta_k(t) + \sum_{k=0}^K \int_0^t h_k Q_k(s) ds, \quad t \geq 0.$$

This is an approximation of the realized cost because the first term on the right-hand side replaces the admission control penalties actually incurred with their means.

In order to derive the approximating Brownian control problem, we consider a sequence of systems indexed by a system parameter $n = 1, 2, \dots$; we attach a superscript of n to various quantities of interest. Following the approach used by Ata [2], we assume that the sequence of systems satisfies the following heavy traffic assumption.

Assumption 1. *For $n \geq 1$, we have that*

$$\lambda^n = n\lambda, \eta_k^n = \eta_k \sqrt{n} \text{ and } \mu_k^n = n\mu_k + \sqrt{n}\beta_k, \quad k = 0, 1, \dots, K,$$

where λ, μ_k, η_k and β_k are nonnegative constants. Moreover, we assume that

$$\lambda = \mu_0 = \frac{\mu_k}{p_k} \text{ for } k = 1, \dots, K.$$

One starts the approximation procedure by defining suitably centered and scaled processes. For $n \geq 1$, we define

$$\begin{aligned} \hat{E}^n(t) &= \frac{E^n(t) - \lambda^n t}{\sqrt{n}} & \text{and} & \quad \hat{\Phi}^n(q) = \frac{\Phi([nq]) - p([nq])}{\sqrt{n}}, \quad t \geq 0, \quad q \geq 0, \\ \hat{A}_k^n &= \frac{Q^n(t)}{\sqrt{n}} & \text{and} & \quad \hat{S}_k^n(t) = \frac{S_k^n(t) - \mu_k^n t}{\sqrt{n}}, \quad t \geq 0, \quad k = 0, 1, \dots, K, \\ \hat{Q}^n(t) &= \frac{Q^n(t)}{\sqrt{n}} & \text{and} & \quad \hat{\xi}^n(t) = \frac{\xi^n(t)}{\sqrt{n}}, \quad t \geq 0. \end{aligned}$$

In what follows, we assume

$$T_k^n(t) = t - \frac{1}{\sqrt{n}} I_k(t) + o\left(\frac{1}{\sqrt{n}}\right), \quad t \geq 0, \quad k = 0, 1, \dots, K, \quad (69)$$

where $I_k(\cdot)$ is the limiting idleness process for server k ; see [18] for an intuitive justification of (69).

Then, defining

$$\begin{aligned} \chi_0^n(t) &= \hat{E}^n(t) + \hat{A}_0^n(t - \Delta_0(t)) + \hat{S}_0^n(T_0^n(t)), \quad t \geq 0, \\ \chi_k^n(t) &= \hat{A}_k^n(t - \Delta_0(t)) + \Phi_k\left(\frac{1}{n} \hat{S}_0^n(T_0^n(t))\right) + p_k \hat{S}_0^n(T_0^n(t)) \\ &\quad - \hat{S}_k^n(T_k^n(t)), \quad t \geq 0, \quad k = 1, 2, \dots, K, \end{aligned}$$

and using Equations (63) - (64) and (69), it is straightforward to derive the following for $t \geq 0$ and $k = 1, \dots, K$:

$$\hat{Q}_0^n(t) = \chi_0^n(t) + (\eta_0 - \beta_0)t - \eta_0 \Delta_0(t) + \mu_0 I_0(t) + o(1), \quad (70)$$

$$\hat{Q}_k^n(t) = \chi_k^n(t) + (\eta_k + p_k \beta_0 - \beta_k)t - \eta_k \Delta_k(t) + \mu_k I_k(t) - p_k \mu_0 I_0(t) + o(1). \quad (71)$$

Moreover, it follows from Equation (67) that $\Delta_k(t)$ is absolutely continuous. We denote its density by $\delta_k(\cdot)$, *i.e.*,

$$\Delta_k(t) = \int_0^t \delta_k(s) ds, \quad t \geq 0, \quad k = 0, 1, \dots, K,$$

where $\delta_k(t) \in [0, 1]$. Using this, we write

$$\hat{\xi}^n(t) = \sum_{k=0}^K \int_0^t v_k \eta_k \delta_k(s) ds + \sum_{k=0}^K \int_0^t h_k \hat{Q}_k^n(s) ds, \quad t \geq 0. \quad (72)$$

Then passing to the limit formally as $n \rightarrow \infty$, and denoting the weak limit of $(\hat{Q}^n, \hat{X}^n, \hat{\xi}^n)$ by

(Z, χ, ξ) , where χ is a $(K + 1)$ -dimensional driftless Brownian motion with covariance matrix (see Appendix C for its derivation)

$$A = \mu_0 \begin{bmatrix} s_0^2 + a^2 & -p_1 s_0^2 & \cdots & \cdots & -p_K s_0^2 \\ -p_1 s_0^2 & p_1(1 - p_1) + p_1^2 s_0^2 + p_1 s_1^2 & p_1 p_2 (s_0^2 - 1) & \cdots & p_1 p_K (s_0^2 - 1) \\ \vdots & p_1 p_2 (s_0^2 - 1) & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & p_{K-1} p_K (s_0^2 - 1) \\ -p_K s_0^2 & p_1 p_K (s_0^2 - 1) & \cdots & \cdots & p_K(1 - p_K) + p_K^2 s_0^2 + p_K s_K^2 \end{bmatrix},$$

we deduce from (70) - (71) and (72) that

$$\begin{aligned} Z_0(t) &= \chi_0(t) + (\eta_0 - \beta_0)t - \int_0^t \eta_0 \delta_0(s) ds + \mu_0 I_0(t), \\ Z_k(t) &= \chi_k(t) + (\eta_k + p_k \beta_0 - \beta_k)t - \int_0^t \eta_k \delta_k(s) ds + \mu_k I_k(t), \quad k = 1, \dots, K, \\ \xi(t) &= \sum_{k=0}^K \int_0^t v_k \eta_k \delta_k(s) ds + \sum_{k=0}^K \int_0^t h_k Z_k(s) ds. \end{aligned}$$

In order to streamline the notation, we make the following change of variables:

$$\begin{aligned} Y_k(t) &= \mu_k I_k(t), \quad k = 0, \dots, K, \\ \theta_0(t) &= \eta_0 \delta_0(t) - (\eta_0 - \beta_0), \quad t \geq 0, \\ \theta_k(t) &= \eta_k \delta_k(t) - (\eta_k + p_k \beta_0 - \beta_k) \end{aligned}$$

and let

$$\begin{aligned} \underline{\theta}_0 &= \beta_0 \quad \text{and} \quad \bar{\theta}_k = \beta_0 - \eta_0, \\ \underline{\theta}_k &= \beta_k - p_k \beta_0 \quad \text{and} \quad \bar{\theta}_k = \beta_k - \eta_k - p_k \beta_0, \quad k = 1, \dots, K. \end{aligned}$$

Lastly, we define the set of negative drift vectors available to the system manager as in Equation (60). As a result, we arrive at the following Brownian system model:

$$Z_0(t) = \chi_0(t) - \int_0^t \theta_0(s) ds + Y_0(t), \quad t \geq 0, \quad (73)$$

$$Z_k(t) = \chi_k(t) - \int_0^t \theta_k(s) ds - Y_k(t) + p_k Y_0(t), \quad k = 1, \dots, K, \quad (74)$$

which can be written as in Equation (2) with $d = K + 1$, where the reflection matrix R is given by Equation (59). Moreover, the processes Y, Z inherit properties in Equation (63) - (73) from their pre-limit counterparts in the queueing model, cf. Equations (65) - (66).

To minimize technical complexity, we restrict attention to stationary Markov control policies as

done in Section 3. That is, $\theta(t) = u(Z(t))$ for $t \geq 0$ for some policy function $u : \mathbb{R}_+^d \rightarrow \Theta$. Then, defining $c = (c_0, c_1, \dots, c_K)^\top$, $h = (h_0, h_1, \dots, h_K)^\top$ and

$$c(z, \theta) = h^\top z + c^\top \theta,$$

as in Equation (61), the cumulative cost incurred over the time interval $[0, t]$ under policy u can be written as in Equation (10). Note that $C^u(t)$ and $\xi(t)$ differ only by a term that is independent of the control. Given $C^u(t)$, one can formulate the discounted control problem as done in Section 3.1. Similarly, the ergodic control problem can be formulated as done in Section 3.2.

Interpreting the solution of the drift control problem in the context of the queueing network formulation. Because the instantaneous cost rate $c(z, \theta)$ is linear in the control, inspection of the HJB equation reveals that the optimal control is of bang-bang nature. That is, $\theta_k(t) \in \{\underline{\theta}_k, \bar{\theta}_k\}$ for all k, t as stated in Equation (62). This can be interpreted in the context of the queueing network displayed in Figure 1 as follows: For $k = 0, 1, \dots, K$, whenever $\theta_k(t) = \bar{\theta}_k$, the system manager turns away the thin stream jobs arriving to buffer k externally, i.e., she shuts off the renewal process $A_k(\cdot)$ at time t . Otherwise, she admits them to the system. Of course, the optimal policy is determined by the gradient $\nabla V(z)$ of the value function through the HJB equation, which we solve for using the method described in Section 5.

6.3 Related example with quadratic cost of control

Çelik and Maglaras [12] and Ata and Barjesteh [3] advance formulations where a system manager controls the arrival rate of customers to a queueing system by exercising dynamic pricing. One can follow a similar approach for the feed-forward queueing networks displayed in Figure 1 with suitable modifications, e.g., the dashed arrows also correspond to arrivals of regular jobs. This ultimately results in a problem of drift control for RBM with the cost of control

$$c(\theta, z) = \sum_{k=0}^K \alpha_k (\theta_k - \underline{\theta}_k)^2 + \sum_{k=0}^K h_k z_k, \quad (75)$$

where $\underline{\theta}$ is the drift rate vector corresponding to a nominal price vector.

6.4 Two parametric families of benchmark policies

Recall the optimal policy can be characterized as

$$u^*(z) = \arg \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\}, \quad z \in \mathbb{R}_+^d. \quad (76)$$

The benchmark policy for the main test problem. In our main test problem (see Section

6.1), we have $c(z, \theta) = h^\top z + c^\top \theta$. Therefore, it follows from (76) that for $k = 0, 1, \dots, K$,

$$u_k^*(z) = \begin{cases} \bar{\theta}_k & \text{if } (\nabla V(z))_k \geq c_k, \\ \underline{\theta}_k & \text{otherwise.} \end{cases}$$

Namely, the optimal policy is of bang-bang type. Therefore, we consider the following linear-boundary policies as our benchmark policies: For $k = 0, 1, \dots, K$,

$$u_k^{\text{lbp}}(z) = \begin{cases} \bar{\theta}_k & \text{if } \beta_k^\top z \geq c_k, \\ \underline{\theta}_k & \text{otherwise,} \end{cases}$$

where $\beta_0, \beta_1, \dots, \beta_K \in \mathbb{R}^{K+1}$ are vectors of policy parameters to be tuned.

In our numerical study, we focus attention on the symmetric case where

$$\begin{aligned} h_0 &> h_1 = \dots = h_K, \\ c_0 &= c_1 = \dots = c_K, \\ p_1 &= \dots = p_K = \frac{1}{K}, \\ \underline{\theta}_1 &= \dots = \underline{\theta}_K, \\ \bar{\theta}_1 &= \dots = \bar{\theta}_K. \end{aligned}$$

Due to this symmetry, the downstream buffers look identical. As such, we restrict attention to parameter vectors of the following form:

$$\begin{aligned} \beta_0 &= (\phi_1, \phi_2, \dots, \phi_2), \text{ and} \\ \beta_i &= (\phi_3, \phi_4, \dots, \phi_4, \phi_5, \phi_4, \dots, \phi_4) \text{ where } \phi_5 \text{ is the } i + 1^{\text{st}} \text{ element of } \beta_i \text{ for } i = 1, \dots, K. \end{aligned}$$

The parameter vector β_0 , which is used to determine the benchmark policy for buffer zero, has two distinct parameters: ϕ_1 and ϕ_2 . In considering the policy for buffer zero, ϕ_1 captures the effect of its own queue length, whereas ϕ_2 captures the effects of the downstream buffers $1, \dots, K$. We use a common parameter for the downstream buffers because they look identical from the perspective of buffer zero. Similarly, the parameter vector β_i ($i = 1, \dots, K$) has three distinct parameters: ϕ_3 , ϕ_4 and ϕ_5 , where ϕ_3 is used as the multiplier for buffer zero (the upstream buffer), ϕ_5 is used to capture the effect of buffer i itself and ϕ_4 is used for all other downstream buffers. Note that all β_i use the same three parameters ϕ_3 , ϕ_4 and ϕ_5 for $i = 1, \dots, K$. They only differ with respect to the position of ϕ_5 , i.e., it is in the $i + 1^{\text{st}}$ position for β_i .

In summary, the benchmark policy uses five distinct parameters in the symmetric case. This allows us to do a brute-force search via simulation on a five-dimensional grid regardless of the number of buffers.

The benchmark policy for the test problem with the quadratic cost of control. In this case, substituting Equation (75) into Equation (76) gives the following characterization of the

optimal policy:

$$u_k^*(z) = \underline{\theta}_k + \frac{(\nabla V(z))_k}{2\alpha_k}, \quad k = 0, 1, \dots, K. \quad (77)$$

Namely, the optimal policy is affine in the gradient. Therefore, we consider the following affine-rate policies as our benchmark policies: For $k = 0, 1, \dots, K$,

$$u_k^{\text{arp}}(z) = \underline{\theta}_k + \beta_k^\top z,$$

where $\beta_0, \beta_1, \dots, \beta_K \in \mathbb{R}^{K+1}$ are vectors of policy parameters to be tuned. We truncate this at the upper bound $\bar{\theta}_k$ if needed.

We focus attention on the symmetric case for this problem formulation too. To be specific, we assume

$$\begin{aligned} h_0 &> h_1 = \dots = h_K, \\ \alpha_0 &= \alpha_1 = \dots = \alpha_K, \\ p_1 &= \dots = p_K = \frac{1}{K}, \\ \underline{\theta}_1 &= \dots = \underline{\theta}_K, \\ \bar{\theta}_1 &= \dots = \bar{\theta}_K. \end{aligned}$$

Due to this symmetry, the downstream buffers look identical. As such, we restrict attention to parameter vectors of the following form:

$$\begin{aligned} \beta_0 &= (\phi_1, \phi_2, \dots, \phi_2), \text{ and} \\ \beta_i &= (\phi_3, \phi_4, \dots, \phi_4, \phi_5, \phi_4, \dots, \phi_4) \text{ where } \phi_5 \text{ is the } i + 1^{\text{st}} \text{ element for } i = 1, \dots, K. \end{aligned}$$

As done for the first benchmark policy above, this particular form of the parameter vectors can be justified using the symmetry as well.

6.5 Parallel-server test problems

In this section, we consider a problem whose solution can be derived analytically by considering a one-dimensional problem. To be specific, we consider the parallel-server network that consists of K identical single-server queues as displayed in Figure 2. Clearly, this network can be decomposed into K separate single-server queues, leading to K separate one-dimensional problem formulations, which can be solved analytically, see Appendix D for details. For this example we have that $R = I_{d \times d}$ and $A = I_{d \times d}$. In addition, we assume that the action space Θ and the cost function $c(z, \theta)$ are the same as above.




Figure 2: A decomposable parallel-server queueing network.

7 Computational results

For the test problems introduced in Section 6, we now compare the performance of policies derived using our method (see Section 5) with the best benchmark we could find. The results show that our method performs well, and it remains computationally feasible up to at least dimension $d = 30$. We implement our method using three-layer or four-layer neural networks with the elu activation function [38] in Tensorflow 2 [1], and using code adapted from that of Han et al. [17] and Zhou et al. [50]; see Appendix E for further details of our implementation.¹

For our main test problem with linear cost of control (introduced previously in Section 6.1), and also for its variant with quadratic cost of control (Section 6.3), the following parameter values are assumed: $h_0 = 2$, $h_k = 1.9$ for $k = 1, \dots, K$, $c_k = 1$ for $k = 0, \dots, K$, and $p_k = 1/K$ for $k = 1, \dots, K$. Also, the reflection matrix R and the covariance matrix A for those families of problems are as follows:

$$R = \begin{bmatrix} 1 & & & & \\ -1/K & 1 & & & \\ \vdots & & \ddots & & \\ -1/K & & & & 1 \end{bmatrix} \text{ and } A = \begin{bmatrix} 1 & 0 & \cdots & \cdots & 0 \\ 0 & 1 & -\frac{1}{K^2} & \cdots & -\frac{1}{K^2} \\ \vdots & -\frac{1}{K^2} & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & -\frac{1}{K^2} \\ 0 & -\frac{1}{K^2} & \cdots & \cdots & 1 \end{bmatrix}.$$

However, as stated previously in Section 6.5, the reflection matrix and covariance matrix for our parallel-server test problems are $R = I_{d \times d}$ and $A = I_{d \times d}$. Problems in that third class have $K = d$ buffers indexed by $k = 1, \dots, K$, and we set $h_1 = 2$ and $h_k = 1.9$ for $k = 2, \dots, K$.

7.1 Main test problem with linear cost of control

For our main test problem with linear cost of control (Section 6.1), we take

$$\underline{\theta}_k = 0 \text{ and } \bar{\theta}_k = b \in \{2, 10\} \text{ for all } k.$$

Also, interest rates $r = 0.1$ and $r = 0.01$ will be considered in the discounted case.

To begin, let us consider the simple case where $K = 0$ (that is, there are no downstream buffers in the queueing network interpretation of the problem) and hence $d = 1$. In this case, one can solve the

¹Our code is available in <https://github.com/nian-si/RBMSolver>.

HJB equation analytically; see Appendix D for details. For the discounted formulation with $r = 0.1$, Figure 3 compares the derivative of the value function computed using the analytical solution, shown in blue, with the neural network approximation for it that we computed using our method, shown in red. (The comparisons for $r = 0.01$ and the ergodic control case are similar.)

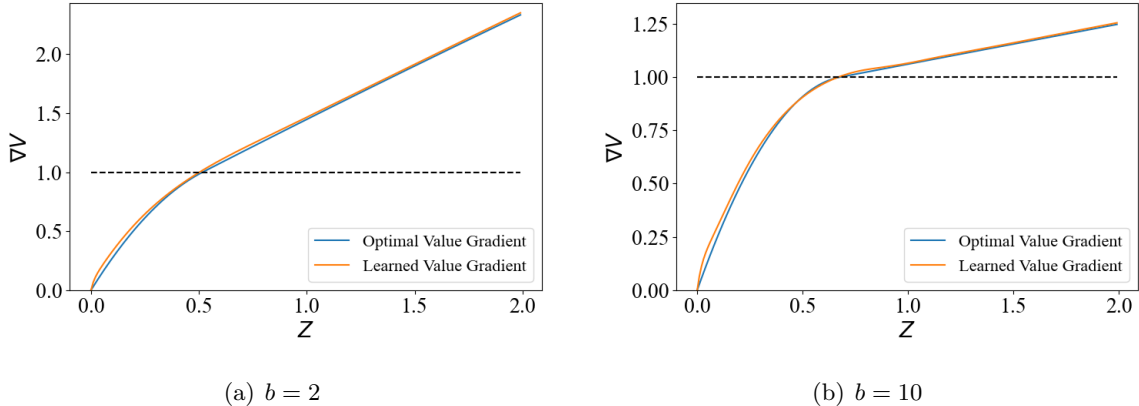


Figure 3: Comparison between the derivative $G_w(\cdot)$ learned from neural networks and the derivative of the optimal value function for the case of $d = 1$ and $r = 0.1$. The dotted line indicates the cost $c_0 = 1$. When the value function gradient is above this dotted line, the optimal control is $\theta = b$, and otherwise it is $\theta = 0$.

Combining Figure 3 with Equation (76), one sees that the policy derived using our method is close to the optimal policy. Table 1 reports the simulated performance with standard errors of these two policies based on four million sample paths and using the same discretization of time as in our computational method. Specifically, we report the long-run average cost under each policy in the ergodic control case, and report the simulated value $V(0)$ in the discounted case. To repeat, the benchmark policy in this case is the optimal policy determined analytically, but not accounting for the discretization of the time scale. Of course, all the performance figures reported in the table are subject to simulation errors. Finally, it is worth noting that our method took less than one hour to compute its policy recommendations using a 10-CPU core computer.

Table 1: Performance comparison of our proposed policy with the benchmark policy in the one-dimensional case ($K = 0$).

		Ergodic	$r = 0.01$	$r = 0.1$
$b = 2$	Our policy	1.455 ± 0.0006	145.3 ± 0.05	14.29 ± 0.004
	Benchmark	1.456 ± 0.0006	145.3 ± 0.05	14.29 ± 0.004
$b = 10$	Our policy	1.375 ± 0.0007	137.2 ± 0.06	13.56 ± 0.005
	Benchmark	1.374 ± 0.0007	137.2 ± 0.06	13.56 ± 0.005

Let us consider now the two-dimensional case ($K = 1$), where the optimal policy is unknown.

Table 2: Performance comparison of our proposed policy with the benchmark policy in the two-dimensional case ($K = 1$).

		Ergodic	$r = 0.01$	$r = 0.1$
$b = 2$	Our policy	2.471 ± 0.0008	246.6 ± 0.08	24.28 ± 0.006
	Benchmark	2.473 ± 0.0008	246.8 ± 0.08	24.29 ± 0.006
$b = 10$	Our policy	2.338 ± 0.0009	233.3 ± 0.09	23.10 ± 0.006
	Benchmark	2.338 ± 0.0009	233.6 ± 0.09	23.10 ± 0.006

Therefore, we compare our method with the best benchmark we could find: the linear boundary policy described in Section 6.3. In the two-dimensional case, the linear boundary policy reduces to the following:

$$\theta_0(z) = b\mathbb{I}\left\{\beta_0^\top z \geq 1\right\} \text{ and } \theta_1(z) = b\mathbb{I}\left\{\beta_1^\top z \geq 1\right\}.$$

Through simulation, we perform a brute-force search to identify the best values of β_0 and β_1 . The policies for $b = 2$ and $b = 10$ are shown in Figures 4 and 5, respectively, for the discounted control case with $r = 0.1$. Our proposed policy sets the drift to b in the red region and to zero in the blue region, whereas the best-linear boundary policy is represented by the white-dotted line. That is, the benchmark policy sets the drift to b in the region above and to the right of the dotted line, and sets it to zero below and left of the line. Table 2 presents the costs with standard errors of the benchmark policy and our proposed policy obtained in a simulation study. The two policies have similar performance. Our method takes about one hour to compute policy recommendations using a 10-CPU core computer.




Figure 4: Graphical representation of the policy learned from neural networks and the benchmark policy for the case $b = 2, d = 2$ and $r = 0.1$




Figure 5: Graphical representation of the policy learned from neural networks and the benchmark policy for the case $b = 10, d = 2$ and $r = 0.1$

Finally, let us consider the six-dimensional case ($K = 5$), where the linear boundary policy reduces to

$$\theta_i(z) = b\mathbb{I}\left\{\beta_i^\top z \geq 1\right\} \text{ for } i = 0, 1, 2, \dots, 5.$$

Although there appear to be 36 parameters to be tuned, recall that we reduced the number of parameters to five in Section 6.4 by exploiting symmetry. This makes the brute-force search computationally feasible. Table 3 compares the performance with standard errors of our proposed policies with the benchmark policies. They have similar performance. In this case, the running time for our method is several hours using a 10-CPU computer.

Table 3: Performance comparison of our proposed policy with the benchmark policy in the six-dimensional case $d = 6$ ($K = 5$).

		Ergodic	$r = 0.01$	$r = 0.1$
$b = 2$	Our policy	7.927 ± 0.001	791.0 ± 0.1	77.83 ± 0.01
	Benchmark	7.927 ± 0.001	791.3 ± 0.1	77.83 ± 0.01
$b = 10$	Our policy	7.565 ± 0.0016	754.8 ± 0.15	74.61 ± 0.01
	Benchmark	7.525 ± 0.0016	751.7 ± 0.15	74.32 ± 0.01

7.2 Test problems with quadratic cost of control

In this section, we consider the test problem introduced in Section 6.3, for which we set $\alpha_k = 1$ and $\underline{\theta}_k = 1$ for all k . As in the previous treatment of our main test example, we report results for the cases of $d = 1, 2, 6$ in Tables 4, 5, 6, respectively, where the benchmark policies are the affine-rate policies discussed in Section 6.4, with policy parameters optimized via simulation and brute-force search. We observe that our proposed policies outperform the best affine-rate policies by very small margins in

all cases. In the one-dimensional ergodic control case ($K = 0$), we obtain analytical solutions to the

Table 4: Performance comparison of our proposed policy with the benchmark policy in the case of quadratic cost of control and $d = 1$.

	Ergodic	$r = 0.01$	$r = 0.1$
Our policy	0.757 ± 0.0004	75.53 ± 0.03	7.415 ± 0.003
Benchmark	0.758 ± 0.0004	75.67 ± 0.03	7.427 ± 0.003

Table 5: Performance comparison of our proposed policy with the benchmark policy in the case of quadratic cost of control and $d = 2$ ($K = 1$)

	Ergodic	$r = 0.01$	$r = 0.1$
Our policy	1.216 ± 0.0005	121.3 ± 0.04	11.94 ± 0.003
Benchmark	1.219 ± 0.0005	121.7 ± 0.05	11.96 ± 0.003

Table 6: Performance comparison of our proposed policy with the benchmark policy in the case of quadratic cost of control and $d = 6$ ($K = 5$)

	Ergodic	$r = 0.01$	$r = 0.1$
Our policy	3.863 ± 0.0008	385.7 ± 0.08	37.92 ± 0.006
Benchmark	3.874 ± 0.0008	386.9 ± 0.08	38.04 ± 0.006

RBM control problem in closed form by solving the HJB equation directly, which reduces to a first-order ordinary differential equation in this case; see Appendix D for details. Figure 6 compares the derivative of the optimal value function (derived in closed form) with its approximation via neural networks in the ergodic case. Combining Figure 6 with Equation (77) reveals that our proposed policy is close to the optimal policy.

In the two-dimensional case, our proposed policy is shown in Figure 7 for the ergodic case, with contour lines showing the state vectors (z_0, z_1) for which the policy chooses successively higher drift rates. The white dotted lines similarly show the states (z_0, z_1) for which our benchmark policy (that is, the best affine-rate policy) chooses the drift rate $\theta_k = 1.5$ (for $k = 0$ in the left panel and $k = 1$ in the right panel).




Figure 6: Comparison of the gradient approximation $G_w(\cdot)$ learned from neural networks with the derivative of the optimal value function for the ergodic control case with quadratic cost of control in the one-dimensional case ($d = 1$).




Figure 7: Graphical representation of the policy learned from neural networks and the benchmark policy for the ergodic case with $d = 2$.

7.3 Parallel-server test problems

This section focuses on parallel-server test problems (see Section 6.5) to demonstrate our method’s scalability. As illustrated in Figure 2, the parallel-server networks are essentially K independent copies of the one-dimensional case. We present the results in Table 7 for $d = 30$ and linear cost of control. When $b = 2$, our policies perform almost equally well as the optimal policy, while for $b = 10$, our policies perform within 1% of the optimal policy. The run-time for our method is about one day in this case using a 20-CPU core computer.

For quadratic cost of control, we are able to solve the test problems up to at least 100 dimensions. The results for $d = 100$ are given in Table 8, where the benchmark policies are the best affine-rate policies (see Section 6.4). The performance of our policy is within 1% of the benchmark performance. The run-time for our method is several days in this case using a 30-CPU core computer.

Table 7: Performance comparison between our proposed policy and the benchmark policy for 30-dimensional parallel-server test problems with linear cost of control.

		Ergodic	$r = 0.01$	$r = 0.1$
$b = 2$	Our policy	42.56 ± 0.003	4247 ± 0.3	417.3 ± 0.02
	Benchmark	42.52 ± 0.003	4244 ± 0.3	417.2 ± 0.02
$b = 10$	Our policy	40.53 ± 0.004	4054 ± 0.4	399.4 ± 0.026
	Benchmark	40.23 ± 0.004	4018 ± 0.4	396.7 ± 0.024

Table 8: Performance comparison between our proposed policy and the benchmark policy for 100-dimensional parallel-server test problems with quadratic cost of control.

		Ergodic	$r = 0.01$	$r = 0.1$
	Our policy	72.74 ± 0.003	7258.3 ± 0.3	712.4 ± 0.02
	Benchmark	72.53 ± 0.003	7237.3 ± 0.3	710.2 ± 0.02

8 Concluding remarks

Consider the general drift control problem formulated in Section 3, assuming specifically that the instantaneous cost rate $c(z, \theta)$ is linear in θ , and further assuming that the set of available drift vectors is a rectangle $\Theta = [0, b_1] \times \cdots \times [0, b_d]$. If one relaxes such a problem by letting $b_i \uparrow \infty$ for one or more i , then one obtains what is called a *singular* control problem, cf. Kushner and Martins [31]. Optimal policies for such problems typically involve the imposition of *endogenous* reflecting barriers, that is, reflecting barriers imposed by the system controller in order to minimize cost, in addition to exogenous reflecting barriers that may be imposed to represent physical constraints in the motivating application.

There are many examples of queueing network control problems whose natural heavy traffic approximations involve singular control; see, for example, Krichagina and Taksar [30], Martins and Kushner [34], and Martins et al. [33]. Given that motivation, we intend to show in future work how the computational method developed in this paper for drift control can be extended in a natural way to treat singular control, and to illustrate that extension by means of queueing network applications.

Separately, the following are three desirable generalizations of the problem formulations pro-
pounded in Section 3 of this paper. Each of them is straightforward in principle, and we expect to see these extensions implemented in future work, perhaps in combination with mild additional restrictions on problem data. (a) Instead of requiring that the reflection matrix R have the Minkowski form (1), require only that R be a completely- \mathcal{S} matrix, which Taylor and Williams [41] showed is a necessary and sufficient condition for an RBM to be well defined. (b) Allow a more general

state space for the controlled process Z , such as the convex polyhedrons characterized by Dai and Williams [13]. (c) Remove the requirement that the action space Θ be bounded.

Appendix A Proof of Proposition 1

Proof. Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}^d$ be right continuous with left limits (rcll). Following Williams [46], we define the oscillation of f over an interval $[t_1, t_2]$ as follows:

$$Osc(f, [t_1, t_2]) = \sup \{|f(t) - f(s)| : t_1 \leq s < t \leq t_2\}, \quad (78)$$

where $|a| = \max_{i=1, \dots, d} |a_i|$ for any $a \in \mathbb{R}^d$. Then for two rccl functions f, g , the following holds:

$$Osc(f + g) \leq Osc(f) + Osc(g). \quad (79)$$

Also recall that the controlled RBM Z satisfies $Z(t) = X(t) + RY(t)$, where

$$X(t) = W(t) - \int_0^t \theta(s) ds, \quad t \geq 0. \quad (80)$$

Then it follows from Theorem 5.1 of Williams [46] that

$$\begin{aligned} Osc(Z, [0, t]) &\leq C Osc(X, [0, t]) \\ &\leq C Osc(W, [0, t]) + C\bar{\theta}t \end{aligned}$$

for some $C > 0$, where $\bar{\theta} = \sum_{l=1}^d (\bar{\theta}_l - \underline{\theta}_l)$ and $\underline{\theta}_l, \bar{\theta}_l$ are the minimal and maximal values on each dimension, and the second inequality follows from (79).

Let $\mathcal{O} = Osc(W, [0, t])$ and recall that we are interested in bounding the expectation $\mathbb{E}[|Z(t)|^n]$. To that end, note that

$$\begin{aligned} |Z(t) - Z(0)|^n &\leq C^n (\mathcal{O} + \bar{\theta}t)^n \\ &= C^n \sum_{k=0}^n \binom{n}{k} \mathcal{O}^k \bar{\theta}^{n-k} t^{n-k}. \end{aligned} \quad (81)$$

To bound $\mathbb{E}[\mathcal{O}^k]$, note that

$$\begin{aligned} \mathcal{O} &= \sup \{|W(t_2) - W(t_1)| : 0 \leq t_1 < t_2 \leq t\} \\ &\leq \sup \{W(s) : 0 \leq s \leq t\} - \inf \{W(s) : 0 \leq s \leq t\} \\ &\leq 2 \sup \{|W(s)| : 0 \leq s \leq t\} \\ &\leq 2 \sup \left\{ \sum_{l=1}^d |W_l(s)| : 0 \leq s \leq t \right\} \\ &\leq 2 \sum_{l=1}^d \sup \{|W_l(s)| : 0 \leq s \leq t\}. \end{aligned}$$

So, by the union bound, we write

$$\begin{aligned}\mathbb{P}(\mathcal{O} > x) &\leq \sum_{l=1}^d \mathbb{P}\left(\sup_{0 \leq s \leq t} W_l(s) > \frac{x}{2d}\right) + \sum_{l=1}^d \mathbb{P}\left(\inf_{0 \leq s \leq t} W_l(s) < -\frac{x}{2d}\right) \\ &\leq 4 \sum_{l=1}^d \mathbb{P}\left(W_l(t) > \frac{x}{2d}\right),\end{aligned}$$

where the last inequality follows from the reflection principle.

Thus,

$$\mathbb{E}[\mathcal{O}^k] = \int_0^\infty x^{k-1} \mathbb{P}(\mathcal{O} > x) dx \leq 4 \sum_{l=1}^d \int_0^\infty x^{k-1} \mathbb{P}\left(W_l(t) > \frac{x}{2d}\right) dx.$$

By change of variable $y = x/d$, we write

$$\begin{aligned}\mathbb{E}[\mathcal{O}^k] &\leq 4 \sum_{l=1}^d (2d)^k \int_0^\infty y^{k-1} \mathbb{P}(W_l(t) > y) dy \\ &= 4 \sum_{l=1}^d (2d)^k \mathbb{E}[|W_l(t)|^k] \\ &= \frac{4 (2d)^k 2^{k/2} t^{k/2} \Gamma\left(\frac{k+1}{2}\right)}{\sqrt{\pi}} \sum_{l=1}^d \sigma_{ll}^k,\end{aligned}$$

where Γ is the Gamma function, and the last equality is a well-known result; see, for example, Equation (12) in [47]. Substituting this into (81) gives the following:

$$\begin{aligned}\mathbb{E}[|Z(t) - Z(0)|^n] &\leq C^n \sum_{k=0}^n \frac{4 (2d)^k \binom{n}{k} 2^{k/2} t^{k/2} \Gamma\left(\frac{k+1}{2}\right) \bar{\theta}^{n-k} t^{n-k}}{\sqrt{\pi}} \left(\sum_{l=1}^d \sigma_{ll}^k\right) \\ &\leq \tilde{C}_n (t^n + 1).\end{aligned}\tag{82}$$

Letting $z = Z(0)$. We write

$$\begin{aligned}|Z(t)|^n &= |Z(t) - z + z|^n \leq (|Z(t) - z| + |z|)^n \\ &\leq \sum_{k=0}^n \binom{n}{k} |Z(t) - z|^k |z|^{n-k}.\end{aligned}$$

Using (82), we can therefore write

$$\begin{aligned}\mathbb{E}[|Z(t)|^n] &\leq \sum_{k=0}^n \binom{n}{k} \tilde{C}_k |z|^{n-k} (t^k + 1) \\ &\leq \hat{C}_n (1 + t^n).\end{aligned}$$

□

Appendix B Validity of HJB Equations

B.1 Discounted Control

Proposition 6. *Let $u \in \mathcal{U}$ be an admissible policy and V^u a C^2 solution of the associated PDE (14)-(15). If both V^u and its gradient have polynomial growth, then V^u satisfies (12).*

Proof. Applying Ito's formula to $e^{-rt} V^u(Z^u(t))$ and using Equation (7), we write

$$\begin{aligned} e^{-rT} V^u(Z^u(T)) - V^u(z) &= \int_0^T e^{-rt} (\mathcal{L}V^u(Z^u(t)) - u(Z^u(t)) \cdot \nabla V^u(Z^u(t)) - rV^u(Z^u(t))) dt \\ &\quad + \int_0^T e^{-rt} \mathcal{D}V^u(Z^u(t)) \cdot dY(t) + \int_0^T e^{-rt} \nabla V^u(Z^u(t)) \cdot dW(t). \end{aligned}$$

Then using (3)-(4) and (14)-(15), we arrive at the following:

$$e^{-rT} V^u(Z^u(T)) - V^u(z) = - \int_0^T e^{-rt} c(Z^u(t), u(Z^u(t))) dt + \int_0^T e^{-rt} \nabla V^u(Z^u(t)) \cdot dW(t). \quad (83)$$

Because ∇V^u has polynomial growth and the action space Θ is bounded, we have that

$$\mathbb{E}_z \left[\int_0^T e^{-rt} \nabla V^u(Z^u(t)) \cdot dW(t) \right] = 0;$$

see, for example, Theorem 3.2.1 of Oksendal [35]. Thus, taking the expectation of both sides of (83) yields

$$V^u(z) = \mathbb{E}_z \left[\int_0^T e^{-rt} c(Z^u(t), u(Z^u(t))) dt \right] + e^{-rT} \mathbb{E}_z [V^u(Z^u(T))].$$

Because V^u has polynomial growth and Θ is bounded, the second term on the right-hand side vanishes by as $T \rightarrow \infty$. Then, because c has polynomial growth and Θ is bounded, passing to the limit as $T \rightarrow \infty$ completes the proof. \square

Proposition 7. *If V is a C^2 solution of the HJB equation (16)-(17), and if both V and its gradient have polynomial growth, then V satisfies (13).*

Proof. First, consider an arbitrary admissible policy u and let V^u denote the solution of the associated PDE (14)-(15). By Proposition 6, we have that

$$V^u(z) = \mathbb{E}_z \left[\int_0^\infty e^{-rt} c(Z^u(t), u(Z^u(t))) dt \right], \quad z \in \mathbb{R}_+^d. \quad (84)$$

On the other hand, because V solves (16)-(17) and

$$u(z) \cdot \nabla V(z) - c(z, u(z)) \leq \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\}, \quad z \in \mathbb{R}_+^d,$$

we conclude that

$$\mathcal{L}V(z) - u(z) \cdot \nabla V(z) + c(z, u(z)) \geq rV(z). \quad (85)$$

Now applying Ito's formula to $e^{-rt}V(Z^u(t))$ and using Equation (7) yields

$$\begin{aligned} e^{-rT}V(Z^u(T)) - V(z) &= \int_0^T (\mathcal{L}V(Z^u(t)) - u(Z^u(t)) \cdot \nabla V(Z^u(t)) - rV(Z^u(t))) dt \\ &\quad + \int_0^T \mathcal{D}V(Z^u(t)) \cdot dY(t) + \int_0^T e^{-rt} \nabla V(Z^u(t)) \cdot dW(t). \end{aligned}$$

Combining this with Equations (3)-(4), (16)-(17) and (85) gives

$$e^{-rT}V(Z^u(T)) - V(z) \geq - \int_0^T e^{-rt} c(Z^u(t), u(Z^u(t))) dt + \int_0^T e^{-rt} \nabla V(Z^u(t)) \cdot dW(t). \quad (86)$$

Because ∇V has polynomial growth and the action space Θ is bounded, we have that

$$\mathbb{E}_z \left[\int_0^T e^{-rt} \nabla V(Z^u(t)) \cdot dW(t) \right] = 0;$$

see, for example, Theorem 3.2.1 of Oksendal [35]. Using this and taking the expectation of both sides of Equation (86) yields

$$V(z) \leq \mathbb{E}_z \left[\int_0^T e^{-rt} c(Z^u(t), u(Z^u(t))) dt \right] + e^{-rT} \mathbb{E}[V(Z^u(T))].$$

Because V has polynomial growth and Θ is bounded, the second term on the right-hand side vanishes as $T \rightarrow \infty$. Then, because c has polynomial growth and Θ is bounded, passing to the limit yields

$$V(z) \leq \mathbb{E}_z \left[\int_0^T e^{-rt} c(Z^u(t), u(Z^u(t))) dt \right] = V^u(z), \quad (87)$$

where the equality holds by Equation (84).

Now, consider the optimal policy u^* , where $u^*(z) = \arg \max_{\theta \in \Theta} \{\theta \cdot \nabla V(z) - c(z, \theta)\}$. For notational brevity, let $Z^* = Z^{u^*}$ denote the RBM under policy u^* . Note from Equation (16) that

$$\mathcal{L}V(z) - u^*(z) \cdot \nabla V(z) + c(z, u^*(z)) = rV(z), \quad z \in \mathbb{R}_+^d. \quad (88)$$

Repeating the preceding steps with u^* in place of u and replacing the inequality with an equality, cf. Equations (85) and (88), we conclude that

$$V(z) = \mathbb{E}_z \left[\int_0^\infty e^{-rt} c(Z^*(t), u^*(Z^*(T))) dt \right] = V^{u^*}(z).$$

Combining this with Equation (87) yields (13). □

B.2 Ergodic Control

Proposition 8. *Let $u \in \mathcal{U}$ be an admissible policy and $(\tilde{\xi}, v^u)$ a C^2 solution of the associated PDE (23)-(24). Further assume that v^u and its gradient have polynomial growth. Then*

$$\tilde{\xi} = \xi^u = \int_{\mathbb{R}_+^d} c(z, u(z)) \pi^u(dz).$$

Proof. Let π^u denote the stationary distribution of RBM under policy u , and let Z^u denote the RBM under policy u that is initiated with π^u . That is,

$$\mathbb{P}(Z^u(0) \in B) = \pi^u(B) \quad \text{for } B \subset \mathbb{R}_+^d.$$

Then applying Ito's formula to $v^u(Z^u(t))$ and using Equation (7) yields

$$\begin{aligned} v^u(Z^u(t)) - v^u(Z^u(0)) &= \int_0^T (\mathcal{L}v^u(Z^u(t)) - u(Z^u(t)) \cdot \nabla v^u(Z^u(t))) dt \\ &\quad + \int_0^T \mathcal{D}v^u(Z^u(t)) \cdot dY(t) + \int_0^T \nabla v^u(Z^u(t)) \cdot dW(t). \end{aligned}$$

Then using Equations (3)-(4) and (23)-(24), we arrive at the following:

$$v^u(Z^u(T)) - v^u(Z^u(0)) = \int_0^T [\tilde{\xi} - c(Z^u(t), u(Z^u(t)))] dt + \int_0^T \nabla v^u(Z^u(t)) \cdot dW(t). \quad (89)$$

Note that the marginal distribution of $Z^u(t)$ is π^u for all $t \geq 0$. Thus, we have that

$$\mathbb{E}_{\pi^u}[v^u(Z^u(T))] = \mathbb{E}_{\pi^u}[v^u(Z^u(0))].$$

Moreover, using Equation (21) and the polynomial growth of ∇v^u , we conclude that

$$\mathbb{E}_{\pi^u} \left[\int_0^T |\nabla v^u(Z^u(t))|^2 dt \right] = T \int_{\mathbb{R}_+^d} |\nabla v^u(z)|^2 \pi^u(dz) < \infty.$$

Consequently, we have that $\mathbb{E} \left[\int_0^T \nabla v^u(Z^u(t)) \cdot dW(t) \right] = 0$; see, for example, Theorem 3.2.1 of Oksendal [35]. Combining these and taking the expectation of both sides of (89) gives

$$\tilde{\xi} = \frac{1}{T} \int_0^T \mathbb{E}_{\pi^u} [c(Z^u(t), u(Z^u(t)))] dt = \int_{\mathbb{R}_+^d} c(z, u(z)) \pi_u(dz) = \xi^u.$$

□

Proposition 9. *Let (v, ξ) be a C^2 solution of the HJB equation (25)-(26), and further assume that both v and its gradient have polynomial growth. Then (27) holds, and moreover, $\xi = \xi^{u^*}$ where the optimal policy u^* is defined by (28).*

Proof. First, consider an arbitrary policy u and note that

$$\xi^u = \int_{\mathbb{R}_+^d} c(z, u(z)) \pi^u(dz),$$

where π^u is the stationary distribution of RBM under policy u . Let Z^u denote the RBM under policy u that is initiated with the stationary distribution π^u . That is,

$$\mathbb{P}(Z^u(0) \in B) = \pi^u(B), \quad B \subset \mathbb{R}_+^d.$$

On the other hand, because (v, ξ) solves the HJB equation and

$$u(z) \cdot \nabla v(z) - c(z, u(z)) \leq \max_{\theta \in \Theta} \{\theta \cdot \nabla v(z) - c(z, \theta)\},$$

we have that

$$\mathcal{L}v(z) - u(z) \cdot \nabla v(z) + c(z, u(z)) \geq \xi \tag{90}$$

Now, we apply Ito's formula to $v(Z^u(t))$ and use Equation (7) to get

$$\begin{aligned} v(Z^u(T)) - v(Z^u(0)) &= \int_0^T (\mathcal{L}v(Z^u(t)) - u(Z^u(t)) \cdot \nabla v(Z^u(t))) dt \\ &\quad + \int_0^T \nabla v(Z^u(t)) \cdot dY(t) + \int_0^T \nabla v(Z^u(t)) \cdot dW(t). \end{aligned}$$

Combining this with Equations (3)-(4), (26) and (90) gives

$$v(Z^u(T)) - v(Z^u(0)) \geq \int_0^T (\xi - c(Z^u(t), u(Z^u(t)))) dt + \int_0^T \nabla v(Z^u(t)) \cdot dW(t). \tag{91}$$

Note that the marginal distribution of $Z^u(t)$ is π^u for all $t \geq 0$. Thus, we have that

$$\mathbb{E}_{\pi^u}[v(Z^u(T))] = \mathbb{E}_{\pi^u}[v(Z^u(0))].$$

Moreover, using Equation (21) and the polynomial growth of ∇v , we conclude that

$$\mathbb{E}_{\pi^u} \left[\int_0^T |\nabla v(Z^u(t))|^2 dt \right] = T \int_{\mathbb{R}_+^d} |\nabla v(z)|^2 \pi^u(dz) < \infty.$$

Consequently, we have that $\mathbb{E} \left[\int_0^T \nabla v(Z^u(t)) dW(t) \right] = 0$; see for example Theorem 3.2.1 of Oksendal [35]. Combining these and taking the expectation of both sides of (91) give

$$\xi T \leq \int_0^T \mathbb{E}_{\pi^u} [c(Z^u(t), u(Z^u(t)))] dt = T \mathbb{E}_{\pi^u} [c(Z^u(0), u(Z^u(0)))],$$

which yields

$$\xi \leq \mathbb{E}_{\pi^u} [c(Z^u(0), u(Z^u(0)))] = \int_{\mathbb{R}_+^d} c(z, u(z)) \pi^u(dz) = \xi^u. \quad (92)$$

Now, consider policy u^* . For notational brevity, let $Z^*(t) = Z^{u^*}(t)$ denote the RBM under policy u^* that is initiated with the stationary distribution π^{u^*} . In addition, note from (25) that

$$\mathcal{L}v(z) - u^*(z) \cdot \nabla v(z) + c(z, u^*(z)) = \xi, \quad z \in \mathbb{R}_+^d. \quad (93)$$

Repeating the preceding steps with u^* in place of u and replacing the inequality with an equality, cf. Equations (90) and (93), we conclude

$$\xi = \mathbb{E}_{\pi^{u^*}} [c(Z^*(0), u(Z^*(0)))] = \int_{\mathbb{R}_+^d} c(z, u^*(z)) \pi^{u^*}(dz) = \xi^{u^*}.$$

Combining this with Equation (92) completes the proof. \square

Appendix C Derivation of the covariance matrix of the feed-forward examples

By the functional central limit theorem for the renewal process [9], we have

$$\hat{E}^n(\cdot) \Rightarrow W_E(\cdot),$$

where $W_E(\cdot)$ is an one-dimensional Brownian motion with drift zero and variance $\lambda a^2 = \mu_0 a^2$. Furthermore, we have

$$\hat{S}_k^n(t) \Rightarrow W_k(\cdot), \text{ for } k = 1, 2, \dots, K,$$

where $W_k(\cdot)$ is an one-dimensional Brownian motion with drift zero and variance $\mu_0 p_k s_k^2$. Now, we turn to $\hat{S}_0^n(t)$ and $\hat{\Phi}^n(t)$. By Harrison [18], we have

$$\text{Cov} \left(\begin{bmatrix} \hat{S}_0^n(t) \\ \hat{\Phi}^n(t) \end{bmatrix} \right) = \mu_0 \Omega^0 + \mu_0 s_0^2 R^0 (R^0)^\top,$$

where $\Omega_{kl}^0 = p_k(\mathbb{I}\{k = l\} - p_l)$ for $k, l = 0, \dots, K$ and $R^0 = [1, -p_1, \dots, -p_K]^\top$. Therefore, we have

$$\text{Cov} \left(\begin{bmatrix} \hat{S}_0^n(t) \\ \hat{\Phi}^n(t) \end{bmatrix} \right) = \mu_0 \begin{bmatrix} s_0^2 & -p_1 s_0^2 & \cdots & \cdots & -p_K s_0^2 \\ -p_1 s_0^2 & p_1(1 - p_1) + p_1^2 s_0^2 & p_1 p_2 (s_0^2 - 1) & \cdots & p_1 p_K (s_0^2 - 1) \\ \vdots & p_1 p_2 (s_0^2 - 1) & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & p_{K-1} p_K (s_0^2 - 1) \\ -p_K s_0^2 & p_1 p_K (s_0^2 - 1) & \cdots & \cdots & p_K(1 - p_K) + p_K^2 s_0^2 \end{bmatrix}.$$

Therefore, the variance of χ is

$$\begin{aligned}
A &= \text{diag}(\lambda a^2, \mu_1 s_1^2, \dots, \mu_K s_K^2) + \\
&\mu_0 \begin{bmatrix} s_0^2 & -p_1 s_0^2 & \cdots & \cdots & -p_K s_0^2 \\ -p_1 s_0^2 & p_1(1-p_1) + p_1^2 s_0^2 & p_1 p_2 (s_0^2 - 1) & \cdots & p_1 p_K (s_0^2 - 1) \\ \vdots & p_1 p_2 (s_0^2 - 1) & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & p_{K-1} p_K (s_0^2 - 1) \\ -p_K s_0^2 & p_1 p_K (s_0^2 - 1) & \cdots & \cdots & p_K(1-p_K) + p_K^2 s_0^2 \end{bmatrix} \\
&= \mu_0 \begin{bmatrix} s_0^2 + a^2 & -p_1 s_0^2 & \cdots & \cdots & -p_K s_0^2 \\ -p_1 s_0^2 & p_1(1-p_1) + p_1^2 s_0^2 + p_1 s_1^2 & p_1 p_2 (s_0^2 - 1) & \cdots & p_1 p_K (s_0^2 - 1) \\ \vdots & p_1 p_2 (s_0^2 - 1) & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & p_{K-1} p_K (s_0^2 - 1) \\ -p_K s_0^2 & p_1 p_K (s_0^2 - 1) & \cdots & \cdots & p_K(1-p_K) + p_K^2 s_0^2 + p_K s_K^2 \end{bmatrix}.
\end{aligned}$$

In particular, if the arrival and service processes are Poisson processes, we have $a = 1$ and $s_k = 1$ for $k = 0, 1, 2, \dots, K$. Then, we have

$$A_{\text{Poisson}} = \mu_0 \begin{bmatrix} 2 & -p_1 & \cdots & \cdots & -p_K \\ -p_1 & 2p_1 & & & \\ \vdots & & \ddots & & \\ \vdots & & & \ddots & \\ -p_K & & & & 2p_K \end{bmatrix}$$

Furthermore, if the service time for server zero is deterministic, i.e., $s_0 = 0$, we have

$$A_{\text{deterministic}} = \mu_0 \begin{bmatrix} a^2 & 0 & \cdots & \cdots & 0 \\ 0 & p_1(1-p_1) + p_1 s_1^2 & -p_1 p_2 & \cdots & -p_1 p_K \\ \vdots & -p_1 p_2 & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & -p_{K-1} p_K \\ 0 & -p_1 p_K & \cdots & \cdots & p_K(1-p_K) + p_K s_K^2 \end{bmatrix}$$

Appendix D Analytical solution of one-dimensional test problems

D.1 Ergodic control formulation with linear cost of control

We consider the one-dimensional control problem with the cost function

$$c(z, \theta) = hz + c\theta \text{ for } z \in \mathbb{R}_+ \text{ and } \theta \in \Theta = [0, b].$$

In the ergodic control case, the HJB equation (25) - (26) is

$$\frac{a}{2}v''(z) - \max_{\theta \in [0, b]} \{\theta \cdot v'(z) - hz - c\theta\} = \xi, \text{ and} \quad (94)$$

$$v'(0) = 0 \text{ and } v'(z) \text{ having polynomial growth rate,} \quad (95)$$

where the covariance matrix $A = a$ in this one-dimensional case. The HJB equation (94) - (95) is equivalent to

$$\frac{a}{2}v''(z) + hz - (v'(z) - c)^+ b = \xi,$$

and the solution is

$$(v^*)'(z) = \begin{cases} \frac{2}{\sqrt{a}} \sqrt{ch + \frac{ah^2}{4b^2}z - \frac{h}{a}z^2} & \text{if } z < z^* \\ \frac{h}{b}z + \frac{ha}{2b^2} - \frac{\sqrt{a}}{b} \sqrt{ch + \frac{ah^2}{4b^2}} + c & \text{if } z \geq z^* \end{cases}, \text{ with}$$

$$z^* = \frac{1}{h} \sqrt{a \left(ch + \frac{ah^2}{4b^2} \right)} - \frac{a}{2b} \text{ and } \xi^* = \sqrt{a \left(ch + \frac{ah^2}{4b^2} \right)},$$

and the optimal control is

$$\theta^*(z) = \begin{cases} 0 & \text{if } z < z^*, \\ b & \text{if } z \geq z^*. \end{cases}$$

D.2 Discounted formulation with linear cost of control

The cost function is still

$$c(z, \theta) = hz + c\theta \text{ for } z \in \mathbb{R}_+ \text{ and } \theta \in \Theta = [0, b],$$

and in the discounted control case, the HJB equation (16) - (17) is

$$\begin{aligned} \frac{a}{2}V''(z) + hz - (V'(z) - c)^+ b &= rV(z), \\ V'(0) &= 0. \end{aligned}$$

The solution is

$$V^*(z) = \begin{cases} V_1(z) & \text{if } z < z^* \\ V_2(z) & \text{if } z \geq z^* \end{cases},$$

and the optimal control is

$$\theta^*(z) = \begin{cases} 0 & \text{if } z < z^* \\ b & \text{if } z \geq z^* \end{cases},$$

where

$$V_1(z) = \frac{h\sqrt{ae}^{-\frac{\sqrt{2}\sqrt{r}z}}{\sqrt{a}}}}{\sqrt{2}r^{3/2}} + \frac{hz}{r} + C_1 e^{\frac{\sqrt{2}\sqrt{r}z}{\sqrt{a}}} + C_2 e^{-\frac{\sqrt{2}\sqrt{r}z}{\sqrt{a}}}, \text{ and}$$

$$V_2(z) = \frac{-bh + brc + hrz}{r^2} + C_2 e^{z\left(\frac{b}{a} - \frac{\sqrt{b^2+2ra}}{a}\right)},$$

for some parameters z^*, C_1, C_2 , to be determined later.

Case 1: $h \leq rc$. Note that if $C_1 = 0$, then we have

$$V_1'(z) = \frac{h}{r} \left(1 - e^{-\frac{\sqrt{2}\sqrt{r}z}{\sqrt{a}}}\right) < \frac{h}{r} \leq c.$$

Therefore, we have

$$V^*(z) = \frac{h\sqrt{a}e^{-\frac{\sqrt{2}\sqrt{r}z}{\sqrt{a}}}}{\sqrt{2}r^{3/2}} + \frac{hz}{r}$$

for the case $h \leq rc$ and the optimal control is always to set $\theta^*(z) = 0$.

Case 2: $h > rc$. We have

$$\begin{aligned} z^* &= \frac{a \log\left(\frac{(h-rc)a}{C_2\lambda(\sqrt{b^2+2ra}-b)}\right)}{b - \sqrt{b^2 + 2ra}}, \\ V_2'(z^*) &= c, \text{ and} \\ V_2''(z^*) &= \frac{(h-rc)\left(\sqrt{b^2 + 2\lambda a} - b\right)}{ra}. \end{aligned}$$

At point z^* , we must have

$$V_1'(z^*) = V_2'(z^*) \text{ and } V_1''(z^*) = V_2''(z^*).$$

Then we can numerically solve for C_1 and C_2 using the following equations:

$$\begin{aligned} V_1'(z^*) &= c, \\ V_1''(z^*) &= \frac{(h-rc)\left(\sqrt{b^2 + 2ra} - b\right)}{ra}. \end{aligned}$$

Table 9 presents numerical values of z^* for different parameter combinations.

Table 9: The numerical values of z^* for different parameter combinations ($a = c = 1$).

		$r = 0.01$	$r = 0.1$
$b = 2$	$h = 2$	0.501671	0.517133
	$h = 1.9$	0.519136	0.535753
$b = 10$	$h = 2$	0.660354	0.674135
	$h = 1.9$	0.678797	0.693707

D.3 Ergodic control formulation with quadratic cost of control

We consider the cost function

$$c(\theta, z) = \alpha(\theta - \underline{\theta})^2 + hz.$$

The HJB equation (25) - (26) then becomes

$$\frac{a}{2}v''(z) - \max_{\theta} \{\theta \cdot v'(z) - hz - \alpha(\theta - \underline{\theta})^2\} = \xi, \text{ and} \tag{96}$$

$$v'(z) = 0 \text{ and } v'(z) \text{ having polynomial growth rate,} \tag{97}$$

which is equivalent to

$$\begin{aligned} hz + \frac{a}{2}v''(z) - \frac{1}{4\alpha} (v'(z))^2 - \underline{\theta}v'(z) &= \xi, \\ v'(0) &= 0. \end{aligned}$$

Let $f(z) = v'(z)$ with $f(0) = 0$. Then we have

$$\xi = hz + \frac{a}{2}f'(z) - \frac{1}{4\alpha} (f(z))^2 - \underline{\theta}f(z), f(0) = 0,$$

which is a Riccati equation. One can solve this equation numerically to find ξ such that $f(\cdot)$ has polynomial growth. For example, if $\alpha = \underline{\theta} = a = 1$, and $h = 2$, we have $\xi^* = 0.8017$.

Appendix E Implementation Details of Our Method

Neural network architecture. We used a three or four-layer, fully connected neural network with 20 - 1000 neurons in each layer; see Tables 10 and 11 for details.

Common hyperparameters. Batch size $B = 256$; time horizon $T = 0.1$, discretization step-size $0.1/64$; see Tables 10 and 11 for details.

Learning rate. The learning rate starts from 0.0005, and decays to 0.0003 and 0.0001 with a rate detailed in Tables 10 and 11.

Optimizer. We used the Adam optimizer [29].

Reference policy. The reference policy sets $\tilde{\theta} = 1$.

Activation function. We use the 'elu' action function [38].

Code. Our code structure follows from that of Han et al. [17] and Zhou et al. [50]. We implement two major changes: First, we have separated the data generation and training processes to facilitate data reuse. Second, we have conducted the RBM simulation. We have also integrated all the features discussed in this section.

Table 10: Hyperparameters used in the test problems with linear costs

Hyperparameters	1-dimensional		2-dimensional		6-dimensional		100-dimensional	
	b=2	b=10	b=2	b=10	b=2	b=10	b=2	b=10
#Iterations	6000		6000		6000		6000	
#Epoches	13	17	15	19	23	27	111	115
Learning rate scheme	0.0005 (0,2000)		0.0005 (0,3000)		0.0005 (0,3000)		0.0005 (0,9500)	
	0.0003 (2000,4000)		0.0003 (3000,6000)		0.0003 (3000,6000)		0.0003 (9500,22000)	
	0.0001 (4000, ∞)		0.0001 (6000, ∞)		0.0001 (6000, ∞)		0.0001 (22000, ∞)	
#Hidden layers	4		4		4		3	
#Neurons in each layer	50		50		50		400	
\tilde{c}_0	—		0.4	7	0.4	7	0.4	7
\tilde{c}_1			800	2400	16000			

Table 11: Hyperparameters used in the test problems with quadratic costs

Hyperparameters	1-dimensional		2-dimensional		6-dimensional		100-dimensional	
	b=2	b=10	b=2	b=10	b=2	b=10	b=2	b=10
#Iterations	6000		6000		6000		12000	
#Epoches	12		14		22		110	
Learning rate scheme	0.0005 (0,3000)		0.0005 (0,3000)		0.0005 (0,3000)		0.0005 (0,9500)	
	0.0003 (3000,6000)		0.0003 (3000,6000)		0.0003 (3000,6000)		0.0003 (9500,22000)	
	0.0001 (6000, ∞)		0.0001 (6000, ∞)		0.0001 (6000, ∞)		0.0001 (22000, ∞)	
#Hidden layers	3		4		4		3	
#Neurons in each layer	20		50		50		1000	

E.1 Decay loss in the test example with linear cost of control

Recall in our main test example with linear cost of control, the cost function is

$$c(z, \theta) = h^\top z + c^\top \theta.$$

In the discounted cost formulation, substituting this cost function into the F function defined in Equation (30) gives the following:

$$F(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t))) = \tilde{\theta} \cdot x + h^\top z - b \sum_{i=1}^d \max(G_{w_2}(\tilde{Z}(t))_i - c, 0). \quad (98)$$

Note that if $G_{w_2}(\tilde{Z}(t)) < c$, we have

$$\frac{\partial F(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t)))}{w_2} = 0,$$

which suggests that the algorithm may suffer from the gradient vanishing problem [25], well-known in the deep learning literature. To overcome this difficulty, we propose an alternative F function

$$\tilde{F}(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t))) = \tilde{\theta} \cdot x + h^\top z - b \sum_{i=1}^d \max(G_{w_2}(\tilde{Z}(t))_i - c, 0) - \tilde{b} \sum_{i=1}^d \min(G_{w_2}(\tilde{Z}(t))_i - c, 0) \quad (99)$$

where \tilde{b} is a decaying function with respect to the training iteration. Specifically, we propose

$$\tilde{b} = \left(\tilde{c}_0 - \frac{\text{iteration}}{\tilde{c}_1} \right)^+,$$

for some positive constants \tilde{c}_0 and \tilde{c}_1 . The specific choices of \tilde{c}_0 and \tilde{c}_1 are shown in Table 10.

We proceed similarly in the ergodic cost case using the function f defined in Equation (40).

E.2 Variance loss function in discounted control

Let us parametrize the value function as $V_{w_1}(z) = \tilde{V}_{w_1}(z) + \xi$. Note that $\partial \tilde{V}_{w_1}(z) / \partial z = \partial V_{w_1}(z) / \partial z$. Therefore, we can rewrite the loss function (52)

$$\begin{aligned} \ell(w_1, w_2) = & \mathbb{E} \left[\left(e^{-rT} (\tilde{V}_{w_1}(\tilde{Z}(T)) + \xi) - (\tilde{V}_{w_1}(\tilde{Z}(0)) + \xi) \right. \right. \\ & \left. \left. - \int_0^T e^{-rt} G_{w_2}(\tilde{Z}(t)) \cdot dW(t) + \int_0^T e^{-rt} F(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t))) dt \right)^2 \right], \end{aligned} \quad (100)$$

By optimizing ξ first, we obtain the following variance loss function:

$$\begin{aligned} \tilde{\ell}(w_1, w_2) = & \text{Var} \left[e^{-rT} \tilde{V}_{w_1}(\tilde{Z}(T)) - \tilde{V}_{w_1}(\tilde{Z}(0)) \right. \\ & \left. - \int_0^T e^{-rt} G_{w_2}(\tilde{Z}(t)) \cdot dW(t) + \int_0^T e^{-rt} F(\tilde{Z}(t), G_{w_2}(\tilde{Z}(t))) dt \right]. \end{aligned}$$

We observe that this trick could accelerate the training speed when r is small. Because when $r > 0$ is small, ξ is of the order $O(1/r)$ and $\tilde{V}_{w_1}(\cdot), G_{w_2}(\cdot)$ are of the order $O(1)$.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283. Savannah, GA, USA, 2016.

- [2] Barış Ata. Dynamic control of a multiclass queue with thin arrival streams. *Operations Research*, 54(5):876–892, 2006.
- [3] Barış Ata and Nasser Barjesteh. An approximate analysis of dynamic pricing, outsourcing, and scheduling policies for a multiclass make-to-stock queue in the heavy traffic regime. *Operations Research*, 71(1):341–357, 2023.
- [4] Baris Ata and Ebru Kasikaralar. Dynamic Scheduling of a Multiclass Queue in the Halfin-Whitt Regime: A Computational Approach for High-Dimensional Problems. 2023.
- [5] Baris Ata, J. Michael Harrison, and Larry A. Shepp. Drift rate control of a Brownian processing system. *Annals of Applied Probability*, 15(2):1145–1160, 2005.
- [6] Baris Ata, Deishin Lee, and Erkut Sonmez. Dynamic volunteer staffing in multicrop gleaning operations. *Operations Research*, 67(2):295–314, 2019.
- [7] Avner Bar-Ilan, Nancy P. Marion, and David Perry. Drift control of international reserves. *Journal of Economic Dynamics & Control*, 31:3110–3137, 2007.
- [8] Christian Beck, Martin Hutzenthaler, Arnulf Jentzen, and Benno Kuckuck. An overview on deep learning-based approximation methods for partial differential equations. *Discrete and Continuous Dynamical Systems - Series B*, 28(6):3697–3746, 2023.
- [9] Patrick Billingsley. *Convergence of probability measures (2nd edition)*. John Wiley & Sons, 1999.
- [10] Jose Blanchet, Xinyun Chen, Nian Si, and Peter W Glynn. Efficient steady-state simulation of high-dimensional stochastic networks. *Stochastic Systems*, 11(2):174–192, 2021.
- [11] Amarjit Budhiraja and Chihoon Lee. Long time asymptotics for controlled diffusions in polyhedral domains. *Stochastic Processes and Their Applications*, 117(8):1014–1036, 2007.
- [12] Sabri Çelik and Costis Maglaras. Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Management Science*, 54(6):1132–1146, 2008.
- [13] Jim G. Dai and Ruth Williams. Existence and uniqueness of semimartingale reflecting brownian motions in convex polyhedrons. *Theory of Probability & Its Applications*, 40(1):1–40, 1996.
- [14] Weinan E, Jiequn Han, and Arnulf Jentzen. Algorithms for solving high dimensional pdes: from nonlinear monte carlo to machine learning. *Nonlinearity*, 35:278–310, 2022.
- [15] Arka P. Ghosh and Ananda P. Weerasinghe. Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Systems*, 55(3):147–159, 2007.
- [16] Arka P. Ghosh and Ananda P. Weerasinghe. Optimal buffer size and dynamic rate control for a queueing system with impatient customers in heavy traffic. *Stochastic Processes and Their Applications*, 120(11):2103–2141, 2010.

- [17] Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [18] J Michael Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications*, pages 147–186. Springer, 1988.
- [19] J Michael Harrison. Brownian models of open processing networks: Canonical representation of workload. *The Annals of Applied Probability*, 10(1):75–103, 2000.
- [20] J Michael Harrison and Vien Nguyen. Brownian models of multiclass queueing networks: Current status and open problems. *Queueing Systems*, 13:5–40, 1993.
- [21] J Michael Harrison and Martin I Reiman. Reflected brownian motion on an orthant. *The Annals of Probability*, 9(2):302–308, 1981.
- [22] J Michael Harrison and Lawrence M Wein. Scheduling networks of queues: heavy traffic analysis of a simple open network. *Queueing Systems*, 5:265–279, 1989.
- [23] J Michael Harrison and Lawrence M Wein. Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Operations research*, 38(6):1052–1064, 1990.
- [24] J Michael Harrison and Ruth J Williams. Brownian models of open queueing networks with homogeneous customer populations. *Stochastics: An International Journal of Probability and Stochastic Processes*, 22(2):77–115, 1987.
- [25] Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998.
- [26] Donald L Iglehart and Ward Whitt. Multiple channel queues in heavy traffic. i. *Advances in Applied Probability*, 2(1):150–177, 1970.
- [27] Donald L Iglehart and Ward Whitt. Multiple channel queues in heavy traffic. ii: Sequences, networks, and batches. *Advances in Applied Probability*, 2(2):355–369, 1970.
- [28] Ioannis Karatzas. A class of singular control problems. *Advances in Applied Probability*, 15(2):225–254, 1983.
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [30] Elena V Krichagina and Michael I. Taksar. Diffusion approximation for GI/G/1 controlled queues. *Queueing Systems*, 12:333–367, 1992.

- [31] Harold J. Kushner and Luiz Felipe Martins. Numerical methods for stochastic singular control problems. *SIAM journal on control and optimization*, 29(6):1443–1475, 1991.
- [32] Harold Joseph Kushner. *Heavy traffic analysis of controlled queueing and communication networks*, volume 28. Springer, 2001.
- [33] Luis Felipe Martins, Steven E. Shreve, and H. Mete Soner. Heavy traffic convergence of a controlled, multiclass queueing system. *SIAM journal on control and optimization*, 34(6):2133–2171, 1996.
- [34] Luiz Felipe Martins and Harold J. Kushner. Routing and singular control for queueing networks in heavy traffic. *SIAM journal on control and optimization*, 28(5):1209–1233, 1990.
- [35] Bernt Oksendal. *Stochastic differential equations: an introduction with applications (sixth edition)*. Springer Science & Business Media, 2003.
- [36] Melda Ormeci Matoglu and John H. Vande Vate. Drift control with changeover costs. *Operations Research*, 59(2):427–439, 2011.
- [37] William P Peterson. A heavy traffic limit theorem for networks of queues with multiple customer types. *Mathematics of operations research*, 16(1):90–118, 1991.
- [38] Andrinandrasana David Rasamoelina, Fouzia Adjailia, and Peter Sinčák. A review of activation function for artificial neural network. In *2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMII)*, pages 281–286. IEEE, 2020.
- [39] Martin I Reiman. Open queueing networks in heavy traffic. *Mathematics of operations research*, 9(3):441–458, 1984.
- [40] Melaine Rubino and Baris Ata. Dynamic control of a make-to-order, parallel-server system with cancellations. *Operations Research*, 57(1):94–108, 2009.
- [41] Lisa Maria Taylor and Ruth J. Williams. Existence and uniqueness of semimartingale reflecting Brownian motions in an orthant. *Probability Theory and Related Fields*, 96(3):283–317, 1993.
- [42] John H. Vande Vate. Average cost Brownian drift control with proportional changeover costs. *Stochastic Systems*, 11(3):218–263, 2021.
- [43] Lawrence M Wein. Brownian networks with discretionary routing. *Operations Research*, 39(2):322–340, 1991.
- [44] Ruth J Williams. On the approximation of queueing networks in heavy traffic. *Stochastic Networks: Theory and Applications*, 4:35–56, 1996.
- [45] Ruth J Williams. Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing systems*, 30:27–88, 1998.

- [46] Ruth J Williams. An invariance principle for semimartingale reflecting brownian motions in an orthant. *Queueing Systems*, 30:5–25, 1998.
- [47] Andreas Winkelbauer. Moments and absolute moments of the normal distribution. *arXiv preprint arXiv:1209.4340*, 2012.
- [48] Kelvin Shuangjian Zhang, Gabriel Peyré, Jalal Fadili, and Marcelo Pereyra. Wasserstein control of mirror langevin monte carlo. In *Conference on Learning Theory*, pages 3814–3841. PMLR, 2020.
- [49] Mo Zhou, Jiequn Han, and Jianfeng Lu. Actor-critic method for high dimensional static hamilton–jacobi–bellman partial differential equations based on neural networks. *SIAM Journal on Scientific Computing*, 43(6):A4043–A4066, 2021.
- [50] Mo Zhou, Jiequn Han, and Jianfeng Lu. Code for “actor-critic method for high dimensional static hamilton–jacobi–bellman partial differential equations based on neural networks”. 2021. URL https://github.com/MoZhou1995/DeepPDE_ActorCritic.