

OPTIMAL CONTROL OF MARKOV JUMP PROCESSES: ASYMPTOTIC ANALYSIS, ALGORITHMS AND APPLICATIONS TO THE MODELING OF CHEMICAL REACTION SYSTEMS*

WEI ZHANG[†], CARSTEN HARTMANN[‡], AND MAX VON KLEIST[§]

Abstract. Markov jump processes are widely used to model natural and engineered processes. In the context of biological or chemical applications one typically refers to the chemical master equation (CME), which models the evolution of the probability mass of any copy-number combination of the interacting particles. When many interacting particles (“species”) are considered, the complexity of the CME quickly increases, making direct numerical simulations impossible. This is even more problematic when one aims at *controlling* the Markov jump processes defined by the CME.

In this work, we study both *open loop* and *feedback* optimal control problems of the Markov jump processes in the case that the controls can only be switched at fixed control stages. Based on Kurtz’s limit theorems, we prove the convergence of the respective control value functions of the underlying Markov decision problem as the copy numbers of the species go to infinity. In the case of the optimal control problem on a finite time-horizon, we propose a hybrid control policy algorithm to overcome the difficulties due to the curse of dimensionality when the copy number of the involved species is large. Two numerical examples demonstrate the suitability of both the analysis and the proposed algorithms.

Keywords. Markov jump process, optimal control problem, large number limit, feedback control policy, hybrid control policy.

AMS subject classifications. 60J27, 93E20.

1. Introduction

In the past decades, discrete-state Markov jump processes have been a major research topic in probability theory receiving much attention in applications like economics, physics, biology and chemistry; see e.g., [1, 16, 19, 27, 58, 60]. For example, in the modelling of chemical reactions, a single state is defined as one possible copy-number combination of the distinct interacting chemical species. After a random waiting time, a reaction occurs and changes this copy-number combination. Since the time and order in which chemical reactions occur is random (referred as *intrinsic* noise), the evolution of the state of the system is random as well. The chemical master equation (CME) models the probability of all possible outcomes over time, giving rise to an extremely large state space (consisting of all copy-number combinations). Consequently, solving the chemical master equation or approximating its solution computationally is a non-trivial, yet unsolved task that has been the objective of intense research over the past decades (see e.g., [47] for a summary).

In many real world applications, one does not only aim at propagating or simulating a process forward in time, but also aims at controlling and optimizing it. In this case, the model equations of a controlled system contain extra terms or parameters that can be manipulated by the decision maker according to some control policy. The latter is chosen so that a given cost functional reaches an optimal (e.g., minimum) value. There are two general approaches to an optimal control task, depending on whether

*Received: September 15, 2016; accepted (in revised form): August 9, 2017. Communicated by David Anderson.

[†]Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany (wei.zhang@fu-berlin.de).

[‡]Institute of Mathematics, Brandenburgische Technische Universität Cottbus-Senftenberg, Platz der Deutschen Einheit 1, 03046 Cottbus, Germany (hartman@c-b-tu.de).

[§]Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany (vkleist@zedat.fu-berlin.de).

the admissible control policies are allowed to depend on the system states (*feedback* or *closed loop* control problem) or not (*open loop* control problem). In the case of an open loop control, the control follows a fixed, deterministic policy regardless of the fact that the underlying dynamics are stochastic. On the other hand, feedback controls are random in the sense that each realization of the process gives rise to a different control that is adapted according to the random states of the system. In principle, one can also consider the case where the control policies depend not only on the current states of the system but also on the past. However, for Markov jump processes, it is known that under certain assumptions the optimal cost value can be achieved by a feedback control policy, which only depends on system's current states (see Section 4.4 of [52] for a precise statement).

For small or moderately sized systems, the underlying optimal control problem can be solved numerically using the dynamic programming principle [5, 14, 52, 61]. However, for large systems, solving the optimal control policy by dynamic programming or related methods becomes difficult without suitable approximations or remodelling steps [8, 51, 56]. Within the area of systems biology or chemical engineering, one such remodelling step that has been extensively exploited by control engineers is to replace the stochastic dynamics by a deterministic system of ordinary differential equations (ODE) that ignores the intrinsic noise (e.g., see [30, 38]). These continuous deterministic reaction rate equations model the concentrations of the interacting chemical species by one ODE per species. The approximation of the stochastic system using the ODE system is mainly based on Kurtz's seminal work [2, 28, 29, 31–34] (also see the recent work on multiscale analysis [10, 49]), which shows that the particle numbers per unit volume of the original Markov jump processes without control can be approximated by the classical reaction rate equations in the large copy-number regime (parameterized by either the total number of particles N or the reaction volume V).

In this article, we investigate the relationship between the optimal control problem for the original Markov jump process and the limiting ODE system. Stochastic control problems for Markov jump processes are also termed “Markov decision processes” (MDP) [5, 25, 52]. We confine our analysis to the situation that the control can only be changed at given discrete points in time (called *control stages*). The key contribution of this paper is twofold. Firstly, applying Kurtz's limit theorem, we prove convergence of the cost value of the controlled Markov jump process to the cost value of the controlled limiting ODE system as $N \rightarrow \infty$, both in the open loop and the feedback case; the convergence results then imply that the optimal open loop control policy for the ODE system can be applied to control the Markov jump process when N is large where the optimal cost is achieved asymptotically. Secondly, based on these theoretical results, we propose a hybrid control policy for the optimal control problem of the Markov jump process on a finite time-horizon; the hybrid control policy not only exploits the information of the optimal control policy for the limiting ODE, but also takes into account the stochasticity of the jump process and thus improves the optimal control policy from the ODE approximation in the pre-limit regime when N is moderately large. In terms of computational complexity, the hybrid algorithm avoids the curse of dimensionality by using an on-the-fly state space truncation. Broadly speaking, the hybrid control policy is related to approximative dynamical programming (ADP) and reinforcement learning that have been extensively studied in the last years [8, 51, 55, 56].

Related work. Although this work is mainly motivated by epidemic, biological, and chemical reaction models, it is important to note that the asymptotic analysis of the related optimal control problems appears relevant in scheduling and queueing theory [22,

59]. In the context of scheduling and queueing problems, the relevant asymptotic regime is the *heavy traffic limit*, under which the stochastic model can be approximated by either a diffusion process or an ODE system. The limit models are named *Brownian network* or *fluid approximation*, depending on whether the limiting differential equation is stochastic or deterministic. Readers interested in the Brownian network approach may consult [23, 36, 37, 40, 42, 59] and references therein. For the fluid approximation of stochastic queueing networks, we refer to [11, 12, 40] (cf. [41] for a discussion of both the fluid and the Brownian network approximation). Optimal control of queueing networks and their fluid approximations has been studied in [3, 4, 13, 39, 43, 44, 50, 57] (see also [35, 48] for an approach using weak convergence techniques).

Despite the vast literature on queueing systems, we emphasize that the models and problems therein are quite different from the ones studied herein. For example, for queueing networks, one is often interested in minimizing the total queue length (or its linear combination) by controlling how each server should allocate the service time to each queue, which explains why many of the rigorous results are confined to linear cost functions or birth-death-processes (e.g., [4, 50]). In the current work, besides the differences of the models, the running cost is allowed to be an arbitrary bounded and (local) Lipschitz function in the system states (see Assumption 2.4 in Section 2) and the jump rates of the process may depend on the controls. A limitation of our work is that the controls are switched only at discrete time points (control stages). However, this assumption allows us to obtain stronger convergence results (with explicit convergence order in some cases) and covers applications in epidemic or chemical reaction networks [6, 14, 24, 54, 61]. Specifically, we will prove the asymptotic optimality of finite and infinite time-horizon open loop policies arising from the deterministic limit equations. Our work complements available results on the asymptotic optimality of the associated closed loop policies or *tracking policies* (e.g., [3, 39, 43]) and gives rise to numerical algorithms that do not require solving the dynamic programming equations on the whole state space (see Section 4).

Outline. The remainder of this paper is organized as follows: In Section 2, we introduce the mathematical problem along with the notations used throughout this paper and two paradigmatic examples. Section 3 is devoted to the extension of Kurtz's limit theorem for Markov jump processes and its application to optimal control problems. Based on this analysis, a hybrid control algorithm is proposed and discussed in Section 4. We present several numerical examples in Section 5, and a technical lemma is recorded in Appendix A.

2. Mathematical setup

In this section, we will first introduce our problem and the notations used, and finally sketch two concrete situations in which the problem is relevant.

2.1. Controlled Markov jump processes. Let \mathbb{X} be a discrete lattice in \mathbb{R}^n and consider the Markov jump process $x(t)$ on it. Suppose that at time $t \geq 0$ and given $x(t) = x \in \mathbb{X}$, the probability for making a transition from x to $x+l$ within the infinitesimal time interval $[t, t+ds)$ is $f(x, l)ds$, $l \in \mathbb{X}$. Denoting by τ the waiting time

$$\tau = \inf_{s > t} \{s - t; x(s) \neq x(t)\}, \quad (2.1)$$

it is known that τ follows an exponential distribution with the rate $\lambda(x) = \sum_{l \in \mathbb{X}} f(x, l)$, i.e., $\tau \sim \text{Exp}(\lambda(x))$.

Jump rates. In this work, we suppose that the jump process $x(t)$ depends on both a parameter $N \gg 1$ and the control $\nu \in \mathcal{A}$, where \mathcal{A} is the *control set*. In applications, N may be related to system's volume or the magnitude of particle numbers, while the control ν may affect the jump rates f . To indicate these dependencies, we denote the jump process as $x^{\nu,N}$ and also introduce the normalized process $z^{\nu,N}(t) = N^{-1}x^{\nu,N}(t)$. It is convenient to think of the normalized variable z as a particle density, which is why we will sometimes refer to $z^{\nu,N}(t)$ as the *normalized density process*. Notice that $z^{\nu,N}$ is a Markov jump process on the *scaled lattice* and, due to its importance in our analysis, we use the notation \mathbb{X}_N and $f_d^{\nu,N} : \mathbb{X}_N \times \mathbb{X}_N \rightarrow \mathbb{R}^+$ for its state space and jump rates, respectively, where \mathbb{R}^+ is the set consisting of non-negative real numbers. \mathbb{X} and $f_o^{\nu,N} : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^+$ will be reserved for the *original process* $x^{\nu,N}$. Notice that the jump rates of the original process may depend on N . The subscripts “d” and “o” which appear in the rate functions simply indicate that they refer to either the normalized density process or the original process. Specifically, we have $\mathbb{X}_N = \{\frac{x}{N} | x \in \mathbb{X}\}$ and $f_d^{\nu,N}(z,l) = f_o^{\nu,N}(Nz,Nl)$ for $z,l \in \mathbb{X}_N$.

Controls. We will discuss the control policies and the controlled Markov jump process in detail. For the sake of simplicity, we will refer to the normalized process $z^{\nu,N}$ only, stressing that all considerations are transferable to the process $x^{\nu,N}$. Suppose that on the time interval $[0,T]$, $K + 1$ time points $0 = t_0 < t_1 < \dots < t_j < t_{j+1} < \dots < t_K = T$ are given and fixed. At each time t_j , $0 \leq j < K$, called a *control stage*, we are allowed to select some control $\nu_j \in \mathcal{A}$ and apply it to the jump process in order to influence its jump rates. Once a control ν_j is selected at time t_j , it will persistently take effect during the time interval $[t_j, t_{j+1})$. When the selection of controls ν_j is allowed to depend on the system's current states at time t_j , the control policy is called a feedback control policy and otherwise it is called an open loop control policy. More generally, we introduce the sets of open loop and feedback control policies on time $[t_k, T]$ for $0 \leq k < K$:

$$\begin{aligned} \mathcal{U}_{o,k} &= \{(\nu_k, \nu_{k+1}, \dots, \nu_{K-1}) \mid \nu_j \in \mathcal{A}, k \leq j < K\}, \\ \mathcal{U}_{f,k} &= \{(\nu_k, \nu_{k+1}, \dots, \nu_{K-1}) \mid \nu_j : \mathbb{X}_N \rightarrow \mathcal{A}, k \leq j < K\}. \end{aligned} \tag{2.2}$$

Notice that in the feedback case, while each policy ν_j is a function of the state, the same notation will be used to denote its value (i.e., the control selected at t_j) when no ambiguity exists. For further simplification, let σ denote either ‘o’ or ‘f’ and we will write $\mathcal{U}_{\sigma,k}$ to refer to either open loop or feedback control policy set.

Given a control policy $u \in \mathcal{U}_{\sigma,k}$, we express the corresponding controlled process in the time interval $[t_k, T]$ as $z^{u,N}(t)$, i.e., the control ν_j is applied during time $t \in [t_j, t_{j+1})$, $k \leq j < K$. The notation $z^{u,N}(t; z)$ will be used to emphasize that the process starts from a fixed initial state $z \in \mathbb{X}_N$ at time t_k (the starting time may be nonzero). Specifically, for a fixed control policy

$$u = (\nu_0, \nu_1, \nu_2, \dots, \nu_{K-1}) \in \mathcal{U}_{\sigma,0},$$

$z^{u,N}(t), t \geq 0$ is a Markov jump process with the property that the probability for system's state to jump from $z^{u,N}(t) = z$ to $z + l$ within the infinitesimal time interval $[t, t + ds)$ at $t \in [t_j, t_{j+1})$, is $f_d^{\nu_j,N}(z,l) ds$ for $l \in \mathbb{X}_N$. That is, application of controls changes the jump rates of the Markov jump process. With the notation

$$j(t) := i, \quad \text{if } t \in [t_i, t_{i+1}), \tag{2.3}$$

we can denote the control policy which is applied to the process $z^{u,N}(t)$ at time t as $\nu_{j(t)}$. Finally, the notation $z^{\nu,N}(t)$ will also be used when we emphasize that the current

control policy at time t is $\nu \in \mathcal{A}$, or when we consider the controlled process on a single stage $[t_j, t_{j+1})$, in which case only the control policy ν applied at time t_j is relevant.

Cost functional. For a control policy $u = (\nu_0, \nu_1, \dots, \nu_{K-1}) \in \mathcal{U}_{\sigma,0}$ and the process $z^{u,N}$, we define the cost functional

$$J_N(z, u) = \mathbf{E}_z^u \left[\sum_{j=0}^{K-1} \left(r(z^{u,N}(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(z^{u,N}(s), \nu_j) ds \right) + \psi(z^{u,N}(T)) \right], \quad (2.4)$$

where \mathbf{E}_z^u denotes the expectation over all realizations of $z^{u,N}$ starting at $z^{u,N}(0) = z$ and evolving under the control policy u . We emphasize that in the feedback case $u \in \mathcal{U}_{f,0}$, we have adopted the convention discussed before, and ν_j in definition (2.4) should be interpreted as $\nu_j = \nu_j(z^{u,N}(t_j))$. The functions $r, \phi: \mathbb{R}^n \times \mathcal{A} \rightarrow \mathbb{R}$, and $\psi: \mathbb{R}^n \rightarrow \mathbb{R}$ correspond to the cost at each control stage t_j , the running cost, and the terminal cost, respectively.

2.2. Limiting process and underlying assumptions. Our analysis in the course of the paper is based on Kurtz’s limit theorems for jump processes [31–34], which state that, for $u \in \mathcal{U}_{o,0}$, the normalized density process $z^{u,N}$ converges to a deterministic limiting process \tilde{z}^u under certain assumptions, and is governed by the ordinary differential equation (ODE)

$$\frac{d\tilde{z}^u(t)}{dt} = F^{\nu_{j(t)}}(\tilde{z}^u(t)), \quad (2.5)$$

or, in integral form,

$$\tilde{z}^u(t) = \tilde{z}^u(0) + \int_0^t F^{\nu_{j(s)}}(\tilde{z}^u(s)) ds. \quad (2.6)$$

Here, the vector field F^ν is defined as the limit of

$$F^{\nu,N}(z) = \sum_{l \in \mathbb{X}_N} l f_a^{\nu,N}(z, l), \quad z \in \mathbb{X}_N, \quad (2.7)$$

as $N \rightarrow \infty$ (see Assumption 2.2), and we have used the notation $j(\cdot)$ which is defined in (2.3). Convergence of $z^{u,N}$ to \tilde{z}^u will be established below in Theorem 3.1.

Limiting control value. We are interested in substituting the optimal control policy for the jump process with an optimal open loop control $u_0 \in \mathcal{U}_{o,0}$ of the limiting process, such that

$$J_N(z, u_0) \approx U_N(z) \triangleq \inf_{u \in \mathcal{U}_{\sigma,0}} J_N(z, u), \quad (2.8)$$

i.e., the infimum (minimum) cost is approximated under the policy u_0 .

The function U_N is called the *value function* or *control value* of the underlying stochastic control problem. It is known that an optimal control $u_{opt,\sigma}^N = \operatorname{argmin}_u J_N(z, u)$ exists when \mathcal{A} is a finite set; see [52] for more details and possible relaxations of the assumptions on the set of admissible controls.

For the related deterministic limiting process \tilde{z}^u satisfying equation (2.5) under some open loop policy $u \in \mathcal{U}_{o,0}$, we define the cost functional by

$$\tilde{J}(z, u) = \sum_{j=0}^{K-1} \left[r(\tilde{z}^u(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(\tilde{z}^u(s), \nu_j) ds \right] + \psi(\tilde{z}^u(T)), \quad (2.9)$$

and the corresponding value function $\tilde{U}(z) = \inf_{u \in \mathcal{U}_{o,0}} \tilde{J}(z, u)$. Note that when \mathcal{A} is a finite set, the minimizer exists since the number of possible open loop control policies u is finite and equal to $|\mathcal{A}|^K$, i.e., $|\mathcal{U}_{o,0}| = |\mathcal{A}|^K$. Convergence of the value function $U_N \rightarrow \tilde{U}$ will be established in the course of the paper.

Standing assumptions. Let Ω be a fixed open subset of the space \mathbb{R}^n . The subsequent analysis rests on the following assumptions:

ASSUMPTION 2.1. *For some fixed $1 < \alpha \leq 2$, we assume that*

$$M_{N,\alpha} := \sup_{\nu \in \mathcal{A}} \sup_{z \in \mathbb{X}_N \cap \Omega} \left(\sum_{l \in \mathbb{X}_N} |l|^\alpha f_d^{\nu,N}(z, l) \right) < \infty, \tag{2.10}$$

and satisfies

$$\lim_{N \rightarrow \infty} M_{N,\alpha} = 0.$$

ASSUMPTION 2.2. *There exist functions $F^\nu : \Omega \rightarrow \mathbb{R}^n$, such that*

$$\omega_N := \sup_{z \in \mathbb{X}_N \cap \Omega, \nu \in \mathcal{A}} |F^{\nu,N}(z) - F^\nu(z)| \tag{2.11}$$

satisfies

$$\lim_{N \rightarrow \infty} \omega_N = 0.$$

ASSUMPTION 2.3. *There exists a constant $L_F \geq 0$, which may depend on the subset Ω , such that*

$$|F^\nu(z') - F^\nu(z)| \leq L_F |z' - z|, \quad \forall z, z' \in \Omega, \nu \in \mathcal{A}.$$

Finally, for the functions related to the cost functional (2.4) of the optimal control problem, we suppose

ASSUMPTION 2.4. *There exist constants $L_r, L_\phi, L_\psi, M_r, M_\phi, M_\psi \geq 0$, which may depend on the subset Ω , such that*

$$\begin{aligned} |r(z_1, \nu) - r(z_2, \nu)| &\leq L_r |z_1 - z_2|, & |\phi(z_1, \nu) - \phi(z_2, \nu)| &\leq L_\phi |z_1 - z_2|, \\ |\psi(z_1) - \psi(z_2)| &\leq L_\psi |z_1 - z_2|, \end{aligned}$$

$\forall z_1, z_2 \in \Omega, \nu \in \mathcal{A}$. Moreover, $|r(z, \nu)| \leq M_r, |\phi(z, \nu)| \leq M_\phi, |\psi(z)| \leq M_\psi, \forall z \in \mathbb{R}^n, \nu \in \mathcal{A}$.

REMARK 2.1. We make some remarks on the above assumptions.

- (1) Although the constants in Assumptions 2.1–2.4 may depend on the subset Ω , we will omit the dependence, since Ω is fixed throughout this paper.
- (2) Instead of utilizing the jump rate function of the density jump process $z^{u,N}$, the quantity in Assumption 2.1 can also be expressed in terms of the original jump process $x^{u,N}$. In fact, using the relation between the functions $f_d^{\nu,N}$ and $f_o^{\nu,N}$, (2.10) is equivalent to

$$M_{N,\alpha} = N^{-\alpha} \sup_{\nu \in \mathcal{A}} \sup_{z \in \mathbb{X}_N \cap \Omega} \left(\sum_{l \in \mathbb{X}} |l|^\alpha f_o^{\nu,N}(Nz, l) \right) < \infty. \tag{2.12}$$

- (3) Assumption 2.2 states that $F^{\nu,N}(z)$ converges to $F^\nu(z)$ uniformly for all $\nu \in \mathcal{A}$ on the subset Ω , while Assumption 2.3 states that the family of the limiting vector fields $F^\nu(z)$ are (local) Lipschitz functions with Lipschitz constant L_F on the set Ω , uniformly for $\nu \in \mathcal{A}$. Similarly, Assumption 2.4 assures that the functions r, ϕ, ψ are Lipschitz on Ω and are bounded on \mathbb{R}^n , uniformly for $\nu \in \mathcal{A}$.

2.3. Applications. Here we consider two prototypical examples of Markov jump processes, which appear relevant in the context of optimal control and to which our results can be applied.

Density dependent Markov chain. The first example is the *density dependent Markov chain* [32], where the jump rates of the original process depend on the density of the system’s states. Specifically, following the notations of Subsection 2.1 and denoting the density dependent Markov chain as $x^{\nu,N}(\cdot)$, it holds that the rate of jumping from state x to $x+l$ under the control $\nu \in \mathcal{A}$ is given by $f_o^{\nu,N}(x,l) = N\eta^\nu(x/N,l)$ for $x,l \in \mathbb{X}$, where $\eta^\nu : \mathbb{R}^n \times \mathbb{X} \rightarrow \mathbb{R}^+$ is a function independent of N . As a consequence,

$$f_d^{\nu,N}(z,l/N) = f_o^{\nu,N}(Nz,l) = N\eta^\nu(z,l)$$

is the rate at which the normalized density process $z^{\nu,N}(\cdot) = N^{-1}x^{\nu,N}(\cdot)$ jumps from $z = x/N$ to $z+l/N = (x+l)/N$. Concrete models of density dependent Markov chains include the predator-prey model, elementary chemical reactions such as $B + C \rightleftharpoons D$, and epidemic models [31, 32].

Notice that if we assume

$$M_\alpha = \sup_{\nu \in \mathcal{A}} \sup_{z \in \Omega} \left(\sum_{l \in \mathbb{X}} |l|^\alpha \eta^\nu(z,l) \right) < \infty, \tag{2.13}$$

then Assumption 2.1 holds, since $M_{N,\alpha} = N^{1-\alpha} M_\alpha$ with $\alpha > 1$. Furthermore, if we define


$$F^\nu(z) = \sum_{l \in \mathbb{X}} l \eta^\nu(z,l), \quad \forall z \in \mathbb{R}^n, \tag{2.14}$$

then equation (2.7) becomes

$$F^{\nu,N}(z) = \sum_{l \in \mathbb{X}_N} l f_d^{\nu,N}(z,l) = \sum_{l \in \mathbb{X}} \frac{l}{N} \cdot N\eta^\nu(z,l) = F^\nu(z), \quad z \in \mathbb{X}_N,$$

where the function $F^\nu(z)$ is independent of N . This implies that Assumption 2.2 trivially holds with $\omega_N \equiv 0$.

Chemical reactions. As a second example, we mention systems of chemical reactions. Consider a reaction network consisting of n chemical species that can undergo m different chemical reactions:



Here, S_i are the different chemical species, κ_k is the rate constant of the k -th reaction, and v_{ki}, v'_{ki} are the molecule numbers of species S_i consumed or generated when the k -th reaction fires. Now let $x^{(i)}(t)$ be the number of molecules of species S_i at time t and define

$$\bar{x}(t) = (x^{(1)}(t), x^{(2)}(t), \dots, x^{(n)}(t))^T \in \mathbb{N}^n \tag{2.16}$$

to be the state of the chemical system at time t . When the k -th reaction fires at time $t > 0$, the system's state jumps from $x(t)$ to $x(t) + (v'_k - v_k)$, where

$$v_k = (v_{k1}, v_{k2}, \dots, v_{kn})^T \in \mathbb{N}^n, \quad v'_k = (v'_{k1}, v'_{k2}, \dots, v'_{kn})^T \in \mathbb{N}^n. \quad (2.17)$$

In order to fully describe the system as a Markov jump process, we still need to specify the Poisson intensity of each reaction (propensity function). Let λ denote a generic propensity function. For simplicity, we will restrict ourself to at most binary reactions, which consume at most two molecules:

- (1) $\emptyset \xrightarrow{\kappa} \text{product}, \quad \lambda = \kappa N$
- (2) $S_i \xrightarrow{\kappa} \text{product}, \quad \lambda = \kappa x^{(i)}$
- (3) $2S_i \xrightarrow{\kappa} \text{product}, \quad \lambda = \frac{\kappa}{N} x^{(i)}(x^{(i)} - 1)$
- (4) $S_i + S_j \xrightarrow{\kappa} \text{product}, \quad \lambda = \frac{\kappa}{N} x^{(i)} x^{(j)},$

where $x = (x^{(1)}, \dots, x^{(n)})$ is the system's state and N is a constant related to the volume of the system (e.g., the total number of molecules or a test tube volume). In the above reactions 1–4, κ is a constant of order one and the scaling of λ with respect to N corresponds to the ‘‘classical scaling’’ considered in [2, 28]. We also refer to [21] for further discussions on the propensity functions. Note that, in general the propensity function is a function of the system state.

For the reaction network described in (2.15), if we denote the propensity functions by $\lambda_k(x)$ when the system is at state x , then the dynamics of $x(t)$ can be written as

$$x(t) = x(0) + \sum_{k=1}^m (v'_k - v_k) Y_k \left(\int_0^t \lambda_k(x(s)) ds \right), \quad (2.18)$$

where $Y_k(\cdot), 1 \leq k \leq m$ are independent Poisson processes with unit intensity. For the system of controlled chemical reactions, we use the notation $\lambda_k^{\nu, N}(x)$ to indicate that the propensities not only depend on N , but also on the control $\nu \in \mathcal{A}$ via the rate constants $\kappa = \kappa_k(\nu)$. From the definition of the reaction events, it is clear that the jump rates introduced before and the propensity functions are related by

$$f_o^{\nu, N}(x, l) = \sum_{\substack{1 \leq k \leq m, \\ v'_k - v_k = l}} \lambda_k^{\nu, N}(x).$$

Notice that if only reactions of type 1, 2, or 4 are involved, the process defined by $f_o^{\nu, N}$ is an instance of the aforementioned density dependent Markov chain. When reactions of type 3 are involved, then the limiting vector field F^ν can be computed from $F^{\nu, N}$ in equation (2.7) by exploiting the fact that $f_d^{\nu, N}(z, l) = N \kappa z^{(i)}(z^{(i)} - N^{-1})$ for $z, l \in \mathbb{X}_N$, where $\kappa = \kappa_k(\nu)$, if $Nl = v'_k - v_k$ for some $1 \leq k \leq m$ (supposing for simplicity only one such index k exists), and $f_d^{\nu, N}(z, l) = 0$ otherwise.

3. Asymptotic analysis of the optimal control problem

In this section, we study optimal control problems in the large number regime based on Kurtz's limit theorem [31–34].

As a first step, given an open loop control $u \in \mathcal{U}_{o,0}$, we establish the approximation result of the Markov jump process $z^{u, N}$ by the ODE limit (2.5). The proof is adapted from Kurtz's argument, in particular [32]. However, for completeness we feel it is necessary to present the proof in detail. As a second step, we confine our attention to

the open loop control problem which is a direct application of Kurtz’s theorem, given that the Assumptions in Subsection 2.2 hold. Specifically, we show that $J_N(z_N, u) \rightarrow \tilde{J}(z_0, u)$ for $u \in \mathcal{U}_{o,0}, z_N \in \mathbb{X}_N, z_N \rightarrow z_0 \in \Omega$ as $N \rightarrow \infty$ (Theorem 3.2). Then, as a third step, we consider the feedback control problem and prove that $U_N(z_N) \rightarrow \tilde{U}(z)$ if $z_N \rightarrow z$, and, in particular, if $u_0 \in \mathcal{U}_{o,0}$ and $\tilde{J}(z, u_0) = \tilde{U}(z)$, then $|J_N(z_N, u_0) - U_N(z_N)| \rightarrow 0$ as $N \rightarrow \infty$ (Theorem 3.3). As we will discuss in detail, an important consequence of Theorem 3.3 is that the optimal (open loop) control policy for the limiting ODE system is almost optimal for the Markov jump process if $N \gg 1$, i.e., it is asymptotically optimal among all *feedback* control policies in $\mathcal{U}_{f,0}$. Finally, we extend the analysis of the finite time-horizon case to discounted optimal control problems on an infinite time-horizon (Theorem 3.4).

3.1. ODE approximation of the normalized Markov jump process. Let $u \in \mathcal{U}_{o,0}$ be some open loop control policy and $z^{u,N}(t) = N^{-1}x^{u,N}(t)$ denote the normalized density Markov jump process. Recall that Ω is the open subset of \mathbb{R}^n introduced in Subsection 2.2. The convergence of the normalized density process as $N \rightarrow \infty$ is described by the following theorem.

THEOREM 3.1. *Let $z^{u,N}(t)$ be the normalized density jump process under the open loop policy $u \in \mathcal{U}_{o,0}$ and suppose the ODE (2.5) has a unique solution $\tilde{z}^u(t)$ on $t \in [0, T]$ starting from $z_0 \in \Omega$. Furthermore, $\exists \gamma > 0$, s.t.*

$$\Omega_{\gamma, z_0, [0, T]}^u := \left\{ z' \in \mathbb{R}^n \mid \inf_{0 \leq t \leq T} |z' - \tilde{z}^u(t)| \leq \gamma \right\} \subseteq \Omega. \tag{3.1}$$

Let τ_N^u be the stopping time for the jump process $z^{u,N}$ to leave the set $\Omega_{\gamma, z_0, [0, T]}^u$, i.e.,

$$\tau_N^u := \inf_{s \geq 0} \left\{ s \mid z^{u,N}(s) \notin \Omega_{\gamma, z_0, [0, T]}^u \right\}. \tag{3.2}$$

(1) *Suppose Assumption 2.3 holds. We have*

$$\mathbf{E}^u \left[\sup_{0 \leq s \leq t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)| \right] \leq \left[\mathbf{E} |z^{u,N}(0) - z_0| + C_{T,N} \right] e^{L_F t}, \tag{3.3}$$

for $0 \leq t \leq T$, where the constant

$$C_{T,N} = T\omega_N + \frac{\alpha}{2(\alpha - 1)} \left(\frac{4TM_{N,\alpha}}{\alpha - 1} \right)^{\frac{1}{\alpha}}, \tag{3.4}$$

with $\alpha \in (1, 2]$, and $\omega_N, M_{N,\alpha}$ are defined in equations (2.11) and (2.10), respectively.

(2) *Suppose Assumptions 2.1–2.3 are satisfied with constant $\alpha \in (1, 2]$ and that*

$$\lim_{N \rightarrow \infty} \mathbf{E} |z^{u,N}(0) - z_0| = 0.$$

Then for any control policy $u \in \mathcal{U}_{o,0}$, we have

$$\lim_{N \rightarrow \infty} \mathbf{E}^u \left[\sup_{0 \leq s \leq t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)| \right] = 0. \tag{3.5}$$

Furthermore, let $\rho > 0$ be given such that $\rho \leq \frac{1}{3}\gamma e^{-L_F T}$. Then $\exists N_0 > 0$ which may depend on ρ , such that

$$\mathbf{P}(\tau_N^u < T) \leq \rho^{-1} \left[\mathbf{E} |z^{u,N}(0) - z_0| + \frac{\alpha}{2(\alpha - 1)} \left(\frac{4TM_{N,\alpha}}{\alpha - 1} \right)^{\frac{1}{\alpha}} \right], \tag{3.6}$$

whenever $N \geq N_0$, where \mathbf{P} is the probability with respect to the process $z^{u,N}$ under the control u . In particular, we have

$$\lim_{N \rightarrow \infty} \mathbf{P}(\tau_N^u < T) = 0.$$

Proof.

(1) Let $w^{u,N}$ be the martingale

$$w^{u,N}(t) = z^{u,N}(t) - z^{u,N}(0) - \int_0^t F^{\nu_{j(s)},N}(z^{u,N}(s)) ds, \tag{3.7}$$

and consider the coupled Markov process $(z^{u,N}(t), w^{u,N}(t))$. For a differentiable function φ of w , Dynkin’s formula [15, 46] entails

$$\begin{aligned} & \mathbf{E}^u [\varphi(w^{u,N}(t \wedge \tau_N^u))] - \mathbf{E}^u [\varphi(w^{u,N}(0))] \\ = & \mathbf{E}^u \left\{ \int_0^{t \wedge \tau_N^u} \left[\sum_{l \in \mathbb{X}_N} \left(\varphi(l + w^{u,N}(s)) - \varphi(w^{u,N}(s)) - l \cdot \nabla \varphi(w^{u,N}(s)) \right) \right. \right. \\ & \left. \left. \times f_d^{\nu_{j(s)},N}(z^{u,N}(s), l) \right] ds \right\}. \end{aligned}$$

In particular, setting $\varphi(z) = |z|^\alpha$, where $\alpha \in (1, 2]$ is the constant in Assumption 2.1, and using Lemma A.1 from Appendix A, we obtain

$$\mathbf{E}^u |w^{u,N}(t \wedge \tau_N^u)|^\alpha \leq \frac{4t}{2^\alpha(\alpha - 1)} \sup_{\nu \in \mathcal{A}} \sup_{z \in \mathbb{X}_N \cap \Omega} \left(\sum_{l \in \mathbb{X}_N} |l|^\alpha f_d^{\nu,N}(z, l) \right) = \frac{4tM_{N,\alpha}}{2^\alpha(\alpha - 1)},$$

which, by Hölder’s inequality and Doob’s maximal inequality, implies that

$$\begin{aligned} \mathbf{E}^u \left[\sup_{0 \leq s \leq t} |w^{u,N}(s \wedge \tau_N^u)| \right] & \leq \left[\mathbf{E}^u \left(\sup_{0 \leq s \leq t} |w^{u,N}(s \wedge \tau_N^u)|^\alpha \right) \right]^{\frac{1}{\alpha}} \\ & \leq \frac{\alpha}{\alpha - 1} \left[\mathbf{E}^u |w^{u,N}(t \wedge \tau_N^u)|^\alpha \right]^{\frac{1}{\alpha}} \\ & \leq \frac{\alpha}{2(\alpha - 1)} \left(\frac{4tM_{N,\alpha}}{\alpha - 1} \right)^{\frac{1}{\alpha}}. \end{aligned} \tag{3.8}$$

Combining equations (3.7) and (2.6) and taking Assumption 2.3 into consideration, it follows that

$$\begin{aligned} & |z^{u,N}(t \wedge \tau_N^u) - \tilde{z}^u(t \wedge \tau_N^u)| \\ \leq & |z^{u,N}(0) - z_0| + L_F \int_0^{t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)| ds \\ & + \int_0^{t \wedge \tau_N^u} |F^{\nu_{j(s)},N}(z^{u,N}(s)) - F^{\nu_{j(s)}}(z^{u,N}(s))| ds + |w^{u,N}(t \wedge \tau_N^u)| \end{aligned}$$

$$\leq |z^{u,N}(0) - z_0| + L_F \int_0^{t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)| ds + t\omega_N + |w^{u,N}(t \wedge \tau_N^u)|.$$

Now let $y^{u,N}(t) = \sup_{0 \leq s \leq t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)|$. Then

$$y^{u,N}(t) \leq y^{u,N}(0) + L_F \int_0^t y^{u,N}(s) ds + T\omega_N + \sup_{0 \leq s \leq T} |w^{u,N}(s \wedge \tau_N^u)|,$$

and Gronwall's inequality implies

$$y^{u,N}(t) \leq \left[y^{u,N}(0) + T\omega_N + \sup_{0 \leq s \leq T} |w^{u,N}(s \wedge \tau_N^u)| \right] e^{L_F t}. \tag{3.9}$$

The estimate (3.3) follows by taking expectations on both sides of the above inequality and using inequality (3.8).

- (2) Assertion (3.5) follows directly from inequality (3.3) by taking the limit $N \rightarrow \infty$ and applying Assumptions 2.1 and 2.2. To prove assertion (3.6), we first choose $N_0 > 0$ such that $T\omega_N \leq \rho$ whenever $N > N_0$. This is possible due to Assumption 2.2.

From the definitions of $y^{u,N}(t)$, the subset $Q_{\gamma, z_0, [0, T]}^u$ in (3.1), and the inequality (3.9), we can deduce that

$$\begin{aligned} y^{u,N}(0) \leq \rho \quad \text{and} \quad \sup_{0 \leq s \leq T} |w^{u,N}(s \wedge \tau_N^u)| \leq \rho \\ \implies y^{u,N}(T) \leq 3\rho e^{L_F T} \leq \gamma \implies \tau_N^u \geq T, \end{aligned}$$

and therefore that

$$\begin{aligned} \mathbf{P}(\tau_N^u < T) &\leq \mathbf{P}(y^{u,N}(0) > \rho) + \mathbf{P}\left(\sup_{0 \leq s \leq T} |w^{u,N}(s \wedge \tau_N^u)| > \rho\right) \\ &\leq \rho^{-1} \left[\mathbf{E}|z^{u,N}(0) - z_0| + \frac{\alpha}{2(\alpha - 1)} \left(\frac{4TM_{N,\alpha}}{\alpha - 1}\right)^{\frac{1}{\alpha}} \right], \end{aligned}$$

where we have used the fact that $y^{u,N}(0) = |z^{u,N}(0) - z_0|$, the inequality (3.8) and the Chebyshev's inequality. □

We conclude this subsection with the following remarks.

REMARK 3.1. From the proof, it is straightforward to see that when $z^{u,N}(0)$ is deterministic and $|z^{u,N}(0) - z_0| \leq \rho \leq \frac{1}{3}\gamma e^{-L_F T}$, estimate (3.6) can be improved to

$$\mathbf{P}(\tau_N^u < T) \leq \rho^{-1} \left[\frac{\alpha}{2(\alpha - 1)} \left(\frac{4TM_{N,\alpha}}{\alpha - 1}\right)^{\frac{1}{\alpha}} \right] \leq \rho^{-1} C_{T,N}. \tag{3.10}$$

REMARK 3.2. For the density dependent Markov chain introduced in Subsection 2.3, it holds that $\omega_N = 0$ and $M_{N,\alpha} = N^{1-\alpha} M_\alpha$, where M_α is given in definition (2.13) with $\alpha \in (1, 2]$. Therefore, the constant in equation (3.4) satisfies

$$C_{T,N} = \mathcal{O}\left(N^{\frac{1}{\alpha} - 1}\right). \tag{3.11}$$

Assuming $\mathbf{E}|z^{u,N}(0) - z_0| \rightarrow 0$ fast enough as $N \rightarrow \infty$, the above implies that the convergence speed in both estimates (3.3) and (3.6) is explicitly of order $N^{\frac{1}{\alpha} - 1}$.

The simplest case is when $z^{u,N}$ is a one-dimensional process and the control set \mathcal{A} is a singleton. For simplicity, we will omit the control u in the notations in the remainder of this paragraph. Suppose that $\eta(z, 1) = 1$ and $\eta(z, l) = 0$ for $l \neq 1, z \geq 0$. Then definition (2.14) implies that $F(z) \equiv 1$, which is Lipschitz continuous with Lipschitz constant $L_F = 0$. For the initial value $z_0 = 0$, equation (2.6) yields $\tilde{z}(t) = t$ and $z^N(t) = N^{-1}P(Nt)$, where $P(\cdot)$ is a Poisson process with unit intensity. We can also choose the subsets $\Omega_{\gamma,0,[0,T]} = \Omega = \mathbb{R}^n$. Further note that Assumption 2.1 holds with $\alpha = 2$ and $M_\alpha = 1$, so that Theorem 3.1 implies

$$\mathbf{E} \left[\sup_{0 \leq s \leq T} \left| \frac{P(Ns)}{N} - s \right| \right] \leq \left(\frac{4T}{N} \right)^{1/2}.$$

3.2. Optimal control on finite time-horizon. In this subsection, we apply the previous approximation result to study both open and closed loop optimal control on a finite time-horizon.

Open loop control. As a straightforward consequence of Theorem 3.1 and Assumptions 2.3–2.4, we have the following result for the open loop control problem.

THEOREM 3.2. *Suppose that Assumptions 2.1–2.4 hold true. Let $z_0 \in \Omega$ and let $u \in \mathcal{U}_{o,0}$ be any open loop control policy of the form $u = (\nu_0, \nu_1, \dots, \nu_{K-1})$ with $\nu_j \in \mathcal{A}, 0 \leq j < K$. Suppose the ODE (2.5) has a unique solution on $[0, T]$ and furthermore the condition (3.1) is satisfied for some $\gamma > 0$. Recall that the cost functionals J_N and \tilde{J} are defined in equations (2.4), (2.9), respectively. Let $z_N \in \mathbb{X}_N \cap \Omega$ and $z_N \rightarrow z_0$ as $N \rightarrow +\infty$. Then $\exists N_0 > 0$, s.t. for $N > N_0$, we have*

$$|J_N(z_N, u) - \tilde{J}(z_0, u)| \leq (|z_N - z_0| + C_{T,N}) \left[L_\phi \frac{e^{L_F T} - 1}{L_F} + (KL_r + L_\psi + \bar{M}) e^{L_F T} \right], \tag{3.12}$$

with the convention $\frac{e^{L_F T} - 1}{L_F} = T$ if $L_F = 0$, and the constant

$$\bar{M} := 6\gamma^{-1}(KM_r + TM_\phi + M_\psi). \tag{3.13}$$

The constant $C_{T,N}$ is defined in equation (3.4) and the other constants are given in Assumptions 2.3–2.4. In particular, when the condition (3.1) is satisfied for all $u \in \mathcal{U}_{o,0}$ for some common $\gamma > 0$, we have

$$\lim_{N \rightarrow \infty} |J_N(z_N, u) - \tilde{J}(z_0, u)| = 0, \tag{3.14}$$

uniformly for all control policies $u \in \mathcal{U}_{o,0}$.

Proof. First of all, let us define the quantity

$$I = \sum_{j=0}^{K-1} \left[r(z^{u,N}(t_j), \nu_j) - r(\tilde{z}^u(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \left(\phi(z^{u,N}(s), \nu_j) - \phi(\tilde{z}^u(s), \nu_j) \right) ds \right] + \psi(z^{u,N}(T)) - \psi(\tilde{z}^u(T)).$$

Then the boundedness conditions in Assumption 2.4 immediately imply $|I| \leq 2(KM_r + TM_\phi + M_\psi)$. Recalling the stopping time τ_N^u in equation (3.2) and the Lipschitz conditions in Assumption 2.4, we also have

$$|I| \leq \sum_{j=0}^{K-1} \left\{ L_r |z^{u,N}(t_j) - \tilde{z}^u(t_j)| + L_\phi \int_{t_j}^{t_{j+1}} |z^{u,N}(s) - \tilde{z}^u(s)| ds \right\}$$

$$+ L_\psi |z^{u,N}(T) - \tilde{z}^u(T)|, \tag{3.15}$$

as long as $\tau_N^u \geq T$. Therefore, using the definitions of the cost functions J_N, \tilde{J} , we have

$$\begin{aligned} & |J_N(z_N, u) - \tilde{J}(z_0, u)| = |\mathbf{E}_{z_N}^u I| \\ & \leq |\mathbf{E}_{z_N}^u (I \cdot \mathbb{1}_{\{\tau_N^u \geq T\}})| + |\mathbf{E}_{z_N}^u (I \cdot \mathbb{1}_{\{\tau_N^u < T\}})| \\ & \leq \mathbf{E}_{z_N}^u (|I| \cdot \mathbb{1}_{\{\tau_N^u \geq T\}}) + 2(KM_r + TM_\phi + M_\psi) \mathbf{P}(\tau_N^u < T), \end{aligned}$$

where $\mathbb{1}$ denotes the indicator function. For the first term above, noticing the fact

$$\mathbf{E}_{z_N}^u \left[\left(\sup_{0 \leq s \leq t} |z^{u,N}(s) - \tilde{z}^u(s)| \right) \mathbb{1}_{\{\tau_N^u \geq T\}} \right] \leq \mathbf{E}_{z_N}^u \left[\sup_{0 \leq s \leq t \wedge \tau_N^u} |z^{u,N}(s) - \tilde{z}^u(s)| \right],$$

using inequality (3.15), and applying Theorem 3.1, we obtain

$$\begin{aligned} & \mathbf{E}_{z_N}^u (|I| \cdot \mathbb{1}_{\{\tau_N^u \geq T\}}) \\ & \leq (|z_N - z_0| + C_{T,N}) \left[\sum_{j=0}^{K-1} \left(L_\phi \int_{t_j}^{t_{j+1}} e^{L_F s} ds + L_r e^{L_F t_j} \right) + L_\psi e^{L_F T} \right] \\ & \leq \left\{ L_\phi \frac{e^{L_F T} - 1}{L_F} + (KL_r + L_\psi) e^{L_F T} \right\} (|z_N - z_0| + C_{T,N}). \end{aligned}$$

Now fix the constant $\rho = \frac{1}{3} \gamma e^{-L_F T}$ and choose N_0 such that $|z_N - z_0| \leq \rho$ when $N > N_0$. The assertion (3.12) then follows after we estimate $\mathbf{P}(\tau_N^u < T)$ by applying Theorem 3.1 (see estimate (3.10) in Remark 3.1). The convergence of the cost function J_N to \tilde{J} follows from estimate (3.12) directly. \square

Feedback control. Now we consider the case of a feedback control problem. In accordance with definition (2.4), we define the cost functional for $u \in \mathcal{U}_{f,k}, z \in \mathbb{X}_N \cap \Omega, 0 \leq k < K$, and the corresponding value function as

$$\begin{aligned} J_N(z, u, k) &= \mathbf{E}_{t_k, z}^u \left[\sum_{j=k}^{K-1} \left(r(z^{u,N}(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(z^{u,N}(s), \nu_j) ds \right) + \psi(z^{u,N}(T)) \right], \\ U_N(z, k) &= \inf_{u \in \mathcal{U}_{f,k}} J_N(z, u, k), \end{aligned} \tag{3.16}$$

with the shorthand $\mathbf{E}_{t_k, z}^u[\cdot] = \mathbf{E}^u[\cdot | z^{u,N}(t_k) = z]$ for the conditional expectation over all realizations of the controlled process starting at $z^{u,N}(t_k) = z$. Notice that, following the convention in Subsection 2.1, we have used the same notation ν_j to denote both the control policy function which depends on system’s state, and the value of the control selected at t_j , i.e., we have $\nu_j = \nu_j(z^{u,N}(t_j))$ in definition (3.16) (see the discussion after (2.2)). By definition, the value function U_N , also called the *optimal cost-to-go*, is the minimum cost value from time t_k to T as a function of the initial data (z, t_k) . In particular, it holds that $U_N(z, K) = \psi(z)$.

Then, in complete analogy with the above definitions, we define

$$\begin{aligned} \tilde{J}(z, u, k) &= \sum_{j=k}^{K-1} \left(r(\tilde{z}^u(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(\tilde{z}^u(s), \nu_j) ds \right) + \psi(\tilde{z}^u(T)), \quad u \in \mathcal{U}_{o,k}, \\ \tilde{U}(z, k) &= \inf_{u \in \mathcal{U}_{o,k}} \tilde{J}(z, u, k), \end{aligned}$$

for $z \in \Omega$, to be the cost functional and the value function of the deterministic limiting process. In what follows, we will omit the dependence of J_N , \tilde{J} and U_N , \tilde{U} on k when $k=0$ so that the notations are consistent with definitions (2.4) and (2.9).

By the dynamic programming principle [52], the necessary conditions for optimality are given in terms of Bellman's equations for the two value functions:

$$\begin{aligned}
 U_N(z_N, k) &= \inf_{\nu \in \mathcal{A}} \mathbf{E}^\nu \left[r(z_N, \nu) + \int_{t_k}^{t_{k+1}} \phi(z^{\nu, N}(s), \nu) ds + U_N(z^{\nu, N}(t_{k+1}), k+1) \right], \\
 \tilde{U}(z, k) &= \inf_{\nu \in \mathcal{A}} \left\{ r(z, \nu) + \int_{t_k}^{t_{k+1}} \phi(\tilde{z}^\nu(s), \nu) ds + \tilde{U}(\tilde{z}^\nu(t_{k+1}), k+1) \right\}, \tag{3.17}
 \end{aligned}$$

with $0 \leq k \leq K-1$, where $z^{\nu, N}(t_k) = z_N \in \mathbb{X}_N$, $\tilde{z}^\nu(t_k) = z \in \Omega$ and the terminal conditions

$$U_N(z_N, K) = \psi(z_N), \quad \tilde{U}(z, K) = \psi(z). \tag{3.18}$$

Notice that in equation (3.17), we have used the notation $\mathbf{E}^\nu = \mathbf{E}_{t_k, z_N}^\nu$ for the conditional expectation and $z^{\nu, N}(t)$, $\tilde{z}^\nu(t)$ for the processes, since the involved quantities and processes only depend on the control ν selected at t_k , rather than the whole control policy.

Before we proceed, we shall first introduce some constants in order to simplify the analysis later on. Let $h = \max\{|t_{j+1} - t_j| : 0 \leq j \leq K-1\}$. In accordance with equation (3.4), we set

$$C_{h, N} = h\omega_N + \frac{\alpha}{2(\alpha-1)} \left(\frac{4hM_{N, \alpha}}{\alpha-1} \right)^{\frac{1}{\alpha}}. \tag{3.19}$$

We also introduce the sequences of numbers a_k, b_k , $0 \leq k \leq K$, satisfying the recursive relations

$$\begin{aligned}
 a_k &= L_r + L_\phi e^{L_F h} h + \bar{M} e^{L_F h} + a_{k+1} e^{L_F h}, \\
 b_k &= L_\phi C_{h, N} e^{L_F h} (t_{k+1} - t_k) + 2\bar{M} C_{h, N} e^{L_F h} + a_{k+1} C_{h, N} e^{L_F h} + b_{k+1}, \tag{3.20}
 \end{aligned}$$

for $0 \leq k \leq K-1$ and $a_K = L_\psi$, $b_K = 0$, where \bar{M} is defined in equation (3.13). The last two expressions can be made more explicit:

$$\begin{aligned}
 a_k &= \left(L_r + L_\phi e^{L_F h} h + \bar{M} e^{L_F h} \right) \frac{e^{L_F h(K-k)} - 1}{e^{L_F h} - 1} + L_\psi e^{(K-k)L_F h}, \\
 b_k &= C_{h, N} e^{L_F h} \left\{ L_\phi (T - t_k) + \left[\frac{L_r + L_\phi e^{L_F h} h + \bar{M} e^{L_F h}}{e^{L_F h} - 1} \left(\frac{e^{L_F h(K-k)} - 1}{e^{L_F h} - 1} - (K-k) \right) \right. \right. \\
 &\quad \left. \left. + L_\psi \frac{e^{L_F h(K-k)} - 1}{e^{L_F h} - 1} \right] + 2\bar{M}(K-k) \right\}, \tag{3.21}
 \end{aligned}$$

for $0 \leq k \leq K$. Notice that under Assumptions 2.1 and 2.2, both $C_{h, N}$ and b_k go to zero as $N \rightarrow \infty$.

Similar to (3.1), we also introduce the set

$$\Omega_{\gamma, z, [t_i, t_j]}^u := \left\{ z' \in \mathbb{R}^n \mid \inf_{t_i \leq t \leq t_j} |z' - \tilde{z}^u(t)| \leq \gamma \right\}, \tag{3.22}$$

between two control stages $t_i < t_j$ where $\tilde{z}^u(t_i) = z$, $u \in \mathcal{U}_{o, i}$. The notation $\Omega_{\gamma, z, [t_i, t_{i+1}]}^u$ will be used when only the control policy $\nu \in \mathcal{A}$ at the control stage t_i is relevant. We have the following approximation result of the value functions.

THEOREM 3.3. *Suppose Assumptions 2.1–2.4 hold. Given $0 \leq k \leq K$ and $z \in \Omega$, s.t. the ODE (2.5) has a unique solution \tilde{z}^u on $[t_k, T]$ for all $u \in \mathcal{U}_{o,k}$ and furthermore, $\exists \gamma > 0$, s.t. $\Omega_{\gamma, z, [t_k, T]}^u \subseteq \Omega$ for all $u \in \mathcal{U}_{o,k}$. Let $z_N \in \mathbb{X}_N \cap \Omega$ be random with $\mathbf{E}|z_N - z| < \infty$. Then $\exists N_k > 0$, s.t.*

$$\mathbf{E}|U_N(z_N, k) - \tilde{U}(z, k)| \leq a_k \mathbf{E}|z_N - z| + b_k, \tag{3.23}$$

for $N > N_k$, with a_k, b_k as given by formulas (3.20) or (3.21). Further suppose that $u_0 \in \mathcal{U}_{o,0}$ is the optimal (open loop) control policy for the process \tilde{z}^u , i.e., $\tilde{J}(z, u_0) = \tilde{U}(z)$, and $z_N \in \mathbb{X}_N \cap \Omega$ is deterministic satisfying $z_N \rightarrow z$ as $N \rightarrow \infty$. Then $\exists N_0 > 0$, s.t. when $N > N_0$,

$$\begin{aligned} & |J_N(z_N, u_0) - U_N(z_N)| \\ & \leq b_0 + a_0 |z_N - z| + \left[L_\phi \frac{e^{L_F T} - 1}{L_F} + (KL_r + L_\psi + \overline{M}) e^{L_F T} \right] \left(C_{T,N} + |z_N - z| \right). \end{aligned} \tag{3.24}$$

In particular, it holds that

$$\lim_{N \rightarrow \infty} |J_N(z_N, u_0) - U_N(z_N)| = 0.$$

Proof. We first prove inequality (3.23) by backward induction from $k=K$ to $k=0$. Let \mathbf{E} denote the expectation with respect to the random variable $z_N \in \mathbb{X}_N \cap \Omega$ and recall that \mathbf{E}^ν is the shorthand of the conditional expectation $\mathbf{E}_{t_k, z_N}^\nu$. For $k=K$, since $z, z_N \in \Omega$, the terminal condition (3.18) and the Lipschitz continuity of the terminal cost ψ in Assumption 2.4 imply that

$$\mathbf{E}|U_N(z_N, K) - \tilde{U}(z, K)| = \mathbf{E}|\psi(z_N) - \psi(z)| \leq L_\psi \mathbf{E}|z_N - z|.$$

Therefore (3.23) holds with $a_K = L_\psi$, $b_K = 0$ and for any $N_K > 0$.

Now suppose inequality (3.23) is true for $k+1 \leq K$. First notice that we have the simple estimate

$$|U_N(z_N, k) - \tilde{U}(z, k)| \leq 2[(K-k)M_r + (T-t_k)M_\phi + M_\psi]$$

under Assumption 2.4. Then, fixing the constant $\rho = \frac{1}{3}\gamma e^{-L_F h}$ and using the Bellman Equation (3.17) for the value function, we can estimate

$$\begin{aligned} & \mathbf{E}|U_N(z_N, k) - \tilde{U}(z, k)| \\ & = \mathbf{E} \left[|U_N(z_N, k) - \tilde{U}(z, k)| \cdot \mathbb{1}_{\{|z_N - z| \leq \rho\}} \right] + \mathbf{E} \left[|U_N(z_N, k) - \tilde{U}(z, k)| \cdot \mathbb{1}_{\{|z_N - z| > \rho\}} \right] \\ & \leq \mathbf{E} \left[|U_N(z_N, k) - \tilde{U}(z, k)| \cdot \mathbb{1}_{\{|z_N - z| \leq \rho\}} \right] \\ & \quad + 6\gamma^{-1} e^{L_F h} [(K-k)M_r + (T-t_k)M_\phi + M_\psi] \mathbf{E}|z_N - z| \\ & \leq \mathbf{E} \left[\left(\sup_{\nu \in \mathcal{A}} \left\{ |r(z_N, \nu) - r(z, \nu)| + \mathbf{E}^\nu \left| \int_{t_k}^{t_{k+1}} \left(\phi(\tilde{z}^\nu(s), \nu) - \phi(z^{\nu, N}(s), \nu) \right) ds \right| \right. \right. \right. \\ & \quad \left. \left. \left. + \mathbf{E}^\nu \left| U_N(z^{\nu, N}(t_{k+1}), k+1) - \tilde{U}(\tilde{z}^\nu(t_{k+1}), k+1) \right| \right\} \right) \cdot \mathbb{1}_{\{|z_N - z| \leq \rho\}} \right] \\ & \quad + \overline{M} e^{L_F h} \mathbf{E}|z_N - z|, \end{aligned} \tag{3.25}$$

where Chebyshev’s inequality has been used and we recall that the constant \overline{M} is defined in equation (3.13).

In the following, let us consider a fixed $z_N \in \mathbb{X}_N$ such that $|z_N - z| \leq \rho$. We consider the process $z^{\nu,N}(s)$ on $[t_k, t_{k+1}]$ with $z^{\nu,N}(t_k) = z_N$ and, similar to definition (3.2), we define the stopping time

$$\tau_N^\nu = \inf_{s \geq t_k} \left\{ s \mid z^{\nu,N}(s) \notin \Omega_{\gamma,z,[t_k,t_{k+1}]}^\nu \right\}.$$

For the notation, see the paragraph following (3.22). Since $\Omega_{\gamma,z,[t_k,T]}^u \subseteq \Omega$ for all $u \in \mathcal{U}_{o,k}$ trivially implies $\Omega_{\gamma,z,[t_k,t_{k+1}]}^\nu \subseteq \Omega$, Theorem 3.1 when considered on the time interval $[t_k, t_{k+1}]$ guarantees that $\exists N' > 0$, s.t. when $N \geq N'$ we have

$$\begin{aligned} \mathbf{E}^\nu \left[\sup_{t_k \leq s \leq t \wedge \tau_N^\nu} |z^{u,N}(s) - \tilde{z}^u(s)| \right] &\leq \left(|z_N - z| + C_{h,N} \right) e^{L_F(t-t_k)}, \quad t \in [t_k, t_{k+1}], \\ \mathbf{P}(\tau_N^\nu < t_{k+1}) &\leq 3\gamma^{-1} e^{L_F h} C_{h,N}, \end{aligned} \tag{3.26}$$

where the second inequality follows from estimate (3.10) in Remark 3.1.

We continue to estimate each of the three terms within the supremum in estimate (3.25). For the first term, noticing that $|z_N - z| \leq \rho < \gamma$ implies $z_N \in \Omega$, and the function r is Lipschitz in Ω ,

$$|r(z_N, \nu) - r(z, \nu)| \leq L_r |z_N - z|.$$

For the second term, using a similar argument as in the proof of Theorem 3.2 and the estimate (3.26), we can obtain, for $N > N'$,

$$\begin{aligned} &\mathbf{E}^\nu \left| \int_{t_k}^{t_{k+1}} (\phi(\tilde{z}^\nu(s), \nu) - \phi(z^{\nu,N}(s), \nu)) ds \right| \\ &\leq \mathbf{E}^\nu \left[\left(\int_{t_k}^{t_{k+1}} |\phi(\tilde{z}^\nu(s), \nu) - \phi(z^{\nu,N}(s), \nu)| ds \right) \mathbb{1}_{\{\tau_N^\nu \geq t_{k+1}\}} \right] + 2hM_\phi \mathbf{P}(\tau_N^\nu < t_{k+1}) \\ &\leq L_\phi \left(|z_N - z| + C_{h,N} \right) e^{L_F h} (t_{k+1} - t_k) + 6h\gamma^{-1} M_\phi e^{L_F h} C_{h,N}. \end{aligned}$$

For the third term, we notice the simple fact that $\Omega_{\gamma,z,[t_k,T]}^u \subseteq \Omega$ for all $u \in \mathcal{U}_{o,k}$ implies $\Omega_{\gamma,z',[t_{k+1},T]}^u \subseteq \Omega$ for all $u \in \mathcal{U}_{o,k+1}$, where $z' = \tilde{z}^\nu(t_{k+1})$, and also that $\tau_N^\nu \geq t_{k+1}$ implies $z^{\nu,N}(t_{k+1}) \in \Omega$. We have

$$\begin{aligned} &\mathbf{E}^\nu \left| U_N(z^{\nu,N}(t_{k+1}), k+1) - \tilde{U}(\tilde{z}^\nu(t_{k+1}), k+1) \right| \\ &\leq \mathbf{E}^\nu \left[\left| U_N(z^{\nu,N}(t_{k+1}), k+1) - \tilde{U}(\tilde{z}^\nu(t_{k+1}), k+1) \right| \cdot \mathbb{1}_{\{\tau_N^\nu \geq t_{k+1}\}} \right] \\ &\quad + 2 \left[(K - k - 1)M_r + (T - t_{k+1})M_\phi + M_\psi \right] \mathbf{P}(\tau_N^\nu < t_{k+1}) \\ &\leq \mathbf{E}^\nu \left[\left| U_N(z^{\nu,N}(t_{k+1}), k+1) - \tilde{U}(\tilde{z}^\nu(t_{k+1}), k+1) \right| \cdot \mathbb{1}_{\{\tau_N^\nu \geq t_{k+1}\}} \right] \mathbf{P}(\tau_N^\nu \geq t_{k+1}) \\ &\quad + 6\gamma^{-1} \left[(K - k - 1)M_r + (T - t_{k+1})M_\phi + M_\psi \right] e^{L_F h} C_{h,N} \\ &\leq a_{k+1} \mathbf{E}^\nu \left[|z^{\nu,N}(t_{k+1}) - \tilde{z}^\nu(t_{k+1})| \cdot \mathbb{1}_{\{\tau_N^\nu \geq t_{k+1}\}} \right] + b_{k+1} + \overline{M} e^{L_F h} C_{h,N} \\ &\leq a_{k+1} (|z_N - z| + C_{h,N}) e^{L_F h} + b_{k+1} + \overline{M} e^{L_F h} C_{h,N}, \end{aligned}$$

for $N > \max\{N', N_{k+1}\}$. In the above, we have used the conclusion for $k+1$ to the conditional expectation $\mathbf{E}^\nu(\cdot | \tau_N^\nu \geq t_{k+1})$.

Substituting the above estimates into (3.25), we conclude

$$\begin{aligned} & \mathbf{E}|U_N(z_N, k) - \tilde{U}(z, k)| \\ & \leq \mathbf{E} \left[L_r |z_N - z| + L_\phi (|z_N - z| + C_{h,N}) e^{L_F h} (t_{k+1} - t_k) + 6h\gamma^{-1} M_\phi e^{L_F h} C_{h,N} \right. \\ & \quad \left. + a_{k+1} (|z_N - z| + C_{h,N}) e^{L_F h} + b_{k+1} + \overline{M} e^{L_F h} C_{h,N} \right] + e^{L_F h} \overline{M} \mathbf{E}|z_N - z| \\ & \leq (L_r + L_\phi e^{L_F h} h + a_{k+1} e^{L_F h} + \overline{M} e^{L_F h}) \mathbf{E}|z_N - z| \\ & \quad + L_\phi C_{h,N} e^{L_F h} (t_{k+1} - t_k) + 2\overline{M} C_{h,N} e^{L_F h} + a_{k+1} C_{h,N} e^{L_F h} + b_{k+1} \\ & = a_k \mathbf{E}|z_N - z| + b_k, \end{aligned}$$

where the recursive relation (3.20) has been used in the last equation. This proves estimate (3.23) for k with $N_k = \max\{N', N_{k+1}\}$.

Equation (3.24) now follows from estimate (3.23) and Theorem 3.2, using the triangle inequality: $\exists N_0 > 0$, s.t. $N > N_0$, we have

$$\begin{aligned} & |J_N(z_N, u_0) - U_N(z_N)| \\ & \leq |J_N(z_N, u_0) - \tilde{J}(z, u_0)| + |\tilde{U}(z) - U_N(z_N)| \\ & \leq b_0 + a_0 |z_N - z| + \left(L_\phi \frac{e^{L_F T} - 1}{L_F} + (K L_r + L_\psi + \overline{M}) e^{L_F T} \right) (C_{T,N} + |z_N - z|). \end{aligned}$$

Convergence $|J_N(z_N, u_0) - U_N(z_N)| \rightarrow 0$ as $N \rightarrow \infty$ readily follows from Assumptions 2.1 and 2.2. \square

REMARK 3.3. As discussed in Remark 3.2, we have $C_{T,N} = \mathcal{O}(N^{\frac{1}{\alpha}-1})$ and thus $b_0 = \mathcal{O}(N^{\frac{1}{\alpha}-1})$ for the density dependent Markov chain introduced in Subsection 2.3. As a consequence, in this case we can explicitly compute the order of convergence in Theorems 3.2 and 3.3. That is, $\exists N_0 > 0$, s.t. when $N > N_0$,

$$|J_N(z_N, u) - \tilde{J}(z_0, u)| \leq C N^{\frac{1}{\alpha}-1}, \quad u \in \mathcal{U}_{0,0},$$

and

$$|J_N(z_N, u_0) - U_N(z_N)| \leq C N^{\frac{1}{\alpha}-1},$$

with $C > 0$ being a generic constant, u_0 being the optimal open loop policy for the limiting process \tilde{z}^u , and U_N being the value function of the stochastic feedback optimal control problem.

3.3. Feedback optimal control on infinite time-horizon with discounted cost. As a final step of our analysis, we consider the discounted optimal control problem on an infinite time-horizon. While the open loop control problem on a finite time horizon that is addressed in Theorem 3.2 will be useful later on in Sections 4 and 5, open loop control on an infinite time-horizon for stochastic processes seems to be less relevant in applications. Therefore, in the following, we consider the feedback optimal control problem with cost functional

$$J_N(z, u) = \mathbf{E}_z \left[\sum_{j=0}^{\infty} e^{-\beta t_j} \left(r(z^{u,N}(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(z^{u,N}(s), \nu_j) ds \right) \right], \quad (3.27)$$

where $\beta > 0$ is a discount factor, $u \in \mathcal{U}_f$ with

$$\mathcal{U}_f = \left\{ (\nu_0, \nu_1, \dots) \mid \nu_j : \mathbb{X}_N \rightarrow \mathcal{A}, \quad 0 \leq j < \infty \right\}, \quad (3.28)$$

and where again the shorthand $\nu_j = \nu_j(z^{u,N}(t_j))$ has been used in equation (3.27).

We assume that the control set \mathcal{A} is finite, which guarantees the existence of the optimal control policy and will simplify the proof of Theorem 3.4 (see below). We emphasize that this assumption is not essential and can be relaxed since we will only consider ϵ -optimal control policies in Theorem 3.4. Also see the related discussions in Subsection 2.2. Furthermore, we only focus on the case when the time stages at which the controls can be changed are uniformly distributed, i.e., $t_j = jh$ for some $h > 0$. This uniformity in time allows us to define value functions which only depend on the system's states and will simplify the discussions below.

It is known (e.g., [52]) that the value function $U_N(z) = \inf_{u \in \mathcal{U}_f} J_N(z, u)$ solves the Bellman equation

$$U_N(z) = \min_{\nu \in \mathcal{A}} \mathbf{E}_z^\nu \left[r(z, \nu) + \int_0^h \phi(z^{\nu,N}(s), \nu) ds + \lambda U_N(z^{\nu,N}(h)) \right], \quad (3.29)$$

where $\lambda = e^{-\beta h} < 1$. Moreover, it is known [52] that there is a map $\pi_N : \mathbb{X}_N \rightarrow \mathcal{A}$, such that $u_{opt} = (\pi_N, \pi_N, \dots) \in \mathcal{U}_f$ is an optimal feedback policy that satisfies $U_N(z) = J_N(z, u_{opt})$ and can be determined by the dynamic programming (i.e., Bellman) equation via

$$\pi_N(z) \in \operatorname{argmin}_{\nu \in \mathcal{A}} \left\{ r(z, \nu) + \mathbf{E}_z^\nu \left[\int_0^h \phi(z^{\nu,N}(s), \nu) ds + \lambda U_N(z^{\nu,N}(h)) \right] \right\}, \quad z \in \mathbb{X}_N.$$

In correspondence with the stochastic control problem, we also consider the optimal control of the deterministic limit dynamics $\tilde{z}^u(\cdot)$ which satisfies ODE (2.5), where

$$u \in \mathcal{U}_o := \left\{ (\nu_0, \nu_1, \dots) \mid \nu_j \in \mathcal{A}, \quad 0 \leq j < \infty \right\}.$$

In this context, it is necessary that the solution $\tilde{z}^u(\cdot)$ exists on $[0, +\infty)$. Recalling the set defined in (3.1), in the following we consider the subset $\Omega_g \subseteq \Omega$ with the following properties:

- (1) $z \in \Omega_g \implies \tilde{z}^u(t) \in \Omega_g, \forall 0 \leq t < \infty, \forall u \in \mathcal{U}_o$, and
- (2) for all $T > 0$, we can find $\gamma > 0$, such that $\Omega_{\gamma, z, [0, T]}^u \subseteq \Omega$ holds for all $u \in \mathcal{U}_o, \forall z \in \Omega_g$.

We emphasize that this (nonempty) subset Ω_g can be easily constructed as long as Ω is large enough and it doesn't have to be unique. In fact, when the solution \tilde{z}^u of the ODE (2.5) starting from $\tilde{z}^u(0) = z$ exists on $[0, +\infty)$ and stays in Ω for all time (without approaching its boundary) for any $u \in \mathcal{U}_o$, it is easy to see that the set $\Omega_g := \{\tilde{z}^u(t) \mid t \geq 0, u \in \mathcal{U}_o\}$ satisfies the above two conditions.

The natural candidate for the deterministic cost functional reads

$$\tilde{J}(z, u) = \sum_{j=0}^{\infty} e^{-\beta jh} \left(r(\tilde{z}^u(jh), \nu_j) + \int_{jh}^{(j+1)h} \phi(\tilde{z}^u(s), \nu_j) ds \right), \quad (3.30)$$

where $z \in \Omega_g$. Notice that again, following the convention in Subsection 2.1, we use the same notation ν_j to denote both the control policy function which depends on system's state, and the value of the control selected at t_j . See the discussion after equation (2.2).

By the dynamic programming principle, the corresponding value function $\tilde{U}(z) = \inf_{u \in \mathcal{U}_o} \tilde{J}(z, u)$ satisfies

$$\tilde{U}(z) = \min_{\nu \in \mathcal{A}} \left\{ r(z, \nu) + \int_0^h \phi(\tilde{z}^\nu(s), \nu) ds + \lambda \tilde{U}(\tilde{z}^\nu(h)) \right\}, \tag{3.31}$$

where $\tilde{z}^\nu(0) = z \in \Omega_g$. We will assume that a map $\pi_\infty : \Omega_g \rightarrow \mathcal{A}$ exists such that

$$\pi_\infty(z) \in \operatorname{argmin}_{\nu \in \mathcal{A}} \left\{ r(z, \nu) + \int_0^h \phi(\tilde{z}^\nu(s), \nu) ds + \lambda \tilde{U}(\tilde{z}^\nu(h)) \right\}, \tag{3.32}$$

where $\tilde{z}^\nu(0) = z \in \Omega_g$.

Assumption 2.4 implies that

$$\tilde{J}(z, u) \leq \sum_{j=0}^{\infty} e^{-\beta j h} \left(M_r + \int_{jh}^{(j+1)h} M_\phi ds \right) = \frac{M_r + M_\phi h}{1 - e^{-\beta h}} =: M_J. \tag{3.33}$$

Similarly, $J_N(z, u) \leq M_J$ and therefore the same upper bound applies to $\tilde{U}(z)$ and $U_N(z)$.

The next theorem provides the relations between the stochastic optimal control problem and the optimal control problem of the limiting ODE.

THEOREM 3.4. *Let the nonempty subset $\Omega_g \subseteq \Omega$ be given.*

- (1) *Suppose that Assumptions 2.3–2.4 hold. For every $\epsilon > 0$, there exists $C_\epsilon > 0$ such that*

$$\sup_{\substack{z, z' \in \Omega_g \\ |z - z'| \leq R}} |\tilde{U}(z) - \tilde{U}(z')| \leq C_\epsilon R + \epsilon, \quad \forall R > 0.$$

- (2) *Suppose that Assumptions 2.1–2.4 hold. Then for all $\epsilon > 0$, there exists $\delta > 0$ and $N' \in \mathbb{N}$, such that when $N \geq N'$,*

$$|U_N(z_N) - \tilde{U}(z)| \leq \epsilon, \tag{3.34}$$

for all $z_N \in \mathbb{X}_N \cap \Omega$, $z \in \Omega_g$, and $|z_N - z| \leq \delta$.

- (3) *Suppose that Assumptions 2.1–2.4 hold. Given $0 < \epsilon' < \epsilon$, $z \in \Omega_g$, and an ϵ' -optimal open loop policy $u = (\nu_0, \nu_1, \dots) \in \mathcal{U}_o$ of the limiting ODE system which satisfies*

$$\tilde{U}(z) \leq \tilde{J}(z, u) \leq \tilde{U}(z) + \epsilon',$$

there exist constants $N' \in \mathbb{N}$ and $\delta > 0$, depending on ϵ, ϵ' and z , such that for $N > N'$, we have

$$J_N(z_N, u) \leq U_N(z_N) + \epsilon$$

for all $z_N \in \mathbb{X}_N \cap \Omega$ and $|z_N - z| \leq \delta$. That is, u is an ϵ -optimal control policy for the feedback optimal control problem (3.27).

Proof.

- (1) Consider two starting points $z, z' \in \Omega_g$ and let $\nu = \pi_\infty(z)$. Let $\tilde{z}^\nu(s; z)$, $\tilde{z}^\nu(s; z')$ be the solutions of the ODE (2.6) on the time interval $[0, h]$ starting from z, z' at $s=0$, respectively. Notice that $z, z' \in \Omega_g$ implies both solutions stay in Ω all time.

By the Lipschitz continuity of the cost functions in Assumption 2.4, and equations (3.31)–(3.32), we have

$$\begin{aligned} \tilde{U}(z') - \tilde{U}(z) &\leq r(z', \nu) + \int_0^h \phi(\tilde{z}^\nu(s; z'), \nu) ds + \lambda \tilde{U}(\tilde{z}^\nu(h; z')) \\ &\quad - r(z, \nu) - \int_0^h \phi(\tilde{z}^\nu(s; z), \nu) ds - \lambda \tilde{U}(\tilde{z}^\nu(h; z)) \\ &\leq L_r |z - z'| + \int_0^h L_\phi |\tilde{z}^\nu(s; z') - \tilde{z}^\nu(s; z)| ds \\ &\quad + \lambda |\tilde{U}(\tilde{z}^\nu(h; z)) - \tilde{U}(\tilde{z}^\nu(h; z'))|. \end{aligned} \quad (3.35)$$

Using Assumption 2.3, the standard ODE theory implies

$$|\tilde{z}^\nu(t; z) - \tilde{z}^\nu(t; z')| \leq e^{L_F t} |z - z'|, \quad 0 \leq t \leq h. \quad (3.36)$$

Now for all $R \geq 0$, we define the function

$$G_1(R) = \sup_{\substack{z_1, z_2 \in \Omega_g, \\ |z_1 - z_2| \leq R}} |\tilde{U}(z_1) - \tilde{U}(z_2)|, \quad (3.37)$$

and it follows from inequality (3.33) that $G_1(R) \leq 2M_J$, $\forall R \geq 0$. Combining estimates (3.35) and (3.36), we find

$$G_1(R) \leq (L_r + L_\phi e^{L_F h}) R + \lambda G_1(e^{L_F h} R),$$

which, upon iterating the above inequality k times, leads to

$$G_1(R) \leq \left(L_r + L_\phi e^{L_F h} \right) \frac{1 - \lambda^k e^{L_F k h}}{1 - \lambda e^{L_F h}} R + 2\lambda^k M_J. \quad (3.38)$$

The first conclusion follows by noticing that $\lambda < 1$.

- (2) Given $\epsilon > 0$ and since $\lambda < 1$, we could first choose $k > 0$ such that $2\lambda^k M_J \leq \frac{\epsilon}{3}$. From the definition of the subset Ω_g , we know $\exists \gamma > 0$, s.t. $\Omega_{\gamma, z, [0, kh]}^u \subseteq \Omega$ is satisfied for all $z \in \Omega_g$ and $u \in \mathcal{U}_o$. Let the constant $0 < \delta < \frac{\gamma}{3} e^{-L_F h}$ and $z_N \in \mathbb{X}_N \cap \Omega$, such that $|z - z_N| \leq \delta$. Given $\nu \in \mathcal{A}$, we consider the stopping time

$$\tau_N^\nu = \inf_{s \geq 0} \left\{ s \mid |z^{\nu, N}(s) - \tilde{z}^\nu(s)| > 3\delta e^{L_F h} \right\} \wedge h, \quad (3.39)$$

where $z^{\nu, N}(0) = z_N$ and $\tilde{z}^\nu(0) = z$. In fact, under Assumptions 2.1–2.3 and using the fact that $\Omega_{\gamma, z, [0, h]}^u \subseteq \Omega$ (see the discussion before Theorem 3.3 on the notations), the same argument in Theorem 3.1 on the time interval $[0, h]$ implies that $\exists N' > 0$, s.t. when $N \geq N'$, we have

$$\mathbf{P}(\tau_N^\nu < h) \leq \delta^{-1} C_{h, N}, \quad (3.40)$$

where the constant $C_{h,N}$ is defined in equation (3.19). Also see estimate (3.10) in Remark 3.1.

More generally, for $R \geq 0$, we define the function

$$G_2(R) = \sup_{\substack{z' \in \mathbb{X}_N \cap \Omega, z \in \Omega_g \\ |z' - z| \leq R}} |U_N(z') - \tilde{U}(z)|,$$

and notice that Assumption 2.4 implies $|G_2(R)| \leq 2M_J, \forall R \geq 0$.

Letting $\nu = \pi_\infty(z) \in \mathcal{A}$, using the dynamic programming equations (3.29) and (3.31), and the estimate (3.40), we can obtain

$$\begin{aligned} & U_N(z_N) - \tilde{U}(z) \\ & \leq \mathbf{E}_{z_N}^\nu \left[\int_0^h \phi(z^{\nu,N}(s), \nu) ds + \lambda U_N(z^{\nu,N}(h)) \right] - \int_0^h \phi(\tilde{z}^\nu(s), \nu) ds - \lambda \tilde{U}(\tilde{z}^\nu(h)) \\ & \leq \mathbf{E}_{z_N}^\nu \left[\left(\int_0^h |\phi(z^{\nu,N}(s), \nu) - \phi(\tilde{z}^\nu(s), \nu)| ds \right) \cdot \mathbb{1}_{\{\tau_N^\nu \geq h\}} \right] \\ & \quad + \lambda \mathbf{E}_{z_N}^\nu \left[|U_N(z^{\nu,N}(h)) - \tilde{U}(\tilde{z}^\nu(h))| \cdot \mathbb{1}_{\{\tau_N^\nu \geq h\}} \right] + 2(hM_\phi + \lambda M_J) \mathbf{P}(\tau_N^\nu < h) \\ & \leq \mathbf{E}_{z_N}^\nu \left[\left(\int_0^h L_\phi |z^{\nu,N}(s) - \tilde{z}^\nu(s)| ds + \lambda |U_N(z^{\nu,N}(h)) - \tilde{U}(\tilde{z}^\nu(h))| \right) \cdot \mathbb{1}_{\{\tau_N^\nu \geq h\}} \right] \\ & \quad + 2(hM_\phi + \lambda M_J) \mathbf{P}(\tau_N^\nu < h) \\ & \leq 3L_\phi \delta e^{L_F h} h + \lambda G_2(3\delta e^{L_F h}) + 2\delta^{-1} (hM_\phi + \lambda M_J) C_{h,N}. \end{aligned}$$

In the above, we have used the facts that

$$\begin{aligned} z \in \Omega_g & \implies \tilde{z}^\nu(h) \in \Omega_g, \\ \tau_N^\nu \geq h & \implies \sup_{0 \leq s \leq h} |z^{\nu,N}(s) - \tilde{z}^\nu(s)| \leq 3\delta e^{L_F h} \leq \gamma \implies z^{\nu,N}(h) \in \Omega_{\gamma, z, [0, h]}^\nu \subseteq \Omega. \end{aligned}$$

Since the same upper bound holds for $\tilde{U}(z) - U_N(z_N)$ as well, taking the supremum over $z_N \in \mathbb{X}_N \cap \Omega, z \in \Omega_g$, such that $|z_N - z| < \delta$, we obtain

$$G_2(\delta) \leq 3L_\phi \delta e^{L_F h} h + \lambda G_2(3\delta e^{L_F h}) + 2\delta^{-1} (hM_\phi + \lambda M_J) C_{h,N},$$

as long as $\delta \leq \frac{\gamma}{3} e^{-L_F h}, N > N'$.

Notice that $\Omega_{\gamma, z, [0, kh]}^u \subseteq \Omega$ implies $\Omega_{\gamma, z', [ih, kh]}^u \subseteq \Omega$ for the same $\gamma > 0$, where $z' = \tilde{z}^u(ih), 0 < i < k, \forall u \in \mathcal{U}_o$. Therefore, iterating the above inequality k times and using the inequality $G_2 \leq 2M_J$, gives

$$\begin{aligned} G_2(\delta) & \leq 3L_\phi \delta e^{L_F h} h \frac{(3\lambda e^{L_F h})^k - 1}{3\lambda e^{L_F h} - 1} + 2\lambda^k M_J \\ & \quad + 2\delta^{-1} (hM_\phi + \lambda M_J) C_{h,N} \frac{3^{-k} \lambda^k e^{-kL_F h} - 1}{3^{-1} \lambda e^{-L_F h} - 1} \\ & \leq 3L_\phi \delta e^{L_F h} h \frac{(3\lambda e^{L_F h})^k - 1}{3\lambda e^{L_F h} - 1} + \frac{\epsilon}{3} \\ & \quad + 2\delta^{-1} (hM_\phi + \lambda M_J) C_{h,N} \frac{3^{-k} \lambda^k e^{-kL_F h} - 1}{3^{-1} \lambda e^{-L_F h} - 1} \end{aligned}$$

for $\delta \leq 3^{-k} e^{-kL_F h} \gamma$, $N > N'$.

Since Assumptions 2.1–2.2 imply that $C_{h,N} \rightarrow 0$ as $N \rightarrow \infty$, we can first choose δ and then N' such that $G_2(\delta) \leq \epsilon$ when $N \geq N'$. The conclusion follows readily.

- (3) We estimate the cost using the definition (3.27). Notice that the constant $\lambda = e^{-\beta h} < 1$ and that the open loop control u is ϵ' -optimal for the deterministic optimal control problem (3.30). For any $k \geq 1$, recalling the stopping time in definition (3.2) and Assumption 2.4, we obtain

$$\begin{aligned}
& J_N(z_N, u) \\
& \leq \mathbf{E}_{z_N}^u \left[\sum_{j=0}^k \lambda^j \left(r(z^{u,N}(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(z^{u,N}(s), \nu_j) ds \right) \right] + \sum_{j=k+1}^{\infty} \lambda^j (M_r + hM_\phi) \\
& \leq \sum_{j=0}^k \lambda^j \left(r(\tilde{z}^u(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(\tilde{z}^u(s), \nu_j) ds \right) + \sum_{j=k+1}^{\infty} \lambda^j (M_r + hM_\phi) \\
& \quad + \mathbf{E}_{z_N}^u \left[\sum_{j=0}^k \lambda^j |r(z^{u,N}(t_j), \nu_j) - r(\tilde{z}^u(t_j), \nu_j)| \right] \\
& \quad + \mathbf{E}_{z_N}^u \left[\sum_{j=0}^k \lambda^j \int_{t_j}^{t_{j+1}} |\phi(z^{u,N}(s), \nu_j) - \phi(\tilde{z}^u(s), \nu_j)| ds \right] \\
& \leq \sum_{j=0}^{\infty} \lambda^j \left(r(\tilde{z}^u(t_j), \nu_j) + \int_{t_j}^{t_{j+1}} \phi(\tilde{z}^u(s), \nu_j) ds \right) + 2 \sum_{j=k+1}^{\infty} \lambda^j (M_r + hM_\phi) \\
& \quad + \mathbf{E}_{z_N}^u \left[\left(\sum_{j=0}^k \lambda^j |r(z^{u,N}(t_j), \nu_j) - r(\tilde{z}^u(t_j), \nu_j)| \right) \cdot \mathbb{1}_{\{\tau_N^u \geq kh\}} \right] \\
& \quad + \mathbf{E}_{z_N}^u \left[\left(\sum_{j=0}^k \lambda^j \int_{t_j}^{t_{j+1}} |\phi(z^{u,N}(s), \nu_j) - \phi(\tilde{z}^u(s), \nu_j)| ds \right) \cdot \mathbb{1}_{\{\tau_N^u \geq kh\}} \right] \\
& \quad + 2(M_r + hM_\phi) \mathbf{P}(\tau_N^u < kh) \sum_{j=0}^k \lambda^j \\
& \leq \mathbf{E}_{z_N}^u \left[\left(\sup_{0 \leq s \leq kh} |z^{u,N}(s) - \tilde{z}^u(s)| \right) \cdot \mathbb{1}_{\{\tau_N^u \geq kh\}} \right] (L_r + hL_\phi) \sum_{j=0}^k \lambda^j \\
& \quad + \tilde{J}(z, u) + 2 \frac{M_r + hM_\phi}{1 - \lambda} \left(\lambda^{k+1} + \mathbf{P}(\tau_N^u < kh) \right) \\
& \leq \tilde{U}(z) + \epsilon' + \mathbf{E}_{z_N}^u \left[\left(\sup_{0 \leq s \leq kh} |z^{u,N}(s) - \tilde{z}^u(s)| \right) \cdot \mathbb{1}_{\{\tau_N^u \geq kh\}} \right] \frac{L_r + hL_\phi}{1 - \lambda} \\
& \quad + 2M_J \left(\lambda^{k+1} + \mathbf{P}(\tau_N^u < kh) \right) \\
& \leq U_N(z_N) + |U_N(z_N) - \tilde{U}(z)| + 2M_J \left(\lambda^{k+1} + \mathbf{P}(\tau_N^u < kh) \right) + \epsilon' \\
& \quad + \mathbf{E}_{z_N}^u \left[\left(\sup_{0 \leq s \leq kh} |z^{u,N}(s) - \tilde{z}^u(s)| \right) \cdot \mathbb{1}_{\{\tau_N^u \geq kh\}} \right] \frac{L_r + hL_\phi}{1 - \lambda}.
\end{aligned}$$

Now for $\epsilon > \epsilon'$, we can first choose $k > 0$ and then obtain $\gamma > 0$ using the property

of the subset Ω_g with $T = kh$. Applying Theorem 3.1 on the time interval $[0, kh]$, inequality (3.34), and Assumptions 2.1–2.2, we can find $N' \in \mathbb{N}$ and $\delta > 0$, such that

$$J_N(z_N, u) \leq U_N(z_N) + \epsilon$$

if $z \in \Omega_g$, $z_N \in \mathbb{X}_N \cap \Omega$, and $|z - z_N| < \delta$. The conclusion follows immediately. \square

4. Algorithms

In this section, we discuss some numerical aspects of the control problems studied in this paper. The main motivation is that, although our previous analysis suggested that the optimal open loop control of the limiting ODE system is a reasonable approximation whenever N is sufficiently large, in applications it is often difficult to verify how large N should be such that the approximation is satisfactory. On the other hand, the optimal feedback control becomes increasingly difficult to compute due to the rapid growth of the state space when N is large. The main purpose of this section is to construct an algorithm which further improves the optimal open loop policy by utilizing the information of the system state (i.e., by adding feedback), while avoiding the curse of dimensionality that is inherent to the dynamic programming approach.

In contrast to the previous sections, this part involves some heuristics, and we confine ourselves to the optimal control problem for a Markov jump process on a finite time-horizon $[0, T]$ with a finite control set \mathcal{A} . To this end, we assume that the parameter N is large, and we remind the reader again that $x^{u, N}$ denotes the original Markov jump process with a control policy u , and $z^{u, N} = N^{-1}x^{u, N}$ stands for the normalized density process. The state spaces on which $x^{u, N}$ and $z^{u, N}$ live are denoted by \mathbb{X} and \mathbb{X}_N , respectively.

4.1. Tau-leaping method. In order to compute the optimal control policy, it is necessary to simulate trajectories of the underlying Markov jump process and to estimate the corresponding cost. The stochastic simulation algorithm (SSA) [17, 18, 21] is a typical Monte Carlo method: At each time step, it determines the waiting time in definition (2.1) as well as the next state according to the jump rates between the current state and the next possible states. When N is large, however, the system becomes numerically stiff because a large number of jump events occur within a short time interval. Since SSA traces every single jump event of the system, the effective step size of the method decreases rapidly, which renders the SSA inefficient.

As a remedy to this problem, the tau-leaping method [9, 20, 21, 26, 53] aims at increasing the effective step size by updating the state vector according to the transitions that may occur within a given time interval. Roughly speaking, instead of computing the waiting time and the next jump, the idea of the tau-leaping method is to answer the question “How many times will each type of jumps occur within a given time interval?” and then update the state vector accordingly. With a proper and carefully chosen step size [9], the tau-leaping method can approximate the SSA quite well and meanwhile reduce the simulation time by up to 1 or 2 orders of magnitude. In our implementation (see the numerical examples in Section 5), we use the explicit tau-leaping method where the leaping time step sizes are determined according to [9].

4.2. State space truncation. The computational complexity for solving the feedback optimal control problem is proportional to the number of states in \mathbb{X} considered (which is of order N^n , with n being the number of species). Therefore,

truncating the state space \mathbb{X} is necessary before numerically solving the optimal feedback control. One such approach to truncating the state space is to consider only states $x = (x^{(1)}, x^{(2)}, \dots, x^{(n)}) \in \mathbb{X}$ that lie within a hypercube defined by $x^{(i)} \in [c_i N, c'_i N]$, $1 \leq i \leq n$, where $0 \leq c_i < c'_i$ are estimations of the lower and upper bounds of the average densities per species. The cut-off values c_i, c'_i could, for example, be determined by launching independent simulations of the jump process controlled by candidate open loop control policies.

Once a truncated state space \mathbb{X}_{cut} has been constructed, then a simple algorithm (Algorithm 1) to compute the optimal feedback control policy can be based on the necessary optimality condition (3.17) with the terminal condition $U_N(\cdot, K) = \psi$ where the expectation value in (3.17) is estimated by a Monte Carlo average. If T is the total simulation time, $\Delta t > 0$ is the average time step size used to generate trajectories (e.g., by SSA or tau-leaping) and we use M independent realizations for each starting state to approximate the expectation value, the overall computational cost of Algorithm 1 is $\mathcal{O}(M \cdot |\mathcal{A}| \cdot |\mathbb{X}_{cut}| \cdot \lceil T/\Delta t \rceil)$.

Algorithm 1 Compute the optimal feedback control policy on truncated state space

- 1: Set $U_N(\cdot, K) = \psi$.
 - 2: **for** $k \leftarrow K - 1$ to 0 **do**
 - 3: **for** each $x \in \mathbb{X}_{cut}$ **do**
 - 4: **for** each $\nu \in \mathcal{A}$ **do**
 - 5: Starting from x at time t_k , generate M trajectories $x_i^{\nu, N}$ till time t_{k+1} , such that $x_i^{\nu, N}(t_{k+1}) \in \mathbb{X}_{cut}$ (generate new realization if $x_i^{\nu, N}(t_{k+1}) \notin \mathbb{X}_{cut}$).
 - 6: Let $z = x/N$, $z_i^{\nu, N} = x_i^{\nu, N}/N$, compute

$$Q(\nu) = \frac{1}{M} \sum_{i=1}^M \left(r(z, \nu) + \int_{t_k}^{t_{k+1}} \phi(z_i^{\nu, N}(s), \nu) ds + U_N(z_i^{\nu, N}(t_{k+1}), k+1) \right).$$
 - 7: **end for**
 - 8: Set $\nu_k(z) = \operatorname{argmin}_{\nu \in \mathcal{A}} Q(\nu)$ and $U_N(z, k) = \min_{\nu \in \mathcal{A}} Q(\nu)$.
 - 9: **end for**
 - 10: **end for**
-

4.3. Hybrid control. Solving the feedback control problem may be computationally infeasible even after truncation of the state space. As already mentioned at the beginning of this section, we will utilize an adaptive state space truncation strategy which exploits information from the (optimal) open loop control policies. The key idea is to assume that the typical states visited by the jump process when an optimal open loop policy is applied are also important states for computing a sufficiently accurate feedback control policy. To this end, the following algorithm generates states (for each control stage) whose densities are scattered around the density values of the system controlled by reasonable open loop control policies.

Adaptive truncation strategy. Let $\mathcal{S}_j \subset \mathbb{X}$ denote the finite state set at the j -th control stage after truncation, $0 \leq j < K$. We construct sets \mathcal{S}_j using the following steps.

- (1) Compute “good” (open loop) candidate policies for the Markov jump process on time $[0, T]$. A control policy $u_k \in \mathcal{U}_{o,0}$ is called “good” if $k < n_{ol}$ and $J_N(u_k) \leq (1 + \epsilon_{ol}) J_N(u_0)$ for appropriately chosen $n_{ol} \in \mathbb{N}$, $n_{ol} \geq 1$ and $\epsilon_{ol} \geq 0$ (in particular, u_0 is

the optimal open loop control policy for the jump process). Sort all “good” control policies $u_k \in \mathcal{U}_{o,0}$ by their costs in non-decreasing order.

- (2) Compute statistics of the controlled jump processes under “good” policies. For each “good” open loop policy u_k , record the average densities $z_{k,j} \in \mathbb{R}^n$ and the standard deviations $\sigma_{k,j} \in \mathbb{R}^n$ of the controlled normalized density process at each stage j , $0 \leq j < K$.
- (3) Compute the truncated sets \mathcal{S}_j . For each “good” open loop policy u_k , generate M_{ol} trajectories and add the states $x \in \mathbb{X}$ of each trajectory at stage j to the set \mathcal{S}_j if

$$x^{(i)}/N \in [z_{k,j}^{(i)} - \zeta\sigma_{k,j}^{(i)}, z_{k,j}^{(i)} + \zeta\sigma_{k,j}^{(i)}], \quad \forall i \in \{1, \dots, n\} \tag{4.1}$$

where $\zeta > 0$ is a pre-selected constant, and $x^{(i)}, z_{k,j}^{(i)}, \sigma_{k,j}^{(i)}$ are the i th components of $x, z_{k,j}, \sigma_{k,j} \in \mathbb{R}^n$.

REMARK 4.1. A few remarks about the above algorithm are in order.

- (1) In the case that the jump process starts from a fixed initial value x , $\mathcal{S}_0 = \{x\}$ is a singleton containing only the initial state.
- (2) Step 1 can be accomplished by enumerating all possible (finite) $u_k \in \mathcal{U}_{o,0}$ and computing the cost $J(u_k)$ by simulating trajectories using SSA or the tau-leaping method. Parameters n_{ol} and ϵ_{ol} are introduced in order to determine the number of “good” open loop policies which might carry important information and will be used to construct the truncated state sets \mathcal{S}_j in Steps 2, 3 above. By the central limit theorem (see [33]), the state distributions of the jump process under “good” open loop policies are approximately Gaussian whenever N is large. Hence the standard statistical estimators for the means and standard deviations computed in Step 2 can capture the distributions to a good approximation.
- (3) Ideally, for every “good” control policy u_k and every control stage j , we would like to record all possible (i.e., reachable) discrete states that satisfy condition (4.1). However, this set may be very large. Therefore, we sample these reachable states in Step 3 with a tunable parameter M_{ol} , which can control the number of states in \mathcal{S}_j . The drawback is that important states may be missing when they are not visited by the M_{ol} trajectories (see below for a patch).

Hybrid control policy. Having the state sets \mathcal{S}_j at hand, the task of computing a feedback control policy is to determine maps $\nu_j: \mathcal{S}_j \rightarrow \mathcal{A}$, $0 \leq j < K$, according to a modification of Algorithm 1. Keeping in mind that the sets \mathcal{S}_j may be only partially sampled, it is quite possible that, at some control stage j , the system fails to reach \mathcal{S}_j under control ν_{j-1} . To remedy this defect, we propose the following strategy: Denote the best available open loop policy as $u_0 = (\nu_0^0, \nu_1^0, \dots, \nu_{K-1}^0)$, and consider the j -th control stage, $0 \leq j < K$ where we suppose that the system has ended up in a state $x \notin \mathcal{S}_j$. Further, let x' be one of the nearest states to x among all states in \mathcal{S}_j , i.e., $x' \in \operatorname{argmin}_{x' \in \mathcal{S}_j} |x - x'|$. Then we apply the control $\nu_j(x')$ if $|x - x'|/N \leq \epsilon_{near}$, where ϵ_{near} is a cut-off parameter, and otherwise we use ν_j^0 . In other words, we replace the original candidate control by the modified control policy $u = (\bar{\nu}_0, \bar{\nu}_1, \dots, \bar{\nu}_{K-1}) \in \mathcal{U}_{f,0}$ with

$$\bar{\nu}_j: \mathbb{X} \rightarrow \mathcal{A}, \quad \bar{\nu}_j(x) = \begin{cases} \nu_j(x'), & \text{if } |x' - x|/N < \epsilon_{near} \\ \nu_j^0, & \text{otherwise.} \end{cases} \tag{4.2}$$

In the following, we keep using ν_j instead of $\bar{\nu}_j$ when no ambiguity exists. This strategy can prevent problems that arise when the feedback policy ν_j at stage j cannot be computed because some rare, but important states are missing due to the insufficient sampling when constructing the set \mathcal{S}_j . Notice that the algorithmic modification can be easily switched off by setting $\epsilon_{near} = 0$. In this case, the feedback policy is applied only when the states belong to \mathcal{S}_j , while open loop policies are applied otherwise. In agreement with the notation used in Sections 1–3, we define

$$\mathcal{U}_{h,k} = \{(\bar{\nu}_k, \bar{\nu}_{k+1}, \dots, \bar{\nu}_{K-1}) \mid \nu_j : \mathcal{S}_j \rightarrow \mathcal{A}, k \leq j < K\}, \quad 0 \leq k < K, \quad (4.3)$$

as the set of all hybrid control policies, where $\bar{\nu}_j$ is defined as in (4.2). The algorithmic task now boils down to finding the optimal hybrid control policy $u \in \mathcal{U}_{h,0}$. In order to solve this task, we consider the cost function $J_N(z, u, k)$ as in (3.16) and define a modified value function as

$$U_N(z, k) = \inf_{u \in \mathcal{U}_{h,k}} J_N(z, u, k), \quad Nz \in \mathcal{S}_k. \quad (4.4)$$

By definition, the value function satisfies the terminal condition $U_N(z, K) = \psi(z)$ and a modified Bellman equation as a necessary optimality condition:

$$U_N(z, k) = \min_{\nu \in \mathcal{A}} \mathbf{E}^\nu \left[\sum_{j=k}^{\tau-1} \left(r(z^{u,N}(t_j), \nu_j(z^{u,N}(t_j))) \right) + \int_{t_j}^{t_{j+1}} \phi(z^{u,N}(s), \nu_j(z^{u,N}(t_j))) ds \right] + U_N(z^{u,N}(t_\tau), \tau), \quad Nz \in \mathcal{S}_k, \quad (4.5)$$

where $z^{u,N}(t_k) = z$, $u = (\nu_k, \nu_{k+1}, \dots, \nu_{K-1})$ with $\nu_k = \nu$ and $(\nu_{k+1}, \dots, \nu_{K-1}) \in \mathcal{U}_{h,k+1}$ is the optimal hybrid control policy starting from stage $k+1$. The terminal index τ is a stopping time, depending on the particular realization, and is either the smallest stage index such that $k < \tau < K$ and $Nz^{u,N}(t_\tau) \in \mathcal{S}_\tau$, or $\tau = K$ otherwise. Notice that in (4.5), only values of $U_N(z, k)$ at states z such that $Nz \in \mathcal{S}_k$ are involved. Based on it, we can compute the optimal hybrid control policy by backward iterations in Algorithm 2 below.

A computational bottleneck in computing the hybrid control policy for $\epsilon_{near} > 0$ is the solution of the minimization problem $\operatorname{argmin}_{x' \in \mathcal{S}_j} |x - x'|$, i.e., to find the nearest neighbor of x in \mathcal{S}_j . The computational complexity of a direct minimization based on a pairwise comparison is $\mathcal{O}(|\mathcal{S}_j|)$, which would increase the computational cost of Algorithm 2 to $\mathcal{O}(M \cdot |\mathcal{A}| \cdot |\mathcal{S}_j|^2 \cdot \lceil T/\Delta t \rceil)$ (assuming $\tau = k+1$ and $|\mathcal{S}_j|$ are constant for simplicity). However, by employing the so-called k - d tree data structure [7] to store the states in \mathcal{S}_j , the computational complexity of finding the nearest neighbor can be reduced to $\mathcal{O}(\ln |\mathcal{S}_j|)$, by which the total computational cost is of the order

$$\mathcal{O}(M \cdot |\mathcal{A}| \cdot |\mathcal{S}_j| \cdot \ln |\mathcal{S}_j| \cdot \lceil T/\Delta t \rceil).$$

In the numerical examples in Section 5 below, our implementation uses the ANN (Approximate Nearest Neighbor) library [45], which provides operations on k - d trees and efficient algorithms for finding the first k -th ($k = 1$ in our case) nearest neighbors.

Algorithm 2 Compute the optimal hybrid control policy

- 1: Set $U_N(\cdot, K) = \psi$.
 - 2: **for** $k \leftarrow K - 1$ to 0 **do**
 - 3: **for** each $x \in \mathcal{S}_k$ **do**
 - 4: **for** each $\nu \in \mathcal{A}$ **do**
 - 5: Set $u = (\nu, \nu_{k+1}, \dots, \nu_{K-1})$, where ν_j is the optimal policy function on the j -th stage, $k < j < K - 1$ (already computed).
 - 6: Generate M trajectories $x_i^{u,N}$ from time t_k to t_{τ_i} where $k < \tau_i$ and t_{τ_i} is either the first time when $x_i^{u,N}(t_{\tau_i}) \in \mathcal{S}_{\tau_i}$ or $\tau_i = K$, $1 \leq i \leq M$.
 - 7: Let $z = x/N$, $z_i^{u,N} = x_i^{u,N}/N$, compute

$$Q(\nu) = \frac{1}{M} \sum_{i=1}^M \left\{ \sum_{j=k}^{\tau_i-1} \left[r(z_i^{u,N}, \nu_j(z_i^{u,N}(t_j))) + \int_{t_j}^{t_{j+1}} \phi(z_i^{u,N}(s), \nu_j(z_i^{u,N}(t_j))) ds \right] + U_N(z_i^{u,N}(t_{\tau_i}), \tau_i) \right\}.$$
 - 8: **end for**
 - 9: Set $\nu_k(z) = \operatorname{argmin}_{\nu \in \mathcal{A}} Q(\nu)$ and $U_N(z, k) = \min_{\nu \in \mathcal{A}} Q(\nu)$.
 - 10: **end for**
 - 11: **end for**
-

5. Numerical examples

In this section, we consider two numerical examples in order to demonstrate the analysis and the algorithms discussed in the previous sections.

5.1. Birth-death process. First, we consider the one-dimensional birth-death process which can be described as

$$x - 1 \xleftarrow{\kappa_-} x \xrightarrow{\kappa_+} x + 1, \tag{5.1}$$

where $x \in \mathbb{N}^+$. We suppose that the process has a density dependent birth rate which is $x \cdot \kappa_+$ when the current state is x , and similarly, $x \cdot \kappa_-$ for the death rate. We fix $T = 3.0$ and $K = 3$, i.e., the control can be switched at time $t = 0.0, 1.0, 2.0$. Two control/parameterization sets $\mathcal{A}_1, \mathcal{A}_2$ shown in Table 5.1 are considered. Each set contains two controls $\nu^{(0)}, \nu^{(1)}$ that affect the jump rates κ_- and κ_+ . For the optimal control problem, let $x^{u,N}(t)$ be system’s state at $t \in [0, T]$ with control $u \in \mathcal{U}_{\sigma,0}$ and set $r(z, \nu) = \psi(z) = 0$, $\phi(z, \nu) = |z - 1.0|$ for $\nu \in \mathcal{A}_i$, $i = 1, 2$, leading to the cost function

$$J_N(z_0, u) = \mathbf{E}_{z_0}^u \left[\int_0^3 |z^{u,N}(t) - 1.0| dt \right], \quad u \in \mathcal{U}_{\sigma,0}, \tag{5.2}$$

with $z^{u,N}(t) = x^{u,N}(t)/N$ and $z^{u,N}(0) = z_0$. Minimizing a cost as in equation (5.2) may arise when one wishes to keep the density of the system (5.1) not far away from 1.0 by controlling the jump rates (depending on the system’s states). Fixing $z_0 = 1.2$ and one of two control sets $\mathcal{A}_1, \mathcal{A}_2$, we shall compare the optimal open loop and feedback control policies for the jump process as N increases, as well as the optimal (open loop) control policy for the related deterministic ODE

$$\frac{d\tilde{z}^u(t)}{dt} = (\kappa_+ - \kappa_-)\tilde{z}^u(t), \quad \tilde{z}^u(0) = 1.2. \tag{5.3}$$

Open loop control. In the case of open loop control, there are $|\mathcal{U}_{o,0}| = 2^3 = 8$ different control policies in total for both the jump process (5.1) and the deterministic ODE (5.3), regardless of the value N , since one of the two controls $\nu^{(0)}, \nu^{(1)}$ can be selected at any of the three control stages. The optimal control is obtained by simply comparing the costs of all 8 possible policies. In Figure 5.1, the evolutions of the means and the standard deviations of the density $z^{u,N}(t)$ are shown for different N . For both control sets $\mathcal{A}_1, \mathcal{A}_2$, it is observed that the standard deviations decrease and the means get closer to that of the ODE controlled by the optimal control policy as N grows larger. For the control set \mathcal{A}_2 , we observe that the suboptimal policy $u_2 = (1, 1, 0)$ leads to a cost which is close to the optimal cost (that is determined by choosing the optimal policy $u_1 = (1, 0, 1)$) of the ODE system. (For the ease of notation, we use the index of the control action to denote the control policy, e.g., $(1, 0, 1)$ means $(\nu^{(1)}, \nu^{(0)}, \nu^{(1)})$.) For the jump processes with $N = 40$ or $N = 100$, u_2 performs even better than u_1 ; cf. Figure 5.3.

Feedback control. Now we turn to the feedback control problem, in which case the optimal control policy can be obtained by iterating the dynamic programming equations (3.17)–(3.18) by backward iterations. As the state space $\mathbb{X} = \mathbb{N}^+$ is infinite, finite state truncation is necessary for Algorithm 1 to work. Based on a rough estimation of the solution of ODE (5.3), and taking account of the form of the cost functional (5.2), the initial condition $z_0 = 1.2$, as well as the jump rates κ_+, κ_- , we truncate the space into the finite subset $\mathbb{X}_{cut} = \{N/2, N/2 + 1, \dots, 2N\} \subset \mathbb{N}^+$ (see discussions in Subsection 4.2).

Figure 5.2 shows the means and the standard deviations of $z^{u,N}(t)$ under the optimal feedback control policy as a function of time for increasing N . Generally, for both control sets \mathcal{A}_1 and \mathcal{A}_2 , the optimal feedback control policies lead to smaller costs as compared to the optimal open loop controls (Figure 5.3). Specifically, we observe in Figure 5.2(a) that, for the control set \mathcal{A}_1 , the standard deviations decrease and the means converge to the densities of the optimally controlled ODE system (by u_2) as N increases. For the control set \mathcal{A}_2 , due to the existence of the competing policy u_2 (in this case, $u_1 = (1, 0, 1)$ is optimal for the ODE system), some states with density close to $z = 1.0$ may select the control $\nu^{(1)}$ at stage $t = 1.0$, while others select $\nu^{(0)}$ (see Table 5.1), which leads to a significant rise in the standard deviation at the next control stage $t = 2.0$ (see Figure 5.2(b)); we moreover notice that the convergence of the empirical means of the controlled jump process at time $t = 2.0$ to the ODE solution is slower than in case of the control set \mathcal{A}_1 as N increases. The last observation is in agreement with Figure 5.4(a) which shows the bimodal probability density function of the optimally controlled process at time $t = 2.0$ that becomes even more pronounced for larger values of N . Nevertheless, Figure 5.3 clearly shows the convergence of the cost values of both open loop and feedback control policies as N increases, in line with the theoretical prediction. Also notice that, in Figure 5.3(b), the optimal costs using feedback and hybrid policies for finite N can be smaller than the optimal cost of the limiting ODE system, i.e., the convergence may be not monotonically decreasing from above. As a final demonstration, Figures 5.3(a) and 5.4(b) show a comparison of the SSA and the tau-leaping methods, with the clear indication that the results using the tau-leaping method are close to the SSA prediction, but at much lower computational cost.

Hybrid control. Finally, we consider the hybrid control policy following the procedure discussed in Subsection 4.3 and we confine our attention to the control set \mathcal{A}_2 . To assess the approximation quality of the hybrid control algorithm, we compute the cost under the open loop control policies for various values of N and with 5000 trajectories for each possible policy. As “good” control policies, we define the suboptimal controls

No.	control	\mathcal{A}_1		\mathcal{A}_2	
		κ_-	κ_+	κ_-	κ_+
0	$\nu^{(0)}$	<u>0.6</u>	1.0	<u>0.8</u>	1.0
1	$\nu^{(1)}$	1.0	0.8	1.0	0.8

TABLE 5.1. Two different control sets $\mathcal{A}_1, \mathcal{A}_2$ for the birth-death jump process. Each set contains two controls where the underlined entries indicate different control actions in \mathcal{A}_1 and \mathcal{A}_2 .

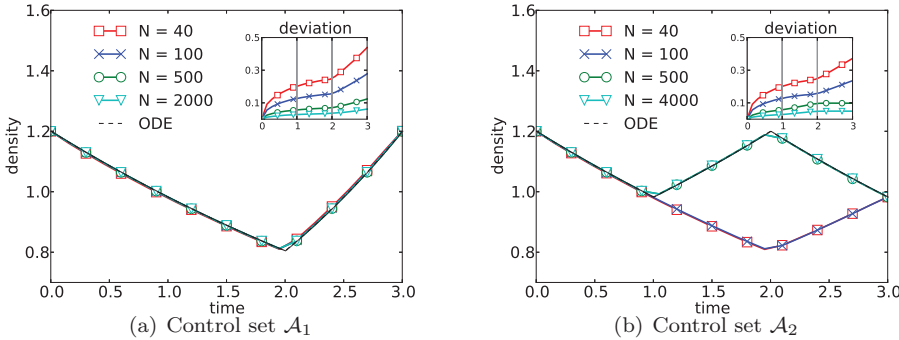


FIGURE 5.1. Birth-death process. Evolution of the empirical means and the standard deviations (inset plot) of the normalized density process $z^{u,N}$ under the optimal open loop control policies in comparison with the ODE solutions. Here N is the scaling number and controls are switched at times $t = 0.0, 1.0, 2.0$. (a) Control Set \mathcal{A}_1 . The optimal policy is $u_2 = (1, 1, 0)$ both for the jump process for all N and for the ODE system. (b) Control Set \mathcal{A}_2 . The optimal policy is $u_2 = (1, 1, 0)$ for the jump process with $N = 40, 100$, but it becomes $u_1 = (1, 0, 1)$ for $N = 500, 4000$ and the ODE system.

with $n_{ol} = 2$ and $\epsilon_{ol} = 0.05$ (see page 316). Sets \mathcal{S}_j are computed from $M_{ol} = 5000$ realizations for each “good” open loop policy according to (4.1) with $\zeta = 2.5$. As Figure 5.4(c) illustrated, the cardinality of the sets \mathcal{S}_1 and \mathcal{S}_2 is much smaller than the cardinality of \mathbb{X}_{cut} used in the feedback control case, which can lead to a tremendous reduction of the computational effort as compared to Algorithm 1 at almost no loss of numerical accuracy (see Figure 5.3).

5.2. Predator-prey model. In this subsection, we consider a two dimensional predator-prey model on the state space $\mathbb{X} = \mathbb{N}^+ \times \mathbb{N}^+$. We call A and B the prey and predator species, and let $x = (x^{(1)}, x^{(2)}) \in \mathbb{X}$ denote the numbers of species A and B . We suppose that both the prey and predator reproduce or decrease naturally, with the predator eating the prey in order to reproduce. Recalling the notations explained in Subsection 2.3, the dynamics of A, B species can be modelled as a jump process on \mathbb{X} according to the rules (see [31])

- (1) $A \xrightarrow{\lambda_1} 2A, \quad A \xrightarrow{\mu_1} \emptyset$
- (2) $B \xrightarrow{\lambda_2} 2B, \quad B \xrightarrow{\mu_2} \emptyset$
- (3) $A + B \xrightarrow{b} B, \quad A + B \xrightarrow{c} A + 2B.$

A control corresponds to a vector $\nu = (\lambda_1, \mu_1, \lambda_2, \mu_2, b, c)$, where each parameter takes positive real values. Now we define the jump vectors $l_1 = (1, 0), l_2 = (0, 1)$ and consider the normalized state vector $z = (z^{(1)}, z^{(2)}) = x/N \in \mathbb{R}^2$ for a fixed scaling parameter $N \gg$

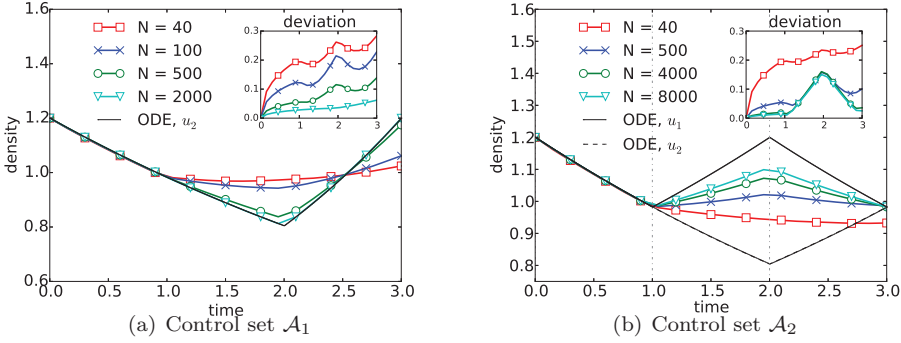


FIGURE 5.2. Birth-death process. Evolution of the empirical means and the standard deviations (inset plot) of the normalized density process $z^{u,N}$ under the optimal feedback control policies in comparison with the ODE solutions. N is the scaling number and controls are switched at times $t=0.0, 1.0, 2.0$. (a) Control set \mathcal{A}_1 : as N increases, the standard deviations decrease and the empirical means get closer to the ODE solution under the optimal policy $u_2 = (1, 1, 0)$. (b) Control set \mathcal{A}_2 : the policies $u_1 = (1, 0, 1)$ and $u_2 = (1, 1, 0)$ are the dominant (sub)optimal control policies for the ODE system.

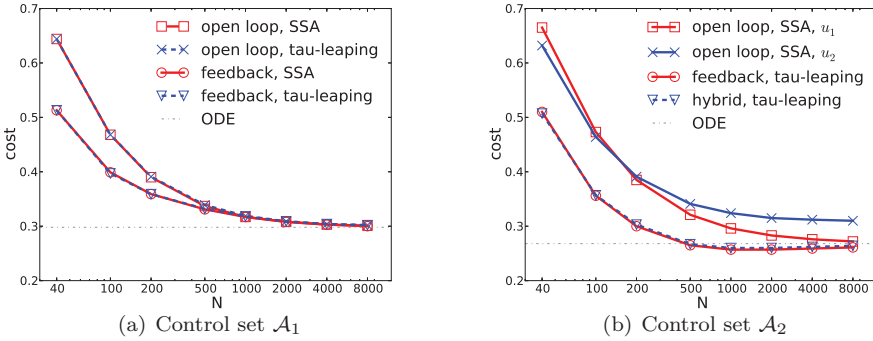


FIGURE 5.3. Birth-death process. Cost values for the jump processes with different scaling number N . Both SSA and tau-leaping methods are used to sample trajectories. For the control set \mathcal{A}_2 , $u_1 = (1, 0, 1)$, $u_2 = (1, 1, 0)$ are the best two open loop policies.

1. The jump rates for the normalized density process are then given by

$$\begin{aligned}
 f_d^{\nu,N}(z, \frac{l_1}{N}) &= \lambda_1 N z^{(1)}, & f_d^{\nu,N}(z, -\frac{l_1}{N}) &= N(\mu_1 + bz^{(2)})z^{(1)}, \\
 f_d^{\nu,N}(z, \frac{l_2}{N}) &= N(\lambda_2 + cz^{(1)})z^{(2)}, & f_d^{\nu,N}(z, -\frac{l_2}{N}) &= \mu_2 N z^{(2)},
 \end{aligned} \tag{5.4}$$

which indicate that the process is density dependent (see Subsection 2.3), with the vector fields $F^{\nu,N}(z)$ in (2.7) given by

$$F^\nu(z) = F^{\nu,N}(z) = \left((\lambda_1 - \mu_1)z^{(1)} - bz^{(1)}z^{(2)}, cz^{(1)}z^{(2)} - (\mu_2 - \lambda_2)z^{(2)} \right). \tag{5.5}$$

Our aim is to study the optimal control problem on a finite time-horizon $[0, T]$, with terminal time $T=5.0$ and $K=5$ control stages at times $t=j \times 1.0$, $0 \leq j \leq 4$. We define




FIGURE 5.4. Birth-death process. Dynamics under the control set \mathcal{A}_2 . (a) The probability distribution of states of the jump process at time $t=2.0$ under the optimal feedback control. (b) The CPU run time (in seconds) for different values of N , where the algorithm is run in parallel using 10 processors in each case. (c) Number of states in the sets \mathbb{X}_{cut} (feedback control) and $\mathcal{S}_1, \mathcal{S}_2$ (hybrid control).

the cost functional as

$$J_N(z_0, u) = \mathbf{E}_{z_0}^u \left[\int_0^{5.0} \left(|z_t^{(1),u,N} - 2z_t^{(2),u,N}| + |z_t^{(1),u,N} - 1.5| \right) dt \right], \quad u \in \mathcal{U}_{\sigma,0}, \quad (5.6)$$

where $z^{u,N}(t) = (z_t^{(1),u,N}, z_t^{(2),u,N}) = N^{-1}x^{u,N}(t)$ is the normalized density jump process with initial condition $z^{u,N}(0) = z_0$. In our numerical experiment, we set $z_0 = (1.0, 0.4)$ and choose $N = 50, 100, 200, 500, 1000, 2000, 4000$.

The particular choice of the cost functional J_N is aimed at maintaining the density of the prey species around $z^{(1)} = 1.5$ over time $[0, 5.0]$, with roughly about two times more prey than predator. The control set \mathcal{A} contains three different controls and is shown in Table 5.2: Observe that, in comparison with $\nu^{(0)}$, the prey reproduces faster under the control $\nu^{(1)}$ and the predators decrease more slowly, while the control $\nu^{(2)}$ has the reverse effect.

Open loop control. We do a brute-force calculation of the optimal open loop control policy based on ranking all possible $3^5 = 243$ policies in $\mathcal{U}_{o,0}$ according to their cost. In each case, 50000 trajectories are sampled using both SSA and tau-leaping methods. From Table 5.3, we conclude that for large N (≥ 500), tau-leaping method outperforms the SSA, as is indicated by the large increment of the effective time step sizes. Except for the system with $N=50$ whose optimal open loop control policy is $u_1 = (0, 2, 1, 0, 2)$ with the corresponding cost 11.26, the optimal policies for other larger N are all $u_2 = (0, 2, 1, 2, 2)$, which is also the optimal policy for the limiting ODE system (for $N=50$, u_2 is the second best policy with cost 11.30), see Figure 5.7. The empirical means and the standard deviations of the normalized density process $z^{u,N}$ are shown in Figure 5.5 for various values of N . As can be expected from the theoretical predictions, we observe that the mean values approach the solution of the limiting ODE, with the standard deviations decreasing as N increases. Convergence of the cost values to the cost value of the limit ODE system is also observed in Figure 5.7.

Hybrid control. We continue to study the hybrid control policy introduced in Subsection 4.3. Firstly, all 243 possible open loop control policies are ordered by their costs, among which we identify all “good” policies with $n_{ol} = 3, \epsilon_{ol} = 0.05$. Then, secondly, we estimate the empirical means and the standard deviations of the process under all “good” policies based on 5000 independent realizations of the process. Thirdly, for

No.	control	λ_1	μ_1	λ_2	μ_2	b	c
0	$\nu^{(0)}$	2.5	0.2	0.2	2.0	2.0	2.0
1	$\nu^{(1)}$	<u>2.7</u>	0.2	0.2	<u>1.5</u>	2.0	2.0
2	$\nu^{(2)}$	2.5	0.2	0.2	<u>2.5</u>	2.0	2.0

TABLE 5.2. Predator-prey model. The control set \mathcal{A} contains three different controls to modify the rates in the predator-prey model. The major differences among the controls are indicated by the underlined rates.

N	50	100	200	500	1000	2000	4000
Δ_{st}	1.8×10^{-3}	9.0×10^{-4}	4.5×10^{-4}	1.8×10^{-4}	9.0×10^{-5}	4.5×10^{-5}	2.2×10^{-5}
$\Delta_{\tau t}$	1.8×10^{-3}	9.0×10^{-4}	4.5×10^{-4}	3.9×10^{-4}	1.1×10^{-3}	2.7×10^{-3}	3.2×10^{-3}

TABLE 5.3. Predator-prey model. Average time step sizes when the SSA (row with label Δ_{st}) or tau-leaping method (row with label $\Delta_{\tau t}$) are used to generate realizations of the predator-prey model.

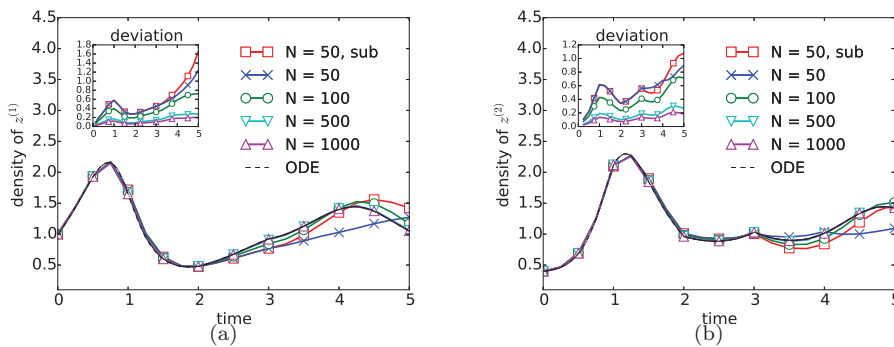


FIGURE 5.5. Predator-prey model. Evolution of the empirical means and the standard deviations (inset plot) of the normalized predator and prey states (densities) under the optimal open loop policy. The curve labeled by “ $N=50$, sub” corresponds to the jump process of size $N=50$ that is controlled by the suboptimal policy u_2 , which becomes the optimal policy for larger N . “ODE” corresponds to the limiting ODE under the optimal policy u_2 .

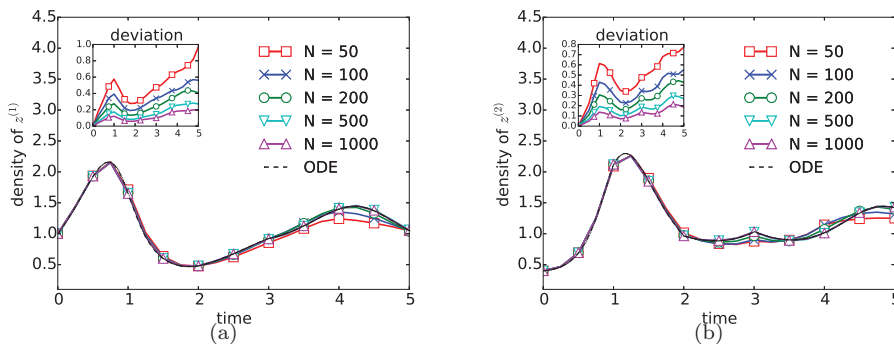


FIGURE 5.6. Predator-prey model. Evolution of empirical means and standard deviations (inset plot) of the normalized predator-prey system under the hybrid control policy.

N	50	100	200	500	1000	2000	4000
N_g	5	5	3	3	3	3	3
M_{ol}	5000	10000	10000	10000	20000	20000	30000
$\min_{1 \leq j \leq 4} \mathcal{S}_j $	4090	8738	12024	11545	23120	26060	40463
$\max_{1 \leq j \leq 4} \mathcal{S}_j $	11420	30572	25784	14587	29369	29597	44513
$9N^2$	22500	90000	360000	2250000	9000000	36000000	144000000

TABLE 5.4. *Predator-prey model with hybrid control.* The row “ $9N^2$ ” shows the estimated state space cardinalities after truncation if a simple cut-off criterion is used. The row “ N_g ” shows the number of the “good” open control policies, and “ M_{ol} ” denotes the number of trajectories generated for each “good” open policy in the calculation of the sets \mathcal{S}_j . The other two rows contain the minimum and maximum numbers of states in the sets \mathcal{S}_j .

ϵ_{near}	N	50	100	200	500	1000	2000	4000
0.0	r_{ol}	13.6%	13.6%	38.1%	66.1%	66.6%	73.2%	74.7%
	time	1.0h	5.3h	5.6h	7.1h	5.0h	5.0h	8.2h
	cost	10.72	9.88	9.58	9.27	9.18	9.13	9.11
0.02	r_{ol}	3.3%	1.1%	0.9%	0.6%	0.3%	0.4%	0.3%
	r_{near}	10.2%	12.0%	36.4%	65.5%	66.3%	72.9%	74.3%
	time	1.1h	5.5h	5.5h	7.0h	5.7h	5.5h	7.2h
	cost	10.60	9.81	9.47	9.25	9.18	9.13	9.11

TABLE 5.5. *Predator-prey model with hybrid control.* The rows “ r_{ol} ” and “ r_{near} ” record the relative frequencies of using an open loop policy or a feedback policy of a nearest neighbor when the hybrid control policy is applied (see Subsection 4.3). The row “time” shows the CPU run time (in hours) needed to compute the optimal hybrid control policy with 20 processors running in parallel for each N .




FIGURE 5.7. *Predator-prey model.* Cost values of the predator-prey model under the optimal open loop (“OL”) control policy, the hybrid control policies with $\epsilon_{near} = 0$ and 0.02, for various values of N . The dotted horizontal line is the optimal cost for the limiting ODE system.

each “good” policy, we generate M_{ol} trajectories once again and collect the accessed states at time t_j in \mathcal{S}_j , $1 \leq j < M$ according to the criterion (4.1) for $\zeta = 3.0$. (Note that \mathcal{S}_0 contains only a single element). The minimum and maximum cardinalities $\min_{1 \leq j \leq M-1} |\mathcal{S}_j|$ and $\max_{1 \leq j \leq M-1} |\mathcal{S}_j|$ of sets \mathcal{S}_j are shown in Table 5.4.

The reader should bear in mind that, if we wanted to compute the optimal feedback

control policy on a globally truncated state space \mathbb{X}_{cut} (see Subsection 4.2), then it would be necessary to include states whose normalized components are within $[0, 3.0] \times [0, 3.0]$ as suggested by the empirical means and standard deviations of the process (see Figure 5.5), which would result in roughly $9N^2$ states in total; even for moderate predator-prey populations, computing the optimal feedback policy on \mathbb{X}_{cut} is therefore extremely costly. Compared to this approach, the adaptive state truncation that gives rise to the sets \mathcal{S}_j is much more efficient in the sense that the overall number of states involved in the computation of the optimal hybrid policy is much smaller; see Table 5.4 and Figure 5.8.

Finally, we compute the optimal hybrid policy using Algorithm 2 and apply it to the predator-prey model in the way explained in Subsection 4.3. The resulting cost values that were estimated based on 50000 independent realizations are shown in Table 5.5, Figure 5.7 and clearly demonstrate the superiority of the hybrid controls over the optimal open loop control policies (in particular, see Table 5.5 for $N = 50, 100, 200$). To explain the observed gain in the numerical speed-up, Table 5.5 also records the relative frequencies r_{ol} of switching to an open loop policy: For $\epsilon_{near} = 0.0$, we observe that the hybrid control frequently switches to the optimal open loop policy, which is an indicator that the sets \mathcal{S}_j are too small as the dynamics often hits an “unknown” state outside \mathcal{S}_j . Yet, for $\epsilon_{near} = 0.02$, we find that r_{ol} decreases significantly which suggests that the sets \mathcal{S}_j contain almost all states that are close to the accessible states under the given control policy. Note moreover that the resulting cost value for $\epsilon_{near} = 0.02$ is slightly improved over the choice $\epsilon_{near} = 0.0$.

Before we conclude, we would like to stress an important observation that the standard deviations of the process are smaller under the hybrid control policy (similarly for the feedback policy) than that under the optimal open loop policy. This effect can be revealed by comparing Figure 5.5 with Figure 5.6 for the same value of N , and it suggests that besides providing smaller costs, both hybrid and feedback control policies have a positive effect on stabilizing the stochastic process.

6. Conclusions and future directions

Due to their wide applicability, Markov Decision Processes have been the subject of intensive research. While the theory is well developed, algorithms for numerically computing optimal controls are restricted to small or moderately sized systems.

The aim of this paper was to analyze optimal control problems for Markov jump process in the large number regime (parameterized by the “particle” number $N \gg 1$), i.e., when the state space is too large to compute the optimal feedback controls using standard algorithms. Based on Kurtz’s limit theorems, we have established convergence results for the value functions of the optimal control problems on finite and infinite time-horizons as $N \rightarrow \infty$. Our results suggest that the optimal open loop control policy for the limiting deterministic system is a good substitute for the controlled Markov jump process, for which the optimal feedback policy may be difficult to compute. Nonetheless, for a given jump process with a possibly large, but finite N , the approximation error induced by replacing the optimal stochastic (feedback) control with the limiting deterministic control is difficult to assess; even for large values of N the stochastic dynamics controlled by a deterministic open loop control policy is not robust under the intrinsic random perturbations, and may hence deviate considerably from the optimal regime. To account for this lack of robustness, we proposed an algorithmic strategy to compute a *hybrid* control policy that is based on a combination of deterministic (open loop) and stochastic (closed loop) controls. The key idea is to truncate the state space adaptively in time, exploiting data gathered from stochastic simulations under near-optimal open




FIGURE 5.8. States selected to construct the hybrid control policy in the predator-prey model for $N = 100$. The value at each grid point $x = (x^{(1)}, x^{(2)})$ counts how many sets \mathcal{S}_j contain the state x , i.e., the value at $x \in \mathbb{X}$ is equal to $\sum_{j=1}^4 \mathbf{1}_{\mathcal{S}_j}(x)$.

loop policies, and then to apply the optimal feedback control policy for all times in which the stochastic realizations resides inside the truncated state space (for all other states, the optimal open loop policy is applied). Both the accuracy and the practicability of the proposed *hybrid* algorithm have been demonstrated numerically with birth-death and predator-prey models.

Before we conclude, it is necessary to mention several related topics which go beyond our current work. Firstly, throughout the article, we have assumed that the cost can be expressed as a function of the normalized density process $z^N(t) = N^{-1}x^N(t)$, which in many cases is the natural variable scaling. In some cases, however, such as complex chemical reaction networks, it might be necessary to consider a more general scaling of the form $z^{(i),N}(t) = N^{-\alpha_i}x^{(i),N}(t)$, $\alpha_i \geq 0$, in which each chemical species comes with its own scaling order. Then, in the limit $N \rightarrow \infty$ it may happen that the limit of $z^{(i),N}(t)$ can be deterministic, stochastic or even hybrid when some of the α_i are equal to zero and others are positive. We emphasize that in these cases determining the correct scaling of the variables is not a trivial task and the convergence analysis is also more involved (see [2,28]). Secondly, besides the large copy-number N , systems in realistic applications may also contain many different species. While our analysis is still valid in this case, it may become computationally challenging to compute the hybrid policy proposed in the current work. One idea to alleviate the difficulty is to first reduce the dimension of the system (in particular when there are both slow and fast reactions or when the quasi-stationary assumption is satisfied), and then utilize the information of the reduced system to design numerical algorithms. Thirdly, it is also interesting to consider the asymptotic analysis of the optimal control problem in the case that the control policy can be switched at any time or when there is uncertainty in the observation of the system's states. We leave these aspects for future work.

Acknowledgement. The authors acknowledge financial support by the Einstein Center of Mathematics (ECMath) and the DFG-CRC 1114 “Scaling Cascades in Com-

plex Systems”.

Appendix A. A technical lemma. The following inequality has been used in the proof of Theorem 3.1.

LEMMA A.1. *Let $\varphi(z) = |z|^\alpha$, where $z \in \mathbb{R}^n$ and $1 < \alpha \leq 2$. We have*

$$0 \leq \varphi(z+w) - \varphi(w) - z \cdot \nabla \varphi(w) \leq \frac{4}{\alpha-1} \varphi\left(\frac{z}{2}\right), \quad \forall z, w \in \mathbb{R}^n. \tag{A.1}$$

Proof. The case $w=0$ can be readily verified. Now assume $w \neq 0$ and consider $z = (z_1, 0, 0, \dots, 0)^T$, $w = (w_1, w')^T$ where $z_1, w_1 \in \mathbb{R}$, $w' \in \mathbb{R}^{n-1}$. Defining $g(r) = r^\alpha$ for $r > 0$, it follows that

$$\begin{aligned} & \varphi(z+w) - \varphi(w) - z \cdot \nabla \varphi(w) \\ &= g\left(\sqrt{(z_1+w_1)^2 + |w'|^2}\right) - g\left(\sqrt{w_1^2 + |w'|^2}\right) - g'\left(\sqrt{w_1^2 + |w'|^2}\right) \frac{w_1 z_1}{\sqrt{w_1^2 + |w'|^2}} \\ &= \int_0^{z_1} \int_0^r \left[g''\left(\sqrt{(s+w_1)^2 + |w'|^2}\right) \frac{(s+w_1)^2}{(s+w_1)^2 + |w'|^2} \right. \\ & \quad \left. + g'\left(\sqrt{(s+w_1)^2 + |w'|^2}\right) \left(\frac{1}{\sqrt{(s+w_1)^2 + |w'|^2}} - \frac{(s+w_1)^2}{((s+w_1)^2 + |w'|^2)^{\frac{3}{2}}} \right) \right] ds dr \\ &= \int_0^{z_1} \int_0^r \left[g''\left(\sqrt{(s+w_1)^2 + |w'|^2}\right) \frac{(s+w_1)^2}{(s+w_1)^2 + |w'|^2} \right. \\ & \quad \left. + \frac{|w'|^2 g'\left(\sqrt{(s+w_1)^2 + |w'|^2}\right)}{((s+w_1)^2 + |w'|^2)^{\frac{3}{2}}} \right] ds dr. \end{aligned}$$

Since $1 < \alpha \leq 2$, we know that $g', g'' \geq 0$, and $\frac{g'(r)}{r} = \frac{g''(r)}{\alpha-1} = \alpha r^{\alpha-2}$ is non-increasing for $r > 0$. We also have the simple inequality $\frac{a+\frac{b}{\alpha-1}}{a+b} \leq \frac{1}{\alpha-1}$, $\forall a, b > 0$. Therefore

$$\begin{aligned} & 0 \leq \varphi(z+w) - \varphi(w) - z \cdot \nabla \varphi(w) \\ &= \int_0^{z_1} \int_0^r g''\left(\sqrt{(s+w_1)^2 + |w'|^2}\right) \frac{(s+w_1)^2 + \frac{1}{\alpha-1}|w'|^2}{(s+w_1)^2 + |w'|^2} ds dr \\ &\leq \frac{1}{\alpha-1} \int_0^{z_1} \int_0^r g''\left(\sqrt{(s+w_1)^2 + |w'|^2}\right) ds dr \\ &\leq \frac{1}{\alpha-1} \int_0^{z_1} \int_0^r g''(|s+w_1|) ds dr \\ &\leq \frac{2}{\alpha-1} \int_0^{|z_1|} \int_0^{\frac{r}{2}} g''(s) ds dr \leq \frac{4}{\alpha-1} g\left(\frac{|z_1|}{2}\right) = \frac{4}{\alpha-1} g\left(\frac{|z|}{2}\right). \end{aligned}$$

For the general case, let A be an $n \times n$ rotation matrix, such that $Az = (z_1, 0, 0, \dots, 0)^T$, $z_1 \in \mathbb{R}$. Then

$$\begin{aligned} & \varphi(z+w) - \varphi(w) - z \cdot \nabla \varphi(w) \\ &= g(|z+w|) - g(|w|) - g'(|w|) \frac{w}{|w|} \cdot z \end{aligned}$$

$$\begin{aligned}
&= g(|Az + Aw|) - g(|Aw|) - g'(|Aw|) \frac{Aw}{|Aw|} \cdot Az \\
&= \varphi(Az + Aw) - \varphi(Aw) - Az \cdot \nabla \varphi(Aw) \\
&\leq \frac{4}{\alpha - 1} g\left(\frac{|Az|}{2}\right) = \frac{4}{\alpha - 1} g\left(\frac{|z|}{2}\right),
\end{aligned}$$

therefore the conclusion also holds for general $z \in \mathbb{R}^n$. \square

REFERENCES

- [1] L.J. Allen, *An Introduction to Stochastic Processes with Application to Biology*, CRC Press, 2011.
- [2] K. Ball, T.G. Kurtz, L. Popovic, and G. Rempala, *Asymptotic analysis of multiscale approximations to reaction networks*, Ann. Appl. Probab., 16:1925–1961, 2006.
- [3] N. Bäuerle, *Asymptotic optimality of tracking policies in stochastic networks*, Ann. Appl. Probab., 10:1065–1083, 2000.
- [4] N. Bäuerle, *Optimal control of queueing networks: an approach via fluid models*, Adv. Appl. Probab., 34:313–328, 2002.
- [5] R.E. Bellman, *Dynamic Programming*, Dover Publications, Incorporated, 2010.
- [6] C.C. Bennett and K. Hauser, *Artificial intelligence framework for simulating clinical decision-making: A Markov decision process approach*, Artif. Intell. Med., 57:9–19, 2013.
- [7] J.L. Bentley, *Multidimensional binary search trees used for associative searching*, Commun. ACM, 18:509–517, 1975.
- [8] D.P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II: Approximate Dynamic Programming*, Athena Scientific, Forth Edition, 2012.
- [9] Y. Cao, D.T. Gillespie, and L.R. Petzold, *Efficient step size selection for the tau-leaping simulation method*, J. Chem. Phys., 124:044109, 2006.
- [10] D. Cappelletti and C. Wiuf, *Elimination of intermediate species in multiscale stochastic reaction networks*, Ann. Appl. Probab., 26:2915–2958, 2016.
- [11] J.G. Dai, *On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models*, Ann. Appl. Probab., 5:49–77, 1995.
- [12] J.G. Dai and S.P. Meyn, *Stability and convergence of moments for multiclass queueing networks via fluid limit models*, IEEE T. Automat. Contr., 40:1889–1904, 1995.
- [13] P. Dupuis, *Explicit solution to a robust queueing control problem*, SIAM J. Control Optim., 42:1854–1875, 2003.
- [14] S. Duwal, S. Winkelmann, C. Schütte, and M. von Kleist, *Optimal treatment strategies in the context of 'treatment for prevention' against HIV-1 in resource-poor settings*, PLoS Comput. Biol., 11:e1004200, 2015.
- [15] E.B. Dynkin, *Markov processes*, Grundlehren der athematischen Wissenschaften, Academic Press Berlin, New York, 1, 1965.
- [16] C. Gardiner, *Stochastic Methods*, Springer, 2010.
- [17] D.T. Gillespie, *A general method for numerically simulating the stochastic time evolution of coupled chemical reactions*, J. Comput. Phys., 22:403–434, 1976.
- [18] D.T. Gillespie, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81:2340–2361, 1977.
- [19] D.T. Gillespie, *Markov Processes: An Introduction for Physical Scientists*, Academic Press, 1991.
- [20] D.T. Gillespie, *Approximate accelerated stochastic simulation of chemically reacting systems*, J. Chem. Phys., 115:1716–1733, 2001.
- [21] D.T. Gillespie, *Stochastic simulation of chemical kinetics*, Annu. Rev. Phys. Chem., 58:35–55, 2007.
- [22] D. Gross, J.F. Shortle, J.M. Thompson, and C.M. Harris, *Fundamentals of Queueing Theory*, Wiley-Interscience, New York, NY, USA, Forth Edition, 2008.
- [23] J.M. Harrison and J.A.V. Mieghem, *Dynamic control of Brownian networks: State space collapse and equivalent workload formulations*, Ann. Appl. Probab., 7:747–771, 1997.
- [24] M. Hauskrecht and H. Fraser, *Planning treatment of ischemic heart disease with partially observable Markov decision processes*, Artif. Intell. Med., 18:221–244, 2000.
- [25] R.A. Howard, *Dynamic Programming and Markov Processes*, The MIT Press, First Edition, 1960.
- [26] Y. Hu and T. Li, *Highly accurate tau-leaping methods with random corrections*, J. Chem. Phys., 130:124109, 2009.
- [27] N.G.V. Kampen, *Stochastic Processes in Physics and Chemistry*, Elsevier, Third Edition, 2007.

- [28] H.W. Kang and T.G. Kurtz, *Separation of time-scales and model reduction for stochastic reaction networks*, Ann. Appl. Probab., 23:529–583, 2013.
- [29] H.-W. Kang, T.G. Kurtz, and L. Popovic, *Central limit theorems and diffusion approximations for multiscale markov chain models*, Ann. Appl. Probab., 24:721–759, 2014.
- [30] D.E. Kirk, *Optimal Control Theory: An Introduction*, Dover Books on Electrical Engineering Series, Dover Publications, 2004.
- [31] T.G. Kurtz, *Solutions of ordinary differential equations as limits of pure jump Markov processes*, J. Appl. Prob., 7:49–58, 1970.
- [32] T.G. Kurtz, *Limit theorems for sequences of jump Markov processes approximating ordinary differential processes*, J. Appl. Prob., 8:344–356, 1971.
- [33] T.G. Kurtz, *Limit theorems and diffusion approximations for density dependent Markov chains*, in Stochastic Systems: Modeling, Identification and Optimization, I, Mathematical Programming Studies, Springer Berlin Heidelberg, 5:67–78, 1976.
- [34] T.G. Kurtz, *Strong approximation theorems for density dependent Markov chains*, Stoch. Proc. Appl., 6:223–240, 1978.
- [35] H.J. Kushner, *Heavy Traffic Analysis of Controlled Queueing and Communication Networks*, Springer, New York, 2001.
- [36] H.J. Kushner and L.F. Martins, *Heavy traffic analysis of a controlled multiclass queueing network via weak convergence methods*, SIAM J. Control Optim., 34:1781–1797, 1996.
- [37] H.J. Kushner and K.M. Ramachandran, *Optimal and approximately optimal control policies for queues in heavy traffic*, SIAM J. Control Optim., 27:1293–1318, 1989.
- [38] S. Lenhart and J.T. Workman, *Optimal Control Applied to Biological Models*, Chapman & Hall/CRC Mathematical and Computational Biology, Chapman & Hall/CRC, 2007.
- [39] C. Maglaras, *Discrete-review policies for scheduling stochastic networks: trajectory tracking and fluid-scale asymptotic optimality*, Ann. Appl. Probab., 10:897–929, 2000.
- [40] A. Mandelbaum, W.A. Massey, and M.I. Reiman, *Strong approximations for Markovian service networks*, Queueing Syst., 30:149–201, 1998.
- [41] A. Mandelbaum and G. Pats, *State-dependent stochastic networks. part I. approximations and applications with continuous diffusion limits*, Ann. Appl. Probab., 8:569–646, 1998.
- [42] L.F. Martins, S.E. Shreve, and H.M. Soner, *Heavy traffic convergence of a controlled, multiclass queueing system*, SIAM J. Control Optim., 34:2133–2171, 1996.
- [43] S.P. Meyn, *Sequencing and routing in multiclass queueing networks part I: Feedback regulation*, SIAM J. Control Optim., 40:741–776, 2001.
- [44] S.P. Meyn, *Control Techniques for Complex Networks*, Cambridge University Press, New York, NY, USA, First Edition, 2007.
- [45] D.M. Mount and S. Arya, *ANN Programming Manual*, 2010.
- [46] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, Springer, Fifth Edition, 2000.
- [47] J. Pahle, *Biochemical simulations: stochastic, approximate stochastic and hybrid approaches*, Brief Bioinform, 10:53–64, 2009.
- [48] G. Pang and M.V. Day, *Fluid limits of optimally controlled queueing networks*, J. Appl. Math. Stoch. Anal., Article ID 68958, 2007.
- [49] P. Pfaffelhuber and L. Popovic, *Scaling limits of spatial compartment models for chemical reaction networks*, Ann. Appl. Probab., 25:3162–3208, 2015.
- [50] A. Piunovskiy and Y. Zhang, *Accuracy of fluid approximations to controlled birth-and-death processes: absorbing case*, Mathematical Methods of Operations Research, 73:159–187, 2011.
- [51] W.B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics)*, Wiley-Interscience, 2007.
- [52] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Inc., New York, NY, USA, First Edition, 1994.
- [53] M. Rathinam, L.R. Petzold, Y. Cao, and D.T. Gillespie, *Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method*, J. Chem. Phys., 119:12784–12794, 2003.
- [54] A.J. Schaefer, M.D. Bailey, S.M. Shechter, and M.S. Roberts, *Modeling Medical Treatment Using Markov Decision Processes*, Springer US, Boston, MA, 593–612, 2004.
- [55] G. Shani, J. Pineau, and R. Kaplow, *A survey of point-based POMDP solvers*, Auton. Agent. Multi Agent Syst., 27:1–51, 2013.
- [56] R.S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 1998.
- [57] M. Čudina and K. Ramanan, *Asymptotically optimal controls for time-inhomogeneous networks*, SIAM J. Control Optim., 49:611–645, 2011.
- [58] J.W. Weibull, *Evolutionary Game Theory*, MIT Press, 1997.
- [59] W. Whitt, *Stochastic-process limits: An introduction to stochastic-process limits and their application to queues*, Springer series in operations research, Springer, New York, Berlin, Paris,

- 2002.
- [60] D.J. Wilkinson, *Stochastic Modelling for Systems Biology*, Chapman & Hall/CRC, 2006.
- [61] S. Winkelmann, C. Schütte, and M. von Kleist, *Markov control processes with rare state observation: Theory and application to treatment scheduling in HIV-1*, Commun. Math. Sci., 12:859–877, 2014.