Splitting Techniques for Rare Event Simulation in Chemical Reaction Networks

by

Michael Snarski

B.Sc., McGill University; Montreal, Canada, 2013

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in The Division of Applied Mathematics at Brown University

PROVIDENCE, RHODE ISLAND

May 2019

This dissertation by Michael Snarski is accepted in its present form
by The Division of Applied Mathematics as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Date_____ _____

Paul Dupuis, Ph.D., Advisor

Recommended to the Graduate Council

Date_____ _____

Jérôme Darbon, Ph.D., Reader

Date_____ _____

Anastasios Matzavinos, Ph.D., Reader

Approved by the Graduate Council

Date_____ _____

Andrew G. Campbell, Dean of the Graduate School

Vita

Michael Snarski graduated from McGill University in 2013 where he received a B. Sc. in Honours Pure Mathematics with First Class Honours. He was supported by a PGS-M grant from NSERC in 2013-2014, an FQRNT Master's scholarship in 2014-2015, and a PGS-D grant from NSERC for 2015-2018. His preprints include [25, 24].

Preface and Acknowledgments

First and foremost, I thank my advisor Paul Dupuis for his guidance throughout my graduate career. He has taken the time to carefully explain the weak convergence approach to the theory of large deviations and its use in the design and analysis of rare event simulation algorithms. He has always taught with great patience and attention to detail and I am grateful for having had the opportunity to learn from him.

Besides my advisor, I would like to thank Jérôme Darbon for his continuous encouragement and his unfailing enthusiasm for mathematics. I have benefitted from his insight into Hamilton-Jacobi-Bellman equations, as well as his practical perspective on the implementation of sophisticated mathematical algorithms. I also extend my deep gratitude to both him and Anastasios Matzavinos for serving on my thesis committee.

I must also thank Cristi Predescu for his exceptional mentorship during my internships at D. E. Shaw Research. He has had a profound impact on my view of applied mathematics, and I hope to continue learning from him over the coming years.

There are many others who have enriched my environment in one way or another. I thank Govind Menon for always being willing to chat about math, Basilis Gidas for his broader perspective on the mathematical community as a whole, and Kavita

Ramanan for starting the Math CoOp. Among my peers, I thank Guo-Jhen Wu, for being an excellent collaborator and an even better friend; my former officemates, Ian Alevy, Michael Burkhart and Sameer Iyer, all of whom were willing to share their daily mathematical puzzles and struggles; and Tingwei Meng, for tolerating my barrage of questions whenever we met.

Last, but certainly not least, I wholeheartedly thank the amazing staff at the Division. Without their work, none of this would be possible. In particular, I thank Rosanna Wertheimer for helping me with all my scholarship issues, Carly Hansen-Decelles for providing the best stationery and attending to any little need, and Elizabeth Fox, for all her organizational help.

Contents

List of Tables

List of Figures

Chapter One

Introduction

1.1 Outline

This thesis is about the splitting technique in rare event simulation and its use in chemical reaction networks. Rare events, when broadly defined, are events which seldom occur but can have a significant impact on the state of a system. They are of importance in many disciplines: physical chemistry, biology, geology, economics and operations research, to name a select few.

In the context of this thesis, what constitutes a "rare" event will always be understood in the sense of *large deviations*, a notion that is precisely defined later in this chapter. The theory of large deviations will be used to study the behavior of the *splitting* technique for rare event simulation. We will then apply the developed techniques on a class of stochastic models known as chemical reaction networks. The particular structure of these reaction networks will allow us to construct splitting schemes which satisfy an appropriate notion of *asymptotic optimality*. Finally, we study the schemes through many numerical examples.

We outline the thesis. In this first chapter, we define some of the mathematical notation used and introduce the basic mathematical results that will be frequently referred to throughout. In Section 1.3 we introduce the *Large Deviations Principle* (LDP) and define the notion of a rare event; in Section 1.4 we introduce Hamilton-Jacobi-Bellman (HJB) equations, as well as the notion of *subsolutions*, which will be used to design the splitting schemes; Section 1.5 describes the relationships between HJB equations, the LDP, and Lyapunov functions.

In Chapter Two, the splitting algorithms are discussed. Section 2.1 introduces the general principle of splitting algorithms and details the particular one we study, known as RESTART. Section 2.2 studies the performance of the RESTART algo-

rithm when the event of interest can occur over long time intervals, while Section 2.3 explores the utility of a "quadratic approximation" around the law of large numbers dynamics, known as a *moderate deviations* scaling.

Chapter Three is dedicated to chemical reaction networks (CRNs). Section 3.1.1 provides the basic background of CRNs; Section 3.2 describes various techniques to identify subsolutions in reaction networks and Section 3.3 uses these techniques to find subsolutions for three classes of problems. These problems are explored from a numerical viewpoint in Chapter Four.

1.2 Preliminaries

The goal of this section is to introduce the mathematical objects used throughout the thesis. An effort is made to keep the notation consistent. We omit the fundamental definitions of probability and refer the reader to an introductory textbook such as [10].

We begin by describing basic notation. We use the following conventions: for any subset B of some topological space, we use $\bar{B}, B^\circ$ and ∂B to denote the closure, interior and boundary of the set, respectively. We use the standard convention for scaling and translating sets. For any set $B \subset \mathbb{R}^d$, constant $c > 0$ and $z \in \mathbb{R}^d$,

$$cB + z = \{cb + z : b \in B\}.$$

We use $\langle \cdot, \cdot \rangle$ to denote the standard Euclidean inner product. For a positive semidefinite matrix A, we use $\| \cdot \|_A^2$ to denote the norm $\langle x, Ax \rangle$. We use $B_\delta(x) \doteq= \{y : |x - y| < \delta\}$ to denote a Euclidean ball of radius $\delta > 0$ around a point $x \in \mathbb{R}^d$.

Dots indicate time derivatives, $\dot{x}(t) = (dx/dt)(t)$, and the symbol ∂_t is used to denote a partial derivative with respect to time when the function has a spatial variable as well. We use D to denote a spatial derivative, often with a subscript to specify the coordinate with respect to which we are differentiating.

We use $\{X^n, n \geqslant 1\}$ to denote a sequence of random variables taking values in a complete separable metric space $\mathcal{X}$, also known as a *Polish space*. We always assume there is an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which the random variables X^n are defined. The σ-algebra $\mathcal{F}$ must contain at least all sets of the form $(X^n)^{-1}(A)$ for Borel sets $A \subset \mathcal{X}$. If $\mathcal{X}$ is a function space and $X^n = \{X^n(t), t \geqslant 0\}$ a stochastic process in either continuous or discrete time, we will usually denote the associated filtration by $\mathcal{F}^n_t$, with $t \in \mathbb{R}_+$ or $t \in \{0, 1, 2, \ldots\}$. In the case of continuous time we assume that the filtration satisfies the usual conditions [29].

For $T > 0$, we use $C([0,T] : \mathbb{R}^d)$ to denote the space of continuous functions defined on $[0,T]$ and taking values in $\mathbb{R}^d$, or simply $C([0,T])$ when the target space is implicitly understood. We always equip $C([0,T] : \mathbb{R}^d)$ with the uniform topology, $\|f - g\|_\infty = \sup_{0 \leqslant t \leqslant T} |f(t) - g(t)|$. If $x \in C([0,T])$ then we use $x(\cdot)$ to denote the object in $C([0,T])$ and $x(t)$ to denote its value in $\mathbb{R}^d$ at a particular point $t \in [0,T]$.

We use $D([0,T] : \mathbb{R}^d)$ to denote the Skorokhod space of "cadlag" or "rcll" functions, that is to say, functions which are **r**ight-**c**ontinuous and have **l**eft-**l**imits. We always use the Skorokhod metric which makes $D([0,T] : \mathbb{R}^d)$ into a Polish space. Note that relativized to $C([0,T] : \mathbb{R}^d)$, the Skorokhod topology coincides with the uniform topology; see [7], p. 124. We use this fact a few times to avoid dealing with the metric on D.

Given an $\mathbb{R}^d$-valued stochastic process $\{X(t), t \geqslant 0\}$ with associated right-

continuous filtration $\{\mathcal{F}_t, t \geqslant 0\}$, for $x \in \mathbb{R}^d$ we use $\mathbb{P}_x$ to denote the probability measure induced by the process starting at x: $\mathbb{P}(X(0) = x) = 1$. Similarly, we use $\mathbb{E}_x$ to denote the associated expectation operator. For a time-homogeneous Markov process, we use $\mathcal{L}$ to denote the infinitesimal generator of the process. That is, for f in a suitable class of functions,

$$(\mathcal{L}f)(x) = \lim_{t \to 0} \frac{\mathbb{E}_x[f(X(t))] - f(x)}{t}.$$

The class of admissible functions f depends on the stochastic process. We provide here the generator of two classes of Markov processes which will be used throughout the thesis: diffusions and pure jump processes.

Example 1.2.1 (Diffusion process)**.** Let $C \in (0, \infty)$, $b : \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma : \mathbb{R}^{d \times k}$ satisfy, for all $x, y \in \mathbb{R}^d$,

$$\|b(x) - b(y)\| + \|\sigma(x) - \sigma(y)\| \leqslant C\|x - y\| \quad \text{and} \quad \|b(x)\| + \|\sigma(x)\| \leqslant C(1 + \|x\|).$$

Let W be a k-dimensional standard Wiener process. Fix any $x_0 \in \mathbb{R}^d$. We define the diffusion process $\{X(t), t \geqslant 0\}$ as the unique strong solution of the stochastic differential equation

$$dX(t) = b(X(t))dt + \sigma(X(t))dW(t), \quad X(0) = x_0.$$

For $f \in C^2(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ with bounded first and second derivatives, the generator $\mathcal{L}f(x)$ is

$$(\mathcal{L}f)(x) = \sum_{i=1}^{d} b_i(x) \frac{\partial f}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^{d} (\sigma\sigma^T)_{i,j}(x) \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

The existence and uniqueness of strong solutions the stochastic differential equation in Example 1.2.1 can be found in Section 5.2 of [49].

Example 1.2.2 (Jump process)**.** Let $\{X(t), t \geqslant 0\}$ be an $\mathcal{X}$-valued Feller pure jump Markov process with finitely many jump directions $v_1, \ldots, v_r$ and associated state-depedent jump rates $\lambda_1(x), \ldots, \lambda_r(x)$. For bounded and continuous $f : \mathcal{X} \to \mathbb{R}$, the generator $\mathcal{L}f(x)$ is

$$(\mathcal{L}f)(x) = \sum_{i=1}^{r} \lambda_i(x)(f(x + v_i) - f(x)).$$

See, for instance, Chapter 3 of [53].

1.3 Large deviations

We state the large deviations principle below, and we provide the rate functions for the two widely used processes in this thesis: pure jump and diffusion processes.

Definition 1.3.1 (Rate function)**.** A function $I : \mathcal{X} \to [0, \infty]$ is called a *rate function* if I has compact sublevel sets, that is, for each $M < \infty$, the set $\{x \in \mathcal{X} : I(x) \leqslant M\}$ is compact.

Recall that a function f defined on $\mathcal{X}$ having closed sublevel sets is equivalent to it being lower semicontinuous, i.e., $\liminf_{n\to\infty} f(x_n) \geqslant f(x)$ for any $x_n \to x$ and $x \in \mathcal{X}$. Moreover, any lower semi-continuous achieves its minimum on a compact set. It follows that a rate function achieves its minimum on any sublevel set.

We now introduce two equivalent formuations of the large deviations principle. The first is in terms of the typical upper and lower bounds [20] that are reminiscent of weak convergence in the Portmanteau theorem ([7], Theorem 2.1), while the second is

called the Laplace principle and is more suitable for the weak convergence approach. We define, for any subset A of $\mathcal{X}$, $\mathcal{X}$,

$$I(A) \doteq \inf_{x \in A} I(x).$$

Definition 1.3.2 (Large deviations principle)**.** A sequence of random variables $\{X^n\}$ taking values in a Polish space $\mathcal{X}$ satisfies a large deviations principle with rate function I if and speed $\{r_n\}$, for any open set $E \subset \mathcal{X}$, X^n satisfies the *large deviations lower bound*

$$\liminf_{n \to \infty} r_n \log \mathbb{P}(X^n \in E) \geqslant -I(E),$$

and, for any closed set $F \subset \mathcal{X}$, X^n satisfies the *large deviations upper bound*

$$\limsup_{n \to \infty} r_n \log \mathbb{P}(X^n \in F) \leqslant -I(F).$$

Remark. Throughout this thesis, we will use a speed of $r_n = 1/n$ unless otherwise indicated.

Definition 1.3.3 (Laplace principle)**.** A sequence of random variables $\{X^n\}$ taking values in a Polish space $\mathcal{X}$ satisfies the Laplace principle with rate function I if, for all bounded continuous functions $h : \mathcal{X} \to \mathbb{R}$, X^n satisfies the Laplace principle lower bound,

$$\liminf_{n \to \infty} \frac{1}{n} \log \mathbb{E}[e^{-nh(X^n)}] \geqslant - \inf_{x \in \mathcal{X}} \{h(x) + I(x)\}$$

and the Laplace principle upper bound,

$$\limsup_{n \to \infty} \frac{1}{n} \log \mathbb{E}[e^{-nh(X^n)}] \leqslant - \inf_{x \in \mathcal{X}} \{h(x) + I(x)\}.$$

We remark that the Laplace principle upper and lower bounds together imply

that the the limit exists and is equal to $\inf_{x\in\mathcal{X}}\{h(x)+I(x)\}$.

An event A is *rare* if it does not contain a minimizer of an associated rate function $I(\cdot)$. Such a definition implies that there is an underlying scaling parameter, denoted either n or ε, and that the probability of the rare event decays exponentially like $\exp(-nI(A))$ or $\exp(-I(A)/\varepsilon)$. All the rare events we will be interested in behave like this.

1.3.1 Law of large numbers

Before computing the rate functions for Examples 1.2.1 and 1.2.2, we need to introduce the appropriate scaling and law of large numbers. The result for small-noise diffusions is straightforward. The law of large numbers for jump processes is a little more complicated and will be used most frequently in this thesis, especially in Chapter 3. We will often refer to this law of large numbers are "Kurtz's theorem", as this is what it is called in the chemical reaction network community. See [51] and [52].

Theorem 1.3.4 (Kurtz)**.** *Let $x_0\in\mathbb{R}^d$, $v_i\in\mathbb{R}^d$ be jump vectors, $\lambda_i:\mathbb{R}^d\to\mathbb{R}_+$, $i=1,\dots,r$ be uniformly bounded Lipschitz functions, and let X^n be the associated pure jump Markov process with the generator in Example 1.2.2. se that $|X^n(0)-x_0|\to 0$ as $n\to\infty$, and let $x(t)$ be the unique solution of*

$$\dot{x}(t)=\sum_{i=1}^{r}v_i\lambda_i(x(t)),\quad x(0)=x_0.$$

Then, for each $T<\infty$, there is $c_1\in(0,\infty)$ and $c_2(\varepsilon)$ with

$$\lim_{\varepsilon\to 0}\frac{c_2(\varepsilon)}{\varepsilon^2}\in(0,\infty),\quad \lim_{\varepsilon\to\infty}\frac{c_2(\varepsilon)}{\varepsilon}=\infty,$$

such that for every $n \in \mathbb{N}$ *and* $\varepsilon > 0$,

$$\mathbb{P}_{x_0}\left(\sup_{0\leqslant t\leqslant T}|X^n(t)-x(t)|\geqslant\varepsilon\right)\leqslant c_1e^{-nc_2(\varepsilon)}.$$

We introduce a transformation which is central to many aspects of large deviations theory. It will allow for the computation of rate functions for Examples 1.2.1 and 1.2.2 when an appropriate law of large numbers scaling is introduced.

Definition 1.3.5. Suppose that $g : \mathbb{R}^d \to (-\infty, \infty]$ is a lower semicontinuous function and assume that $\{y : g(y) < \infty\}$ is non-empty. Its Legendre-Fenchel transform is defined as

$$g^*(p) = \sup_{y\in\mathcal{Y}}\{\langle p, y\rangle - g(y)\}.$$

The asterisk notation will typically denote a Legendre-Fenchel transform of a function unless otherwise noted.

The Legendre-Fenchel transform is a cornerstone of convex analysis, and many excellent treatises exist – see for instance the classical text by Rockafeller [56].

To see how the Legendre-Fenchel transform relates to the theory of large deviations, consider a discrete-time Markov process of the form

$$X^n_{i+1} = X^n_i + \frac{1}{n}v_i(X^n_i), \quad X^n_0 = x_0,$$

where the v_i are i.i.d. vector fields. Let $X^n(\cdot)$ denote the linearly-interpolated continuous time process with $X^n(i/n) = X^n_i$ for $i = 0\ldots, n$, so that $X^n(t)$ is defined for $t \in [0, 1]$.

Definition 1.3.6. We denote the cumulant generating function of the process by

$$H(x,\alpha) \doteq \log \mathbb{E}[\exp(\langle \alpha, v_i(x)\rangle)] \tag{1.3.1}$$

and frequently refer to $H(x,\alpha)$ as the *Hamiltonian* associated to the stochastic process X^n. The *local rate* is the Legendre-Fenchel transform of the Hamiltonian and is denoted by

$$L(x,\beta) \doteq H(x,\cdot)^*(\beta).$$

In general, we will assume the following for all pairs Legendre pairs (H, L). The conditions will be stated precisely in the statements of theorems.

Condition 1.3.7. *For all $x \in \mathbb{R}^d$ there is a unique $b(x)$ such that $L(x, b(x)) = 0$, and that $b(\cdot)$ is Lipschitz continuous. H and L are continuous in (x,α) and (x,β) (respectively) on their domains of finiteness, and both are strictly convex in the second variable.*

We use the Hamiltonian to define the "zero cost" dynamics as follows.

Definition 1.3.8. The trajectory $\varphi \in C([0,T])$ which satisfies

$$\dot{\varphi}(t) = D_p H(\varphi(t), 0)$$

for all $t \in [0,T]$ is called the *law of large numbers* trajectory and is said to satisfy the *noiseless dynamics*.

1.3.2 Examples of LDPs

We now state the LDPs for both jump and diffusion processes. While the conditions we impose are quite broad, they can be relaxed if necessary. In particular, the case of degenerate diffusion matrices can be readily treated. See for instance [35].

Theorem 1.3.9. *Suppose there is $C \in (0,\infty)$ such that $b : \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ satisfy, for all $x, y \in \mathbb{R}^d$,*

$$\|b(x) - b(y)\| + \|\sigma(x) - \sigma(y)\| \le C\|x - y\|, \quad \text{and} \quad \|b(x)\| + \|\sigma(x)\| \le (1 + \|x\|).$$

Let X^ε denote the unique strong solution to the stochastic differential equation (see [49])

$$dX(t) = b(X(t))dt + \sigma(X(t))dW(t), \quad X(0) = x_0, t \in [0,T].$$

Assume that $\sigma(x)$ is invertible for all x. The sequence $\{X^\varepsilon\}$ satisfies an LDP with rate function

$$I_T(\varphi) \doteq \int_0^T L(\varphi(s), \dot\varphi(s))ds,$$

where L is the local rate obtained as the Legendre-Fenchel transform of H, where

$$H(x,\alpha) = \langle \alpha, b(x)\rangle + \frac{1}{2}\|\alpha\|^2_{\sigma(x)}, \quad L(x,\beta) = \|\beta - b(x)\|^2_{\sigma^{-1}(x)}.$$

The LDP for jump processes is next. Once again, the conditions can be significantly relaxed. In particular, rather than the specific $n\lambda_i$ scaling of the jump rates, one can take general sequence λ^n such that $n^{-1}\lambda^n$ converges uniformly to λ as $n \to \infty$. See, e.g., [59]

Theorem 1.3.10. *Let X^n be an $\mathbb{R}^d$-valued jump process as in Example 1.2.2 with*

generator

$$(\mathcal{L}^n f)(x) = \sum_{i=1}^{r} n\lambda_i(x)\left(f\left(x + \frac{v_i}{n}\right) - f(x)\right),$$

where $f : \mathbb{R}^d \to \mathbb{R}$ *is bounded and continuous, the* v_i, $i = 1, \ldots, r$ *are jump directions in* $\mathbb{R}^d$, *and* $\lambda_i : \mathbb{R}^d \to [0, \infty)$ *are state-dependent jump rates. Furthermore, assume that for each* i, $\log \lambda_i(x)$ *is bounded and Lipschitz continuous. Then* $\{X^n, n \geq 1\}$ *satisfies an LDP with rate function*

$$I_T(\varphi) \doteq \int_0^T L(\varphi(s), \dot{\varphi}(s))ds,$$

where L *is the local rate obtained as the Legendre-Fenchel transform of* H, *where*

$$H(x, \alpha) = \sum_{i=1}^{r} \lambda_i(x)(e^{\langle \alpha, v_i \rangle} - 1)$$
$$L(x, \beta) = \inf\left\{\sum_{i=1}^{r} \lambda_i(x)\ell\left(\frac{u_i}{\lambda_i(x)}\right) : \sum_{i=1}^{r} u_i v_i = \beta, u_i \geq 0\right\}$$

and

$$\ell(x) \doteq x \log x - x + 1.$$

1.4 Hamilton-Jacobi-Bellman equations

In this section we compile a few facts about a type of partial differential equation (PDE) known as the *Hamilton-Jacobi-Bellman* (HJB) equation. These facts will be referred to throughout the thesis and can be omitted until called for. The main point is to understand that there exists a relationship between the HJB PDE and a family of underlying calculus of variations problems, and that for the purposes of rare event estimation the exact relationships can be weakened to inequalities.

We first discuss the relation to rare events and introduce some mathematical notation in Section 1.4.1. We then introduce the notion of *subsolutions* in the time-dependent and time-independent cases in Section 1.4.2.

1.4.1 Rare event simulation and setup

In the context of rare event simulation, we will find that the solution to the calculus of variations problem can guide us in "directing" trajectories towards the rare event set of interest. In the case of importance sampling, this is done by altering the dynamics of the underlying process according to the most likely trajectory, that is, the solution to the calculus of variations problem. For splitting, this is done by branching the trajectory when it moves in a "promising" direction.

For both importance sampling and splitting, the information can be encoded through the *value function*, which can be described as a solution to an associated partial differential equation. The value function is typically denoted by V and can be either a function of two arguments (t, x) in the time-dependent setting, or a single spatial variable x in the time-independent setting. In the case of importance sampling, the dynamics are steered according to the spatial gradient $D_x V(t, x)$ or $D_x V(x)$, whereas for splitting the thresholds are determined by level sets of $V(t, x)$ or $V(x)$. In other words, importance sampling requires gradient information, whereas for splitting simply knowing the value function is sufficient. Observe that less regularity is required from the value function in the case of splitting, and in particular one need not insist on "classical-sense" solutions. The boundary conditions that are imposed for the rare event problems we study will make it appropriate to refer to the value function as the "cost-to-go" of the trajectory.

In general, we have in mind a discrete time Markov process $\{X_i^n, i = 0, \ldots, nT\}$ defined by

$$X_{i+1}^n = X_i^n + \frac{1}{n} v_i(X_i^n), \quad i = 0, \ldots, nT, X_0^n = x,$$

where the $v_i(\cdot)$ are i.i.d. vector fields given on some probability space. By $\{X^n(t), t \in [0, T]\}$ we mean the continuous time piecewise-linear interpolation with interpolation points $X^n(i/n) = X_i^n$, i.e.

$$X^n(t) = X_i^n + (X_{i+1}^n - X_i^n)(nt - i), \quad t \in \left[\frac{i}{n}, \frac{i+1}{n}\right], i = 0, 1, \ldots, nT.$$

We recall that the Hamiltonian $H(x, \alpha)$ is determined by the process model as the (state-dependent) cumulant generating function of the increments (see (1.3.1)). We place the following conditions on H.

Condition 1.4.1. *We assume that the cumulant generating function is finite everywhere, i.e.* $H(x, \alpha) < \infty$ *for all* $x \in \mathbb{R}^d$ *and* $\alpha \in \mathbb{R}^d$.

The Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) is

$$\partial_t V(t, x) - H(x, D_x V(t, x)) = 0.$$

Suppose we formally impose the boundary conditions $V(T, x) = 0$ if $x = y$ and $V(T, x) = \infty$ if $x \neq y$ for some fixed $y \in \mathbb{R}^d$. Then, formally, a solution to the time-dependent HJB PDE is

$$V(t, x) = \inf \left\{ \int_t^T L(\varphi(s), \dot{\varphi}(s)) ds : \varphi(t) = x, \varphi(T) = y \right\}.$$

The time-independent or *steady-state* HJB PDE is

$$H(x, -D_x V(x)) = 0,$$

and the analogous variational characterization will be explored in Section 1.5.

The boundary conditions will depend on the rare event problem but are typically taken so that the function V is non-positive when the trajectory escapes a domain of interest. The notions of existence and uniqueness for HJB PDEs are related to the notion of viscosity solutions [12], but for our purposes the definitions of subsolutions outlined below will be sufficient.

1.4.2 Subsolutions

It turns out that for lower bounds on the asymptotic performance (defined later in Definition 2.1.2), it suffices to work with lower bounds on the cost. The content of this section is to make this heuristic precise by defining varions notions of subsolution and to see how they interrelate in the context of rare event simulation. Throughout, we use the convention that an overbar refers to a subsolution, whereas a function with the same letter without an overbar refers to a solution.

We now discuss time-dependent and time-independent subsolutions for the problem of rare-event simulation. The problems of interest can always be phrased as *escape* from a certain domain D, though depending on the problem it may be more natural to speak of entrance into a set.

1.4.2.1 Time-dependent case

Fix some domain $D \subset \mathbb{R}^d$, and let $H(\cdot,\cdot) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ be a Hamiltonian which satisfies Condition 1.4.1. The subsolutions in this section are related to the problem

of estimating

$$\mathbb{P}_x(X^n(T) \notin D).$$

For sets of the form $\{X^n(t) \notin D, \text{ for some } t \in [0,T]\}$, the definition is entirely analogous, with the only change being in the boundary conditions.

Definition 1.4.2. A function $\bar{V} : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ is a *classical sense subsolution* if it is continuously differentiable in both variables and if

$$\partial_t \bar{V}(t,x) - H(x, -D_x \bar{V}(t,x)) \geqslant 0, \quad (t,x) \in [0,T) \times \mathbb{R}^d, \tag{1.4.1}$$

and $\bar{V}(T,x) \leqslant 0$ for $x \in D$.

A weaker notion of subsolution is as follows. Let

$$K_{t,x;T} = \{\varphi \in \mathrm{AC}([0,T] : \mathbb{R}^d) : \varphi(t) = x, \varphi(T) \notin D\}$$

Definition 1.4.3. A function $\bar{V} : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ is a *weak sense subsolution* if it is continuous, bounded from below and

$$\bar{V}(t,x) \leq \inf_{\varphi \in K_{t,x;T}} \{I_T(\varphi) + \bar{V}(T, \varphi(T))\}, \quad (t,x) \in [0,T) \times \mathbb{R}^d, \tag{1.4.2}$$

and $\bar{V}(T,x) \leqslant 0$ for $x \notin D$.

Proposition 1.4.4. *If $\bar{V}$ is a classical sense subsolution as in Definition 1.4.2, then it is a weak sense subsolution as in Definition 1.4.3.*

Proof. Suppose that $\bar{V}$ satisfies the classical definition of a subsolution, as in (1.4.1), and let $\varphi \in K_{t,x;T}$. For any $(x,\beta) \in \mathbb{R}^d \times \mathbb{R}^d$ and $\alpha \in \mathbb{R}^d$,

$$L(x,\beta) \geqslant \langle \alpha, \beta \rangle - H(x,\alpha).$$

Using $\varphi(s), \dot{\varphi}(s)$ and $-D\bar{V}(s,\varphi(s))$ in place of x, β and α, and integrating from t to T, we obtain

$$\begin{aligned}\int_t^T L(\varphi(s),\dot{\varphi}(s))\mathrm{ds} &\geqslant \int_t^T \langle \varphi(s), -D\bar{V}(s,\varphi(s))\rangle - H(\varphi(s), -D\bar{V}(s,\varphi(s)))\mathrm{ds}. \\ &= \int_t^T -\frac{d}{ds}\bar{V}(s,\varphi(s)) + \partial_s \bar{V}(s,\varphi(s)) - H(\varphi(s), -D\bar{V}(s,\varphi(s)))\mathrm{ds} \\ &\geqslant -(\bar{V}(T,\varphi(T)) - \bar{V}(t,\varphi(t))).\end{aligned}$$

In the first equality of the above display we used the chain rule on $\bar{V}$, and in the second inequality we used the classical subsolution property (1.4.1). Using $\varphi(t) = x$ and re-arranging, we obtain the desired inequality for an arbitrary $\varphi \in K_{t,x;T}$, whereupon (1.4.2) follows by infimising over $\varphi \in K_{t,x;T}$. The boundary condition is satisfied by assumption. □

1.4.2.2 Time-independent case

In this section we consider subsolutions which are related to the problem of estimating

$$\mathbb{P}_x(X^n(t) \notin D, \text{ for some } t \in [0,T]).$$

Definition 1.4.5. A function $\bar{V} : \mathbb{R}^d \to \mathbb{R}$ is a *classical sense subsolution* if it is continuously differentiable, bounded below, as well as

$$H(x, -D_x\bar{V}(x)) \leq 0$$

for all $x \in D$ and

$$\bar{V}(x) \leq 0, \quad x \notin D.$$

Just as for the time-independent case, we can formulate a weaker notion of subsolution in terms of a variational problem. Let

$$C_{x;T} = \{\varphi \in \mathrm{AC}([0,T] : \mathbb{R}^d) : \varphi(0) = x, \varphi(t) \in D, \varphi(T) \notin D, t \in (0,T), T < \infty\}$$

Definition 1.4.6. A continuous function $\bar{V} : \mathbb{R}^d \to \mathbb{R}$ is a *weak sense subsolution* if it is bounded from below and

$$\bar{V}(x) \leq \inf_{\varphi \in C_{x;T}} \left\{ I_T(\varphi) + \bar{V}(\varphi(T)) \right\}, \tag{1.4.3}$$

for all $x \in D$, and $\bar{V}(z) \leq 0$ for $z \notin D$.

1.5 Freidlin-Wentzell Quasipotential

In this section we will use the local rate (Definition 1.3.6) to define an object known as the *quasipotential* which was originally introduce by Freidlin and Wentzell [35] in the context of small noise analysis in dynamical systems. The quasipotential plays a central role in this thesis and is used to overcome important obstacles in rare event simulation. Of particular importance are the many different ways of identifying the quasipotential and its relationship to the HJB.

In order to define the quasipotential, it is necessary to assume that the underlying domain contains a unique attracting fixed point. We first define these notions in Section 1.5.1. Relationships with the quasipotential are then categorized as follows: in Section 1.5.2 we express the quasipotential as a scaled limit of stationary distributions; in Section 1.5.3 we express it as a solution to the steady-state HJB PDE, and demonstrate how it can be used to construct an appropriate subsolution; in Section

1.5.3.1 we further develop the PDE characterization in the one-dimensional case. Finally, Section 1.5.4 shows that the quasipotential is always a Lyapunov function for the law of large numbers dynamics.

1.5.1 Definition of the quasipotential

We define a notion of stability which will be used throughout the thesis, and subsequently we define the quasipotential relative to a stable point.

Definition 1.5.1. Let $b : \mathbb{R}^d \to \mathbb{R}^d$ be a Lipschitz continuous function, and suppose that c is a fixed point or *steady state* of b, i.e. $b(c) = 0$. Let B be an open set containing c. We say that c is *stable* on B if, for any $\varepsilon > 0$ and $x \in B$, there is $T < \infty$ such that $\varphi(0) = x$ and $\dot{\varphi}(t) = b(\varphi(t))$ imply $|\varphi(t) - c| < \varepsilon$ for all $t \geqslant T$. The *domain of attraction* of c is the largest such open set B. We will say that the domain B is *attracted to* c.

We remark that this notion of stability is sometimes referred to as *local* stability. When we wish to refer to *global* stability in this thesis, we will do so explicitly.

We now arrive at the definition of the quasipotential.

Definition 1.5.2. Let $L(x, \beta)$ be the local rate of some stochastic process (see Definition 1.3.6). The quasipotential is the function $Q : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty]$ defined for $y, x \in \mathbb{R}^d$ by

$$Q(y, x) \doteq \inf_{\varphi(0)=y, \varphi(T)=x, T<\infty} \left\{ \int_0^T L(\varphi(s), \dot{\varphi}(s)) ds \right\}. \tag{1.5.1}$$

where the infimum is over all absolutely continuous trajectories. Let $c \in \mathbb{R}^d$ be a

stable steady state for the noiseless dynamics (Definition 1.3.8). We call

$$Q(x) \doteq Q(c, x)$$

the quasipotential relative to c.

1.5.2 Quasipotential and stationary distributions

The only purpose of this section is to explicitly separate out a well-known result of Freidlin and Wentzell's which relates a specific scaling of a stationary distribution to the quasipotential. The statement of the theorem is based on the statement in [35], Theorem 4.3 in Chapter 4. It is stated there for diffusion processes, but holds more generally with minor modifications to the proof.

Theorem 1.5.3. *Let c be the unique stable equilibrium of the noiseless dynamics $\dot{\varphi}(t) = D_p H(\varphi(t), 0)$, and assume that c is globally attracting, i.e., the domain of attraction is all of $\mathbb{R}^d$. Let $\{X^n\}$ be a stochastic process with local rate $L(x, \beta)$, the Legendre-Fenchel transform of $H(x, \alpha)$ (see Definition 1.3.6), both satisfying Condition 1.3.7, and suppose that for each $n \geq 1$, $\{X^n(t), t \geq 0\}$ has a unique invariant measure π^n. We have*

$$\lim_{n\to\infty} \frac{1}{n} \log \pi^n(A) = -\inf_{x\in A} Q(c, A) \tag{1.5.2}$$

for any open set $A \subset \mathbb{R}^d$ with compact boundary ∂A and which satisfies $\partial(\bar{A}) = \partial A$.

1.5.3 Quasipotential and HJB

In this section we focus on the most relevant characterization of the quasipotential for our purposes, that is, through the steady state Hamilton-Jacobi-Bellman equation. Throughout the section we assume that we are given a sequence of stochastic processes $\{X^n\}$ which satisfies a large deviations principle with speed $r_n = 1/n$ and rate function $I_T(\varphi)$. We let D be a domain containing a steady state c, and we assume that all of D is attracted to c.

The following theorem is the main result of the section; see Theorem 4.3 in Chapter 5 of [35].

Theorem 1.5.4. *Let $H(x,\alpha) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ be a Hamiltonian defined as a cumulant generating function, as in Definition 1.3.6, and suppose $H(x,\cdot)$ is strictly convex and continuously differentiable for all $x \in \mathbb{R}^d$. Let $U : D \to [0,\infty)$ be a continuously differentiable function which is continuous on $\bar{D}$, with $U(0) = 0$, $U(x) > 0$ and $DU(x) \neq 0$ for $x \neq 0$, and which satisfies the HJB PDE*

$$H(x, DU(x)) = 0,$$

Then $U(x) = Q(c,x)$ on $\{U(\cdot) \leqslant \min_{y\in\partial D} U(y)\}$.

Proof. Fix some x in $\{U \leqslant \min_{y\in\partial D} U(y)\}$. Let φ be any absolutely continuous trajectory satisfying $\varphi(0) = 0$, $\varphi(T) = x$ and $\varphi(t) \in D$ for all $t \in [0,T]$. As in the proof of Proposition 1.4.4, for any $(x,\beta) \in \mathbb{R}^d \times \mathbb{R}^d$ and $\alpha \in \mathbb{R}^d$,

$$L(x,\beta) \geqslant \langle \alpha, \beta\rangle - H(x,\alpha).$$

Using $\varphi(s), \dot{\varphi}(s)$ and $DU(\varphi(s))$ in place of x, β and α, and integrating from t to T,

we obtain

$$\begin{aligned}\int_t^T L(\varphi(s),\dot{\varphi}(s))\mathrm{ds} &\geqslant \int_t^T \langle \varphi(s), DU(\varphi(s))\rangle - H(\varphi(s), DU(\varphi(s)))\mathrm{ds}. \\ &= \int_t^T \frac{d}{ds}U(\varphi(s)) - H(\varphi(s), DU(\varphi(s)))\mathrm{ds} \\ &= U(\varphi(T)) - U(\varphi(0)) \\ &= U(x).\end{aligned}$$

Thus we find $I_T(\varphi) \geqslant U(x)$. Since $U(x) \leqslant \min_{y\in\partial D} U(y)$, there is no advantage in taking curves φ which pass through the boundary ∂D, and so $Q(x) \geqslant U(x)$.

To complete the proof we must show that $Q(x) = U(x)$ for $x \in \{U(x) \leq \min_{y\in\partial D} U(y)\}$. To do this, we will construct a trajectory φ which achieves the minimum in the variational definition of Q, equation (1.5.1). This is a somewhat delicate point, since when starting at a rest point it can take an "infinite" amount of time to leave it with optimal cost. To account for this, we first note that an equivalent definition of the quasipotential is

$$Q(c,x) = \inf\left\{\int_{T_1}^{T_2} L(\varphi(s),\dot{\varphi}(s))ds : \varphi(T_1) = c, \varphi(T_2) = x, -\infty \leq T_1 < T_2 < \infty\right\}$$

Let φ be a solution to

$$\dot{\varphi}(t) = D_\alpha H(\varphi(t), DU(\varphi(t)))$$

with the terminal condition $\varphi(T_2) = x$. A solution exists so long as $\varphi(t) \in D$ since H is continuous in (x,α) (see for instance [41], Chapter 1). By definition of φ, U

and the Legendre-Fenchel transformation, we have

$$\begin{aligned}\frac{d}{dt}U(\varphi(t)) &= \langle \dot{\varphi}(t), DU(\varphi(t))\rangle \\ &= L(\varphi(t), \dot{\varphi(t)}) + H(\varphi(t), DU(\varphi(t)) \\ &= L(\varphi(t), \dot{\varphi(t)}),\end{aligned}$$

where the last equality follows since $H(\varphi(t), DU(\varphi(t))) = 0$ so long as $\varphi(t) \in D$. The local rate must be positive, since $L(x, \beta) = 0$ if and only if $\beta = D_\alpha H(x, 0)$. In particular, U decreases for decreasing t, and since $U(x) \leq \min_{y\in\partial D}$ it must be that $\varphi(t) \leq \min_{y\in\partial D}$ for all $t \in [T_1, T_2]$, and hence $\varphi(t) \in D$ for all such t. More can be said: if it were not the case that $U(\varphi(t)) = c$ for some $T_1 < T_2$ (allowing for the possibility that $T_1 = -\infty$), then $U(\varphi(t))$ stays outside a neighbourhood of $\{c\}$ and we have $(d/dt)U(\varphi(t)) = L(\varphi(t), \dot{\varphi}(t)) \geq \text{const} > 0$ for some constant. This would contradict the finiteness of $U(\varphi(t))$ as $t \to -\infty$.

It remains only to verify that the trajectory φ achieves the desired minimum. Indeed, we have

$$\begin{aligned}\int_{T_1}^{T_2} L(\varphi(t), \dot{\varphi}(t))dt &= \int_{T_1}^{T_2} \frac{d}{dt}U(\varphi(t))dt \\ &= U(\varphi(T_2)) - U(\varphi(T_1)) = U(x).\end{aligned}$$

Such a trajectory can be realized as a trajectory on $[0, \infty)$, and so the Theorem is proved. □

The quasipotential can be used to construct a time-independent subsolution in the sense of Definition 1.4.6 as follows.

Lemma 1.5.5. *For any constant $c \leq \inf_{x\notin\mathcal{D}} Q(x)$, $\bar{V}(x) = -Q(x)+c$ is a subsolution*

in the sense of Definition 1.4.6.

Proof. First consider the case $x \notin \mathcal{D}$. Note that the infimum of Q on the boundary of $\mathcal{D}$ agrees with its infimum on the complement of $\mathcal{D}$, since the positivity of $L(x, \beta)$ ensures that any point in $\mathcal{D}^c$ cannot have lower cost than a point in $\mathcal{D}$. If $c \leq \inf_{x \notin \mathcal{D}} Q(x)$, then for $x \notin \mathcal{D}$, $c - Q(x) \leq 0$ and therefore $c - Q(x)$ satisfies the boundary condition.

For $x \in \mathcal{D}$, observe that Definition 1.4.6 for the choice $\bar{V}(x) = -Q(x) + c$ is equivalent to

$$-Q(x) \leq \inf_{\varphi(0)=x;\varphi(T)\notin\mathcal{D};\varphi(s)\in\mathcal{D},s\in(0,T),T<\infty} \{I_T(\varphi) - Q(\varphi(T))\}.$$

If we temporarily assume the infimum is achieved and write $Q(\varphi(T)) \leq Q(x)+I_T(\varphi)$, we can heuristically interpret this as saying that the cost to exit starting at 0 must be less than the cost to go from 0 to x and then exit starting from x.

To make this heuristic precise, we note that $Q(x, y) \leq I_T(\varphi)$ whenever $\varphi(0) = x, \varphi(T) = y$. Let φ be any function appearing in the characterization of the subsolution property, and let $y = \varphi(T) \notin \mathcal{D}$. It follows from the definition of Q and a standard dynamic programming argument that for any $x, y, z \in \mathbb{R}^d$

$$Q(z, y) \leq Q(z, x) + Q(x, y). \tag{1.5.3}$$

Rearranging gives

$$-Q(z, x) + c \leq -Q(z, y) + c + Q(x, y).$$

Letting $z = 0$ and using $Q(x, y) \leq I_T(\varphi)$ when $\varphi(0) = x, \varphi(T) = y$, we have

$$\bar{V}(x) \leq \bar{V}(y) + I_T(\varphi).$$

Since $y \notin \mathcal{D}$ implies $\bar{V}(y) \leq 0$ it follows that $\bar{V}(x) \leq I_T(\varphi)$. Since φ is any function appearing in the definition of a subsolution, (1.4.3) holds, i.e., $\bar{V}$ is a subsolution. □

1.5.3.1 One-dimensional systems

Owing to Theorem 1.5.4, it is possible to identify the quasipotential on an appropriate neighbourhood of a stable steady state by solving a partial differential equation which depends on the quasipotential only through its gradient. In the one-dimensional case any function has a potential and it is straightforward to identify the quasipotential having only gradient information. This section fleshes out the details of such a construction.

Remark. The results of this section will only be used in Chapter 3, where we denote the second variable of the Hamiltonian by p rather than α. We will also use p instead of α in this section.

Corollary 1.5.6. *Consider a one-dimensional domain $D = (a, b) \subset \mathbb{R}$ which contains a single steady state $c \in D$ and assume c is stable. Let $p(x)$ be a continuous increasing function with its unique zero at $x = c$ and which satisfies $H(x, p(x)) = 0$ for all $x \in D$. Then*

$$U(x) = \int_c^x p(y)dy \tag{1.5.4}$$

satisfies the conditions of Theorem 1.5.4 and $U(x) = Q(c, x)$ for all $x \in D$.

Proof. First we verify the conditions of Theorem 1.5.4. Since $p(x)$ is continuous on

D, $U(x)$ is continuously differentiable on D and continuous on the closure of D. From (3.2.6), we have $DU(x) = p(x)$, whence $H(x, DU(x)) = 0$ for all $x \in D$, as well as $U(c) = 0$. Since p is increasing and has its unique zero at $x = c$, $p(x)$ be positive for $x > c$ and negative for $x < c$. Thus, for $x \in D \setminus \{c\}$, $U(x)$ is positive and $DU(x) \neq 0$. We must also have $DU(c) = p(c) = 0$. Thus U satisfies all the conditions of Theorem 1.5.4.

It only remains to show that in fact $U(x) = Q(c, x)$ for all $x \in D$, rather than only for $x \in \{x : U(x) \leq \min_{y \in \partial D} U(y)\}$. This constraint is used in the proof of Theorem 1.5.4 while establishing the inequality $Q(c, x) \geq U(x)$ to restrict the infimization problem in the quasipotential to curves which lie entirely within D. In the present one-dimensional case it is clear this argument extends to all $x \in D$. Indeed, suppose that $D = (a, b)$, $U(a) < U(b)$, and $x > c$. Any trajectory $\{\varphi(t), [0, T]\}$ which leads from c to x at time T can be assumed without loss of generality to satisfy $\varphi(t) \geq c$ for all $t \in [0, T]$. Indeed, let $\tau = \sup\{t > 0 : \varphi(t) \leq c\}$. We must have $\tau < T$ since $\varphi(T) = x$. Setting $\phi(t) = \varphi(t + \tau)$ produces a path ϕ with lower cost which satisfies $\phi(T - \tau) = x$ and $\phi(0) = c$. A similar argument can be applied to the case $x < c$. This establishes that (3.2.6) defines the quasipotential on the whole domain. □

Next we show that such a function p always exists.

Lemma 1.5.7. *Let $H(x, p) : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$ be continuously differentiable, strictly convex in p for each fixed x, and which satisfies $H(x, 0) = 0$ for all $x \in \mathbb{R}_=+$. Let $c > 0$ be a stable steady state and let D denote its domain of attraction. Then there is a continuous function $p : D \to \mathbb{R}$ which satisfies the conditions of Corollary 1.5.6.*

Proof. Let B denote the domain of attraction of the stable steady state c. We first show that for any $x \in B \setminus \{c\}$ there is a unique $p(x) \neq 0$ such that $H(x, p(x)) = 0$.

We then prove that $\{p(x) : x \in B\}$ is a continuous curve which passes through $(c, 0)$.

Let $x \in B\backslash\{c\}$ and suppose without loss of generality that $x < c$; the case $x > c$ is similar. Since any point in B is attracted to c we must have $D_pH(x, 0) = \Phi(x) > 0$. Since H is continuously differentiable in (x, p), there must be $p' < 0$ such that $H(x, p') < 0$. For each fixed $x \in \mathbb{R}_{>0}$ the function $p \to H(x, p)$ is strictly convex, so we must have $H(x, p) \to \infty$ as $p \to \pm\infty$. By the intermediate value theorem applied to some $p'' > p'$ for which $H(x, p'') > 0$, there is $p(x) < p'$ for which $H(x, p(x)) = 0$. Moreover, the convexity of $H(x, \cdot)$ implies that there are at most two roots to $H(x, p) = 0$. Since $H(x, 0) = 0$ for all $x \in \mathbb{R}_{>0}$, $p(x)$ is the only other root and is therefore the unique p which satisfies $p \neq 0$ and $H(x, p) = 0$.

Now set $p(c) = 0$. It remains to show that $p(x)$ is continuous. We first establish continuity on $D\backslash\{c\}$. We claim that for $x \in D\backslash\{c\}$, we must have $D_pH(x, p(x)) \neq 0$. Since $H(x, \cdot)$ is strictly convex, $D_pH(x, p(x)) = 0$ would imply $p(x)$ is the global minimum of H, and since $p(x) \neq 0$ we must have $H(x, 0) > 0$, a contradiction to the assumptions on H. Fix any particular $x' \in D \setminus \{c\}$ and set $p' = p(x')$. Then $D_pH(x, p(x)) \neq 0$, and by continuity of D_pH, the inequality holds on a neighbourhood of (x', p'). By the implicit function theorem, there is a neighbourhood E of x' and a continuous function $g : E\mathbb{R}$ such that $g(x') = p'$ and $H(x, g(x)) = 0$ for all $x \in E$. The uniqueness of $p(x)$ established above implies that in fact $g(x) = p(x)$ and hence that p is continuous on E. Since $x' \in D \setminus \{c\}$ was arbitrary, this establishes the uniqueness of p on $D \setminus \{c\}$.

Finally, since $D_pH(c, 0) = 0$ and $H(c, 0) = 0$, the strict convexity of $H(c, \cdot)$ implies that $p = 0$ is the unique global minimum of $H(c, \cdot)$ and that there exist no other roots. Consequently, we must have $p(x) \to 0$ whenever $x \to c$, since continuity of H implies $0 = \lim_{x\to c} H(x, p(x)) = H(c, p(c))$.

□

1.5.4 Relation to Lyapunov functions

Since the quasipotential provides a measure of the "cost" to move away from a stable steady state, it is natural to ask whether it can serve as a Lyapunov function for the noiseless dynamics (Definition 1.3.8). The answer is affirmative, and this fact is implicit in many Freidlin-Wentzell style arguments. In this section we provide some of the details. A much deeper (and related) result is that relative entropy is a Lyapunov function for the Kolmogorov forward equation associated to the prelimit Markov process, and the fact that the quasipotential is a Lyapunov function can be seen as a "double limit" version of this result, where both the number of particles and the time horizon are sent to infinity; see [23, 8] for further details.

Definition 1.5.8. A (strict) *Lyapunov function* for the system $\dot{x} = b(x)$ is a function $U : \mathbb{R}^d_+ \to [0, \infty)$ such that $U(x) = 0$ if $x = c$ and $U(x) > 0$ otherwise, which satisfies

$$\frac{d}{dt}U(x(t)) < 0$$

for $x(t) \neq c$.

It is straightforward to show that if a steady state c admits a strict Lyapunov function, then it is stable in the sense of Definition 1.5.1.

Proposition 1.5.9. *Consider a Hamiltonian H, and suppose that c is a stable steady state for the noiseless dynamics $\dot{x} = D_\alpha H(x, 0)$. Let D denote the domain of attraction of c, so that in particular c is the unique steady state in D. Let $x(t)$ satisfy the noiseless dynamics $\dot{x} = D_\alpha H(x, 0)$ with $x(0) \in D$, and let $Q(x) = Q(c, x)$ denote the*

quasipotential relative to c. Then

$$\frac{d}{dt}Q(x) < 0, \quad x \neq c$$

and Q is a Lyapunov function for the system $\dot{x} = D_\alpha H(x, 0)$.

Proof. First, note that $Q(c) = 0$ and $Q(x) > 0$ for all $x \neq c$, since there is a unique $b(x)$ such that $L(x, b(x)) = 0$ and since c is assumed to be a stable steady state, we have $b(c) = 0$.

It remains only to show that $Q(x(t+h)) - Q(x(t)) \leqslant 0$ for all $t \geqslant 0$ and $h > 0$. Let $\varepsilon > 0$ be arbitrary and let $T < \infty$ and $\varphi \in \mathrm{AC}([0, T])$ be such that $\varphi(0) = 0$, $\varphi(T) = x(t)$, and $I_T(\varphi) < Q(x(t)) + \varepsilon$. Define a new trajectory ψ by concatenating it with φ on $[0, T]$ and with the trajectory $\{x(s) : s \in [t, t+h]\}$ on $(T, T+h]$. Since $x(\cdot)$ follows the noiseless dynamics it has zero cost and $I_T(\psi) = I_T(\varphi)$. Since $\psi(0) = 0$ and $\psi(T + h) = x(t + h)$, it follows that $Q(x(t + h)) \leqslant Q(x(t)) + \varepsilon$. Since $\varepsilon > 0$ was arbitrary, $Q(x(t + h)) \leqslant Q(x(t))$.

To rule out the case of equality, note that if there are $t \geqslant 0$ and $h > 0$ such that $Q(x(t + h)) = Q(x(t))$, then by the mean value theorem there is $\xi \in (t, t + h)$ such that $DQ(x(\xi)) \cdot \dot{x}(\xi)h = 0$. Since c is the unique stable steady state in D, $\dot{x}(\xi) \neq 0$ and $DQ(x(\xi)) \neq 0$ unless $x(\xi) = c$. Thus equality cannot hold unless $x = c$. The proof is complete. □

Chapter Two

Splitting

2.1 Introduction

This chapter studies an algorithm for the estimation of rare event probabilities of the form

$$\mathbb{P}(X^n(t) \notin D, \text{ for some } t \in [0,T]) \quad \text{or} \quad \mathbb{P}(X^n(T) \notin D), \quad n \geq 1 \tag{2.1.1}$$

where $D \subset \mathbb{R}^d$ is an open set satisfying appropriate conditions, and $\{X^n\}$ is a Markov process given by

$$X^n_{i+1} = X^n_i + \frac{1}{n} v_i(X^n_i), \quad X^n_0 = x_0, n \geq 1. \tag{2.1.2}$$

We will place conditions on the i.i.d. vector fields $v_i(\cdot)$ which ensure $\{X^n\}$ satisfies some form of large or moderate deviations principle with a rate function I and speed $r(n)$.

We will largely call these *escape* problems and the probabilities in (2.1.1) the *escape probabilities*. Of course, one could equivalently consider estimating the "hitting time" of the complement of the set D.

The outline of the Chapter is as follows. We first discuss in Section 2.1.1 the general computational obstacle that rare event simulation techniques seek to overcome and in Section 2.1.2 we describe the general idea of *splitting*. In Section 2.1.3 we define a specific variant of splitting known as RESTART; Section 2.1.5 defines notions of performance and asymptotic optimality; Section 2.1.4 relates these notions of performance to the subsolution property from Section 1.4.2 in the previous chapter. Finally, Section 2.1.6 describes upper bounds on the second moment of the RESTART algorithm which will be used in the later sections of this chapter to establish asymptotic performance bounds on RESTART.

Section 2.2 analyzes the use of time-independent subsolutions for the problem of escape over a long time interval. When a uniform large deviations principle is available, the process can be allowed to escape a domain D over increasingly long time intervals $[0, T(n)]$, where $T(n)$ grows polynomially with n. We establish conditions under which notions of asymptotic optimality from Section 2.1.5 hold.

When a large deviations principle is not available, or the subsolutions are too difficult to construct, one can instead resort to moderate deviations for events which are not "too rare". The moderate deviations scaling amounts to a quadratic expansion of the cost around the law of large numbers dynamics. In Section 2.3 we establish optimality results for the splitting algorithm in the moderate deviations framework.

2.1.1 The problem of estimating rare events

It is insightful to start with the question of whether or not a brute force method might be "good enough" for estimating probabilities of rare events. In the present context, the de facto brute force method is Monte Carlo simulation: one generates independent trials of an event and computes the success rate. Such a method, which we call *standard Monte Carlo* is hopeless for estimating rare events when the LDP scaling parameter n becomes sufficiently large. This is due to the rapid exponential decay of the probabilities and the comparatively weak rate decay of the variance of the average, which goes only inversely proportional to n.

To illustrate, suppose we are interested in estimating $p_A^n \doteq \mathbb{P}_{x_0}(X^n \in A)$ for some $A \subset D([0,T])$. Suppose $X_1^n, \ldots, X_K^n$ are independent samples, all with the

same distribution as X^n, and consider the standard Monte Carlo estimator

$$\hat{\theta}^n_{\mathrm{MC}} \doteq \frac{1}{K}\sum_{i=1}^{K} 1_A(X^n_i).$$

Assume that $\{X^n\}$ satisfies a large deviations principle with rate function $I(\cdot)$. If $I(A) > 0$, then $\mathbb{P}(X^n \in A) \approx e^{-nI(A)}$ as $n \to \infty$. The variance of the estimator is

$$\mathrm{Var}(\hat{\theta}^n_{\mathrm{MC}}) = \frac{1}{K^2}\sum_{i=1}^{K} \mathrm{Var}(1_A(X^n_i)) = \frac{p^n_A(1-p^n_A)}{K}.$$

Hence, the relative error is

$$\frac{\sqrt{\mathrm{Var}(\hat{\theta}^n_{\mathrm{MC}})}}{p^n_A} \approx \frac{1}{\sqrt{Kp^n_A}}.$$

In order to obtain a bounded relative error, the number K of samples must grow exponentially as n increases, which is infeasible in practice.

The relatively high variance of the estimator $\hat{\theta}_{\mathrm{MC}}$ is due to the 0 or 1 nature of its summands $1_A(\cdot)$. An ideal estimator would give the correct estimate p^n_A every time, and a "good" estimator would return values close to the estimate p^n_A as frequently as possible. For $n \geq 1$ large the summands of the Monte Carlo estimator return values of 0 all but an exponentially small number of times, and the only non-zero values returned are 1.

2.1.2 Multi-level splitting

The basic premise of splitting algorithms in the context of estimating event probabilities associated to a process $\{X_i, i \geq 0\}$ is to evolve *particles* according to the

law of $\{X_i\}$ and split the weight of a particle at times where it seems more likely the event of interest will occur. As mentioned previously, we will focus on escape probabilities for which the trajectory is more likely to be within a specific region over a time interval or at a specific time, and we wish to estimate the probability that it deviates from this typical behavior.

A splitting algorithm is by nature a branching process; in splitting, the branching locations are called *thresholds*. The standard implementation of the algorithm specifies the thresholds as a sequence of sets $C_0 \subset \cdots \subset C_J$ which can be subsets of $\mathbb{R}^d$ or $[0,T] \times \mathbb{R}^d$. The temporal information is relevant when escape must occur at a specific time T rather than any time $t \in [0,T]$. The algorithm also specifies a *splitting rate*, which we take to be fixed across the thresholds and which will be denoted by R.

In the standard version of multi-level splitting, a single particle is started at an initial location $x_0 \in C_J \setminus C_{J-1}$ and evolved until either the terminal time T is reached or, if allowed, escape has occurred prior to time T. If a particle ends up in C_0 then it has escaped, though escape may be possible at a higher threshold. When a particle enters a set C_j, $0 \leq j < J$ for the first time, it produces $R-1$ offspring which are then evolved independently along with the original particle. The R resulting trajectories are identically distributed and conditionally independent given the splitting time and location. Figure 2.1 shows a sample of the resulting branching process.

While splitting can certainly decrease the variance of the underlying estimator, the design of the thresholds is critical to good performance. There are at least two ways the method can exhibit poor behavior. Perhaps the most important is that the branching process could be supercritical, i.e., R is too large relative to the spacing of the C_j's and exponentially many particles are generated. The second is that there

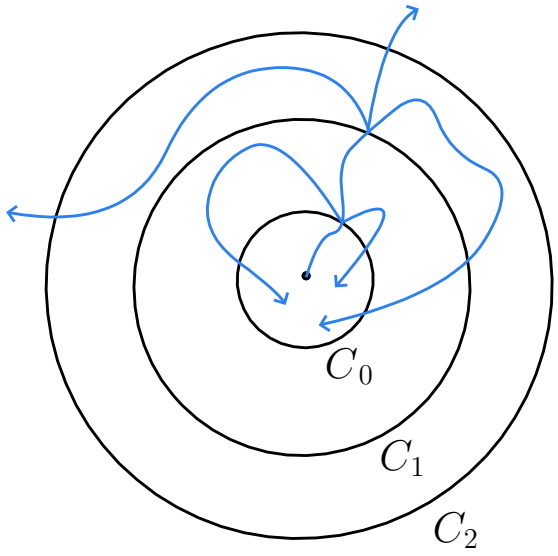

Figure 2.1: A sample trajectory of a process $\{X_i, i \geq 0\}$ and two splitting events with a splitting rate $R = 3$.

could be high variance of the weights, which could be due to particles exiting at different regions of D or too many simulations which produce no particles that end up exiting D. Both of these issues will be resolved by an appropriately constructed *importance function*, which is a function that specifies the location of the thresholds, and this importance function will end up being based on a subsolution; see Section 2.1.4.

There is another important pitfall which is more computational in nature. In the context of escape problems, the probability of escape decays exponentially with the LDP scaling parameter n, and the limiting trajectories $X^\infty(t) = \lim_{n\to\infty} X^n(t)$ never escape. For large (but finite) n, particles typically do not escape and much computational effort is spent simulating trajectories which end up in a higher threshold than the one they were created in. These particles ultimately contribute only a zero

value to the final estimate. The RESTART algorithm is a variant of the splitting algorithm which limits this computational cost.

2.1.3 RESTART Splitting algorithm

The computational issue raised in the context of escaping a domain D is that ordinary splitting methods spend time evolving trajectories which contribute only a value of 0 to the final estimate. The same contribution could be obtained by simply stopping the simulation of a trajectory if it does not look "promising". The idea of the RESTART algorithm (Repetitive Simulation Trials After Reaching Thresholds) is to discard particles which go above the threshold in which they were created. In order to ensure unbiasedness of the estimate, particles are split *every time* they cross into a higher threshold. The proof of unbiasedness is not trivial and can be found in [19].

The RESTART algorithm is formulated in terms of an importance function which determines the thresholds C_j, and the importance function may or may not be time-dependent. The algorithm will be formulated in terms of a time-independent importance function $U(x)$ and the definitions can be extended to the time-dependent case by viewing time as a spatial variable, i.e. $x = (t, y) \in [0, T] \times \mathbb{R}^d \subset \mathbb{R}^{d+1}$. In both the time-dependent and time-independent cases, the importance functions will be based on an associated HJB PDE. For the remainder of the section, we focus on the time-independent case.

An *importance function* is a bounded continuous function $U : \mathbb{R}^d \to \mathbb{R}$. Only the relative values of $U(x)$ at different points matter, and so we assume for simplicity that $U(x) \geq 0$. To guarantee this assumption, we will later see that $U(x)$ will be an appropriately chosen additive shift and scaling of the subsolution $\bar{V}(x)$.

The splitting thresholds are based on closed sets C_j, for $0 \leq j \leq J-1 \doteq \lceil U(0)/\Delta \rceil - 1$, which are defined by

$$C_j \doteq \{x \in D : U(x) \leq j\Delta\}, \tag{2.1.3}$$

and $C_J \doteq D$, where $\Delta \doteq \log(R)$. We define the function $\sigma(x) \doteq j$ when $x \in C_j \backslash C_{j-1}$. We define a piecewise constant approximation to the importance function U by

$$\bar{U}(x) \doteq j\Delta, \quad \text{for } x \in C_j \setminus C_{j-1},$$

for $0 \leq j \leq J$, where $C_{-1} \doteq \emptyset$ by convention.

Splitting occurs whenever a threshold with a lower index is entered. It may happen that several thresholds are crossed at one step, even though this is unlikely. In that case, it is important to know which killing thresholds to assign to the newly created particles. As in [19], we let $q_\ell(j,k)$ be the number of new thresholds that are given killing threshold ℓ, assuming the particle moves from $x \in C_j \setminus C_{j-1}$ to $x \in C_k \setminus C_{k-1}$ for $k < j$.

Suppose we have a splitting rate R, i.e., the number of descendants from each particle is $R-1$ when crossing one threshold. If the particle moves from $x \in C_j \backslash C_{j-1}$ to $x \in C_k \setminus C_{k-1}$ we define

$$q_\ell(j,k) \doteq \begin{cases} 0 & \text{if } j \leq k, \\ (R-1)R^{j-\ell-1} & \text{if } k \leq j-1 \text{ and } k \leq \ell \leq j-1. \end{cases}$$

Before defining the estimator and the RESTART algorithm, we need to introduce some more notation. Let N_i be the number of particles at time i, $Z_{i,j}$ the position of the j^{th} particle at time i, $C_{i,j}$ the current threshold of the j^{th} particle at time i and

$K_{i,j}$ the killing threshold of the j^{th} particle at time i. We define the unnormalized empirical measure $\bar{\delta}_{Z_i}$ by

$$\bar{\delta}_{Z_i} \doteq \sum_{j=1}^{N_i} \delta_{Z_{i,j}},$$

and the estimator γ then equals

$$\gamma \doteq \sum_{i=0}^{\infty} \int_{\mathbb{R}^d} 1_{D^c}(x) e^{\bar{U}(x)-\bar{U}(0)} \bar{\delta}_{Z_i}(dx). \tag{2.1.4}$$

Note that the weight $\exp(\bar{U}(x) - \bar{U}(0))$ in (2.1.4) equals R^{j-J} when $x \in D_j = C_j \setminus C_{j-1}$. The state-dependence of $1_{D^c}(x)e^{\bar{U}(x)-\bar{U}(0)}$ in (2.1.4) is to account for the fact that the process can exit D at different thresholds.

The RESTART Algorithm can be found in Algorithm 1, and can also be found in e.g. [17] and [19]. Note that the pseudocode represents a "parallel" version of the algorithm, since the particles for a given threshold are split and then simulated until it is terminated, which occurs when the particle reaches the next threshold, the target set, or is killed. One could also implement a "sequential" version in which a particle is simulated until it either reaches the target set or is killed, while recording the times and locations of splitting. Once the particle is terminated, the algorithm returns to the highest threshold for which particles remain to be simulated and starts anew.

We also remark that RESTART is particularly appropriate for the problem of estimating escape from a domain over a long time interval. This is captured in the problem of Section 2.2, where the simulation is over increasing time intervals

$[0, T(n)]$, where n is the LDP parameter.

Initialization;
$N_0 = 1$, $Z_{0,1} = 0$, $C_{0,1} = J$, $K_{0,1} = J$, $\gamma = 0$, $i = 0$;
while $N_i \neq 0$ **do**
  $N_{i+1} = 0$;
  for $j = 1, \ldots, N_i$ **do**
    If the particle did not reach the boundary;
    if $Z_{i,j} \in D$ **then**
      Generate dummy variable $Y_{i,j}$ according to the law of X_{i+1};
      If the particle is not killed;
      if $\sigma(Y_{i,j}) \leq K_{i,j}$ **then**
        $N_{i+1} = N_{i+1} + 1$;
        $Z_{i+1,N_{i+1}} = Y_{i,j}$;
        $C_{i+1,N_{i+1}} = \sigma(Y_{i,j})$;
        $K_{i+1,N_{i+1}} = K_{i,j}$;
        If the particle needs to be branched;
        if $\sigma(Y_{i,j}) < C_{i,j}$ **then**
          for $k = 1, \ldots, J$ **do**
            for $\ell = 1, \ldots q_k(C_{i,j}, \sigma(Y_{i,j}))$ **do**
              $N_{i+1} = N_{i+1} + 1$;
              $Z_{i+1,N_{i+1}} = Y_{i,j}$;
              $C_{i+1,N_{i+1}} = \sigma(Y_{i,j})$;
              $K_{i+1,N_{i+1}} = k$;
            end
          end
        end
      end
    end
    If the particle reached the boundary;
    if $Z_{i,j} \notin D$ **then**
      $\gamma = \gamma + e^{\bar{U}(Z_{i,j})}$;
    end
  end
  i = i+1;
end
$\gamma = e^{-\bar{U}(0)}\gamma$;

Algorithm 1: RESTART Algorithm

2.1.4 Subsolutions as importance functions

The proposed method for generating importance functions is through a function $\bar{V}$ which will satisfy a subsolution property associated to the underlying dynamics of the sequence of processes $\{X^n\}$. We recall the main assumption that $\{X^n\}$ satisfies an LDP with rate function I and speed $r(n)$.

To motivate the use of a subsolution, one can ask what kind of properties one seeks in a general importance function U. Recall that U specifies the splitting thresholds via (2.1.3). If U decreases too slowly then splitting occurs infrequently and the scheme is not very helpful, while if U decreases too rapidly then splitting occurs too often and there is an exponential growth in the number of particles, which makes the scheme impossible to simulate.

Let π_j denote the entrance distribution on the threshold C_j, and let p denote the probability of a particle distributed according to π_j reaching C_{j-1} before being killed. Out of R independent particles distributed according to π_j, we expect Rp particles to make it to C_{j-1}, so choosing $Rp \leq 1$ avoids an explosion in the number of particles. When X^n satisfies a large deviations principle, for fixed $n \geq 1$ we have the heuristic approximation $p \approx e^{-r(n)I_T(\varphi_j)}$, where φ_j is a trajectory satisfying $\varphi_j(0) \in \partial C_j$ and $\varphi_j(T) \in C_{j-1}$ with minimal cost. If we scale the splitting rate as $R = e^{r(n)\Delta}$, then we want $e^{r(n)\Delta - r(n)I_T(\varphi_j)} \leq 1$. The definition of threshold implies $U(\varphi_j(0)) - U(\varphi_j(T)) \approx \Delta$. Taking logarithms in $Rp \leq 1$ and dividing by $r(n)$, we obtain

$$U(\varphi_j(0)) - U(\varphi_j(T)) \leq I_T(\varphi_j).$$

This rough estimate conveys why subsolutions as defined in Definition 1.4.6 may be appropriate importance functions. It also suggests the appropriate scaling of

thresholds which are introduced below.

We introduce the objects associated with the splitting scheme with an appropriate n-dependent scaling in the time-dependent and time-independent cases. In both cases, we require the speed $r(n)$ of the associated large deviations principle and the splitting rate $R \geq 1$ which determines $\Delta \doteq \log(R)$ and the scaled quantity $\Delta_n \doteq \Delta / r(n)$.

2.1.4.1 Time-independent subsolutions

We consider the problem of estimating

$$\mathbb{P}_{x_0}(X^n(t) \notin D, \text{ for some } t \in [0,T]), x_0 \in D$$

using a time-independent subsolution. In principle, one could also use a time-dependent subsolution for the same problem, but these can be difficult to construct even in a one dimensional setting, and we shall see that time-independent subsolutions are suitable when the time interval $[0,T]$ is large.

The appropriate definition of a weak-sense, time-independent subsolution for this problem is the one from Definition 1.4.6: a continuous function $\bar{V} : \mathbb{R}^d \to \mathbb{R}$ is a subsolution if it is bounded below and

$$\bar{V}(x) \leq \inf_{\varphi(0)=x;\varphi(T)\notin D;\varphi(s)\in D, s\in(0,T), T<\infty} \left\{ \int_0^T L(\varphi(s), \dot{\varphi}(s)) ds + \bar{V}(\varphi(T)) \right\}, \quad (2.1.5)$$

for all $x \in D$, and $\bar{V}(z) \leq 0$ for $z \notin D$.

Given a subsolution $\bar{V}$, we define the thresholds as follows:

$$C_j^n = \{x : \bar{V}(0) - \bar{V}(x) \geq (J^n - j)\Delta_n\},$$

for $j = 0, \ldots, J^n$, where J^n is the smallest positive integer such that $J^n\Delta_n \geq \bar{V}(0) - \bar{V}(x)$ for all $x \in D$. For any subsolution $\bar{V}$, $J^n\Delta_n$ approximates the maximum of $\bar{V}(0) - \bar{V}(x)$ on D:

$$J^n\Delta_n \to M \doteq \max_{x\in\bar{D}}\{\bar{V}(0) - \bar{V}(x)\}. \tag{2.1.6}$$

For simplicity we define the difference sets $D_j^n \doteq C_j^n \setminus C_{j-1}^n$, where $C_{-1}^n = \emptyset$ by convention.

The n-dependent piecewise continuous importance functions associated to a subsolution $\bar{V}$ will be denoted by $\bar{U}^n$, but we will speak of the splitting scheme associated to $\bar{V}$. For $x \in \mathbb{R}^d$ the piecewise constant function $\bar{U}^n$ is defined by

$$\bar{U}^n(x) \doteq \sum_{j=0}^{J^n} j\Delta_n 1_{D_j^n}(x). \tag{2.1.7}$$

For ease of notation we write $\bar{U}_j^n = \bar{U}^n(x)$ for $x \in D_j^n$. Observe that for $x \in D_j^n$, we have by (2.1.6) and (2.1.7)

$$|\bar{U}^n(x) - \bar{U}^n(0) - (\bar{V}(x) - \bar{V}(0))| \leq \Delta_n. \tag{2.1.8}$$

Thus, the differences $\bar{U}^n(x) - \bar{U}^n(y)$ converge uniformly to the differences $\bar{V}(x) - \bar{V}(y)$ as $n \to \infty$.

Roughly speaking, the quality of a subsolution is determined by how closely it approximates the corresponding calculus of variations problem, which will be com-

monly referred to as the "cost to exit D" starting at a given point. More precisely, define the function $W : \mathbb{R}^d \to [0, \infty)$ by

$$W(x) \doteq \inf \{I_T(\varphi) : \varphi(0) = x, \varphi(T) \notin D, T < \infty\}, \tag{2.1.9}$$

for $x \in D$ and $W(x) \doteq 0$ otherwise. Note that a subsolution $\bar{V}$ is never greater than W. It will turn out that, if $\bar{V}(x_0) = W(x_0)$, then the scheme is optimal in a specific sense described in Definition 2.1.2

It turns out that we can make a particularly convenient choice of subsolution which is appropriate for the problem of escape from D over a long time interval when D contains a unique stable equilibrium in the sense of Definition 1.5.1. For convenience, we take this unique equilibrium point to be $0 \in D$. We recall the notion of the quasipotential from Definition 1.5.2, which is the function $Q : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty]$ defined for $y, x \in \mathbb{R}^d$ by

$$Q(y, x) \doteq \inf_{\varphi(0)=y, \varphi(T)=x, T<\infty} \left\{ \int_0^T L(\varphi(s), \dot{\varphi}(s)) ds \right\}.$$

where the infimum is over all absolutely continuous trajectories. The quasipotential with respect to the unique attracting point 0 is $Q(x) \doteq Q(0, x)$.

We will later place conditions on Q which guarantee that Q is continuous and $Q(x) < \infty$ for $x \in D$; see Condition 2.2.2. It follows from the definitions of Q and W that

$$\min_{x \in \partial D} Q(x) = W(0). \tag{2.1.10}$$

From Lemma 1.5.5 that the negative of the quasipotential is always a subsolution, and $-Q(x) + c$ is a subsolution so long as $c \leq \inf_{x \notin D} Q(x)$. Figure 2.2 gives an illustration of choosing the subsolution as the negative of the quasipotential plus a

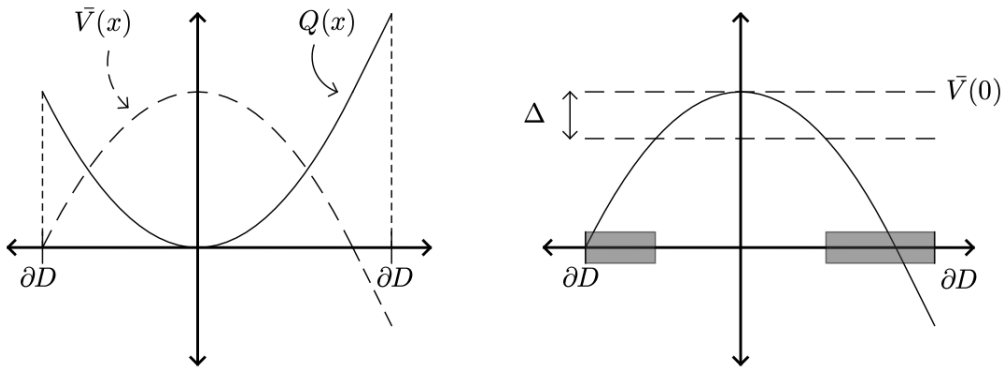

Figure 2.2: Left figure: The quasipotential $Q(x)$ for a simple one-dimensional domain, and $\bar{V}(x) = -Q(x) + \min_{x\in\partial D} Q(x)$. Right figure: The dashed level lines indicate a spacing of Δ, starting from $\bar{V}(0) = J\Delta$. The shaded region indicates the set of $x \in C_{J-1} = \{x : \bar{V}(0) - \bar{V}(x) \geq \Delta\}$.

constant, and this can be interpreted as follows. $Q(x)$ provides a lower bound on the cost of going from 0 to x, while a subsolution provides a lower bound on the cost to exit starting at x. If we were to choose the point x on the boundary which minimizes $Q(x)$, we would obtain a lower bound on the cost to escape starting at 0. Every other point in D can be related to the origin by the dynamic programming inequality (1.5.3).

We denote the minimum cost to exit via

$$c_Q \doteq \min_{x\notin D} Q(x),$$

and we remark that continuity of Q implies

$$c_Q = \min_{x\in\partial D} Q(x).$$

The constant c_Q will be referred to a few times in the rest of the section, but we emphasize that it plays no role in the determination of the RESTART algorithm since the thresholds are determined by *differences* of the subsolution. Rather, we will see that the optimality of a subsolution $\bar{V}$ is determined by whether or not

$\bar{V}(0) = c_Q$.

A final remark is that while we only consider estimating probabilities in this thesis, one could very well consider the estimation of so-called "risk-sensitive functionals" of the form $\mathbb{E}_x[e^{-nF(X^n(\rho^n))}1_{\{\rho^n \leq T(n)\}}]$, where $F : \mathbb{R}^d \to [0, \infty)$ is measurable. The same notion of subsolution would yield asymptotically optimal estimators, though one would have to appropriately modify the boundary conditions. For more information, see e.g. [18] and [21].

2.1.4.2 Time-dependent subsolutions

We will consider the problem of estimating

$$\mathbb{P}_{x_0}(X^n(T) \notin D)$$

using a time-dependent subsolution.

The appropriate definition of a weak-sense, time-dependent suboslution for this problem is the one from Definition 1.4.3: a continuous function $\bar{V} : [0, T] \times \mathbb{R}^d \to \mathbb{R}$ is a subsolution if it is bounded below and

$$\bar{V}(t, y) \leq \inf_{\varphi(t)=y, \varphi(T) \notin D} \left\{ \int_t^T L(\varphi(s), \dot{\varphi}(s))ds + \bar{V}(T, \varphi(T)) \right\}. \tag{2.1.11}$$

The corresponding thresholds are

$$C_j^n = \{(t, y) \in [0, T] \times \mathbb{R}^d : \bar{V}(0, 0) - \bar{V}(t, y) \geq (J^n - j)\Delta_n\}, \tag{2.1.12}$$

and we define the difference sets $D_j^n = C_j^n \setminus C_{j-1}^n$ for $j = 0, \ldots, J^n$, where J^n is the smallest positive integer such that $J^n\Delta_n \geq \bar{V}(0,0) - \bar{V}(t,y)$ for all $(t,y) \in [0,T] \times \mathbb{R}^d$. As in Section 2.1.4.1

$$M \doteq \lim_{n\to\infty} J^n\Delta_n. \tag{2.1.13}$$

We denote the associated n-dependent piecewise continuous importance function by $\bar{V}^n$, and we define it as

$$\bar{V}^n(t,y) = \sum_{j=0}^{J^n} j\Delta_n 1_{D_j^n}(t,y). \tag{2.1.14}$$

Observe that for $(t,y) \in D_j^n$, we have

$$|\bar{V}^n(t,y) - \bar{V}^n(0,0) - (\bar{V}(t,y) - \bar{V}(0,0))| \leq \Delta_n. \tag{2.1.15}$$

Thus, the differences $\bar{V}^n(t,y) - \bar{V}^n(0,0)$ converge uniformly to the differences $\bar{V}(t,y) - \bar{V}(0,0)$ as $n \to \infty$.

2.1.5 Performance measures and the rare event setting

In this section we discuss performance measures for unbiased estimators of rare events associated to a sequence of processes $\{X^n, n \geq 1\}$ with $X^n \in D([0,T])$ which satisfy a large deviations principle with rate function I and speed $r(n)$.

Let $A \subset D([0,T])$ be measurable and suppose that $I(A) > 0$. Throughout, we will assume that A is an I-continuity set, i.e. $I(\bar{A}) = I(A^\circ)$. This is not necessary and simply ensures that certain limits exist, which simplifies some definitions. Let $\hat{\theta}^n$ be an unbiased estimator of $\mathbb{P}(X^n \in A)$ so that $\mathbb{E}[\hat{\theta}^n] = \mathbb{P}(X^n \in A)$. A natural performance measure used to judge the quality of an estimator is its variance. If we restrict to unbiased estimators, then minimizing the variance is equivalent to

minimizing the second moment.

Given an estimator $\hat{\theta}$, let $\mathcal{S}(\hat{\theta})$ denote its second moment,

$$\mathcal{S}(\hat{\theta}^n) \doteq \mathbb{E}[(\hat{\theta}^n)^2], \tag{2.1.16}$$

and note that by Jensen's inequality,

$$\mathcal{S}(\hat{\theta}^n) \geq 2\mathbb{P}(X^n \in A).$$

Owing to the large deviations scaling of the escape probability, it is convenient to frame the minimization of the second moment as the maximization of the scaled quantity $-r(n)^{-1} \log \mathcal{S}^n(\bar{V})$. The large deviation principle then implies

$$\limsup_{n\to\infty} -\frac{1}{r(n)} \log \mathcal{S}(\hat{\theta}^n) \leq 2I(A). \tag{2.1.17}$$

This will be the main quantity of interest for us.

Definition 2.1.1. The *decay rate* of an unbiased estimator $\hat{\theta}^n$ is given by

$$\liminf_{n\to\infty} -\frac{1}{r(n)} \log \mathcal{S}(\hat{\theta}^n).$$

Observe that (2.1.17) places an upper bound on the decay rate of any unbiased estimator: it can be no better than twice the decay rate of the probability. When this upper bound is achieved, we say that the estimator is *asymptotically optimal*.

Definition 2.1.2. Fix $T > 0$ and $A \subset D([0,T])$ be measurable. An estimator is said to be *asymptotically optimal* for the problem of estimating $\mathbb{P}_{x_0}(X^n \in A)$ if

$$\liminf_{n\to\infty} -\frac{1}{r(n)} \log \mathcal{S}(\hat{\theta}^n) \geq 2I(A). \tag{2.1.18}$$

The notion of optimality is a little subtle, however, when the time intervals $[0, T]$ are allowed to depend on the large deviations scaling parameter n. As discussed in Section 2.2.3, asymptotic optimality of an estimator over fixed time intervals $[0, T]$ does not imply the optimality over time intervals $[0, T(n)]$ when $T(n)$ increases polynomially with n. For this reason we provide a few definitions of asymptotic optimality below depending on the problem of interest. Another issue to consider is the computational expense of simulating the many trajectories of a splitting algorithm. Indeed, if exponentially many particles are created for a single run of a splitting algorithm, the comparison with standard Monte Carlo is unfair. We address this point now by introducing the notion of work and "work-normalized"-error.

Let N_i^n denote the number of particles at time i, where N_i is defined for the unscaled algorithm in Section 2.1.3. The *work* of a single run of RESTART is the quantity

$$w^n = \sum_{i=1}^{\infty} N_i^n. \tag{2.1.19}$$

The work-normalized error is given by (see [38])

$$\lim_{n\to\infty} \frac{1}{n} \log \frac{\mathcal{S}(\hat{\theta}^n)\mathbb{E}_0[w^n]}{\mathbb{P}_{x_0}(X^n \in A}.$$

Once again by Jensen's inequality,

$$\lim_{n\to\infty} \frac{1}{n} \log \frac{\mathcal{S}(\hat{\theta}^n)\mathbb{E}_0[w^n]}{\mathbb{P}_{x_0}(X^n \in A)^2} \geq \frac{1}{n} \log \mathbb{E}[w^n].$$

It has been shown in [19], Theorem 5.7, that over finite time intervals the RESTART scheme does subexponential work when based on a subsolution. That is, when the simulation occurs over a time interval $[0, T]$,

$$\lim_{n\to\infty} \frac{1}{n} \log \mathbb{E}_0[w^n] = 0.$$

An argument similar to the one in [19] shows that the same result holds for work over increasing time intervals $[0, T(n)]$ when $T(n)$ is polynomial in n. The work-normalized error achieves its optimal value of 0 if and only if the work is subexponential and the scheme is asymptotically optimal in the sense of Definition 2.1.2 (or Definition 2.2.14 in Section 2.2 when T depends on n and A is of a specific form). Since the work is subexponential when the RESTART scheme is based on a subsolution, to demonstrate that the work-normalized error decays to its optimal value of zero, it suffices to establish asymptotic optimality. In our numerical examples we do record the work w^n.

For splitting algorithms in general, the decay rate of the estimator inherently depends on the importance functions used, which in turn are built using subsolutions. We introduce the following notation: let $\gamma^n \doteq \gamma^n(\bar{V})$ denote the RESTART estimator based on a subsolution $\bar{V}$, and let $\mathcal{S}^n(\bar{V}) \doteq \mathcal{S}(\gamma^n(\bar{V}))$ denote the associated pre-asymptotic decay rate.

2.1.6 RESTART inequalities

We have from Theorem 17.6 in [17] the following pre-asymptotic upper bound for the second moment of RESTART. The inequality first appeared in [19], Theorem 4.5, and is presented here in a slightly modified form, i.e. in [19] the RESTART algorithm is not limited to estimating probabilities and hence their notation is different. The statement concerns the estimator without explicit dependence on a scaling parameter n, and so we denote the second moment of the RESTART estimator by $\mathcal{S}(\bar{U})$.

Theorem 2.1.3. *Suppose that A is open, B is closed with $A \cap B = \emptyset$, Y_i is a discrete time Markov chain, and that the stopping time $H \doteq \inf\{i \geq 1 : Y_i \in A \cup B\}$ is almost*

surely finite. Then if $\bar{U}$ is used to design the splitting scheme, for all $y_0 \in (A \cup B)^c$ we have

$$\mathcal{S}(\bar{U}) \leq e^{-\bar{U}(y_0)}\mathbb{E}_{y_0}\left[\sum_{i=1}^{H} e^{\bar{U}(Y_{i-1})}[\mathbb{P}_{Y_i}(Y_H \in B)]^2\right] + e^{-\bar{U}(y_0)}\mathbb{E}_{y_0}\left[e^{\bar{U}(Y_H)}1_B(Y_H)\right]. \tag{2.1.20}$$

We will apply this theorem to deduce an analogous result for the present situation of exiting the domain D within some time $T(n)$. We do so by introducing time as a state variable and judiciously choosing A and B.

Corollary 2.1.4. *For an open domain D, fix some $n \in \mathbb{N}$ and let $\rho^n \doteq \inf\{i \geq 0 : X_i^n \notin \mathcal{D}\}$. If $\bar{V}$ is used to determine the importance functions $n\bar{U}^n$ used to construct the splitting scheme, we have*

$$\mathcal{S}^n(\bar{V}) \leq e^{-n\bar{U}^n(0)}\mathbb{E}_0\left[\sum_{i=1}^{nT(n)} e^{n\bar{U}^n(X_{i-1}^n)}[\mathbb{P}_{X_i^n}(\rho^n \leq T(n) - i/n)]^2\right] + e^{-n\bar{U}^n(0)}\mathbb{E}_0\left[e^{n\bar{U}^n(X_{\rho^n}^n)}1\{\rho^n \leq T(n)\}\right].$$

Proof. Fix any $n \geq 1$ and set $A = D \times (T(n), \infty)$ and $B = D^c \times [0, T(n)]$. Note that A is open, B is closed and $A \cap B = \emptyset$. We will apply Theorem 2.1.3 to the Markov chain $Y_i = (X_i^n, i/n)$, with $\bar{U}(x,t) = n\bar{U}^n(x)$ being the importance function used to design the splitting scheme. The stopping time $H \doteq \inf\{i \geq 1 : Y_i \in (A \cup B)\}$ is finite almost surely by finiteness of $T(n)$ for each $n \geq 1$.

Given $y = (x, k/n) \in D \times [0, T(n)]$, $\mathbb{P}_y(Y_H \in B)$ is the probability that there is some $i \in \{k+1, \ldots, nT(n)\}$ for which $X_i^n \notin D$, conditioned on $X_k^n = x$. By the

Markov property

$$\mathbb{P}_{Y_i}(Y_H \in B) = \mathbb{P}_{(X_i^n, i/n)}(\rho^n \leq T(n) - i/n),$$

where ρ^n on the right-hand side above is the first entrance time for the process started at X_i^n. Applying this observation to the first term in (2.1.20), we obtain with $y_0 = (0,0)$

$$e^{-n\bar{U}(y_0)}\mathbb{E}_{y_0}\left[\sum_{i=1}^{H} e^{n\bar{U}(Y_{i-1})}[\mathbb{P}_{Y_i}(Y_H \in B)]^2\right]$$
$$\leq e^{-n\bar{U}^n(0)}\mathbb{E}_{(0,0)}\left[\sum_{i=1}^{nT(n)} e^{n\bar{U}^n(X_{i-1}^n)}[\mathbb{P}_{(X_i^n, i/n)}(\rho^n \leq T(n) - i/n)]^2\right].$$

The inequality follows because $H \leq nT(n)$, so one is taking a sum over a larger number of positive terms. The Markov chain X_i^n is time-homogeneous, so with an abuse of notation,

$$\mathbb{P}_{(X_i^n, i/n)}(\rho^n \leq T(n) - i/n) = \mathbb{P}_{X_i^n}(\rho^n \leq T(n) - i/n).$$

For the second term in (2.1.20), note that

$$1_B(Y_H) = 1\{\rho^n \leq T(n)\},$$

and $Y_H = X_{\rho^n}^n$. By again using $y_0 = (0,0)$ and time-homogeneity of X_i^n, we conclude the proof. □

When we consider the moderate deviations scaling in Section 2.3, we will instead

consider the problem of estimating

$$\mathbb{P}(Y^n(T) \notin D^*)$$

for some $T < \infty$ and appropriately scaled set $D^* \subset \mathbb{R}^d$. We assume that the process Y^n satisfies an LDP with speed $r(n) = a_n^2$. We apply Theorem 2.1.3 to the open set $A = (T, \infty) \times D$ and the closed set $B = [T, \infty) \times (D^*)^c$ to obtain the following corollary, the proof of which is entirely analogous to the proof of Corollary 2.1.4 and is therefore omitted.

Corollary 2.1.5. *Let $\bar{V}$ be a subsolution in the sense of Definition 1.4.3 from Chaper 1 with the appropriate boundary conditions, and let $\bar{V}^n(t, y)$ denote the piecewise constant importance function generated by $\bar{V}$ via (2.1.7). Recall that $\mathcal{S}^n(\bar{V})$ denotes the second moment of the estimator γ^n when $\bar{V}^n$ is used to specify the splitting thresholds. Then*

$$\mathcal{S}^n(\bar{V}) \leq \mathbb{E}_{0,0}\left[\sum_{i=1}^{nT} e^{a_n^{-2}(\bar{V}((i-1)/n, Y^n_{i-1}) - \bar{V}(0,0))}(\mathbb{P}_{i/n, Y^n_i}(Y^n_{nT} \notin D^*))^2\right]$$
$$+ \mathbb{E}_{0,0}[e^{a_n^{-2}(\bar{V}((nT-1)/n, Y^n_{nT-1}) - \bar{V}(0,0))} 1_{D_*^c}(Y^n_{nT})].$$

2.2 Large-deviations based splitting over long time intervals

Remark. This section is joint work with A. Buijsrogge, a PhD candidate at the University of Twente.

In this section we study *exit probabilities* for discrete time stochastic processes,

where the process escapes from some neighborhood of an attractor prior to a given time. In contrast to existing work on exit probabilities, see for example [35, 59], we allow the time by which the process has to exit to grow polynomially with the large deviations scaling parameter. One motivation for considering this scaling is that when the time interval is large and the escape probability is small, the probability closely approximates the inverse of the mean escape time, and in particular the exponential decay rate for the probability and growth rate for the escape time coincide. Although one can easily conjecture an expression for the decay rate for such probabilities, it does not follow from standard sample path large deviation estimates, which apply to bounded time intervals. A first contribution is to apply Freidlin-Wentzell type arguments to rigorously determine the decay rate.

The main focus is on Monte Carlo estimation of the probability using *splitting* methods; see e.g. [37]. In rare event simulation the two most commonly used methods are *importance sampling* and those based on interacting particles, which include splitting as a particular case. For a discussion on some of the differences and similarities of importance sampling and splitting we refer to [14]. For both importance sampling and splitting it turns out that one can design and analyze efficient Monte Carlo schemes using *subsolutions* to a Hamilton-Jacobi-Bellman equation that is naturally associated to the process through the large deviation rate function: see for example [18, 19, 21]. This method has been successfully applied in various settings, for example queueing theory [22], and will also be used in the current chapter. However, it was shown in [16] that importance sampling has some shortcomings when applied to the problem of estimating probabilities to escape from the neighborhood of an attractor. For example, when dealing with such problems one is tempted, owing to the fact that the time interval is large, to consider subsolutions to the corresponding time independent Hamilton-Jacobi-Bellman equation as the basis for

algorithm design. It is sometimes the case (e.g., for reversible systems) that useful subsolutions can be found much more easily for the time independent version. However, as discussed in [16], since the attractor is inside the interior of the domain of interest, importance sampling schemes based on such time independent subsolutions generically degrade as the time interval gets large. We will consider the design of splitting schemes, and show that no such degradation occurs.

Within the setting of splitting schemes there are those which allow killing to improve efficiency without introducing bias, with one of the most well known of such schemes being the RESTART method [63]. As will be discussed later in the chapter, the improvement in efficiency of RESTART over ordinary splitting increases in proportion to the time interval, and so in this chapter we focus on this version of splitting, though the theoretical results can be proved for ordinary splitting as well, and in fact with simpler proofs.

The outline of the remainder of the section is as follows. In Section 2.2.1 we introduce the model, state conditions that will hold throughout and develop some preliminary results. Then in Section 2.2.2 we determine the decay rate for the probability to escape from a domain within some time $T(n)$, where n is the large deviation parameter and $T(n)$ is allowed to grow polynomially in n. In Section 2.1.3 we introduce RESTART and determine the asymptotic decay rate.

2.2.1 Process model

The problem of interest is to estimate exit probabilities of a discrete time process $\{X_i^n\}$ from a bounded open set $\mathcal{D} \subset \mathbb{R}^d$ over a time interval $[0, T(n)]$. Previously

The index n serves a dual purpose: we assume that X^n satisfies a large deviations principle (LDP) with rate n, and we also assume $T(n) \to \infty$ as $n \to \infty$. We will require that the LDP be uniform with respect to the initial condition, as characterized in Definition 2.2.3.

A large class of processes satisfying our assumptions can be obtained from recursive chains of the form

$$X_{i+1}^n = X_i^n + \frac{1}{n} v_i(X_i^n), \tag{2.2.1}$$

where $v_i(\cdot)$ are independent and identically distributed random vectors fields whose distribution on $\mathbb{R}^d$ is given by a stochastic kernel $\mu(\cdot|x)$, $x \in \mathbb{R}^d$, having mean $b(x)$ and with $b(\cdot)$ Lipschitz continuous. We will use this as a canonical model for discussion. However, the results will be valid in greater generality, with the key assumption being that an appropriate uniform LDP is available. This covers models such as Markov-modulated processes.

The process X^n naturally induces a family of probability measures $\mathbb{P}_x$, where $\mathbb{P}_x(X^n(0) = x) = 1$. We let $\mathbb{E}_x[\cdot]$ denote the expected value with respect to $\mathbb{P}_x$. Recall the cumulant generating function of the vector fields,

$$H(x, \alpha) = \log \int_{\mathbb{R}^d} e^{\langle \alpha, u \rangle} \mu(du|x).$$

We assume that the following conditions are satisfied by H and the stochastic kernel μ.

Condition 2.2.1. *1. For all $\alpha \in \mathbb{R}^d$, we have $\sup_{x \in \mathbb{R}^d} H(x, \alpha) < \infty$;*

2. The map $x \to \mu(\cdot|x)$ is continuous in the topology of weak convergence.

The cumulant generating function $H(x, \alpha)$ is used to define the local rate associ-

ated to the system (2.2.1) via the Legendre-Fenchel transformation. For $(x, \beta) \in \mathbb{R}^d$, let

$$L(x, \beta) \doteq \sup_{\alpha \in \mathbb{R}^d} \{\langle \alpha, \beta \rangle - H(x, \alpha)\}. \tag{2.2.2}$$

We first formulate conditions on the local rate L. For models of the form (2.2.1), since L is the Legendre-Fenchel transform of a cumulant generating function it is automatic that $L(x, \cdot)$ is convex and that there is a unique $b(x)$ such at $L(x, b(x)) = 0$ for each $x \in \mathbb{R}^d$. We will assume the following properties of L. These can all be related to properties of the kernel $\mu(\cdot|x)$ for models of the form (2.2.1), see [17], and can also apply to other types of models.

Condition 2.2.2. *We assume that $L : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty]$ satisfies the following properties:*

1. *for every compact $K \subset \mathbb{R}^d$, $\lim_{c \to \infty} \inf_{x \in K, |\beta| = c} L(x, \beta) = \infty$;*
2. *for every $x \in \mathbb{R}^d$ there is a unique $b(x)$ such at $L(x, b(x)) = 0$, and $b : \mathbb{R}^d \to \mathbb{R}^d$ is Lipschitz continuous;*
3. *for every $x \in \mathbb{R}^d$ $L(x, \cdot)$ is convex;*
4. *for every compact set K, there is $r > 0$ such that $L(x, \beta) \leq 1/r$ if $x \in K$, $|\beta| \leq r$.*

The most significant of these is the last property, which is used to establish that the infimum of the rate function (defined below) over the interior and closure of sets of trajectories coincide. This is not needed, but without it both the statement of results and their proofs become much more complicated.

To state the LDP for X_i^n, it is convenient to interpolate it as a piecewise constant

process. For fixed $T \in (0, \infty)$, define

$$X^n(t) \doteq X_i^n, \quad t \in (i/n, (i+1)/n], \ i = 0, \ldots, \lfloor nT \rfloor,$$

where $\lfloor a \rfloor$ is the integer part of a. We consider X^n as taking values in $D([0,T]) \doteq D([0,T] : \mathbb{R}^d)$, with the usual Skorohod topology [29]. We could also have considered a piecewise linear interpolation that takes values in $C([0,T]) \doteq C([0,T] : \mathbb{R}^d)$ and obtain an LDP with the same rate function, but then we would have to justify the use of the Markov property at only the interpolation times and this seems to us to lead to a more complicated proof. We will make use of the fact that the Skorokhod topology on $D([0,T])$ relativized to $C([0,T])$ coincides with the uniform topology given by the metric

$$d(\varphi_1, \varphi_2) \doteq \|\varphi_1 - \varphi_2\|_\infty = \sup_{0 \le t \le T} |\varphi_1(t) - \varphi_2(t)|.$$

We omit the implicit dependence of the metric d on the time interval $[0,T]$ for notational simplicity, as it will be clear from the context what the interval is. For any set $A \subset C([0,T])$ and $\varphi \in C([0,T])$, we define $d(\varphi, A) \doteq \inf_{\phi \in A} d(\varphi, \phi)$. For objects living in $\mathbb{R}^d$ we use absolute values $|\cdot|$ to denote the standard Euclidean norm.

The rate function associated with the process X^n over the interval $[0,T]$ is

$$I_T(\varphi) \doteq \int_0^T L(\varphi(s), \dot{\varphi}(s)) ds,$$

if $\varphi(t)$ is absolutely continuous, and $I_T(\varphi) = \infty$ otherwise. We phrase the uniform large deviations principle in terms of level sets of I_T, which are defined by

$$\Phi_{x,T}(s) \doteq \{\varphi \in D([0,T]) : I_T(\varphi) \le s, \varphi(0) = x\}.$$

For the purposes of this section, we use a uniform version of the large deviations principle, which is taken from [35, p. 74].

Definition 2.2.3 (Uniform large deviations principle)**.** The sequence X^n satisfies a uniform large deviations principle if:

1. the functional I_T is lower semicontinuous on $D([0,T])$ and for each $T \in [0,\infty)$, compact $K \subset \mathbb{R}^d$ and $s < \infty$, the set $\cup_{x\in K}\Phi_{x,T}(s)$ is compact;

2. for any $\delta > 0$, $\gamma > 0$, $s_0 < \infty$ and any compact $K \subset \mathbb{R}^d$, there is $N \in \mathbb{N}$ such that
$$\mathbb{P}_x(\|X^n - \phi\|_\infty < \delta) \geq \exp(-n(I_T(\phi) + \gamma))$$
for all $n \geq N$, all $x \in K$ and all $\phi \in \Phi_{x,T}(s_0)$;

3. for any $\delta > 0$, $\gamma > 0$, $s_0 < \infty$ and any compact $K \subset \mathbb{R}^d$, there exists $N \in \mathbb{N}$ such that
$$\mathbb{P}_x(d(X^n, \Phi_{x,T}(s_0)) \geq \delta) \leq \exp(-n(s - \gamma))$$
for all $n \geq N$, $s \leq s_0$ and $x \in K$.

This definition allows us to phrase the final assumption on the process X_i^n.

Condition 2.2.4. *For each $T \in (0,\infty)$, the sequence $\{X^n, n \in \mathbb{N}\}$ satisfies the uniform LDP in Definition 2.2.3 with a rate function $I_T(\varphi)$ of the form*

$$I_T(\varphi) = \int_0^T L(\varphi(s), \dot{\varphi}(s))ds,$$

for some L satisfying Condition 2.2.2.

Having posed all conditions on this process, we are interested in the probability that this process escapes from a domain $\mathcal{D}$ before some time $T(n)$. We denote the

first exit time of the process X^n from the set $\mathcal{D}$ by

$$\rho^n \doteq \inf\{t \geq 0 : X^n(t) \notin \mathcal{D}\}.$$

We assume that the process exits the set $\mathcal{D}$ in finite time with probability 1, i.e. for all $x \in \mathcal{D}$,

$$\mathbb{P}_x(\rho^n < \infty) = 1.$$

Condition 2.2.5. *We impose the following conditions on the set $\mathcal{D} \subset \mathbb{R}^d$:*

1. *$\mathcal{D}$ is a bounded open subset of $\mathbb{R}^d$;*
2. *$\mathcal{D}$ satisfies a regularity condition at all points of its boundary: for any $\delta > 0$ and $p \in \partial\mathcal{D}$, there is a point q in the interior of $\mathcal{D}^c$ with $|p - q| < \delta$;*
3. *there is a unique asymptotically stable equilibrium point $O \in \mathcal{D}$ for the noiseless dynamics $\dot{\phi} = b(\phi)$ in the sense of Lyapunov: for any $\mu > 0$, there is $\delta > 0$ such that if $|\phi(0) - O| < \delta$, then $|\phi(t) - O| < \mu$ for all $t \geq 0$. Without loss of generality, we take the unique equilibrium point O to be the origin $0 \in \mathbb{R}^d$.*

The assumption that $\mathcal{D}$ is bounded is not necessary for the results on the decay rate of RESTART to hold, since one can always restrict to an appropriately chosen bounded subset of the domain. Such a generalization is straightforward but cumbersome and is therefore omitted. With regard to $T(n)$ we assume the following.

Condition 2.2.6. *$T(n)$ grows polynomially in n, by which we mean*

$$\lim_{n\to\infty} T(n) = \infty,$$

and

$$\lim_{n\to\infty} \frac{1}{n} \log T(n) = 0.$$

Throughout the section many results hold asymptotically as $n \to \infty$, and such statements are proved by establishing appropriate estimates for all n greater than some threshold value $N(\cdot)$, where $(\cdot)$ represents parameters in the proof. Our convention is to explicitly write this dependence once and thereafter simply write N. We do this for all limit statements, not just those involving n.

2.2.1.1 Preliminary results

In this section we will state some results that will be used in Section 2.2.2 to estimate the rate of decay of the exit probability from $\mathcal{D}$ by some time $T(n)$. The proofs of the lemmas in this section follow more or less directly from the conditions that were presented in Section 2.2.1 and hence these proofs can be found in Section 2.2.4.

The uniform bound on the cumulant generating function $H(x, \alpha)$ allows us to establish an asymptotic bound on the maximum jump size of the process, which is the content of Lemma 2.2.7.

Lemma 2.2.7. *Assume Condition 2.2.1. Then for all $\delta > 0$,*

$$\limsup_{n\to\infty} \sup_{x\in\mathbb{R}^d} \frac{1}{n} \log \mathbb{P}_x \left(\max_{0\le i\le nT(n)} |X^n_{i+1} - X^n_i| \ge \delta \right) = -\infty.$$

Recall the definition of $W(\cdot)$ from (2.1.9),

$$W(x) = \inf \left\{ I_T(\varphi) : \varphi(0) = x, \varphi(T) \notin \mathcal{D}, T < \infty \right\}.$$

The function $W(\cdot)$ has the following properties.

Lemma 2.2.8. *Under Condition 2.2.2, the function $W(\cdot)$ satisfies*

1. *$W(\cdot)$ is continuous on $\bar{\mathcal{D}}$,*
2. *for every $x \in \mathcal{D}$, $W(x) \leq W(0)$.*

The last lemma in this section gives an upper bound on the probability that the process does not enter the μ-neighborhood of the origin within some finite time T, starting at some point x that is not in the μ-neighborhood of the origin. The proof of this lemma can be found in [35], Lemma 2.2 in Chapter 4.

Lemma 2.2.9. *For any $\mu > 0$ and any $M < \infty$, there are $T < \infty$ and $N < \infty$ such that any trajectory starting at $x \notin \mathcal{E}_\mu = \{x : |x| < \mu\}$ satisfies*

$$\mathbb{P}_x(X^n(t) \notin \mathcal{E}_\mu, 0 \leq t \leq T) \leq e^{-nM},$$

for all $n \geq N$.

2.2.2 Long time estimates

The main result of this section is the following estimate.

Theorem 2.2.10. *Suppose Conditions 2.2.1, 2.2.2, 2.2.4, 2.2.5 and 2.2.6 are met. Then for any $x \in \mathcal{D}$ we have*

$$\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}_x(\rho^n \leq T(n)) = -W(x). \tag{2.2.3}$$

Since $T(n) \to \infty$, this is not covered by an LDP of the form in Definition 2.2.3, and in fact requires a Freidlin-Wentzell analysis appropriate for large time problems as in [35]. The proof is divided into upper and lower bounds for the decay of $\mathbb{P}_x(\rho^n \leq T(n))$. We start with the upper bound when the process starts in a small neighborhood of the attracting point 0.

Lemma 2.2.11. *Assume the conditions of Theorem 2.2.10 are met. For any $\varepsilon > 0$, there is a $\mu > 0$ and $N < \infty$ such that for all $n \geq N$ and all $y \in \mathcal{E}_\mu \doteq \{x : |x| < \mu\}$,*

$$\frac{1}{n} \log \mathbb{P}_y(\rho^n \leq T(n)) \leq -W(y) + \varepsilon. \tag{2.2.4}$$

Proof. Let $\varepsilon > 0$. The function $y \to W(y)$ is continuous on $\mathcal{D}$ by Lemma 2.2.8, hence uniformly continuous on any closed neighborhood of 0. For the given $\varepsilon > 0$, there is $\mu > 0$ such that

$$|W(x) - W(y)| < \varepsilon/8 \quad \text{whenever } |x - y| < 2\mu. \tag{2.2.5}$$

We choose such a $\mu > 0$ and any $\eta \in (0, \mu/2)$. We also assume μ small enough that $\mathcal{E}_{\mu/2} \subset \mathcal{D}$. Both μ and η will remain fixed for the rest of the proof. For readability we have divided the proof in 5 steps.

Step 1. *A priori estimates.*

Define the sets

$$\begin{aligned} \Gamma_\mu^\eta &\doteq \{x : \mu - \eta \leq |x| \leq \mu\}, \\ \gamma_\mu^\eta &\doteq \left\{x : \frac{\mu}{2} - \eta \leq |x| \leq \frac{\mu}{2}\right\}. \end{aligned}$$

We will choose a time T_1 such that the probability that the process does not enter the

$\mu/2$-neighborhood of the attracting point 0 is superexponentially small. By Lemma 2.2.9, there are $T_1(\mu)$ and $N_1(\mu)$ such that $T \geq T_1$ and $n \geq N_1$ together imply, for any $x \in \mathcal{D} \setminus \mathcal{E}_{\mu/2}$, that

$$\mathbb{P}_x(X^n(t) \in \mathcal{D} \setminus \mathcal{E}_{\mu/2}, t \in [0, T_1]) \leq e^{-n(\bar{W}+1)}, \tag{2.2.6}$$

where

$$\bar{W} \doteq \inf_{x \in \Gamma_\mu^\eta} W(x).$$

Without loss of generality, we may assume that $\bar{W} > \varepsilon$.

We claim there is $\delta(\varepsilon) > 0$ such that any path φ which satisfies $\varphi(0) \in \Gamma_\mu^\eta$ and $\varphi(t) \notin \mathcal{D}$ for some $t \in [0, T_1]$ also satisfies

$$d(\varphi, \Phi_{\varphi(0),T_1}(\bar{W} - \varepsilon/4)) \geq \delta > 0.$$

If not, for any $\delta > 0$ we could find a path φ^δ with $\varphi^\delta(0) \in \Gamma_\mu^\eta$ and $\varphi^\delta(t_*) \notin \mathcal{D}$ for some $t_* \in [0, T_1]$, and $d(\varphi^\delta, \Phi_{\varphi^\delta(0),T_1}(\bar{W} - \varepsilon/4)) < \delta$. By Lemma 2.2.17, on the compact set $K = \{x : \inf_{z \in \mathcal{D}} |x - z| \leq 1\}$ there is $c > 0$ such that for all $x, y \in K$, we can construct φ_{xy} satisfying $\varphi_{xy}(0) = x$, $\varphi_{xy}(\tau) = y$ with $\tau \in (0, \infty)$ and $I_\tau(\varphi_{xy}) \leq c|x - y|$.

Choose $\delta \in (0, \varepsilon/8c)$, and without loss of generality assume $\delta < 1$. We obtain a path φ^δ starting in Γ_μ^η for which $\varphi^\delta(t_*) \notin \mathcal{D}$ at some $t_* \in [0, T_1]$. Under the assumption that $d(\varphi^\delta, \Phi_{\varphi^\delta(0),T_1}(\bar{W} - \varepsilon/4)) < \delta$, we can find $\phi \in \Phi_{\varphi^\delta(0),T_1}(\bar{W} - \varepsilon/4)$ such that $\|\phi - \varphi^\delta\|_\infty < \delta$. Since $\phi(0) = \varphi^\delta(0) \in \Gamma_\mu^\eta$, the definition of $\bar{W}$ implies that $\phi(t_*) \in \mathcal{D} \subset K$, while $\delta < 1$ ensures $\varphi^\delta(t_*) \in K$. Applying Lemma 2.2.17 to the points $x = \phi(t_*)$ and $y = \varphi^\delta(t_*)$, we obtain a path φ_{xy} connecting x and y in time τ with cost $I_\tau(\varphi_{xy}) \leq c\delta < \varepsilon/8$. Concatenating ϕ on $[0, t_*]$ with φ_{xy} on $(t_*, t_* + \tau]$, we obtain a path starting in Γ_μ^η and ending outside of $\mathcal{D}$ at time $t_* + \tau < \infty$, with cost

less than $I_{t_*}(\phi) + I_\tau(\varphi_{xy}) \le \bar{W} - \varepsilon/8$. This contradicts the definition of $\bar{W}$. Thus $\delta > 0$ must exist as claimed.

Observe that we have the inclusion

$$\{\rho^n \le T_1, X^n(0) = y\} \subset \{d(X^n, \Phi_{y,T_1}(\bar{W} - \varepsilon/4)) \ge \delta\}.$$

By the finite time uniform large deviations upper bound on the interval $[0, T_1]$, there is $N_2(\varepsilon, T_1) < \infty$ for which $n \ge N_2$ implies, for any $y \in \Gamma_\mu^\eta$,

$$\mathbb{P}_y(d(X^n, \Phi_{y,T_1}(\bar{W} - \varepsilon/4)) \ge \delta) \le e^{-n(\bar{W} - \varepsilon/4 - \varepsilon/4)}.$$

Hence, for $n \ge N_2$,

$$\mathbb{P}_y(\rho^n \le T_1) \le e^{-n(\bar{W} - \frac{\varepsilon}{2})}. \tag{2.2.7}$$

With T_1 chosen, we introduce the stopping times $\sigma_0 \doteq 0$, and for $j \ge 1$,

$$\begin{aligned} \tau_j^n &\doteq \inf\{t > \sigma_{j-1}^n : X^n(t) \in (\gamma_\mu^\eta \cup \mathcal{D}^c)\} \wedge (T_1 + \sigma_{j-1}^n), \\ \sigma_j^n &\doteq \inf\{t > \tau_j^n : X^n(t) \in \Gamma_\mu^\eta\}. \end{aligned}$$

The stopping times τ_j^n and σ_j^n also depend on η and μ, but we do not include this dependence to avoid an overload of notation. Figure 2.3 shows a sample path with its associated stopping times.

Step 2. *The probability of an excursion to $\partial\mathcal{D}$ within time $T(n)$ is approximately the probability of excursion over a single time interval $(\sigma_{j-1}^n, \tau_j^n)$ times the expected number of such intervals.*

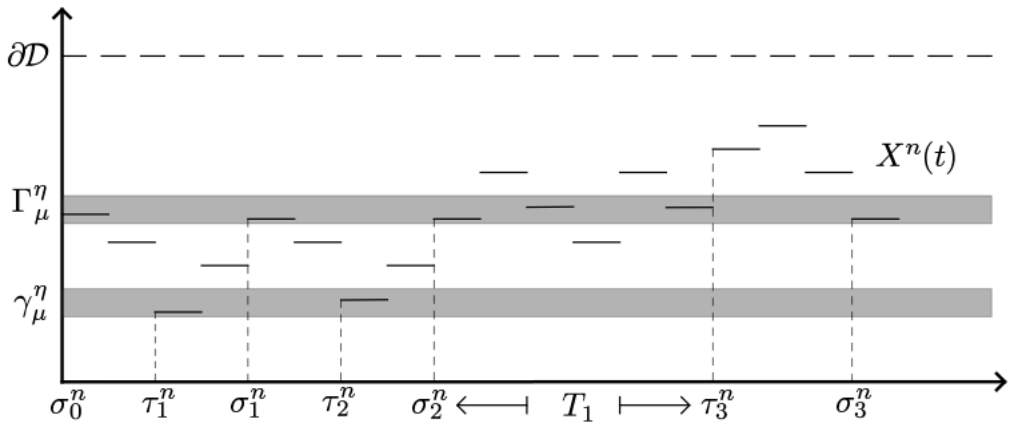

Figure 2.3: The piecewise constant process $X^n(t)$ moving between the sets Γ_μ^η and γ_μ^η. The process in the figure has not exited the domain $\mathcal{D}$ nor hit γ_μ^η at time τ_3^n, but nevertheless σ_3^n is triggered when the process returns to Γ_μ^η since it wandered around for time longer than T_1.

Define the random variables

$$B_j^n \doteq \begin{cases} 1 & \text{if} \quad X^n(\tau_j^n) \notin \mathcal{D} \quad \text{or} \quad \tau_j^n - \sigma_{j-1}^n \geq T_1, \\ 0 & \text{else.} \end{cases}$$

For any $y \in \Gamma_\mu^\eta$ it holds that

$$\begin{aligned} &\mathbb{P}_y(\rho^n \leq T(n)) \\ &\quad \leq \mathbb{P}_y\left(B_j^n = 1 \text{ for some } j \text{ with } \tau_j^n \leq T(n)\right) \\ &\quad = \mathbb{P}_y\left(B_j^n = 1 \text{ for some } j \text{ with } \left(\tau_1^n - \sigma_0^n + \sum_{i=2}^{j} \tau_i^n - \sigma_{i-1}^n + \sigma_{i-1}^n - \tau_{i-1}^n\right) \leq T(n)\right) \\ &\quad \leq \mathbb{P}_y\left(B_j^n = 1 \text{ for some } j \text{ with } \left(\sum_{i=1}^{j} \tau_i^n - \sigma_{i-1}^n\right) \leq T(n)\right) \\ &\quad \leq \mathbb{E}_y\left[\sum_{i=1}^{M^n} B_i^n\right], \end{aligned} \tag{2.2.8}$$

where $M^n \doteq \inf\left\{j \geq 1 : \sum_{i=1}^{j} \tau_i^n - \sigma_{i-1}^n > T(n)\right\}$. We can bound (2.2.8) by

$$\mathbb{E}_y\left[\sum_{i=1}^{M^n} B_i^n\right] \leq \mathbb{E}_y\left[M^n\right] \sup_{x\in\Gamma_\mu^\eta} \mathbb{E}_x\left[B_1^n\right],$$

for any $y \in \Gamma_\mu^\eta$. The proof of this statement can be found in Section 2.2.4, Lemma 2.2.18. Therefore, for any $y \in \Gamma_\mu^\eta$,

$$\mathbb{E}_y\left[\sum_{i=1}^{M^n} B_i^n\right] \leq \mathbb{E}_y[M^n]\left(\sup_{x\in\Gamma_\mu^\eta} \mathbb{E}_x[B_1^n]\right). \tag{2.2.9}$$

Step 3. *The expected number of intervals $(\sigma_{j-1}^n, \tau_j^n)$ in time $T(n)$ is approximately $T(n)$ divided by the expected duration of each interval. A lower bound on the expected duration gives an upper bound on the expected number of intervals.*

For any $y \in \Gamma_\mu^\eta$, we have

$$\begin{aligned}
\mathbb{E}_y[M^n] \inf_{x\in\Gamma_\mu^\eta} \mathbb{E}_x\left[\tau_1^n\right] &\leq \mathbb{E}_y\left[\sum_{i=1}^{M^n} \tau_i^n - \sigma_{i-1}^n\right] \\
&\leq T(n) + \tau_{M^n}^n - \sigma_{M^n-1}^n \\
&\leq T(n) + T_1.
\end{aligned}$$

See Lemma 2.2.18 for the proof of the first inequality.

It takes some positive time to travel from Γ_μ^η to either γ_μ^η or to escape $\mathcal{D}$, so $\mathbb{E}_x[\tau_1^n] \geq s > 0$ for some fixed s independent of $y \in \Gamma_\mu^\eta$ and i. A precise proof can be found in Lemma 2.2.19. Hence, for all $y \in \Gamma_\mu^\eta$,

$$\mathbb{E}_y[M^n] \leq \frac{T(n) + T_1}{s}. \tag{2.2.10}$$

Step 4. *Combine estimates for* $y \in \Gamma_\mu^\eta$.

From equations (2.2.8) and (2.2.9), we find for any $y \in \Gamma_\mu^\eta$

$$\mathbb{P}_y(\rho^n \le T(n)) \le \mathbb{E}_y[M^n] \left(\sup_{x \in \Gamma_\mu^\eta} \mathbb{E}_x[B_1^n] \right). \tag{2.2.11}$$

In (2.2.10) the first factor was bounded above. For the second factor, note that for any $y \in \Gamma_\mu^\eta$,

$$\mathbb{E}_y\left[B_1^n\right] = \mathbb{P}_y(B_1^n = 1) = \mathbb{P}_y(\{X^n(\tau_1) \notin \mathcal{D}, \tau_1 \le T_1\} \cup \{\tau_1 \ge T_1\}). \tag{2.2.12}$$

We consider the two events separately and then bound the probability of the union by the sum of the probabilities. The probability of the first event is upper bounded in (2.2.7) for $n \ge N_2$. For the second event, note that

$$\{\tau_1 \ge T_1\} \subset \{X^n(t) \in \mathcal{D} \setminus \mathcal{E}_{\mu/2}, t \in [0, T_1]\} \cup \{\left|X_{i+1}^n - X_i^n\right| \ge \eta\},$$

for if the process did enter the $\mu/2$-neighborhood of the origin but did not trigger the event τ_1, then it must have jumped over the set γ_μ^η. From Lemma 2.2.7, we can choose $N_3(\mu, \eta) < \infty$ for which $n \ge N_3$ implies

$$\mathbb{P}_y(\left|X_{i+1}^n - X_i^n\right| \ge \eta) \le e^{-n(\bar{W}+1)}. \tag{2.2.13}$$

The event $\{X^n(t) \in \mathcal{D} \setminus \mathcal{E}_{\mu/2}, t \in [0, T_1]\}$ is also superexponentially small, as the constant T_1 was originally chosen so that any $y \in \Gamma_\mu^\eta \subset (\mathcal{D} \setminus \mathcal{E}_{\mu/2})$ satisfies (2.2.6) for all $n \ge N_1$.

Set $N_4 = \max(N_1, N_2, N_3)$. Then, by (2.2.6), (2.2.7), (2.2.12) and (2.2.13), for

any $n \geq N_4$ and any $y \in \Gamma_\mu^\eta$

$$\begin{aligned}\mathbb{E}_y\left[B_1^n\right] &\leq \mathbb{P}_y(\rho^n \leq T_1) + \mathbb{P}_y(X^n(t) \in \mathcal{D} \setminus \mathcal{E}_{\mu/2}, t \in [0, T_1]) + \mathbb{P}_y(\left|X_{i+1}^n - X_i^n\right| \geq \eta) \\ &\leq e^{-n\left(\bar{W} - \frac{\varepsilon}{2}\right)} + 2e^{-n(\bar{W}+1)}.\end{aligned}$$

From (2.2.11) it follows that

$$\mathbb{P}_y(\rho^n \leq T(n)) \leq \frac{T(n) + T_1}{s}(e^{-n(\bar{W} - \frac{\varepsilon}{2})} + 2e^{-n(\bar{W}+1)}).$$

We can now take logarithms and scale by n. By using Condition 2.2.6 we can choose $N_5 \in (N_4, \infty)$, so that $n \geq N_5$ implies

$$\frac{1}{n}\log \mathbb{P}_y(\rho^n \leq T(n)) \leq -\bar{W} + \frac{\varepsilon}{2} + \frac{\varepsilon}{8}.$$

The additional error term of $\varepsilon/8$ appears to account for the polynomial factor and the additional exponential term which decays at rate $\bar{W} + 1$.

Since $|x - y| < 2\mu$ whenever $x, y \in \Gamma_\mu^\eta$, (2.2.22) ensures that $-\bar{W} \leq -W(y) + \varepsilon/8$. We conclude that whenever $n \geq N_5$ and $y \in \Gamma_\mu^\eta$, we have

$$\frac{1}{n}\log \mathbb{P}_y(\rho^n \leq T(n)) \leq -\bar{W} + \frac{\varepsilon}{2} + \frac{\varepsilon}{8} \leq -W(y) + \frac{3\varepsilon}{4}. \tag{2.2.14}$$

In particular, (2.2.4) holds for $y \in \Gamma_\mu^\eta$.

Step 5. *Extend estimate to $y \in \mathcal{E}_\mu$.*

The estimate (2.2.14), which holds for all $y \in \Gamma_\mu^\eta$, can be harnessed to obtain

(2.2.4) for all $y \in \mathcal{E}_\mu$. Let $z \in \mathcal{E}_\mu \setminus \Gamma_\mu^\eta$, and define the stopping time

$$H_\Gamma^n \doteq \inf\{t \geq 0 : X^n(t) \in \Gamma_\mu^\eta\}.$$

In order to escape $\mathcal{D}$, a trajectory starting at z must pass through Γ_μ^η, or else jump over it. In the latter case we have

$$\mathbb{P}_z(\rho^n < H_\Gamma^n) \leq \mathbb{P}_z\left(\max_{i=0,\dots,nT(n)} |X_{i+1}^n - X_i^n| \geq \eta\right).$$

Owing to Lemma 2.2.7, the probability of jumping over Γ_μ^η is superexponentially small: we can choose $N_6(\mu,\eta)$ so large that

$$\mathbb{P}_z\left(\max_{i=0,\dots,nT(n)} |X_{i+1}^n - X_i^n| \geq \eta\right) \leq e^{-n(\bar{W}+1)}, \tag{2.2.15}$$

for all $n \geq N_6$. By the law of total probability,

$$\begin{aligned}\mathbb{P}_z(\rho^n \leq T(n)) &= \mathbb{P}_z(\rho^n \leq T(n) \mid \rho^n < H_\Gamma^n)\, \mathbb{P}_z(\rho^n < H_\Gamma^n) \\ &\quad + \mathbb{P}_z(\rho^n \leq T(n) \mid \rho^n > H_\Gamma^n)\, \mathbb{P}_z(\rho^n > H_\Gamma^n) \\ &\leq e^{-n(\bar{W}+1)} + \mathbb{P}_{X^n(H_\Gamma^n)}(\rho^n \leq T(n) - H_\Gamma^n),\end{aligned} \tag{2.2.16}$$

where the last step follows from the strong Markov property applied at H_Γ^n. Since $X^n(H_\Gamma^n) \in \Gamma_\mu^\eta$ by definition of H_Γ^n, and since $H_\Gamma^n \geq 0$, it also holds that

$$\mathbb{P}_{X^n(H_\Gamma^n)}(\rho^n \leq T(n) - H_\Gamma^n) \leq \sup_{y \in \Gamma_\mu^\eta} \mathbb{P}_y(\rho^n \leq T(n)).$$

Recall that the choice of $\mu > 0$ in (2.2.22) ensures $W(z) \leq W(y) + \varepsilon/8$. Using (2.2.14) in (2.2.16), we obtain for the given $z \in \mathcal{E}_\mu \setminus \Gamma_\mu^\eta$ and $n \geq \max(N_5, N_6)$ that

$$\mathbb{P}_z(\rho^n \leq T(n)) \leq e^{-n(\bar{W}+1)} + e^{-n(W(y)-\frac{3\varepsilon}{4})} \leq 2e^{-n(W(z)-7\varepsilon/8)}.$$

By choosing n large enough and applying the logarithmic scaling, we can absorb the coefficient of 2 into the remaining $\varepsilon/8$ of room for error. This establishes the desired estimate for $z \in \mathcal{E}_\mu \setminus \Gamma_\mu^\eta$ and thus for all $y \in \mathcal{E}_\mu$. □

The next lemma extends the asymptotic upper bound to include points in $\mathcal{D}\setminus\{0\}$.

The next lemma extends the asymptotic upper bound to all points in $\mathcal{D}$. It is stated in its uniform version as required for the proof of Theorem 2.2.15. For simplicity, we do not prove the corresponding lower bound in its uniform version, though it uniformly there too.

Lemma 2.2.12. *Under the conditions of Theorem 2.2.10,*

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}_z(\rho^n \le T(n)) \le -W(z) \tag{2.2.17}$$

uniformly in $z \in \mathcal{D}$.

Proof. Given an open set U, a sequence of real numbers $\{p_n(z)\}_{z\in U}$ converges uniformly to some number p on any compact subset $K \subset U$ if and only if it converges locally uniformly on U, i.e. for any $z \in U$ and $\varepsilon > 0$, there is a neighborhood $B(z) \subset U$ of z and $N(\varepsilon, z) < \infty$ such that $|p_n(z) - p| < \varepsilon$ for all $z \in B(z)$. Since $\overline{\mathcal{D}}$ is compact, to establish the uniform convergence of the upper bound (2.2.17) on $\mathcal{D}$, it suffices to establish the locally uniform convergence on an open set U containing $\overline{\mathcal{D}}$.

Let U be any open set containing $\mathcal{D}$, and let $z \in U$ and $\varepsilon > 0$ be given. We will find $\alpha(z, \varepsilon) > 0$ and $N(z, \varepsilon) < \infty$ such that

$$\frac{1}{n} \log \mathbb{P}_y(\rho^n \le T(n)) \le -W(y) + \varepsilon \tag{2.2.18}$$

for all $n \geq N$ and $y \in B_\alpha(z) \doteq \{y : |z - y| < \delta\}$.

The particular case of $z = 0$ is covered by Lemma 2.2.11, for which one can simply take $\alpha = \mu$, so it suffices to consider $z \in U \setminus \{0\}$.

Next, we eliminate the case where $W(z) = 0$. Since W is continuous on $\mathbb{R}^d$, there is $\alpha(z, \varepsilon) > 0$ such that $0 < -W(y) + \varepsilon$ for all $y \in B_\alpha(z)$. Since $\mathbb{P}_y(\rho^n \leq T(n)) \leq 1$, it automatically follows that (2.2.18) holds for all $y \in B_\alpha(z)$, as required. Furthermore, we observe that $W(z) = 0$ for any $z \notin \mathcal{D}$, and so we may focus on points $z \in D \setminus \{0\}$.

Let then $z \in \mathcal{D} \setminus \{0\}$ with $W(z) > 0$, and let $\varepsilon > 0$ be given. Without any loss of generality, we may assume that $0 < \varepsilon < W(z)$. By Lemma 2.2.11, there is $\mu_0(\varepsilon) > 0$ and $N_1(\varepsilon) < \infty$ such that $n \geq N_1$ implies, for all $y \in \mathcal{E}_{\mu_0} = \{x : |x| < \mu_0\}$, that

$$\frac{1}{n} \log \mathbb{P}_y(\rho^n \leq T(n)) \leq -W(y) + \frac{\varepsilon}{4}. \tag{2.2.19}$$

We claim that we can choose $\mu(z, \varepsilon) \in (0, \mu_0)$ and $\alpha(z, \varepsilon)$ such that: (1) $B_\alpha(z) \subset \mathcal{D} \setminus \mathcal{E}_\mu$ and (2) we have the inequality

$$W(y) \leq \inf_{x \in \mathcal{E}_\mu} W(x) + \frac{\varepsilon}{2} \tag{2.2.20}$$

for all $y \in B_\alpha(z)$.

We first choose μ. By part 2 of Lemma 2.2.8 we have $W(z) \leq W(0)$, and by the first part of the same lemma W is continuous at 0. Thus we can choose $\mu > 0$ such that $W(0) < W(x) + \varepsilon/4$ for all $x \in \mathcal{E}_\mu$, and without loss of generality $\mu < \min(|z|/2, \mu_0)$. In particular,

$$W(z) \leq \inf_{x \in \mathcal{E}_\mu} W(x) + \frac{\varepsilon}{4}. \tag{2.2.21}$$

Next, since $\mu < |z|/2$ we can take $\alpha(z, \varepsilon, \mu) > 0$ small enough that $B_\alpha(z) \cap \mathcal{E}_\mu = \emptyset$ and moreover we may insist that $B_\alpha(z) \subset \mathcal{D}$ because $z \in \mathcal{D}$ and $\mathcal{D}$ is open. Since W is continuous at z, α can be taken even smaller to guarantee

$$|W(y_1) - W(y_2)| < \varepsilon/4 \tag{2.2.22}$$

for all $y_1, y_2 \in B_\alpha(z)$. Note that μ depends only on z and ε, so α inherently depends only on z and ε as well. Combined with (2.2.21), an application of the triangle inequality shows that this choice of α guarantees (2.2.20).

Define

$$\zeta_\mu^n \doteq \inf\{t > 0 : X^n(t) \notin \mathcal{D} \setminus \mathcal{E}_\mu\},$$

which is the first time the process enters the μ-neighborhood of the attracting point 0 or escapes the set $\mathcal{D}$. By Lemma 2.2.9, there are $T_1(\mu) < \infty$ and $N_2(\mu) < \infty$ so large that $n \geq N_2$ implies, for any $x \notin \mathcal{E}_\mu$

$$\mathbb{P}_x(\zeta_\mu^n > T_1) \leq e^{-n(W(x)+1)}. \tag{2.2.23}$$

Then, for any $y \in B_\alpha(z)$,

$$\begin{aligned}\mathbb{P}_y(\rho^n \leq T(n)) = \mathbb{P}_y(\rho^n \leq T_1) + \mathbb{P}_y(T_1 < \rho^n \leq T(n), \zeta_\mu^n \leq T_1) \\ + \mathbb{P}_y(T_1 < \rho^n \leq T(n), \zeta_\mu^n > T_1).\end{aligned} \tag{2.2.24}$$

We estimate each term separately. For the first term, we use the finite time uniform large deviations upper bound, part 3 of Definition 2.2.3. Let $\bar{W} \doteq \inf_{y \in B_\alpha(z)} W(y)$. Since the closure of $B_\alpha(z)$ is compact, an argument similar to the one in Step 1 of Lemma 2.2.11 shows that for the given $\varepsilon > 0$ there is $\delta(\varepsilon) > 0$ such

that, for all $y \in B_\alpha(z)$,

$$\{\rho^n \le T_1, X^n(0) = y\} \subset \{d(X^n, \Phi_{y,T_1}(\bar{W} - \varepsilon/8)) \ge \delta\}.$$

Applying the large deviations upper bound on the compact set $\overline{B_\alpha(z)}$ with $s_0 = \bar{W} - \varepsilon/8$, $\gamma = \varepsilon/8$ and $\delta > 0$ as above, we find $N_3(\varepsilon, T_1) < \infty$ for which $n \ge N_3$ implies, for all $y \in B_\alpha(z)$,

$$\mathbb{P}_y(\rho^n \le T_1) \le e^{-n(\bar{W}-\varepsilon/4)}.$$

By the choice of α in (2.2.22) we have $W(y) \le \bar{W} + \varepsilon/4$, and so

$$\mathbb{P}_y(\rho^n \le T_1) \le e^{-n(\bar{W}-\varepsilon/4)} \le e^{-n(W(y)-\varepsilon/2)} \tag{2.2.25}$$

for all $y \in B_\alpha(z)$.

For the second term, note that on the event $\{T_1 < \rho^n \le T(n)\} \cap \{\zeta_\mu^n \le T_1\}$, the process has entered the μ-neighborhood of 0, which is contained in the μ_0-neighborhood of 0 by the choice of μ. After entering $\mathcal{E}_\mu$ at time ζ_μ^n, it has $T(n) - \zeta_\mu^n$ time remaining to exit $\mathcal{D}$. By allowing $T(n)$ time we increase the probability of exiting. Thus, by (2.2.19) and the strong Markov property, $n \ge N_1$ implies

$$\begin{aligned} \mathbb{P}_y(T_1 < \rho^n \le T(n), \zeta_\mu^n < T_1) &\le \sup_{x \in \mathcal{E}_\mu} \mathbb{P}_x(\rho^n \le T(n)) \\ &\le e^{-n(\inf_{x \in \mathcal{E}_\mu} W(x) - \varepsilon/4)}. \end{aligned} \tag{2.2.26}$$

For the last term in (2.2.24), we use the estimate (2.2.23),

$$\mathbb{P}_y(T_1 < \rho^n \le T(n), \zeta_\mu^n > T_1) \le \mathbb{P}_y(\zeta_\mu^n > T_1) \le e^{-n(W(y)+1)}, \tag{2.2.27}$$

which holds for all $y \in B_\alpha(z) \subset \mathcal{D} \setminus \mathcal{E}_\mu$ whenever $n \geq N_2$.

Set $N_4 \doteq \max(N_1, N_2, N_3)$. Using (2.2.25), (2.2.26) and (2.2.27) in (2.2.24), we find for all $n \geq N_4$ and $y \in B_\alpha(z)$

$$\mathbb{P}_y(\rho^n \leq T(n)) \leq e^{-n(W(y)-\varepsilon/2)} + e^{-n(\inf_{x\in\mathcal{E}_\mu} W(x)-\varepsilon/4)} + e^{-n(W(y)+1)}. \tag{2.2.28}$$

From (2.2.20), for any $y \in B_\alpha(z)$,

$$e^{-n(\inf_{x\in\mathcal{E}_\mu} W(x)-\varepsilon/4)} \leq e^{-n(W(y)-3\varepsilon/4)},$$

for all $n \geq N_4$. It therefore follows from (2.2.28) that, for all $y \in B_\alpha(z)$,

$$\frac{1}{n}\log \mathbb{P}_y(\rho^n \leq T(n)) \leq -W(y) + \frac{3\varepsilon}{4} + \frac{3}{n},$$

and choosing $N_4 \leq N(z,\varepsilon) < \infty$ large enough that $3/n < \varepsilon/4$, we obtain (2.2.18) for all $y \in B_\alpha(z)$, as required. □

Next we prove the lower bound for the decay of $\mathbb{P}_x(\rho^n \leq T(n))$.

Lemma 2.2.13. *Assume the conditions of Theorem 2.2.10. For any $x \in \mathcal{D}$, we have*

$$\liminf_{n\to\infty} \frac{1}{n}\log \mathbb{P}_x(\rho^n \leq T(n)) \geq -W(x).$$

Proof. It suffices to show that for any $x \in \mathcal{D}$ and any $\varepsilon > 0$, there is $N(\varepsilon) < \infty$ such that $n \geq N$ guarantees

$$\frac{1}{n}\log \mathbb{P}_x(\rho^n \leq T(n)) \geq -W(x) - \varepsilon.$$

To establish the estimate, we will take a trajectory which is within ε of the infimum in the definition of $W(x)$, and use part 4 of Condition 2.2.2 in conjunction with part 2 of Condition 2.2.5 to extend the trajectory a little past the boundary. This will allow us to apply the large deviations lower bound, which is stated in its uniform version as part 2 of Definition 2.2.3.

Fix $x \in \mathcal{D}$ and let $\varepsilon > 0$. By the variational definition of $W(x)$, there is a trajectory φ and $T < \infty$ such that $\varphi(0) = x$, $\varphi(T) \notin \mathcal{D}$ and $I_T(\varphi) \leq W(x) + \varepsilon/4$. Define the compact set $K = \{y : \inf_{z \in \mathcal{D}} |y - z| \leq 1\}$. By Lemma 2.2.17, there is $c > 0$ such that for all $p, q \in K$, we can find φ_{pq} and τ which satisfy $\varphi_{pq}(0) = p$, $\varphi_{pq}(\tau) = q$, $\tau \in (0, \infty)$ and $I_\tau(\varphi_{pq}) \leq c|p - q|$. Let $a \in (0, \varepsilon/4c)$ and assume without loss of generality that $a < 1$.

$I_T(\varphi) < \infty$ implies φ must be absolutely continuous, so there is a time $s \in (0, T]$ such that $\varphi(s) \in \partial\mathcal{D}$. Set $p \doteq \varphi(s)$ and note that $p \in K$. By the regularity property in part 2 of Condition 2.2.5, there is a point q in the interior of $\mathcal{D}^c$ which satisfies $|p - q| < a$, and since $a < 1$ we also have $q \in K$. From Lemma 2.2.17 we obtain a trajectory φ_{pq} and $\tau > 0$ as described above with cost

$$I_\tau(\varphi_{pq}) \leq ca < \varepsilon/4.$$

Let ϕ denote the concatenation of the paths φ on $[0, s]$ and φ_{pq} on $(s, s + \tau]$, that is,

$$\phi(t) = \begin{cases} \varphi(t) & t \in [0, s], \\ \varphi_{pq}(t - s) & t \in (s, s + \tau], \end{cases}$$

and note that ϕ is continuous at s. Moreover,

$$I_{s+\tau}(\phi) \leq I_T(\varphi) + I_\tau(\varphi_{pq}) \leq W(x) + \varepsilon/2.$$

Finally, since q lies in the interior of $\mathcal{D}^c$, there is $\delta > 0$ such that $|q - z| < \delta$ implies $z \in \mathcal{D}^c$. In particular, if $\|\phi' - \phi\|_\infty < \delta$, then $\phi'(s+\tau) \in \mathcal{D}^c$. If n is large enough to guarantee $T(n) \geq s + \tau$, then

$$\{\|X^n - \phi\|_\infty < \delta\} \subset \{\rho^n \leq T(n)\}.$$

Applying part 2 of Definition 2.2.3 with $s_0 = W(x) + 1$, $\gamma = \varepsilon/2$, and the given $\delta > 0$, we find $N < \infty$ such that for all $n \geq N$,

$$\mathbb{P}_x(\|X^n - \phi\|_\infty < \delta) \geq \exp(-n(I_T(\phi) + \varepsilon/2)).$$

Taking logarithms and using $I_T(\phi) \leq W(x) + \varepsilon/2$, we obtain the desired result. □

We can now combine the upper and lower bounds obtained above to deduce Theorem 2.2.10.

Proof of Theorem 2.2.10. Lemma 2.2.12 guarantees that for any $x \in \mathcal{D}$

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}_x(\rho^n \leq T(n)) \leq -W(x),$$

while Lemma 2.2.13 guarantees that for the same x

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}_x(\rho^n \leq T(n)) \geq -W(x).$$

Thus the limit exists and (2.2.3) in Theorem 2.2.10 holds. □

2.2.3 Analysis of asymptotic performance

The main result of this section is to establish a lower bound on the decay rate of the RESTART estimator for the problem of escape from a domain $\mathcal{D}$ over increasing time intervals $[0, T(n)]$. The lower bound is stated in terms of the value of the subsolution $\bar{V}$ at the origin. Recalling that $W(0)$ is the minimal cost to exit $\mathcal{D}$ starting at 0, we show that $\bar{V}(0) = W(0)$ ensures asymptotic optimality in the sense of the following definition.

Definition 2.2.14. If the estimator γ^n satisfies

$$\liminf_{n\to\infty} -\frac{1}{n} \log \mathcal{S}(\gamma^n) = 2W(0), \tag{2.2.29}$$

we say that γ^n is asymptotically optimal.

Note that this definition is similar to the one in Definition 2.1.2, but has the important difference that A is not a subset of $D([0, T])$, but rather a sequence $\{A_n\}$ of subsets in $D([0, T(n)])$.

When the subsolution is defined in terms of the quasipotential as in Lemma 1.5.5 using the constant $c_Q = W(0)$, the RESTART estimator is asymptotically optimal. An upper bound on the decay rate is automatic from Jensen's inequality, as discussed at the beginning of Section 2.1.5.

It would also be possible to establish upper and lower bounds on the decay rate in terms of a variational problem, and with such a formulation one could obtain a limit

on the asymptotic decay rate (the upper and lower bounds would agree). Instead of formulating the analysis in terms of escape from $\mathcal{D}$, one would have to consider the problem of moving from any point $x \in \mathcal{D}$ to any other point $y \in \mathcal{D}$. While limits for the asymptotic decay rate (as opposed to just lower bounds) are useful for the comparison of estimators, the main goal of this section is to demonstrate that the decay rate of splitting algorithms does not degrade over long time intervals, and this is accomplished with only lower bounds.

Before stating and proving the main result, we first point out that the increasing time intervals $[0, T(n)]$ can indeed ruin asymptotic optimality when using importance sampling instead of splitting. It has been shown in [16], where stochastic differential equation models are considered, that when importance sampling is based on a time-independent subsolution, the following lower bound on the second moment holds:

$$\mathcal{S}^n(\bar{V}) \geq e^{C_1(T-K)} e^{-nC_2}.$$

Here T is the time interval for escape, C_1 and C_2 are positive constants, and $K < T$ is a fixed constant. If T were n dependent, e.g., $T(n) = n^2$, then

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathcal{S}^n(\bar{V}) = \infty.$$

Not only is the estimator not asymptotically efficient, the decay rate of this importance sampling scheme is worse than ordinary Monte Carlo simulation.

For one-dimensional stochastic differential equations, it is possible to construct time-dependent subsolutions which ensure that the importance sampling schemes do not degrade over long time intervals. These time-dependent subsolutions are difficult to construct and it is not known how to construct them in any generality for higher

dimensions.

We show that these problems do not arise with splitting in Theorem 2.2.15. As discussed in [16], the lower decay rate is due to the exponential growth of the likelihood ratio when the trajectories stay near the origin, which it can do with large enough probability over a long time interval. In contrast, such trajectories do not affect the decay rate of splitting schemes since they are not multiplied by a likelihood ratio; all that matters is that on average, out of R particles born in a threshold, one of them reaches a threshold with a lower index.

The analysis is carried out for the RESTART scheme, but for ordinary splitting the result will be the same and the analysis will be easier. By using techniques similar to the ones in [19] it is straightforward to get estimates on the second moment.

The main result is stated in the following theorem. It places a lower bound on the asymptotic decay rate in terms of the value of the subsolution at the origin.

Theorem 2.2.15. *Suppose Conditions 2.2.1, 2.2.2, 2.2.4, 2.2.5 and 2.2.6 are met. Let $\bar{V}$ be a subsolution in the sense of Definition 1.4.6, and let $\mathcal{S}^n(\bar{V})$ denote the second moment of the estimator γ^n when $\bar{V}$ is used to determine the piecewise constant importance functions $\bar{U}^n$. We have the following lower bound:*

$$\liminf_{n\to\infty} -\frac{1}{n}\log \mathcal{S}^n(\bar{V}) \geq \bar{V}(0) + W(0).$$

Proof. We first consider a general subsolution $\bar{V}$ used to construct the importance

function $\bar{U}^n(x)$ defined by (2.1.7). From Corollary 2.1.4,

$$\mathfrak{S}^n(\bar{V}) \le \mathbb{E}_0\left[\sum_{i=1}^{nT(n)} e^{n(\bar{U}^n(X^n_{i-1})-\bar{U}^n(0))}[\mathbb{P}_{X^n_i}(\rho^n \le T(n) - i/n)]^2\right] + \mathbb{E}_0\left[e^{n(\bar{U}^n(X^n(\rho^n))-\bar{U}^n(0)))}1\{\rho^n \le T(n)\}\right]. \tag{2.2.30}$$

There are two terms in (2.2.30), and both are positive and n-dependent. If two sequences of positive numbers $\{a^1_n\}$ and $\{a^2_n\}$ satisfy

$$\liminf_{n\to\infty} -n^{-1}\log(a^i_n) \ge a,$$

then

$$\liminf_{n\to\infty} -n^{-1}\log(a^1_n + a^2_n) \ge a. \tag{2.2.31}$$

It therefore suffices to check the limit on both terms individually. We start with the second, easier term.

Let $\varepsilon > 0$ be arbitrary. From (2.1.8) it follows that for all n large enough we have

$$\bar{U}^n(x) - \bar{U}^n(0) \le \bar{V}(x) - \bar{V}(0) + \varepsilon.$$

If $x \notin \mathcal{D}$ then $\bar{V}(x) \le 0$ by the boundary condition for the subsolution property in Definition **??**. Since $X^n(\rho^n) \notin \mathcal{D}$, for all n sufficiently large we have

$$\mathbb{E}_0\left[e^{n(\bar{U}^n(X^n(\rho^n))-\bar{U}^n(0))}1\{\rho^n \le T(n)\}\right] \le e^{n(-\bar{V}(0)+\varepsilon)}\mathbb{P}_0\left(\rho^n \le T(n)\right).$$

Using the last display and Theorem 2.2.10 we find

$$\liminf_{n\to\infty} -\frac{1}{n}\log \mathbb{E}_0\left[e^{n(\bar{U}^n(X^n(\rho^n))-\bar{U}^n(0))}1\{\rho^n \le T(n) - i/n\}\right] \ge \bar{V}(0) + W(0) - \varepsilon,$$

which handles the second term. We remark that the second term contributes to the overall rate of decay even though it might seem exponentially negligible compared to the first term in (2.2.30). This is because we expect the highest contributions to the variance to come from the particles which split closest to the escape set, which is what the term represents.

Next we handle the first term. Again, let $\varepsilon > 0$. Owing to the polynomial growth of $T(n)$ and the logarithmic scaling in the statement of the theorem, it suffices to bound each term in the sum (similarly to (2.2.31)). Observe that for any $i \in \{0, \dots, nT(n)\}$,

$$\begin{aligned}
&\mathbb{E}_0\left[e^{n(\bar{U}^n(X^n_{i-1})-\bar{U}^n(0))}\mathbb{P}_{X^n_i}(\rho^n \le T(n) - i/n)^2\right] \\
&= \sum_{j=0}^{J^n} \mathbb{E}_0\left[1_{\{X^n_i \in D^n_j\}} e^{n(\bar{U}^n(X^n_{i-1})-\bar{U}^n(0))}\mathbb{P}_{X^n_i}(\rho^n \le T(n) - i/n)^2\right].
\end{aligned}$$

Since J^n grows linearly with n, owing to the logarithmic scaling – similarly to (2.2.31) – we may estimate the expected value individually on each set D^n_j. To do this, we first resolve the time index disparity between the indicator $1_{\{X^n_i \in D^n_j\}}$ and the point X^n_{i-1} appearing in the argument of $\bar{U}^n$. We will show that the probability of X^n_i and X^n_{i-1} not lying in the same set D^n_j is negligible as $n \to \infty$. Define

$$\delta^n_i \doteq [\bar{U}^n(X^n_{i-1}) - \bar{U}^n(0)] - [\bar{U}^n(X^n_i) - \bar{U}^n(0)] = \bar{U}^n(X^n_{i-1}) - \bar{U}^n(X^n_i).$$

Since $\mathcal{D}$ is bounded and $\bar{U}^n$ is piecewise constant on $\mathcal{D}$,

$$B \doteq \sup_{x,y\in\mathcal{D}} |\bar{U}^n(x) - \bar{U}^n(y)| < \infty. \tag{2.2.32}$$

For any $\varepsilon_1 > 0$ and each $i \in \{0, \ldots, nT(n)\}$, we have

$$\begin{aligned}
&\mathbb{E}_0\left[1_{\{X_i^n \in D_j^n\}} e^{n(\bar{U}^n(X_{i-1}^n) - \bar{U}^n(0))} \mathbb{P}_{X_i^n}(\rho^n \le T(n) - i/n)^2\right] \\
&\quad = \mathbb{E}_0\left[1_{\{X_i^n \in D_j^n\}} e^{n(\bar{U}^n(X_i^n) - \bar{U}^n(0))} e^{n\delta_i^n} \mathbb{P}_{X_i^n}(\rho^n \le T(n) - i/n)^2\right] \\
&\quad \le e^{n\varepsilon_1} \mathbb{E}_0\left[1_{\{X_i^n \in D_j^n\}} e^{n(\bar{U}^n(X_i^n) - \bar{U}^n(0))} \mathbb{P}_{X_i^n}(\rho^n \le T(n) - i/n)^2\right] \\
&\qquad + \mathbb{E}_0\left[e^{n(B+\delta_i^n)} 1\{|\delta_i^n| \ge \varepsilon_1\}\right].
\end{aligned} \tag{2.2.33}$$

Inequality (2.1.8) implies that $\bar{U}^n(0) - \bar{U}^n(x) \to \bar{V}(0) - \bar{V}(x)$ uniformly in x, so for the given $\varepsilon_1 > 0$ there is $N_1(\varepsilon_1) < \infty$ such that $n \ge N_1$ and $|\delta_i^n| \ge \varepsilon_1$ implies

$$|\bar{V}(X_{i-1}^n) - \bar{V}(X_i^n)| \ge \varepsilon_1/2.$$

By the uniform continuity of $\bar{V}$ on $\overline{\mathcal{D}}$, there is $\eta(\varepsilon_1) > 0$ such that $|\bar{V}(z) - \bar{V}(y)| \ge \varepsilon_1/2$ implies $|z - y| \ge \eta$. Thus for the given $\varepsilon_1 > 0$ we have $\eta > 0$ such that

$$\{|\delta_i^n| \ge \varepsilon_1\} \subset \{|X_{i-1}^n - X_i^n| > \eta\}.$$

Lemma 2.2.7 implies that for any $\eta > 0$, we can find $N_2(\eta) < \infty$ large enough that, for all $i \in \{0, ..., nT(n)\}$ and $n \ge N_2$,

$$\mathbb{P}_x(|X_i^n - X_{i-1}^n| \ge \eta) \le e^{-n(2B + \bar{V}(0) + W(0) + 1)}.$$

From equation (2.2.32) we have $|\delta_i^n| \leq B$, so the last term in (2.2.33) satisfies

$$\begin{aligned}\mathbb{E}_0[e^{n(\delta_i^n+B)}1\{|\delta_i^n| \geq \varepsilon_1\}] &\leq e^{n2B}e^{-n(2B+\bar{V}(0)+W(0)+1)} \\ &= e^{-n(\bar{V}(0)+W(0)+1)}, \end{aligned} \tag{2.2.34}$$

whenever $n \geq N_3 = \max(N_1, N_2)$. For large enough n, this exponential decay is strictly greater than the claimed decay of $\bar{V}(0) + W(0)$, hence it will become clear later that this term can be ignored when we establish the lower bound on the decay rate of the second moment of the estimator.

For fixed i and j, the third term in (2.2.33) satisfies

$$\begin{aligned} & e^{n\varepsilon_1}\mathbb{E}_0\left[1_{\{X_i^n \in D_j^n\}}e^{n(\bar{U}^n(X_i^n)-\bar{U}^n(0))}\mathbb{P}_{X_i^n}(\rho^n \leq T(n) - i/n)^2\right] \\ & \leq e^{n\varepsilon_1}\sup_{z\in D_j^n}\mathbb{P}_z(\rho^n \leq T(n)) \\ & \quad \times \mathbb{E}_0\left[1_{\{X_i^n \in D_j^n\}}e^{n(\bar{U}_j^n-\bar{U}^n(0))}\mathbb{P}_{X_i^n}(\rho^n \leq T(n) - i/n)\right]. \end{aligned} \tag{2.2.35}$$

In the inequality above, we used the shorthand notation $\bar{U}_j^n = \bar{U}^n(x)$ for $x \in D_j$, and we enlarged the set $\{\rho^n \leq T(n) - i/n\}$ to $\{\rho^n \leq T(n)\}$. Since we are interested in taking limits as $n \to \infty$, it suffices to demonstrate the upper bound (or lower bound on the rate) for any sequence $\{j_n\}_{n\geq 1}$ such that, as $n \to \infty$,

$$j_n\Delta_n \to \theta \in [0, M],$$

where we recall from (2.1.6) that $M = \max_{x\in\bar{\mathcal{D}}}\{\bar{V}(0) - \bar{V}(x)\}$ and $J^n\Delta_n \to M$.

The limits for j_n and J^n determine a limit for $U_{j_n}^n - U^n(0)$ as follows. Let

$$L(\theta) \doteq \{x : \bar{V}(0) - \bar{V}(x) = M - \theta\}.$$

If $x_n \in D^n_{j_n} = C^n_{j_n} \setminus C^n_{j_n-1}$ for all $n \geq 1$, then

$$x_n \in \{x : (J_n - j_n + 1)\Delta_n > \bar{V}(0) - \bar{V}(x) \geq (J_n - j_n)\Delta_n\}.$$

By passing to a convergent subsequence we may assume without loss of generality that $x_n \to \bar{x}$. Then $\bar{x} \in L(\theta)$ and $\bar{V}(x_n) - \bar{V}(0) \to \theta - M$. By (2.1.8) we have for the given $\varepsilon > 0$ that

$$\bar{U}^n_{j_n} - \bar{U}^n(0) \leq \theta - M + \varepsilon/2,$$

for all sufficiently large n. This will allow us to estimate the $\bar{U}^n_j - \bar{U}^n(0)$ term in (2.2.35) when j is replaced by j_n. Furthermore, the supremum over $z \in D^n_{j_n}$ in (2.2.35) will be replaced in the limit by one over $z \in L(\theta)$. Finally, the Markov property for $\{X^n_i\}$ implies that

$$\mathbb{E}_0[1_{\{X^n_i \in D^n_{j_n}\}} \mathbb{P}_{X^n_i}(\rho^n \leq T(n) - i/n)] \leq \mathbb{P}_0(\rho^n \leq T(n)).$$

Thus an upper bound on (2.2.35) with $j = j_n$ is

$$e^{n(\theta - M + \varepsilon_1 + \varepsilon/2)} \sup_{z \in D^n_{j_n}} \mathbb{P}_z(\rho^n \leq T(n))\ \mathbb{P}_0(\rho^n \leq T(n)), \tag{2.2.36}$$

for all sufficiently large n.

To summarize, (2.2.36) provides an upper bound on the first term in (2.2.33) via (2.2.35), and the second term in (2.2.33) will turn out to be exponentially negligible compared to the first term owing to the estimate (2.2.34). It remains to determine the asymptotic behavior of (2.2.36) when the large deviations scaling is applied.

We first show that

$$\liminf_{n\to\infty} -\frac{1}{n}\log\left(\sup_{z\in D^n_{j_n}} \mathbb{P}_z(\rho^n \le T(n))\right) \ge \inf_{z\in L(\theta)} W(z). \tag{2.2.37}$$

The limit inferior of a sequence of real numbers can be defined as the infimum over all its convergent subsequences, so it suffices to establish the inequality in (2.2.37) along any convergent subsequence $\{n_k\}_{k\ge1}$. Let $\{n_k\}_{k\ge1}$ denote such a subsequence, and let $\varepsilon_2 > 0$ be given. By Lemma 2.2.12, there is $N_4(\varepsilon_2) < \infty$ such that, for all $z \in \overline{\mathcal{D}}$ and $n \ge N_4$,

$$-\frac{1}{n}\log \mathbb{P}_z(\rho^n \le T(n)) \ge W(z) - \varepsilon_2. \tag{2.2.38}$$

Let $K_1(N_4) < \infty$ be such that $n_k \ge N_4$ for all $k \ge K_1$. For each $k \ge 1$, choose $z_{n_k} \in D^{n_k}_{j_{n_k}}$ such that

$$\sup_{z\in D^{n_k}_{j_{n_k}}} \frac{1}{n}\log \mathbb{P}_z(\rho^n \le T(n)) < \frac{1}{n}\log \mathbb{P}_{z_{n_k}}(\rho^n \le T(n)) + \varepsilon_2. \tag{2.2.39}$$

By passing to a further subsequence if necessary, we may assume without any loss that $z_{n_k} \to \bar{z} \in L(\theta)$. If $z_{n_k} \in \overline{\mathcal{D}}$, (2.2.38) applies with $z = z_{n_k}$ and $n = n_k$ for all $k \ge K_1$, while for $z_{n_k} \notin \overline{\mathcal{D}}$ we have $W(z_{n_k}) = 0$ and the same estimate holds for any $k \ge 1$. Moreover, since W is continuous and $z_{n_k} \to \bar{z}$, there is $K_2(\varepsilon_2) < \infty$ such that $W(z_{n_k}) > W(\bar{z}) - \varepsilon_2$ for all $k \ge K_2$. Combining this with (2.2.39) and using monotonicity of the logarithm, we find, for all $k \ge \max(K_1, K_2)$,

$$\begin{aligned}
-\frac{1}{n_k}\log\left(\sup_{z\in D_{n_k}} \mathbb{P}_z(\rho^{n_k} \le T(n_k))\right) &> -\frac{1}{n_k}\log \mathbb{P}_{z_{n_k}}(\rho^{n_k} \le T(n_k)) - \varepsilon_2 \\
&\ge W(z_{n_k}) - 2\varepsilon_2 \\
&> W(\bar{z}) - 3\varepsilon_2.
\end{aligned}$$

Since $\bar{z} \in L(\theta)$, we automatically have $W(\bar{z}) \geq \inf_{z \in L(\theta)} W(z)$. Since $\{n_k\}_{k \geq 1}$ and $\varepsilon_2 > 0$ were arbitrary, we conclude that (2.2.37) holds.

We now return to (2.2.36). We first take logarithms, scale by $-n$, and take the limit inferior. Using superadditivity of the limit inferior, (2.2.37), and another application of Lemma 2.2.12 for the probability starting at 0, we obtain

$$\begin{aligned}\liminf_{n\to\infty} -\frac{1}{n} \log \left(e^{n(\theta - M + \varepsilon_1 + \varepsilon/2)} \sup_{z \in D_{j_n}^n} \mathbb{P}_z(\rho^n \leq T(n)) \, \mathbb{P}_0(\rho^n \leq T(n)) \right) \\ \geq M - \theta - \varepsilon_1 - \varepsilon/2 + \inf_{z \in L(\theta)} W(z) + W(0).\end{aligned}$$

To conclude the proof, we use the subsolution property of $\bar{V}$, which implies that $\bar{V}(x) \leq W(x)$ for all x. For $z \in L(\theta)$, $\bar{V}(z) = \bar{V}(0) - M + \theta$, so

$$\inf_{z \in L(\theta)} W(z) \geq \bar{V}(0) - M + \theta.$$

Thus,

$$M - \theta - \varepsilon_1 - \varepsilon/2 + \inf_{z \in L(\theta)} W(z) + W(0) \geq \bar{V}(0) + W(0) - \varepsilon_1 - \varepsilon/2.$$

Since both ε and ε_1 were arbitrary, we send them to zero and arrive at the desired result for a general subsolution $\bar{V}$. □

Corollary 2.2.16. *If the subsolution is defined through $\bar{V}(x) = -Q(x) + c_Q$, then $\bar{V}(0) = c_Q = W(0)$ and the estimator is asymptotically optimal.*

Proof. In the case where $\bar{V}(x) = -Q(x) + c_Q$, we recall that Lemma 1.5.5 shows $-Q(x) + c_Q$ is indeed a subsolution. From Theorem 2.2.15, we have the lower bound

$$\liminf_{n\to\infty} -\frac{1}{n} \log \mathcal{S}^n(\bar{V}) \geq c_Q + W(0),$$

while Jensen's inequality implies the corresponding upper bound. Equation (2.1.10) implies $c_Q = W(0)$, which yields asymptotic optimality in the sense of Definition 2.2.14. □

2.2.4 Proofs of various Lemmas

Proof of Lemma 2.2.7. Fix $\delta > 0$ and note that by a union bound,

$$\mathbb{P}_x \left(\max_{0 \le i \le nT(n)} |X_{i+1}^n - X_i^n| \ge \delta \right) \le nT(n) \max_{0 \le i \le nT(n)} \mathbb{P}_x \left(|X_{i+1}^n - X_i^n| \ge \delta \right),$$

and since $|\cdot|$ denotes the Euclidean norm on $\mathbb{R}^d$, we also have for each $0 \le i \le nT(n)$ and standard unit vector $e_k \in \mathbb{R}^d$, $k = 1, \ldots, d$,

$$\mathbb{P}_x \left(|X_{i+1}^n - X_i^n| \ge \delta \right) \le d \max_{1 \le k \le d} \mathbb{P}_x \left(|\langle X_{i+1}^n - X_i^n, e_k \rangle| \ge \delta / \sqrt{d} \right).$$

Since d is fixed and $n^{-1} \log(nT(n)) \to 0$ as $n \to \infty$ by Condition 2.2.6, it suffices to establish that for each fixed $1 \le k \le d$, and any $\delta > 0$,

$$\limsup_{n \to \infty} \sup_{x \in \mathbb{R}^d} \max_{0 \le i \le nT(n)} \frac{1}{n} \log \mathbb{P}_x(|\langle X_{i+1}^n - X_i^n, e_k \rangle| \ge \delta) = -\infty.$$

Recall that $X_{i+1}^n - X_i^n = n^{-1} v_i(X_i^n)$, where the $v_i(\cdot)$ are i.i.d. state-dependent vector fields with stochastic kernel $\mu(\cdot|x)$ for $x \in \mathbb{R}^d$. Thus

$$\max_{0 \le i \le nT(n)} \mathbb{P}_x(|\langle X_{i+1}^n - X_i^n, e_k \rangle| \ge \delta) \le \sup_{x \in \mathbb{R}^d} \int_{\mathbb{R}^d} 1_{\{|\langle y, e_k \rangle| \ge n\delta\}} \mu(dy|x).$$

For any $x \in \mathbb{R}^d$, Chernoff's inequality implies that for any $\alpha_k > 0$,

$$\begin{aligned}\int_{\mathbb{R}^d} 1_{\{|\langle y,e_k\rangle| \geq n\delta\}} \mu(dy|x) &= \int_{\mathbb{R}^d} \left(1_{\{\langle y,e_k\rangle \geq n\delta\}} + 1_{\{-\langle y,e_k\rangle \geq n\delta\}}\right) \mu(dy|x) \\ &\leq \int_{\mathbb{R}^d} \left(e^{\alpha_k(\langle y,e_k\rangle - n\delta)} + e^{\alpha_k(-\langle y,e_k\rangle - n\delta)}\right) \mu(dy|x) \\ &= e^{-n\delta\alpha_k}(e^{H(x,\alpha')} + e^{H(x,-\alpha')}),\end{aligned}$$

where $\alpha' = \alpha_k e_k$. Let $f(\alpha) = \sup_{x\in\mathbb{R}^d} e^{H(x,\alpha)}$, and recall from Condition 2.2.1 that $f(\alpha) < \infty$ for all $\alpha \in \mathbb{R}^d$. We conclude from the above display that for every $\alpha_k > 0$ with $\alpha' = \alpha_k e_k$,

$$\max_{0\leq i\leq nT(n)} \mathbb{P}_x(|\langle X_{i+1}^n - X_i^n, e_k\rangle| \geq \delta) \leq e^{-n\delta\alpha_k}(f(\alpha') + f(-\alpha')).$$

We will now show that, after taking logarithms and scaling by n, the left-hand side in the above display can be made arbitrarily negative for sufficiently large n. Let $M < \infty$ be given and choose $\alpha_k = (M+1)\delta^{-1}$. Choose $N < \infty$ such that $n \geq N$ implies $n > \log(f(\alpha') + f(-\alpha'))$, and note that monotonicity of the logarithm allows us to commute the maximum over $0 \leq i \leq nT(n)$ to obtain

$$\begin{aligned}\max_{0\leq i\leq nT(n)} \frac{1}{n} \log \mathbb{P}_x(|\langle X_{i+1}^n - X_i^n, e_k\rangle| \geq \delta) &\leq -M - 1 + \frac{\log(f(\alpha') + f(-\alpha'))}{n} \\ &\leq -M.\end{aligned}$$

Since the upper bound is independent of x, holds for all $n \geq N$, and $M < \infty$ was arbitrary, the proof is complete. □

Lemma 2.2.17. *Under Condition 2.2.2, for any compact set K there is $c \in (0,\infty)$ such that for any x and $y \in K$, there exists $\tau \in (0,\infty)$ and a smooth function φ, with $\varphi(0) = x$, $\varphi(\tau) = y$, $\tau = |x-y|/r$ for which $I_\tau(\varphi) < c\,|x-y|$.*

Proof of Lemma 2.2.17. Let the compact set K be given, and let $r \in (0, \infty)$ be as in property 4 of Condition 2.2.2 Take $\tau = |x - y| / r$. Then

$$\varphi(t) = x + \frac{y - x}{|y - x|} rt.$$

Then $\varphi(0) = 0$, $\varphi(\tau) = y$ and $L(\varphi(t), \dot{\varphi}(t)) \leq r^{-1}$ for all $t \in [0, T]$. Thus,

$$\int_0^\tau L(\varphi(t), \dot{\varphi}(t)) dt \leq \frac{\tau}{r} = c\,|x - y|\,,$$

where $c = 1/r^2$. □

Proof of Lemma 2.2.8. First we prove that W is continuous at any point $x \in \mathcal{D}$. By Lemma 2.2.17 applied to the compact set $\bar{\mathcal{D}}$, there is $c \in (0, \infty)$ such that for any $x, y \in \mathcal{D}$ satisfying $|x - y| < \varepsilon/2c$, the cost of going from x to y in some finite time $\tau \in (0, \infty)$ is less than $\varepsilon/2$. By concatenating this path from x to y with a near-minimizing one from y to some point not in $\mathcal{D}$, and using additivity of I_T on disjoint time segments, we obtain $|W(x) - W(y)| < \varepsilon$.

If $x \in \overline{\mathcal{D}} \setminus \mathcal{D}$ then $W(x) = 0$, and the above proof can be repeated on a neighborhood N of x, separating into the cases $y \in N \cap \mathcal{D}$ and $y \notin N \cap \mathcal{D}$.

Next we show that for any $\varepsilon > 0$ and any $x \in \overline{\mathcal{D}}$, $W(x) \leq W(0) + \varepsilon$, which implies that $W(x) \leq W(0)$ for all $x \in \mathcal{D}$. Note that we might have $W(x) = 0$, if the noiseless dynamics starting at x pass through the complement of $\mathcal{D}$ before reaching any sufficiently small neighborhood of the origin, in which case $W(x) \leq W(0)$ automatically since $W(0) \geq 0$. Otherwise, since W is continuous at 0, there is $\delta > 0$ such that any $y \in \mathcal{E}_\delta = \{x : |x| < \delta\}$ satisfies $W(y) < W(0) + \varepsilon$. Since 0 is the unique attracting fixed point, for any $x \in \mathcal{D}$ there is $T > 0$ and a trajectory

$\varphi \in C([0,T])$ with $\varphi(0) = x$ such that $I_T(\varphi) = 0$ and $\varphi(T) \in \mathcal{E}_\delta$. Thus,

$$W(x) \leq I_T(\varphi) + W(\varphi(T)) < W(0) + \varepsilon,$$

as claimed. □

Lemma 2.2.18. *For all* $y \in \Gamma_\mu^\eta$,

$$\mathbb{E}_y \left[\sum_{i=1}^{M^n} B_i^n \right] \leq \mathbb{E}_y \left[M^n \right] \sup_{x \in \Gamma_\mu^\eta} \mathbb{E}_x \left[B_1^n \right],$$

and

$$\mathbb{E}_y \left[\sum_{i=1}^{M^n} \tau_i^n - \sigma_{i-1}^n \right] \geq \mathbb{E}_y \left[M^n \right] \inf_{x \in \Gamma_\mu^\eta} \mathbb{E}_x \left[\tau_1^n \right].$$

Proof of Lemma 2.2.18. Consider

$$S_k^n = \sum_{i=1}^{k} \tau_i^n - \sigma_{i-1}^n.$$

Let $\mathcal{F}_i^n = \sigma(X_j^n, 0 \leq j \leq i)$, $\mathcal{F} = \sigma(\cup_{i=1}^\infty \mathcal{F}_i^n)$ and $\mathcal{G}_k^n = \mathcal{F}_{\sigma_k^n}^n$, where

$$\mathcal{F}_{\sigma_k^n}^n = \{A \in \mathcal{F} : A \cap \{\sigma_k^n \leq m\} \in \mathcal{F}_m^n, m \geq 1\},$$

is the sigma-algebra of the stopping time σ_k^n. Since $\sigma_i^n \geq \tau_i^n$, S_k^n is $\mathcal{G}_k^n$-measurable. Let

$$M^n \doteq \inf\{k \geq 1 : S_k^n > T(n)\}.$$

Then M^n is a stopping time with respect to the filtration generated by S_k^n and hence

$\mathcal{G}_k^n$. In particular, $\{M^n \geq k\} = \{M^n \leq k-1\}^c$ is $\mathcal{G}_{k-1}^n$-measurable. We have

$$\begin{aligned}
\mathbb{E}_y\left[\sum_{i=1}^{M^n} \mathbb{E}_{X^n_{\sigma^n_{i-1}}} B_1^n\right] &= \mathbb{E}_y\left[\sum_{i=1}^{M^n} \mathbb{E}_y[B_i^n | X^n_{\sigma^n_{i-1}}]\right] \\
&= \mathbb{E}_y\left[\sum_{i=1}^{M^n} \mathbb{E}_y[B_i^n | \mathcal{G}^n_{i-1}]\right] \\
&= \sum_{i=1}^{\infty} \mathbb{E}_y[1_{\{M^n \geq i\}} \mathbb{E}_y[B_i^n | \mathcal{G}^n_{i-1}]] \\
&= \sum_{i=1}^{\infty} \mathbb{E}_y[1_{\{M^n \geq i\}} B_i^n] \\
&= \mathbb{E}_y\left[\sum_{i=1}^{M^n} B_i^n\right].
\end{aligned}$$

In the above display, the first two equalities follow from the strong Markov property conditioned on $X^n_{\sigma^n_{i-1}}$, and pulling out the infinite sum from the expectation sign is permitted by Tonelli's theorem because all terms in the sum are positive. It immediately follows that

$$\mathbb{E}_y\left[\sum_{i=1}^{M^n} B_i^n\right] \leq \mathbb{E}_y M^n (\sup_{x \in \Gamma_\mu^\eta} \mathbb{E}_x B_1^n).$$

Similarly, we have

$$\mathbb{E}_y\left[\sum_{i=1}^{M^n} \tau_i^n - \sigma_{i-1}^n\right] = \mathbb{E}_y\left[\sum_{i=1}^{M^n} \mathbb{E}_{X^n_{\sigma^n_{i-1}}} \tau_1^n\right],$$

and so

$$\mathbb{E}_y\left[\sum_{i=1}^{M^n} \tau_i^n - \sigma_{i-1}^n\right] \geq \mathbb{E}_y M^n \inf_{x \in \Gamma_\mu^\eta} \mathbb{E}_x \tau_1.$$

□

Lemma 2.2.19. *There exist $s > 0$ and $N < \infty$ such that for all $n \geq N$ and any $y \in \Gamma_\mu^\eta$, we have $\mathbb{P}_y(\tau_1^n \leq s) \leq \frac{1}{2}$.*

Proof. We may assume that $s \leq 1$ without any loss, since the probability is decreased by decreasing s. The probability $\mathbb{P}_y(\tau_1^n \leq s)$ is the probability that X^n exits $\mathcal{D}$ or enters γ_μ^η before time s. If instead of the piecewise constant interpolation of X^n we had taken the piecewise continuous one, it would not be possible for X^n to jump over the set γ_μ^η and the probability $\mathbb{P}_y(\tau_1^n \leq s)$ would be increased. Consequently, it suffices to establish the desired inequality when X^n is the piecewise continuous interpolation. For the purposes of this proof it is convenient to do so, and by an abuse of notation we use X^n to denote this process, so that $X^n \in C([0,1])$ $\mathbb{P}_y$-almost surely for the remainder of the proof.

Let $\Delta = \min(\mu/2 - \eta, \operatorname{dist}(\mathcal{E}_\mu, \partial\mathcal{D})) > 0$ where $\mathcal{E}_\mu = \{x : |x| < \mu\}$. Any trajectory which starts in Γ_μ^η and exits $\mathcal{D}$ or enters γ_μ^η must travel a distance of at least $\Delta > 0$ at some point in time. Let

$$A_s = \{\varphi \in C([0,s]) : |\varphi(t) - \varphi(0)| \geq \Delta \text{ for some } t \in [0,s], \varphi(0) \in \Gamma_\mu^\eta\}.$$

Then $\{\tau_1^n \leq s\} \subset \{X^n \in A_s\}$. To get an upper bound on the probability of $\{X^n \in A_s\}$ we place a lower bound on the cost of any $\varphi \in A_s$. To do this, we are free to choose $s > 0$ as small as we like. Moreover, it suffices to establish the lower bound for φ which are absolutely continuous, as otherwise $I_s(\varphi) = +\infty$.

To bound the cost from below, first let $M \geq 8/\Delta$. By Lemma 2.2.2 applied on the compact set $\bar{\mathcal{D}}$, there is $\beta_M < \infty$ such that for any β satisfying $|\beta| \geq \beta_M$, we have $L(x,\beta) \geq M|\beta|$ for all $x \in \bar{\mathcal{D}}$. This will allow us to place a lower bound on the cost $I_s(\varphi)$ via the magnitude of the derivative $|\dot{\varphi}|$ for absolutely continuous φ.

To use the estimate $L(\varphi, \dot{\varphi}) \geq M|\dot{\varphi}|$, we must place a lower bound on the Lebesgue measure of the set $\{|\dot{\varphi}| \geq \beta_M\}$, which we view as a subset of $[0,s]$. For

any $s > 0$ and $\varphi \in A_s \subset C([0,s])$, there is $t \in [0,s]$ such that $|\varphi(t) - \varphi(0)| \geq \Delta$. By the fundamental theorem of calculus,

$$\begin{aligned} \Delta &\leq \int_{\{|\dot{\varphi}| \geq \beta_M\}} |\dot{\varphi}| \mathrm{dt} + \int_{\{|\dot{\varphi}| < \beta_M\}} |\dot{\varphi}| \mathrm{dt} \\ &\leq \int_{\{|\dot{\varphi}| \geq \beta_M\}} |\dot{\varphi}| \mathrm{dt} + s\beta_M. \end{aligned}$$

Choosing $s < \Delta/2\beta_M$ we find

$$\int_{\{|\dot{\varphi}| \geq \beta_M\}} |\dot{\varphi}| \mathrm{dt} \geq \Delta/2.$$

so for such $s \in (0, \Delta/2\beta_M)$, any $\varphi \in A_s$ satisfies

$$I_s(\varphi) \geq \int_{\{|\dot{\varphi}| \geq \beta_M\}} L(\varphi, \dot{\varphi}) \mathrm{dt} \geq \int_{\{|\dot{\varphi}| \geq \beta_M\}} M|\dot{\varphi}| \mathrm{dt} \geq M\frac{\Delta}{2} = 4.$$

We claim there are $s \in (0, \Delta/2\beta_M)$ and $\delta > 0$ such that $d(\varphi, \Phi_{\varphi(0),s}(2)) > \delta$ for all $\varphi \in A_s$. Assume without loss of generality that $\Delta/2\beta_M < 1$ and let

$$F \doteq \cup_{x \in \overline{\Gamma_\mu^\eta}} \Phi_{x,1}(2).$$

Since the closure of the set Γ_μ^η is compact, the uniform LDP implies that F is compact in $C([0,1])$; see part 1 of Definition 2.2.3. For $\phi \in C([0,1])$, define the modulus of continuity

$$\omega(\phi, \sigma) \doteq \sup_{|t_2 - t_1| \leq \sigma} |\phi(t_2) - \phi(t_1)|, \quad 0 \leq t_1 \leq t_2 \leq 1.$$

The compactness of $F \subset C([0,1])$ implies there is a uniform bound on the modulus of continuity for functions $\phi \in F$. Let

$$\bar{\omega}(\sigma) \doteq \sup_{\phi \in F} \omega(\phi, \sigma),$$

and note that $\bar{\omega}(\sigma) \to 0$ as $\sigma \to 0$. Choose $\sigma > 0$ small enough that $\bar{\omega}(\sigma) \leq \Delta/2$, and for such $\sigma > 0$, choose $s = \min(\sigma, \Delta/2\beta_M)$. Note that $s \in (0, 1)$.

For this fixed s, consider any $\varphi \in A_s$. We claim that $d(\varphi, \Phi_{\varphi(0),s}(2)) \geq \Delta/4$. If not, there is $\phi \in \Phi_{\varphi(0),s}(2)$ for which $d(\varphi, \phi) < \Delta/4$. The function $\phi \in C([0, s])$ can be naturally identified with a function in $F \subset C([0, 1])$ by concatenating it with the zero-cost trajectory $\dot{\phi} = b(\phi)$ on $(s, 1]$; see part 2 of Condition 2.2.2. In particular, $I_s(\phi) \leq 2$ implies that $\omega(\phi, s) \leq \Delta/2$. Therefore, for all $t \in [0, s]$,

$$|\varphi(t) - \varphi(0)| \leq |\varphi(t) - \phi(t)| + |\phi(t) - \varphi(0)| \leq \frac{\Delta}{4} + \frac{\Delta}{2},$$

where we used the fact that $\phi(0) = \varphi(0)$ and $\omega(\phi, s) \leq \Delta/2$ to bound the second term. It follows that $|\varphi(t) - \varphi(0)| \leq 3\Delta/4$ for all $t \in [0, s]$, which contradicts $\varphi \in A_s$. We conclude that the choice $\delta = \Delta/4$ indeed guarantees $d(\varphi, \Phi_{\varphi(0),s}(2)) > \delta$ for all $\varphi \in A_s$.

Finally, the large deviations upper bound with $\gamma = 1$, $s_0 = 2$ and the given $\delta > 0$ implies there exists N such that for all $n \geq N$,

$$\mathbb{P}_x(d(X^n, \Phi_{\varphi(0),s}(2)) > \delta) \leq e^{-n}.$$

Since $\{\tau_1 \leq s\} \subset \{X^n \in A_s\} \subset \{d(X^n, \Phi_{\varphi(0),s}(2)) > \delta\}$, the proof is complete.

□

2.3 Moderate deviations-based splitting

In this section we investigate splitting schemes based on moderate deviations asymptotics for $\mathbb{R}^d$-valued discrete-time processes of the form

$$X_{i+1}^n = X_i^n + \frac{1}{n} b(X_i^n) + \frac{1}{n} u_i(X_i^n), \quad X_0^n = x_0$$

where b is a "nice" drift term and u_i are mean-zero, independent and identically distributed random state-dependent vector fields. We denote by $X^n(t)$ the linear interpolated process which satisfies $X^n(i/n) = X_i^n$ (see (2.3.3)). The goal will is to estimate probabilities of the form

$$\mathbb{P}(X^n(T) \in D). \tag{2.3.1}$$

We call these *terminal escape probabilities*.

Under conditions on X^n stronger than the ones assumed in this section, **true?** large deviations theory provides insight into how to construct numerical algorithms for the estimation of probabilities in (2.3.1); see [19][18]. The two most common approaches are splitting and importance sampling, both of which require the design of a function which will determine the performance of the algorithm. For splitting this function determines the splitting thresholds, while for importance sampling it determines the state-dependent change of measure. In both cases, the function can be characterised as a subsolution to a suitable partial differential equation (PDE) called the Hamilton-Jacobi-Bellman (HJB) equation.

In many cases the HJB equation and its associated control problem can be intractable or impossible to solve explicitly, and even the construction of subsolutions

can be difficult, especially in a high-dimensional setup. It is natural to ask whether it is valid to approximate the control problem and design a scheme based on the approximated solution.

It turns out that this is the case for both importance sampling and splitting, but it is mildly surprising how one justifies this. As the parameter n tends to infinity, the probabilities in (2.3.1) decay exponentially fast and one needs to consider whether the scheme is even usable. For events which are "rare, but not too rare", one can embed the sequence X^n into an amplified sequence Y^n which satisfies a large deviations principle and for which the event of interest is truly rare. This is proved in [47] and is the main result on which the present section is based. The amplified process Y^n has a Gaussian-style rate function, so the associated control problem is the classical linear-quadratic regulator and explicit solutions are available. The cost one obtains for Y^n is exactly what one would obtain by linearizing around the noiseless dynamics for X^n and making a quadratic approximation of the cost, but this can only be justified after one has gone through the trouble of placing the target probability in a large deviations framework.

2.3.1 Process model

In this section we specify the process model more precisely and place conditions on the relevent objects. We use the same assumptions as in [47] and follow their notation closely. The process model is

$$X_{i+1}^n = X_i^n + \frac{1}{n} b(X_i^n) + \frac{1}{n} u_i(X_i^n), \quad X_0^n = x_0,$$

where $u_i(\cdot)$ are independent and identically distributed vector fields with state-dependent distribution $\mu(\cdot|x)$, that is, $\mathbb{P}(u_i(x) \in A) = \mu(A|x)$ for all Borel sets $A \subset \mathbb{R}^d$. We assume that the vector fields have mean zero, i.e. $\mathbb{E}u_i(x) = 0$ for all $x \in \mathbb{R}^d$.

Define the centered logarithmic moment generating function

$$H_c(x, \alpha) \doteq \log \left(\int_{\mathbb{R}^d} \exp(\langle \alpha, u \rangle) \mu(\mathrm{d}u|x) \right), \quad \alpha \in \mathbb{R}^d.$$

The subscript c serves to distinguish H_c from the usual uncentered log moment-generating function $H(x, \alpha) = H_c(x, \alpha) + \langle \alpha, b(x) \rangle$. We will impose the following conditions on μ and b:

Condition 2.3.1. *We have the following conditions:*

1. *there exists $\lambda > 0$ and $K_{\mathrm{mgf}} < \infty$ such that*

$$\sup_{x \in \mathbb{R}^d} \sup_{\|\alpha\| \leqslant \lambda} H_c(x, \alpha) \leqslant K_{\mathrm{mgf}};$$

2. *$x \to \mu(\cdot|x)$ is continuous with respect to the topology of weak convergence;*

3. *$x \to b(x)$ is continuously differentiable, and both $b(x)$ and its derivative $Db(x)$ are uniformly bounded by some constant $K_b < \infty$.*

Define the matrix A component-wise by

$$A_{ij}(x) \doteq \int_{\mathbb{R}^d} u_i u_j \mu(\mathrm{d}u|x),$$

and observe that it can be derived from the centered log moment generating function,

$$\frac{\partial H_c(x,0)}{\partial \alpha_i} = \int_{\mathbb{R}^d} u_i \mu(\mathrm{du}|x) = 0, \quad \frac{\partial^2 H_c(x,0)}{\partial \alpha_i \partial \alpha_j} = \int_{\mathbb{R}^d} u_i u_j \mu(\mathrm{du}|x) = A_{ij}(x). \quad (2.3.2)$$

Note that the matrix A_{ij} is nonnegative definite and symmetric.

It may turn out that the matrix A_{ij} is singular, but we will still need to refer to the inverse of A as well as its square root $A^{1/2}$. We copy Definition 2.1. of [47] below with some minor modifications to specify how the inverse should be interpreted.

Definition 2.3.2. For a symmetric, nonnegative definite matrix A we can write $A = Q\Lambda Q^T$, where Q is an orthogonal matrix whose columns are the eigenvectors of A and Λ is the diagonal matrix consisting of the nonnegative eigenvalues of A. Throughout the paper we use the condition that $A^{1/2} = Q\Lambda^{1/2}Q^T$, where $\Lambda^{1/2}$ is the diagonal matrix consisting of the positive square roots of the nonnegative eigenvalues of A. In addition, let $\|\alpha\|_A^2 = \langle Q\alpha, \Lambda Q\alpha\rangle$, and let $\|\alpha\|_{A^{-1}}^2 = \langle Q\alpha, \Lambda^{-1}Q\alpha\rangle$, where Λ^{-1} is a diagonal matrix consisting of the inverse of the eigenvalues for the positive eigenvalues and ∞ for the zero eigenvalues. Consequently, $\langle Q\alpha, \Lambda^{-1}Q\alpha\rangle$ has a well-defined and finite value for α in the span of the positive eigenvalues and is equal to ∞ outside of that linear span.

Define the continuous-time linear interpolation of X_i^n for all $i \geq 0$ by

$$X^n(t) = (i + 1 - nt)X_i^n + (nt - i)X_{i+1}^n, \quad t \in [i/n, i/n + 1/n). \quad (2.3.3)$$

The noiseless process $X_i^{n,0}$ defined for $i \geq 0$ by

$$X_{i+1}^{n,0} = X_i^{n,0} + \frac{1}{n} b(X_i^{n,0}), \quad X_0^{n,0} = x_0, \quad (2.3.4)$$

and let $X^{n,0}(t)$ denote the analogously defined continuous-time linear interpolation of $X^{n,0}$. It is clear that $X^{n,0}(\cdot) \to X^0(\cdot)$ in $C([0,T] : \mathbb{R}^d)$ equipped with the supremum norm, where

$$X^0(t) = x_0 + \int_0^t b(X^0(s))ds.$$

Moreover, since $\mathbb{E}u_i(x) = 0$ for all $x \in \mathbb{R}^d$, we have $X^n \to X^0$ in probability in $C([0,T] : \mathbb{R}^d)$ ([46], Theorem 3.3.1). Under stronger assumptions than in Condition 2.3.1, including the assumption that Condition 2.3.1.1 holds with $\lambda = \infty$, X^n satisfies a large deviations principle with speed $r(n) = n$ and rate function defined for absolutely continuous functions by

$$I_L(\varphi) \doteq \inf\left\{\int_0^T L_c(\varphi(s), u(s))ds : \varphi(t) = x_0 + \int_0^t b(\varphi(s)) + u(s)ds, t \in [0,T]\right\},$$

and $I_L(\varphi) = +\infty$ if φ is not absolutely continuous. Here L_c is the Fenchel-Legendre transform of H_c,

$$L_c(x,\beta) \doteq \sup_{\alpha \in \mathbb{R}^d}\{\langle \alpha, \beta\rangle - H_c(x,\alpha)\}.$$

Under appropriate assumptions on the centered local rate L_c, for a suitably regular closed set $D \subset \mathbb{R}^d$ we have the asymptotic approximation

$$\lim_{n\to\infty} -\frac{1}{n}\log \mathbb{P}(X^n(T) \notin D) = \inf\{I_L(\varphi) : \varphi(0) = x_0, \varphi(T) \notin D\}. \tag{2.3.5}$$

As mentioned in the introduction, the control problem which appears in (2.3.5) and determines the rate of decay of the target probability can often be impossible to solve explicitly. In such a case, one might attempt a diffusion approximation of the process, which would result in a simpler (quadratic) control problem. However, after performing such an approximation, the limit in (2.3.5) is no longer valid, since in general L_c need not be quadratic. The moderate deviations scaling provides a rescaled process for which the quadratic expansion of the cost is valid and an

analogue of (2.3.5) holds with a modified rate function. The scaling and what is meant by the MDP approximation can be a little confusing, so the next section is dedicated to stating exactly how the two approximation problems relate to each other, and one how "embeds" this problem into another.

2.3.2 Problem embedding

In this section we will consider the escape problem for X^n in terms of a scaled and centred process Y^n. Recall our convention for the scaling and translation of sets: for any set $B \subset \mathbb{R}^d$, constant $c > 0$ and $z \in \mathbb{R}^d$,

$$cB + z = \{cb + z : b \in B\}. \tag{2.3.6}$$

In order to obtain a Gaussian-style rate function, we will have to amplify the noise in the process X^n. Without an amplification of the noise, events which fall in the large deviations regime will generally have much more rigid constraints on how they must occur and lead to the control problem (2.3.5). (**fair?**). To separate the noise from the driving dynamics, we first shift X^n by its noiseless version $X^{n,0}$; see equation (2.3.4). Using the notation in (2.3.6), the terminal escape problem can then be recast as estimating

$$\mathbb{P}(X^n(T) - X^{n,0}(T) \in D - X^{n,0}(T)). \tag{2.3.7}$$

We place the following assumptions on the sequence $\{a_n\}$ which will be used to scale the process, and which we will sometimes refer to as the *embedding sequence*. The embedding sequence must go to 0, but not too fast and not too slow.

Condition 2.3.3. *The embedding sequence a_n satisfies the following: $a_n > 0$,*

$$1.\ \lim_{n\to\infty} \sqrt{n} a_n = \infty, \quad 2.\ \lim_{n\to\infty} \sqrt{\log n} a_n = 0.$$

The assumption that a_n does not decay too fast is necessary for the moderate deviations scaling in [47]. The assumption that a_n does not decay too slow is used for the asymptotic optimality results. We will see later this imposes no restriction on the splitting algorithm discussed in Section 2.1.3.

For any sequence a_n satisfying Condition 2.3.3, define the scaled, centered process Y^n by

$$Y_i^n = \sqrt{n} a_n (X_i^n - X_i^{n,0}),$$

and let $Y^n(t)$ denote the linearly interpolated process (as in (2.3.3)). Note that

$$Y_{i+1}^n = Y_i^n = \frac{a_n}{\sqrt{n}}(b(X_i^n) - b(X_i^{n,0})) + \frac{a_n}{\sqrt{n}} u_i(X_i^n), \quad Y_0^n = 0.$$

Let $D_x b(x)$ denote the matrix whose $(i,j)^{\text{th}}$ entry is

$$(D_x b(x))_{i,j} = \frac{\partial b_i(x)}{\partial x_j}.$$

In [47], the authors show that Y^n satisfies a large deviations principle on $C([0,T] : \mathbb{R}^d)$ with speed $r(n) = a_n^{-2}$ and rate function

$$I_M(\varphi) \doteq \inf \left\{ \int_0^T \|u(t)\|^2 dt : \varphi(t) = \int_0^t D_x b(X^0(s))\varphi(s) ds + \int_0^t A^{1/2}(X^0(s)) u(s) ds, t \in [0,T] \right\}, \tag{2.3.8}$$

where $A^{1/2}$ is specified in Definition 2.3.2. This is known as a *moderate deviations*

principle for X^n. Observe that in the definition of I_M, the cost incurred is quadratic in the control u. This makes the control problem much simpler to solve explicitly.

To relate the process Y^n back to the original problem of terminal escape proabilities, we must first fix a particular value of n for which to estimate the terminal escape probability (2.3.7). For clarity, we distinguish this fixed value of n by labelling it n_*. The above probability can be rewritten in terms of the Y^n process as

$$\mathbb{P}(Y^{n_*}(T) \notin \sqrt{a_{n_*} n_*}(X^n(T) - X^{n,0}(T)) = \mathbb{P}(Y^{n_*}(T) \notin D_*),$$

where $D_* \doteq \sqrt{n_*} a_{n_*} D$. Notice that D_* does not depend on n; it is a fixed escape set for the Y^n process. Under mild assumptions on the set D (which carry over to the set D_*), the large deviations principle for Y^n identifies the decay rate of the probabilities:

$$\lim_{n\to\infty} -a_n^2 \log \mathbb{P}(Y^n(T) \notin D_*) = \inf_{\varphi:\varphi(0)=0,\varphi(T)\notin D_*} I_M(\varphi). \tag{2.3.9}$$

The hope is the value of n_* is large enough that the above asymptotic approximation for Y^n is valid. We note that the same issue arises when one uses the large deviations approximation for the design of rare event simulation algorithms, and in those cases the approach has reasonable success.

Another issue to address is the implicit dependence of the approximation on the choice of embedding sequence a_n. It is useful to remember that the embedding sequence serves as a way of zooming in on a neighbourhood of the path $X^0(t)$ and studying its fluctuations further into the tails than a diffusion approximation. However, the limiting behavior (that is, the rate function I_M) is independent of the choice of embedding sequence, and therefore the information provided through the large de-

viations principle for Y^n should have no inherent dependence on a_n. We demonstrate in Section 2.3.4 that the choice of a_n will not affect the design of splitting algorithms, just as it did not affect the design of the importance sampling scheme in [47].

2.3.3 Subsolutions to the HJB

This section identifies subsolutions associated to terminal escape problems for the Y^n process. Recall that the exponential rate of decay of the terminal escape probability $\mathbb{P}(Y^n(T) \notin D_*)$ is identified by solving the underlying control problem in (2.3.9). As pointed out in the introduction, one of the advantages of the moderate deviations scaling is that one is able to solve some of control problems explicitly. We first write down the value of the *point-to-point cost*. Then, in order to produce a subsolution that is computationally tractable, we impose the terminal constraint that $z \notin D_*$ by including a terminal cost which is affine in the spatial variable.

The point-to-point cost is defined as follows: for $(t, y), (T, z) \in [0, T) \times \mathbb{R}^d$, let

$$C(t, y, T, z) \doteq \inf \left\{ \frac{1}{2} \int_t^T \|u(s)\|^2 ds : \varphi(T) = z, \text{ and for } s \in [t, T], \right. \tag{2.3.10}$$
$$\left. \varphi(s) = y + \int_t^s D_x b(X^0(q))\varphi(q)dq + A^{1/2}(X^0(q))u(q)dq \right\}.$$

To express the solution, we first need a control-theoretic definition.

Definition 2.3.4. The *state transition function* associated to the control problem (2.3.10) is the matrix-valued function $\Psi(s, t)$ which satisfies $\Psi(t, t) = I$ for all $t \in [0, T]$, where I is the d-dimensional identity matrix, and

$$\partial_t \Psi(s, t) = D_x b(X^0(t))\Psi(s, t)$$

for all $0 \leqslant s < t$.

The solution of the variational problem (2.3.10) can now be expressed using the state transition function.

Theorem 2.3.5. *The variational problem (2.3.10) admits the following explicit solution:*

$$C(t,y,T,z) = \frac{1}{2}\left\|z - \Psi(t,T)y\right\|^2_{M^{-1}_{t,T}}, \tag{2.3.11}$$
$$M_{t,T} \doteq \int_t^T \Psi(q,T)A(X^{(0)}(q))\Psi(q,T)^T dq.$$

While the explicit solution (2.3.11) is readily computable, to identify the value of the control problem (2.3.9) one needs to further infimize $C(t,y,T,z)$ over $z \notin D_*$. This is not in general available in closed form and can be computationally expensive. As mentioned, we will instead solve an alternative control problem with an affine terminal constraint. To this end, let G be a lower-semicontinuous function that is bounded below and define

$$V(t,y) \doteq \inf\left\{\frac{1}{2}\int_t^T \|u(s)\|^2 ds + G(\varphi(T)) : \text{ for } s \in [t,T], \right. \tag{2.3.12}$$
$$\left. \varphi(s) = y + \int_t^s D_x b(X^0(q))\varphi(q)dq + A^{1/2}(X^0(q))u(q)dq \right\}.$$

The control problem in (2.3.9) can then be recovered by taking $G(y) = \infty \cdot 1_{D_*}(y)$; since D_* is open, then G is lower-semicontinuous. Moreover, V is a weak-sense solution to the HJB

$$\partial_t V(t,y) = \frac{1}{2}\|D_y V(t,y)\|^2_{A(X^0(t))} - \langle D_y V(t,y), D_x b(X^0(t))y\rangle \tag{2.3.13}$$

with $V(T,y) = G(y)$. Of course, for $G(y) = \infty \cdot 1_{D_*}(y)$, V can also be expressed as

$$V(t,y) = \inf_{z \notin D_*} C(t,y,T,z).$$

Recall from Definition 1.4.3 that a weak sense subsolution $\bar{V}$ to the same equation satisfies

$$\partial_t \bar{V}(t,y) \geq \frac{1}{2}\|D_y \bar{V}(t,y)\|^2_{A(X^0(t))} - \langle D_y \bar{V}(t,y), D_x b(X^0(t))y \rangle \tag{2.3.14}$$

and $\bar{V}(T,y) \leq G(y)$. It turns out that for a specific form of G, namely

$$G^{\xi,c}(y) = \langle \xi, y \rangle + c, \tag{2.3.15}$$

it is possible to construct an explicit weak-sense *solution* $V^{\xi,c}$ to the HJB PDE (2.3.13). Moreover, it is clear from the variational definition of a time-dependent subsolution (see Definition 1.4.3) that the pointwise minimum of two subsolutions $\bar{V}_1$ and $\bar{V}_2$ with corresponding terminal costs G_1 and G_2 is also a subsolution with terminal cost $\min(G_1(y), G_2(y))$. Putting these facts together, we can therefore construct subsolutions to (2.3.13) as follows:

1. Approximate G from below by the pointwise minimum of affine functions $\bar{G}_1, \ldots, \bar{G}_K$ of the form (2.3.15), so that $\bar{G}(y) = \min_{i=1,\ldots,K} \bar{G}_i(y) \leq G(y)$;
2. Construct explicit *solutions* V_i to (2.3.13) with terminal cost G_i; each V_i is a subsolution to the same equation with terminal cost G since $G_i \leq G$;
3. Take the pointwise minimum of the V_i to form the subsolution $\bar{V}(t,y) = \min_{i=1,\ldots,K} V_i(t,y)$.

The hardest part of this construction is the approximation of G from below, which must be done on a case-by-case basis. The goal is to guarantee asymptotic optimality in the sense of Definition 2.1.2, which from Theorem 2.3.7 is achieved if $\bar{V}(0,0) = V(0,0)$. More details are provided in Section 3.2.3.2, and Section 3.3 discusses explicit constructions for specific examples.

We end the section with the explicit solution to (2.3.13) with an affine terminal condition of the form (2.3.15); see for instance Theorem 4.1 in [47] for a proof.

Theorem 2.3.6. *Given $c \in \mathbb{R}$ and $\xi \in \mathbb{R}^d$, the function $V^{\xi,c}$ defined by*

$$V^{\xi,c}(t,y) = \langle \xi, \Psi(t,T)y \rangle + c - \frac{1}{2}\left\langle \xi, \int_t^T \Psi(s,T)A(X^0(s))\Psi(s,T)^T ds \right\rangle \tag{2.3.16}$$

is a classical sense solution to (2.3.13).

Remark. With the appropriate scaling, one can rewrite the solution to the Hamilton-Jacobi equation in terms of the x variable using the relationship

$$y^n(t,x) \doteq \sqrt{n}a_n(x - X^{n,0}(t)).$$

For any $n \geq 1$, $x \in \mathbb{R}^d$, $t \in [0,T]$, let

$$W(t,x) = \frac{1}{na_n^2}V\left(t, \sqrt{n}a_n(x - X^{n,0}(t))\right) = \frac{1}{na_n^2}V(t, y^n(t,x)). \tag{2.3.17}$$

Then

$$\sqrt{n}a_n D_x W(t,x) = D_y V(t, y^n(t,x)), \quad \frac{dy^n(t,x)}{dx} = \sqrt{n}a_n x,$$

and

$$\partial_t W(t,x) = \frac{1}{na_n^2}\Big(\frac{1}{2}\|D_y V(t,y^n(t,x))\|^2_{A(X^0(t))} - \langle D_y V(t,y^n(t,x)), D_x b(X^0(t))y^n(t,x)\rangle\Big) = \frac{1}{2}\|D_x W(t,x)\|^2_{A(X^0(t))} - \langle D_y W(t,x), D_x b(X^0(t))x\rangle. \tag{2.3.18}$$

so W satisfies a HJB of its own, with the appropriate modification of boundary conditions.

2.3.4 Invariance of embedding

In this section, we show that the splitting algorithm is independent of the choice of embedding sequence $\{a_n\}$. This means that once a selection $n = n_*$ has been made, the (x,t) values which fall in threshold C_j^n defined in (2.1.12) do not depend on $\{a_n\}$.

For concreteness, we consider the problem of estimating

$$\mathbb{P}(X^n(T) \notin D), \quad X^n(0) = x_0 \tag{2.3.19}$$

for some fixed $n = n_*$ and fixed $D \subset \mathbb{R}^d$. To demonstrate the invariance of embedding, we will write down the solution $V(t,y)$ to the HJB equation (2.3.13) in terms of x using $y(x,t) = \sqrt{n_*}a_{n_*}(x - X^{n_*,0}(t))$. In doing so, we find that the thresholds determined by the solution V in (x,t) coordinates do not depend on a_n, and only on n_*.

From Theorem 2.3.5, the optimal cost for the scaled process Y^n to move from y

at time t to z at time T is

$$C(t,y,T,z) = \frac{1}{2}\|z - \Psi(t,T)y\|^2_{M^{-1}_{t,T}},$$

where we recall that the matrix $M_{t,T}$ defined in equation (2.3.11) does not depend on either y or z.

For the given n_*, set $D_* = \sqrt{n_*}a_{n_*}(D - X^{n,0}(T))$. The solution $V(t,y)$ of the control problem associated to (2.3.19) is given by

$$V(t,y) = \inf\{C(t,y,T,z) : z \notin D_*\}.$$

From (2.1.12) and the relation $\Delta_n = \Delta/r(n) = \Delta a_n^2$, the thresholds are defined as

$$C_j^n = \{(t,y) : V(0,0) - V(t,y) \geq (J^n - j)\Delta a_n^2\}. \tag{2.3.20}$$

For any (t,y) there is x such that

$$y = \sqrt{n_*}a_{n_*}(x - X^{n_*,0}(t)),$$

and any $z \in D_*$ can be written as $z = \sqrt{n_*}a_{n_*}(w - X^{n_*,0}(T))$ for $w \in D$. Thus,

$$\begin{aligned} C(t,y,T,z) &= \frac{1}{2}\|z - \Psi(t,T)y\|^2_{M^{-1}_{t,T}} \\ &= \frac{1}{2}\left\|\sqrt{n_*}a_{n_*}\left[w - X^{n_*,0} - \Psi(t,T)(x - X^{n_*,0})\right]\right\|^2_{M^{-1}_{t,T}} \\ &= n_* a_{n_*}^2 C(t, x - X^{n_*,0}(t), T, w - X^{n_*,0}). \end{aligned}$$

It follows that the solution $V(t,y)$ can be rewritten with a scaling of $n_* a_{n_*}^2$ in front:

$$V(t,y) = n_* a_{n_*}^2 \inf\left\{C(t, x - X^{n_*,0}(t), T, w - X^{n_*,0}) : w \notin D, y = \sqrt{n_*}a_{n_*}(x - X^{n_*,0}(t))\right\}.$$

Setting

$$W(t,x) = \inf_{w \in D} C(t, x - X^{n_*,0}(t), T, w - X^{n_*,0}), \tag{2.3.21}$$

the thresholds take the form

$$C_j^{n_*} = \{(t,x) : W(0,x_0) - W(t,x) \geq (J^{n_*} - j)\Delta/n_*\}.$$

It may be preferable to see the invariance directly from the change of variables of the HJB PDE, as in (2.3.17). Repeating the calculations in (2.3.18) with inequalities instead of equalities, one finds that $\bar{W}$ is a subsolution to an appropriately scaled HJB with appropriately scaled boundary conditions. The thresholds C_j^n in (t,x) coordinates become

$$\begin{aligned} C_j^{n_*} &= \{(t,y) : \bar{V}(0,0) - \bar{V}(t,y) \geq (J^{n_*} - j)\Delta a_{n_*}^2\} \\ &= \{(t,x) : \bar{W}(0,x_0) - \bar{W}(t,x) \geq (J^{n_*} - j)\Delta/n_*\}. \end{aligned}$$

The takeaway is that the thresholds one constructs for the x variables have the same $1/n$ spacing one would have if X^n satisfied an LDP, except that we use a "zoomed in" rate function which sees the cost only up to second order.

2.3.5 Analysis of asymptotic performance

In this section we establish lower bounds on the decay rate of moderate-deviations based splitting schemes for the terminal escape problem $\mathbb{P}(X^n(T) \notin D)$. We recall that in Section 2.3.2 the problem was embedded into a sequence Y^n which satisfies a large deviations principle with the rate function I_M in (2.3.8). To estimate the terminal escape probability for a specific value of $n = n_*$, the domain D to escape is

scaled and translated by

$$D_* = \sqrt{n} a_{n_*}(D - X^{n,0}(t)).$$

This set D_* is fixed and independent of n. The function $V(t,y)$ as defined in (2.3.12) with $G = \infty \cdot 1_{D_*}(y)$ identifies the decay rate of the probabilities, i.e.

$$\lim_{n\to\infty} -a_n^2 \log \mathbb{P}(Y^n(T) \notin D_*) = V(0,0).$$

Recall the notation from the end of Section 2.1.5: when the splitting scheme is based on a subsolution $\bar{V}$ (see Definition 1.4.3 or (2.1.11)), $\gamma^n(\bar{V})$ denotes the associated RESTART estimator and the second moment of $\gamma^n(\bar{V})$ is denoted by $\mathcal{S}^n(\bar{V})$. We will establish that the decay rate is bounded below by $V(0,0) + \bar{V}(0,0)$. In particular, if $\bar{V}(0,0) = V(0,0)$, then the estimator $\gamma^n(\bar{V})$ will be asymptotically optimal in the sense of Definition 2.1.2.

Theorem 2.3.7. *Let $\bar{V}$ be a subsolution in the sense of Definition 1.4.3 for the problem of estimating $\mathbb{P}(Y^n(T) \notin D_*)$. We have the following lower bound:*

$$\liminf_{n\to\infty} -a_n^2 \log \mathcal{S}^n(\bar{V}) \geq \bar{V}(0,0) + V(0,0). \tag{2.3.22}$$

Proof. Let a subsolution $\bar{V}$ satisfying (2.1.11) and $V(T,y) \leq 0$ for $y \notin D_*$ be given and let $\bar{V}^n$ denote the associated piecewise constant importance functions defined by (2.1.14). Corollary 2.1.5 gives the following pre-asymptotic upper bound on the second moment:

$$\mathcal{S}^n(\bar{V}) \leq \mathbb{E}_{0,0}\left[\sum_{i=1}^{nT} e^{a_n^{-2}(\bar{V}((i-1)/n, Y^n_{i-1}) - \bar{V}(0,0))} (\mathbb{P}_{i/n, Y^n_i}(Y^n_{nT} \notin D^*))^2\right]$$
$$+ \mathbb{E}_{0,0}[e^{a_n^{-2}(\bar{V}((nT-1)/n, Y^n_{nT-1}) - \bar{V}(0,0))} 1_{D_*^c}(Y^n_{nT})].$$

To simplify notation, define for $i = 0, 1, \ldots, nT$ and $y \in \mathbb{R}^d$,

$$A_i^n(y) \doteq \exp\left[\frac{1}{a_n^2}\left(\bar{V}^n\left(\frac{i}{n}, y\right) - \bar{V}^n(0,0)\right)\right], \tag{2.3.23}$$

$$B_i^n(y) \doteq \mathbb{P}_{i/n,y}(Y_{nT}^n \notin D_*). \tag{2.3.24}$$

With this new notation, we have

$$\begin{aligned}\mathcal{S}^n(\bar{V}) \leqslant \mathbb{E}_{0,0}\left[\sum_{i=1}^{n(\rho^n \wedge T)} A_{i-1}^n(Y_{i-1}^n) B_i^n(Y_i^n)^2\right] \\ + \mathbb{E}_{0,0}\left[A_{n\rho^n}^n(Y_{nT}^n) 1_{D_*^c}(Y_{nT}^n)\right].\end{aligned} \tag{2.3.25}$$

There are two terms in (2.3.25), and both are positive and n-dependent. Owing to the logarithmic large deviations scaling, to establish (2.3.22) it suffices to check the limit on both terms individually. We start with the second, easier term.

Let $\varepsilon > 0$ be arbitary. From (2.1.15), $\bar{V}^n(t,y) - \bar{V}^n(0,0) \to \bar{V}(t,y) - \bar{V}(0,0)$ uniformly in (t,y), so for all n large enough we have

$$\bar{V}^n(t,y) - \bar{V}^n(0,0) \leqslant \bar{V}(t,y) - \bar{V}(0,0) + \varepsilon.$$

If $y \notin D_*$ then $\bar{V}(T,y) \leqslant 0$ by the boundary condition of a subsolution. Thus, for all n sufficiently large we have

$$\mathbb{E}_{0,0}[A_{n\rho^n}^n(Y^n(\rho^n)) 1_{D_*^c}(Y_{nT}^n)] \leqslant \exp\left[\frac{1}{a_n^2}(-\bar{V}(0,0) + \varepsilon)\right] \mathbb{P}_{0,0}(Y_{nT}^n \notin D_*).$$

Applying the large deviations upper bound for the process $\{Y_i^n, i \geq 0\}$, we obtain

$$\limsup_{n\to\infty} a_n^2 \log \mathbb{P}_{0,0}(Y_{nT}^n \notin D_*) \leqslant e^{-a_n^{-2} V(0,0)}.$$

Combining the two above displays and taking care of the negative signs, we obtain

$$\liminf_{n\to\infty} -a_n^2 \log \mathbb{E}_{0,0}\left[A^n_{n\rho^n}(Y^n(\rho^n))1_{D_*^c}(Y^n_{nT})\right] \geq \bar{V}(0,0) + V(0,0) - \varepsilon,$$

which handles the second term. We remark that the second term contributes to the overall rate of decay even though it might seem negligible compared to the first term in (2.3.25). This is because we expect the highest contributions to the variance to come from the particles which split closest to the escape set.

Next we handle the first term in (2.3.25). Owing to the linear growth of T, it suffices to bound each term in the sum. Indeed, if the desired limit holds for each i, then along any subsequence i_n, $0 \leqslant i_n \leqslant nT$ we have

$$\begin{aligned}
&\liminf_{n\to\infty} -a_n^2 \log \mathbb{E}_{0,0}\left[\sum_{i=1}^{nT} A^n_{i-1}(Y^n_{i-1})B^n_i(Y^n_i)^2\right] \\
&\quad \geq \liminf_{n\to\infty} -a_n^2 \log\left(nT\ \mathbb{E}_{0,0}[A^n_{i_n-1}(Y^n_{i_n-1})B^n_{i_n}(Y^n_{i_n})^2]\right) \\
&\quad \geq \bar{V}(0,0) + V(0,0) - \lim_{n\to\infty} a_n^2 \log n,
\end{aligned}$$

and this last term goes to 0 because of Condition 2.3.3.2. Thus it suffices to establish the limit for each fixed i.

Similarly, since J^n grows linearly with a_n^2 we may estimate the expected value individually on each set $D_j^n \doteq C_j^n \setminus C_{j-1}^n$, where the thresholds C_j^n are as in (2.3.20). Thus it suffices to establish the limit

$$\liminf_{n\to\infty} -a_n^2 \log \mathbb{E}_{0,0}[1_{D^n_{j_n}}(i_n/n, Y^n_{i_n})A^n_{i_n-1}(Y^n_{i_n-1})B^n_{i_n}(Y^n_{i_n})^2] \geq \bar{V}(0,0) + V(0,0) \tag{2.3.26}$$

along subsequences $\{j_n\}$ and $\{i_n\}$.

Before taking limits we address the time index disparity between the indicator $1_{D_j^n}(i/n, Y_i^n)$ and the point $\big((i-1)/n, Y_{i-1}^n\big)$ appearing in the argument of $\bar{V}^n$ in the expression for A_i^n from (2.3.23). Condition 2.3.1.1 states that the moment generating function of u_i is finite in a λ- neighbourhood of the origin uniformly in x, so in particular the distribution of $u_i(x)$ has moments of all orders for all $x \in \mathbb{R}^d$. Using the uniform convergence of $\bar{V}^n$ to $\bar{V}$ from (2.1.15), it is straightforward to argue that the probability of Y_i^n and Y_{i-1}^n not lying in the same set D_j^n is superexponentially small as $n \to \infty$, as was done in the proof of Theorem 2.2.15 (see equations (2.2.33) and (2.2.34)). Consequently, if we replace A_{i-1}^n by A_i^n in (2.3.26) it will not affect the limit. We do so for the remainder of the proof.

Returning to (2.3.26) with i instead of $i-1$, observe that for fixed i and j,

$$\begin{aligned}
&\mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)(i/n, Y_i^n)A_i^n(Y_i^n)B_i^n(Y_i^n)^2] \\
&\quad\leqslant \sup_{(t,y)\in D_j^n} \mathbb{P}_{t,y}(Y^n(T) \notin D_*)\mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)A_i^n(Y_i^n)B_i^n(Y_i^n)] \\
&\quad\leqslant e^{a_n^{-2}(\bar{V}_j^n - \bar{V}(0,0))} \sup_{(t,y)\in D_j^n} \mathbb{P}_{t,y}(Y^n(T) \notin D_*)\ \mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)B_i^n(Y_i^n)], \quad (2.3.27)
\end{aligned}$$

where we used the notation $\bar{V}_j^n = \bar{V}^n(t,y)$ for $(t,y) \in D_j^n$. From the definition of B_i^n and the Markov property,

$$\begin{aligned}
\mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)B_i^n(Y_i^n)] &= \mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)\mathbb{P}_{i/n,Y_i^n}(Y^n(T) \notin D_*)] \\
&\leqslant \mathbb{P}_{0,0}(Y^n(T) \notin D_*). \quad (2.3.28)
\end{aligned}$$

Suppose now $0 \leqslant j_n \leqslant J^n$ is any subsequence such that $j_n\Delta_n$ converges to some number θ,

$$j_n\Delta_n \to \theta \in [0, M],$$

and define the level set

$$L(\theta) = \{(t, y) : \bar{V}(0, 0) - \bar{V}(t, y) = M - \theta\},$$

where we recall that $M = \lim_{n\to\infty} J^n \Delta_n$ was defined in (2.1.13). If $(t_n, y_n) \in D^n_{j_n}$ for all $n \geq 1$ then

$$(J_n - j_n + 1)\Delta_n \geq \bar{V}^n(0, 0) - \bar{V}^n(t_n, y_n) \geq (J_n - j_n)\Delta_n,$$

which together with the uniform convergence of $\bar{V}^n$ to $\bar{V}$ implies

$$\liminf_{n\to\infty} -(\bar{V}^n_{j_n} - \bar{V}(0, 0)) = M - \theta. \tag{2.3.29}$$

This identifies the limit of the first factor in (2.3.27) when we apply the large deviations scaling. For the second factor, let (t_n, y_n) denote a convergent subsequence, so that $(t_n, y_n) \to (t, y)$. From the constraint $(t_n, y_n) \in D^n_{j_n}$, we must have $(t, y) \in L(\theta)$. The large deviation upper bound for the process Y^n_i yields, for all $(t, y) \in L(\theta)$,

$$\liminf_{n\to\infty} -a_n^2 \log \mathbb{P}_{t,y}(Y^n(T) \notin D_*) \geq V(t, y). \tag{2.3.30}$$

Finally, the large deviations applied for the initial condition $(0, 0)$ gives

$$\liminf_{n\to\infty} -a_n^2 \log \mathbb{P}_{0,0}(Y^n(T) \notin D_*) \geq V(0, 0). \tag{2.3.31}$$

Combining (2.3.29), (2.3.30) and (2.3.31) with the estimate (2.3.28), we apply the

large deviations scaling in (2.3.27) to deduce

$$\begin{aligned}\liminf_{n\to\infty} -a_n^{-2}\mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)A_i^n(Y_i^n)B_i^n(Y_i^n)^2] \\ \geq M - \theta + \inf_{t,y\in L(\theta)} V(t,y) + V(0,0).\end{aligned}$$

By the subsolution property, $V(t,y) \geq \bar{V}(t,y)$ for all $(t,y) \in [0,T]\times\mathbb{R}^d$. Thus,

$$\inf_{(t,y)\in L(\theta)} V(t,y) \geq \bar{V}(0,0) - M + \theta.$$

It follows that

$$\begin{aligned}\liminf_{n\to\infty} -a_n^{-2}\mathbb{E}_{0,0}[1_{D_j^n}(i/n, Y_i^n)A_i^n(Y_i^n)B_i^n(Y_i^n)^2] \\ \geq \bar{V}(0,0) + V(0,0),\end{aligned}$$

as claimed. □

Chapter Three

Chemical Reaction Networks

3.1 Introduction

This chapter is dedicated to a class of mathematical models known as chemical reaction networks (CRNs). The ultimate goal of the chapter is to construct subsolutions which determine good splitting schemes for the purposes of rare event simulation, as investigated in the previous chapter. These subsolutions will be tailored depending on the rare event of interest and can be constructed owing to the particular structure of CRNs. The techniques and examples developed in this chapter will then be tested numerically in the next one.

We outline the chapter, and then the remainder of the section. This first section introduces the theory of chemical reaction networks and defines the basic dynamics and objects used throughout the chapter. The next section, Section 3.2, describes techniques for identifying subsolutions using the specific structure of CRNs. These techniques are then used in Section 3.3 to construct subsolutions associated to various rare events arising in CRNs.

The remainder of the section proceeds as follows. We first provide some background on the history of CRNs in Section 3.1.1. Next, we formally define the CRN as a directed graph in Section 3.1.2 and make a few remarks about the definition. Finally, Sections 3.1.3 and 3.1.4 are dedicated to defining the two main models associated to the graph underlying a CRN: first, as a system of coupled ordinary differential equations, and second, as a jump Markov process.

Throughout, the object of interest will be a sequence $\{X^n, n \geq 1\}$ of scaled jump Markov processes which satisfy a law of large numbers $n \to \infty$. Each X^n will be a stochastic description of the CRN with some number of particles depending on n, and it is this sequence which will exhibit the "rare events" we are interested in.

3.1.1 Background

Broadly speaking, CRN models attempt to describe in one way or another the behavior of chemical species interacting in a well-stirred homogeneous medium, so that every particle of the same chemical species experiences identical conditions. In particular, CRNs do not track the spatial location of any individual species and provide only macroscopic information. For us, this will be either a discrete number of a certain type of species, or a concentration value, which is a non-negative real number. The origins of CRN theory can be traced back to the "law of mass action" and the principle of detailed balance for kinetic systems, developed towards the end of the 19th century. In an introductory course on chemistry, one typically encounters relationships of the form

$$A \rightleftarrows B$$

which one can vaguely interpret as saying that A reacts to B and vice-versa. The law of mass action states that at "equilibrium", the "rate" of the forward reaction $A \to B$ is proportional to the concentration of A, and the rate of the reverse reaction $B \to A$ is proportional to the concentration of B. The words equilibrium and rate here are in quotation marks because they mean something quite specific in chemistry and are often misused in mathematical treatises of the subject. The law can be formulated in terms of a coupled system of differential equations, often referred to as mass-action kinetics. The perspective we take on the subject might be said to originate from the pioneering work of R. Horn, F. Jackson and M. Feinberg in the early 70s; see [31, 44, 45]. The work of Horn, Jackson and Feinberg showed how graph theoretic notions could be useful in addressing questions about systems of differential equations (namely, mass-action kinetics).

We will not give a full historical account here, but we must note that CRNs have

experienced an increasing amount of attention from a wide variety of mathematical and scientific communities in the past few decades. As is often the case, depending on the community they may be known by a different name. They are sometimes called "stochastic Petri nets" in computer science and are related to "vector addition systems" or "population protocols". Reaction networks are also naturally related to queuing theory, and if one were to consider more general reaction rates than we do, a large class of queuing models could be formulated as reaction networks. Chemical engineers would be familiar with the term "master equation" (the Kolmogorov forward equation) and "Gillespie's algorithm" or "kinetic Monte Carlo", which is just the algorithmic description of a jump Markov process. In this thesis an important class of networks are "complex-balanced" networks, which are better known as "toric varieties" in algebraic geometry. We are able to solve the time-independent Hamilton-Jacobi-Bellman equation for a class of networks owing to the algebraic structure of these toric varieties. The original connection to reaction networks is from Anne Shiu's PhD thesis; see [60] or [61]. Work on stochastic models of reaction networks also dates back to at least the 1950s (see [5, 6, 15]), and in the 70s Kurtz established the law of large numbers for jump processes, which relates the usual stochastic and deterministic models [51, 52]; see also Theorem 1.3.4 in Chapter 1. The style of analysis here has been widely used in the physics community; see for instance [26, 62, 9, 50, 55, 48] or [54, 1] for some on the large deviations perspective. The arxiv preprint [2] in particular treats many of the same objects we discuss here, though not for the same purposes.

3.1.2 Definition of CRNs

In this section we introduce some notation used throughout the chapter and define chemical reaction networks formally.

Notation 3.1.1. Given $x \in \mathbb{R}^d_{\geqslant 0}$ and $\zeta \in \mathbb{Z}^d$, we set

$$x^\zeta \doteq x_1^{\zeta_1} \cdots x_d^{\zeta_d}.$$

For $a \in \mathbb{Z}^d$, we define $|a| = a_1 + \cdots + a_d$.

Definition 3.1.2. A chemical reaction network (CRN) in d species is a directed graph $G_\kappa = (\mathcal{C}, \mathcal{R}, \kappa)$, where the vertices $a \in \mathcal{C}$ are d-dimensional vectors whose entries are non-negative integers. Each vertex $a \in \mathcal{C}$ is identified with a monomial x^a in d variables. The directed edges are defined by ordered pairs $(a, b) \in \mathcal{R} \subset \mathcal{C} \times \mathcal{C}$ and are labeled with positive rate constants $\kappa_{a,b} > 0$.

The reason for associating each vertex to the monomial x^a will become clearer once the deterministic dynamics are introduced, and can be thought of as arising from a counting argument (see the beginning of Section 3.1.4). We also note that the rate constants $\kappa_{a,b}$ may or may not be important depending on the properties of the CRN one is interested in. When they are not important, we will simply use G to denote the pair $(\mathcal{C}, \mathcal{R})$, with the assumption that $\kappa_{a,b} = 1$ for all $(a, b) \in \mathcal{R}$.

Remark. It is natural to identify the vertices $a \in \mathcal{C}$ with their labels, which are in turn entirely described by the vector of exponents of the monomials. Thus $a \in \mathcal{C}$ is always identified with some vector in $\mathbb{Z}^d$. The vector of exponents a is called a *complex* in the literature. A *reaction* is a directed edge $(a, b) \in \mathcal{R}$ and is identified with the vector $\zeta_{a,b} = b - a$, where the difference is interpreted in $\mathbb{Z}^d$. We define the *rate* of a reaction $(a, b) \in \mathcal{R}$ to be $\lambda_\zeta(x) = \lambda_{a,b}(x) = \kappa_{a,b} x^a$, where $\kappa_{a,b} > 0$ is the

rate constant taken to be 1 unless otherwise specified, and we call x^a the *unweighted rate*. Note that we might have $(a, b) \neq (a', b')$ but $\zeta_{a,b} = \zeta_{a',b'}$.

Remark. Our definition of CRN is slightly less general than the usual one, which allows for the labels of the vertices to be general rational functions of $(x_1, \ldots, x_d)$ rather than just monomials. We do not consider the general case in this thesis, though we expect much of the work could be extended to rates $\lambda(x)$ which are Lipschitz continuous, satisfy subexponential growth and decay conditions, and take values in $[0, \infty)$.

Example 3.1.3. Let $\mathcal{C} = \{[2, 0], [1, 1], [0, 2]\} = \{a, b, c\}$ and $\mathcal{R} = \mathcal{C} \times \mathcal{C}$, the set of all ordered pairs in $\mathcal{C}$. The graph G corresponding to the CRN $(\mathcal{C}, \mathcal{R})$ is

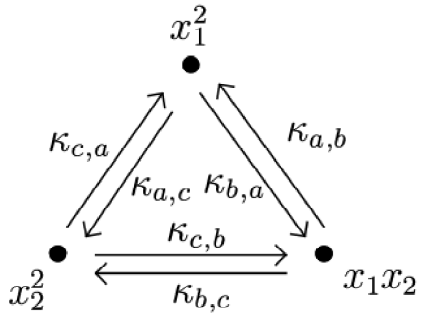

Figure 3.1

The graphical depiction in Example 3.1.3 will often be expressed in terms of the standard notation in chemistry, as in the expression $A \to B$. For instance, we could write the same network without the explicit κ labels as

$$2A \rightleftarrows A + B$$

$$A + B \rightleftarrows 2B$$

$$2B \rightleftarrows 2A.$$

One clear disadvantage of this way of depicting the network is that the vertices are repeated, so at first glance one might think $|\mathcal{C}| = 6$ even though $|\mathcal{C}| = 3$. As such, we will only use this format when no such confusion will arise.

We first present the deterministic description in terms of a system of coupled ordinary differential equations, the "law of mass action". The second description is stochastic and is an analogue of the "law of mass action" since it has a natural law of large number scaling limit attached to it which converges to the system of differential equations, as described by Theorem 1.3.4.

3.1.3 Deterministic dynamics

In this section we explore the deterministic dynamics associated to any CRN G_κ known as "mass-action kinetics". We first define the dynamics as a system of differential equations, and subsequently we explore the structure of the dynamics in terms of the underlying graph G_κ.

System of autonomous ODEs $\dot{\boldsymbol{x}} = \boldsymbol{\Phi}(\boldsymbol{x})$. To any d-dimensional CRN G_κ, we associate the deterministic dynamics $\Phi : \mathbb{R}_+^d \to \mathbb{R}^d$ given by

$$\Phi(x) \doteq \sum_{(a,b)\in\mathcal{R}} \zeta_{a,b}\lambda_{a,b}(x) = \sum_{(a,b)\in\mathcal{R}} (b-a)\kappa_{a,b}x^a.$$

What is the range of $x(t)$ for given $x(0)$? Recall that the $\zeta_{a,b}$ are vectors in $\mathbb{Z}^d$ indexed by the reactions $(a,b) \in \mathcal{R}$ and determined by the differences $b - a$. By the

definition of a solution,

$$x(t) - x(0) = \sum_{(a,b)\in\mathcal{R}} \zeta_{a,b} \left(\int_0^t \lambda_{a,b}(x(s))\mathrm{ds} \right),$$

so $x(t) - x(0)$ is a linear combination of the vectors $\zeta_{a,b}$. In general, the domain of the trajectory will be the intersection an affine subset of $\mathbb{R}^d$ with the positive orthant $\mathbb{R}^d_{\geqslant 0}$.

We note that existence and uniqueness for the system $\dot{x} = \Phi(x)$ is guaranteed for any initial condition $x(0)$. This follows from the Lipschitz continuity of $\lambda_{a,b}(x)$ (see for instance [41]).

Definition 3.1.4. Let $G = (\mathcal{C}, \mathcal{R})$ be a CRN. The subspace

$$\mathcal{S}_0 \doteq \left\{ \sum_{(a,b)\in\mathcal{R}} \alpha_{a,b}\zeta_{a,b} : \alpha_{a,b} \in \mathbb{R}, (a,b) \in \mathcal{R} \right\}$$

is called the *stoichiometric subspace* of the network. For a given initial condition $x(0) \in \mathbb{R}^d_{\geqslant 0}$, we let

$$\mathcal{S}_{x(0)} \doteq \mathbb{R}^d_{\geqslant 0} \cap (\mathcal{S}_0 + x(0)) = \mathbb{R}^d_{\geqslant 0} \cap \{s + x(0) : s \in \mathcal{S}_0\}. \tag{3.1.1}$$

The set $\mathcal{S}_{x(0)}$ is called the *stoichiometric compatibility class at* $x(0)$. When the initial condition $x(0)$ is implied or irrelevant, we simply write $\mathcal{S}$ rather than $\mathcal{S}_{x(0)}$ and refer to $\mathcal{S}$ as *the* stoichiometric compatibility class or simply the *domain* of G.

We will also use $\partial\mathcal{S}$ to denote the boundary of the set $\mathcal{S}$, and $\mathcal{S}^\circ$ to denote its interior. A *facet* of $\mathcal{S}$ is a subset of $\partial\mathcal{S}$ with co-dimension 1. Any facet is of the form

$$\{x \in \mathbb{R}^d_+ : x_i = 0\} \cap \partial\mathcal{S} \tag{3.1.2}$$

for a unique $i \in \{1, \ldots, d\}$. Finally, note that $x(t) \in \mathcal{S}_{x(0)}$ for all $t \geqslant 0$ whenever $\dot{x}(t) = \Phi(x(t))$.

Remark. In the definition of a reaction network, we allow for complex $a = 0$, which will always be associated with the d-dimensional zero vector. We call this complex the *empty complex*. It corresponds to a constant inflow of species since $x^0 = 1$. As defined, CRNs do not allow for constant outflow.

Graphical structure. The underlying structure of the system $\dot{x} = \Phi(x)$ depends on the relationships between the various species which are captured in the graph G_κ. Observe that $\Phi(x)$ in Example 3.1.3 can be written as

$$\Phi(x) = Y L_\kappa^T \psi(x),$$

where $\psi(x) = (x_1^2, x_1 x_2, x_2^2)$,

$$Y = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \end{pmatrix}, \quad L_\kappa = \begin{pmatrix} -\kappa_{a,b} - \kappa_{a,c} & \kappa_{a,b} & \kappa_{a,c} \\ \kappa_{b,a} & -\kappa_{b,a} - \kappa_{b,c} & \kappa_{b,c} \\ \kappa_{c,a} & \kappa_{c,b} & -\kappa_{c,a} - \kappa_{c,b} \end{pmatrix}.$$

This is a general decomposition that is useful since it separates out three aspects of the deterministic dynamics:

1. all of the nonlinearity in the dynamics is stored in the map $\psi(x)$;
2. the orientations of the edges in the graph G_κ are captured in the matrix L_κ, and
3. there is a difference between the rate of change of concentration of each species, and the rate of change of concentration of each complex.

We define these objects more generally.

Definition 3.1.5. Let G_κ be a CRN. The deterministic dynamics $\Phi(x)$ can be written as

$$\Phi(x) = Y L_\kappa^T \psi(x),$$

where the superscript T denotes a transpose.

1. The map $\psi : \mathbb{R}^d \to \mathbb{R}^{\mathcal{C}}$ denotes the vector *concentration to complexes* map which takes the vector of concentrations x to the vector of unweighted rates indexed by the complexes $a \in \mathcal{C}$, i.e.

$$(\psi(x))_a \doteq x^a.$$

   We often abuse notation and assume a labeling on the complexes to identify $\psi(x)$ with a vector in $\mathbb{R}_+^{|\mathcal{C}|}$. Similarly, we use (a, b) to refer to a point either in $\mathcal{C} \times \mathcal{C}$ or $|\mathcal{C}| \times |\mathcal{C}|$.

2. The map $L_\kappa : \mathbb{R}^{|\mathcal{C}|} \to \mathbb{R}^{|\mathcal{C}|}$ is the matrix whose $(a, b)^{\text{th}}$ entry is

$$(L_\kappa)_{a,b} \doteq \begin{cases} \kappa_{a,b} & \text{if} \quad a \neq b \\ -\sum_{b \neq a} \kappa_{a,b} & \text{if} \quad a = b \end{cases}.$$

   We call L_κ the weighted Laplacian on the graph G.

3. The "complexes to concentrations" map $Y : \mathbb{R}^{|\mathcal{C}|} \to \mathbb{R}^d$ is the matrix whose columns are given by the complexes $a \in \mathbb{Z}_+^d$. We have

$$\Phi(x) = Y L_\kappa^T \psi(x),$$

   where right multiplication by the column vector $\psi(x)$ is interpreted as standard

matrix-vector multiplication.

The decomposition in Definition 3.1.5 encodes the graphical structure of the CRN.

Definition 3.1.6. A CRN $G = (\mathcal{C}, \mathcal{R})$ is said to be *weakly reversible* if for any $(a, b) \in \mathcal{R}$, there is a sequence of edges $(b_0, b_1), (b_1, b_2), \ldots, (b_{k-1}, b_k) \in \mathcal{R}$ such that $b_0 = b$ and $b_k = a$.

We remark that weak reversibility is determined entirely by the form of the Laplacian L_κ in Definition 3.1.5. The rate constants $\kappa_{a,b}$ are irrelevant for weak reversibility, so long as they are positive. In Example 3.1.3, the Laplacian L_κ with $\kappa_{a,b} = 1$ for all $(a, b) \in \mathcal{R}$ is

$$L_\kappa = \begin{pmatrix} -2 & 1 & 1 \\ 1 & -2 & 1 \\ 1 & 1 & -2 \end{pmatrix}.$$

Identifying complexes $a, b \in \mathcal{C}$ with coordinates on $\mathbb{R}^{|\mathcal{C}|}$, it is clear that

$$(L_\kappa)_{a,b} = 1\{\text{edge } a \to b\}.$$

This allows us to verify the reachability condition in Definition 3.1.6.

The decomposition $\Phi(x) = \psi(x) L_\kappa^T Y$ of Definition 3.1.5 also allows a deeper understanding of the structure of *steady states* of the network. A *steady state* is a point $c \in \mathbb{R}^d_{>0}$ such that $\Phi(c) = 0$. A natural question to ask is: does a given CRN admit multiple steady states? In some cases, the question can be answered directly from the graph $G = (\mathcal{C}, \mathcal{R})$ without regard to the rate constants – see Section 3.2.1

on "complex-balanced" networks. In other cases, the rate constants $\kappa_{a,b}$ are relevant, and a more detailed analysis is needed. For an algebraic perspective in which one considers polynomials in the rate constants, see [60].

3.1.4 Stochastic dynamics

In this section we define the stochastic dynamics associated to a CRN G_κ. These are intimately tied to the deterministic dynamics introduced in the previous section via the law of large numbers. We first motivate the mass-action rates used in the stochastic dynamics, which also explain the form of the rates $\lambda_{a,b}(x) = \kappa_{a,b}x^a$. We then define a scaled jump Markov process and consider its domain. Finally, we reintroduce the objects from Chapter 2 used for the construction of subsolutions in the context of CRNs.

Counting arguments for mass action rates. We remind ourselves of the underlying physical process we are trying to model: chemical species of a certain type interacting in a well-stirred medium, with any two particles of the same species experiencing identical conditions. Under such assumptions, what is a natural way of associating a "rate" to a reaction? For instance, if there are X_A particles of species A and X_B particles of species B, what should be the rate of the reaction $A \to B$?

Perhaps the most straightforward is to count the number of ways two particles can associate, in this case, $X_A X_B$. If we had a reaction of the form $A + A \to B$, or $2A \to B$ in our notation, the number of ways of choosing two particles out of X_A without regard to order is $X_A(X_A-1)/2$. In general, if a reaction requires n_A particles of species A, its rate should be proportional to $X_A!/(X_A - n_A)!$ if $X_A \geqslant n_A$ and 0

otherwise. This heuristic counting argument motivates the stochastic dynamics.

Jump Markov Processes. To any CRN G_κ we associate the continuous time jump Markov process $X^n(t)$ with generator

$$(\mathcal{L}^n f)(x) = \sum_{(a,b)\in\mathcal{R}} \lambda^n_{a,b}(x)\left(f\left(x+\frac{\zeta_{a,b}}{n}\right)-f(x)\right), \quad x\in\frac{1}{n}\mathbb{Z}^d_+,$$

where $\lambda^n_{a,b}:\mathbb{Z}^d\to[0,\infty)$ is defined by

$$\lambda^n_{a,b}(x) \doteq \frac{n\kappa_{a,b}}{n^{|a|}}\prod_{i=1}^d \frac{x_i!}{(x_i-a_i)!}1\{x_i\geqslant a_i\}.$$

Observe that for $x\in n^{-1}\mathbb{Z}^d$ and n large,

$$\begin{aligned}\lambda^n_{a,b}(nx) &= n\kappa_{a,b}\prod_{i=1}^d \frac{nx_i(nx_i-1)\cdots(nx_i-a_i+1)}{n^{a_i}}1\{nx_i\geqslant a_i\}\\ &\approx n\kappa_{a,b}\prod_{i=1}^d x_i^{a_i}\\ &= n\lambda_{a,b}(x).\end{aligned}$$

In [51], Kurtz established the functional law of large numbers for the process X^n. By Theorem 1.3.4 in Chapter 1, we have $X^n(\cdot)\to x(\cdot)$ in probability in the uniform topology, so one can say that $X^n(\cdot)$ "approximates" the deterministic dynamics $x(\cdot)$ for large n. In general, we will assume that $\lambda_{a,b}(x)=\kappa_{a,b}x^a 1_{\{x\geq 0\}}$ and call these *mass-action* rates.

We define the state space of the scaled process X^n passing through the point $c\in\mathbb{R}^d_{>0}$ by

$$\mathcal{S}^n_c = \mathcal{S}_c\cap\frac{1}{n}\mathbb{Z}^d = \{x\in\mathcal{S}_c : x=n^{-1}z, z\in\mathbb{Z}\},$$

where the stoichiometric compatibility class $\mathcal{S}_c$ was defined in Definition 3.1.4. When no particular point c needs to be distinguished, we simply write $\mathcal{S}^n$. Thus the sequence of processes $\{X^n\}$ approximates the deterministic dynamics as $n \to \infty$. Note that the state space $\mathcal{S}_c^n$ is only a sublattice of the affine space $\mathcal{S}_c$ (see Definition 3.1.4), and that in general there can be multiple sublattices of the same affine space. Moreover, the different sublattices can exhibit significantly different dynamics. For instance, Example 3.1.7 demonstrates that an extinction event can depend on the initial configuration.

Example 3.1.7. Consider the CRN

$$\begin{aligned} A_1 + A_2 &\to 2A_3 \\ A_1 + A_3 &\to 2A_2 \\ A_2 + A_3 &\to 2A_1. \end{aligned}$$

This reaction network is not weakly reversible. The dynamics of this network depend on the initial configuration $X(0) = (X_1(0), X_2(0), X_3(0))$.

To see this, let $\vec{v} = (1, 1, 1)$, the normal vector to the family of planes in which the reaction network takes values. For a given initial configuration $X(0)$, the only absorbing states are at the vertices, $(|X(0)|, 0, 0)$, $(0, |X(0)|, 0)$ and $(0, 0, |X(0)|)$, where $|X(0)| = \sum_{i=1}^{3} X_i(0)$.

Let n_1, n_2 and n_3 be the total number of times the first, second and third reactions occur to reach an absorbing state from $X(0)$. Then, for instance,

$$X_1(t) = X_1(0) + 2n_1 - n_2 - n_3.$$

By equating the values of the two extinct species in each of the vertices, we obtain

three pairs of equations,

$$X_i(0) - X_j(0) = 3(n_i - n_j).$$

In other words, if $X(0)$ can lead to an absorbing state, then at least one of the differences $X_i(0) - X_j(0), i \neq j$ must be divisible by 3. Thus,

$$X(0) = (400, 397, 399)$$

will eventually reach an absorbing state, whereas

$$\tilde{X}(0) = (400, 396, 400)$$

will not (even though both live on the same plane: $(X(0) - \tilde{X}(0)) \cdot \vec{v} = 0$.

While it might seem that this kind of behavior could impact the rare event behavior, the scaling in the law of large numbers renders the large deviations analysis insensitive to this kind of combinatorial issue. We will ignore the sublattices where extinction could occur, and we note that it is straightforward to verify that the dynamics generated by a weakly reversible reaction network (see Definition 3.1.6) are ergodic and that for each such such sublattice $\mathcal{S}^n$ there is a unique stationary distribution π^n. We do so throughout.

Subsolutions. Since $\{X^n\}$ is a Markov process satisfying a law of large numbers, it is natural to conjecture $\{X^n\}$ might also satisfy a corresponding large deviations principle. Based on the results of Chapter 2, subsolutions can be used as importance functions for the design of splitting schemes. In this section we define the basic objects arising from the large deviations framework which will be used to define and

analyze subsolutions for rare event estimation.

Let $G_\kappa = (\mathcal{C}, \mathcal{R}, \kappa)$ be a d-dimensional reaction network. We define the Hamiltonian $H(x,p) : \mathbb{R}^d_+ \times \mathbb{R}^d \to \mathbb{R}$ by

$$H(x,p) = \sum_{(a,b)\in\mathcal{R}} \kappa_{a,b} x^a (e^{\langle \zeta_{a,b}, p\rangle} - 1). \tag{3.1.3}$$

Note that in this chapter, we will use p to denote the second variable of H, as opposed to α in earlier chapters. It is not difficult to see that Hamiltonians of the form (3.1.3) are in one-to-one correspondence with d-dimensional CRNs with mass-action rates $\lambda_{a,b}(x) = \kappa_{a,b} x^a$.

The local rate $L(x,\beta)$ is the Legendre transform of the cumulant generating function of $H(x,p)$, and can be expressed in terms of the convex optimization problem

$$L(x,\beta) = \inf\left\{ \sum_{(a,b)\in\mathcal{R}} \lambda_{a,b}(x)\ell\left(\frac{u_{a,b}}{\lambda_{a,b}(x)}\right) : \sum_{(a,b)\in\mathcal{R}} u_{a,b}\zeta_{a,b} = \beta \right\}. \tag{3.1.4}$$

For reaction networks, the pair of Hamilton's ODEs is given by

$$\begin{aligned} \dot{x}(t) &= \sum_{(a,b)\in\mathcal{R}} \zeta_{a,b}\lambda_{a,b}(x(t)) e^{\langle p(t), \zeta_{a,b}\rangle} \\ \dot{p}(t) &= \sum_{(a,b)\in\mathcal{R}} \nabla_x \lambda_{a,b}(x(t))(1 - e^{\langle p(t), \zeta_{a,b}\rangle}), \end{aligned} \tag{3.1.5}$$

and these are expected to be naturally related to minimizers of the variational problem

$$\inf_{\varphi}\{I_T(\varphi) : \varphi(0) = y, \varphi(T) = z\},.$$

We recall that a defining property of a continuously differentiable time-independent

subsolution $\bar{V}(x)$ is to satisfy

$$H(x, -D\bar{V}(x)) \leq 0,$$

with the boundary conditions determined by the rare event. The time-dependent version for a continuously differentiable subsolution $\bar{V}(t, x)$ is

$$\partial_t \bar{V}(t, x) - H(x, -D\bar{V}(t, x)) \geq 0,$$

with the appropriate boundary conditions.

3.2 Techniques for identifying subsolutions

In this section we present three techniques to identify subsolutions for the purpose of designing splitting schemes; see Sections 1.4.2 and 2.1.4 for the definitions of subsolutions and their relation to splitting schemes, respectively. These techniques will then be applied in Section 3.3 to the design of subsolutions for estimating rare event probabilities via splitting.

The identification of subsolutions is in general a hard problem and there is no catch-all method. Moreover, one typically hopes to save on computational effort, so any numerical methods used to evaluate (or approximate) the subsolutions at a point should be significantly less expensive than the simulation cost if a subsolution were not used. With this in mind, we note that there are significant assumptions needed for the use of each technique. These are described below.

The first one, described in Section 3.2.1 requires that the network have a special

structure known as being "complex-balanced". This structure implies the uniqueness of a steady state in each stoichiometric compatibility class (see Definition 3.1.4), and from the steady state alone one can identify a smooth solution to the HJB PDE $H(x, DU(x)) = 0$. Moreover, non-trivial complex-balancing networks of arbitrarily large dimension can fairly easily be constructed. However, while complex-balanced networks are ubiquitous and well-studied, they have a rigid structure. Since there is exactly one steady state per compatibility class, complex-balancing networks cannot exhibit metastability. Instead, the rare events studied with this method are escape from rest point problems, and issues of hitting the boundary.

The second technique is restricted to one-dimensional CRNs. Each one-dimensional CRN corresponds to a one-dimensional Hamiltonian which is time-independent. In particular, the Hamiltonian itself (the "energy") is a conserved quantity and it is straightforward to compute the quasipotential numerically by quadrature. Since this can be done for a general one-dimensional CRN, we are free to explore all types of rare events, including metastability issues.

Finally, we describe the moderate deviations approximation to CRNs. This technique can be applied to general d-dimensional reaction networks without making any assumptions on the structure of the network. The drawback is that the moderate deviations approximation cannot capture more complex rare-event behavior, since it is essentially a quadratic approximation around the law of large numbers. We will illustrate how they can be used to estimate rare events which are not "too rare" in the context of escaping a rest point and leaving a boundary.

3.2.1 Subsolutions for complex-balanced networks

In this section we introduce the notion of complex-balanced networks and describe some of their basic properties. The main result of the section is that one can explicitly identify the quasipotential for any complex-balanced CRN. Owing to the discussion of Section 2.1.4.1, we can use the quasipotential to construct time-independent subsolutions for the problem of estimating escape probabilities over increasing time intervals $[0, T(n)]$. Theorem 2.2.15 shows that the corresponding splitting scheme will be asymptotically optimal in the sense of Definition 2.2.14.

In the first part of this section we define the notion of complex-balanced networks and state a few well-known consequences. In the second part, we provide a simple condition for being complex-balanced, known as the Deficiency Zero theorem (see Theorem 3.2.4). In the third part we exhibit an explicit solution to the steady-state HJB PDE.

3.2.1.1 Complex-balanced networks

Let $G_\kappa = (\mathcal{C}, \mathcal{R}, \kappa)$ be a d-dimensional CRN, and recall the maps L_κ and ψ from Definition 3.1.5.

Definition 3.2.1. The CRN G_κ is said to be *complex balanced* at a point $c \in \mathbb{R}^d_{>0}$ if

$$\psi(c) L_\kappa = 0.$$

We call c a *complex balancing equilibrium*.

The definition of complex balanced implies that in every stochiometric compati-

bility class, there is c such that for any complex $z \in \mathcal{C}$,

$$\sum_{a:(a,z)\in\mathcal{R}} \kappa_{a,z} c^a = \sum_{b:(z,b)\in\mathcal{R}} \kappa_{z,b} c^z \tag{3.2.1}$$

One can interpret the left-hand side of (3.2.1) as the "inflow" into the complex z, and the right-hand as the "outflow" from z.

The study of complex balanced networks goes back to the seminal work of Feinberg, Horn and Jackson [31, 32, 33, 34, 44, 45]. Of particular interest has been the global attractor conjecture, which asks whether every complex balanced system has a globally attracting fixed point within each stoichiometric compatibility class.

There are several important facts about complex balanced networks. Much of the modern phrasing we use here is due to the expository work of Anderson, though many of the results are originally due to Horn and Jackson:

1. if $c \in \mathbb{R}^d_{>0}$ is a complex-balancing equilibrium, then every $\bar{c} \in \mathbb{R}^d_{>0}$ satisfying $\Phi(\bar{c}) = 0$ must be complex-balancing, i.e., $\psi(\bar{c})L_\kappa = 0$;

2. for any stoichiometric compatibility class $\mathcal{S}(x_0)$, there is a unique complex-balancing equilibrium $c \in \mathcal{S}(x_0)$;

3. the network is necessarily weakly reversible.

In light of properties 1 and 2 above, any CRN which admits a complex balancing equilibrium has only complex balancing equilibria, and they exist in each stoichiometric compatibility class. In such a case we simply refer to the network as being *complex balanced*, without regard to the specific equilibrium c.

3.2.1.2 Deficiency Zero Theorem

There exists sufficient (though far from necessary) conditions which guarantee that a network is complex balanced. What we call here the Deficiency Zero theorem is not exactly what it is called by other authors (see e.g. [44], Theorem 4A, or Theorem 3.1 in [4]). However, the result as stated here is what will be most useful for our purposes of constructing reaction networks which are complex-balanced. First, we define the *deficiency* of a CRN.

Definition 3.2.2. The deficiency δ of a network $G_\kappa = (\mathcal{C}, \mathcal{R}, \kappa)$ is given by

$$\delta \doteq \dim(\ker(Y) \cap \text{image}(L_\kappa^T)).$$

It is easiest to calculate the deficiency for weakly reversible reaction networks. The following is a simplified version of Proposition 5.1 in [40].

Proposition 3.2.3. *If $G = (\mathcal{C}, \mathcal{R})$ is weakly reversible, then*

$$\delta = |\mathcal{C}| - \#\text{cc}(G) - \dim(\mathcal{S}_0), \tag{3.2.2}$$

where $|\mathcal{C}|$ is the number of complexes, $\#\text{cc}(G)$ is the number of connected components of the graph G, and $dim(\mathcal{S}_0)$ is the dimension of the stoichiometric subspace (see Definition 3.1.4).

Though the notion of deficiency is widely used in the study of the dynamics reaction networks, we do not investigate the issue further and simply take advantage of the simple formula (3.2.2) for the deficiency in the case of a weakly reversible reaction network. For a linear-algebraic motivation for the notion of deficiency,

see Definition 3.2 in [40] and its subsequent arguments; for the algebraic geometry perspective, see [60].

We now state the Deficiency Zero theorem.

Theorem 3.2.4 (Deficiency Zero)**.** *Suppose that $G_\kappa = (\mathcal{C}, \mathcal{R})$ is a CRN with a deficiency of $\delta = 0$ and $\{\kappa_{a,b}\}$ is any choice of positive rate constants. The network is weakly reversible if and only if it is complex balanced.*

Recall that weak reversibility (Definition 3.1.6) is a graph-theoretic condition that is easily checked by inspection (at least for small graphs), whereas the deficiency of a network is related to the dimension of the span of the reaction vectors $\zeta_{a,b}$. It is therefore straightforward to construct reaction networks which are complex balanced, at least for "low" dimensions (say $d < 20$).

Anderson et al. determine in [4] that any network which is complex balanced has the following stationary distribution:

$$\pi^n(nx) = \frac{1}{Z^n} \prod_{i=1}^{d} e^{-nc_i} \frac{(nc_i)^{nx_i}}{(nx_i)!}, \quad x \in \mathcal{S}_c^n,$$

where Z^n is a positive normalising constant. If $\mathcal{S}_c^n = n^{-1}\mathbb{Z}_+^d$, then $Z^n = 1$ and the distribution is of product Poisson form. Having determined the stationary distribution, in [3] they perform a standard calculation via Stirling's approximation to obtain

$$\lim_{n\to\infty} -\frac{1}{n} \log \pi^n(x) = \sum_{i=1}^{d} c_i \ell \left(\frac{x_i}{c_i} \right),$$

where $c = (c_1, \ldots, c_d)$ satisfies $\Phi(c) = 0$. They then recognise the right-hand side as a well-known Lyapunov function for the system $\dot{x}(t) = \Phi(x(t))$, which was used by [45] to show that the steady state is locally asymptotically stable in its stoichiometric

compatability class $\mathcal{S}_c$. Based on the theory of Chapter 1 (Theorems 1.5.3, 1.5.4 and Proposition 1.5.9) we should be able to connect this Lyapunov function with a solution of the time-independent HJB. This is the content of the next section. A more complete description of the relationships between Lyapunov functions, solutions of the HJB PDE and large deviations can be found in [23, 8].

3.2.1.3 Solution of HJB

We now identify the quasipotential for complex-balanced CRNs using only information about the steady state. As mentioned earlier, the quasipotential will allow us to construct time-independent subsolutions for escape problems, and the associated splitting schemes will be asymptotically optimal over long time intervals (see Section 2.1.4.1).

Theorem 3.2.5. *Let G_κ be a complex-balanced d-dimensional CRN with rate constants $\{\kappa_{a,b}\}_{(a,b)\in\mathcal{R}}$, let $c \in \mathbb{R}^d_{>0}$ be a complex balancing steady state, and let D denote the domain of attraction of c. The quasipotential relative to c is given on $\mathcal{S}_c \cap \{x : U(x) \leq \min_{y\in\partial D} U(y)\}$ by*

$$U(x) = \sum_{i=1}^{d} c_i \ell\left(\frac{x_i}{c_i}\right) \tag{3.2.3}$$

where $\ell(x) = x\log x - x + 1$.

Proof. First, note that the Hamiltonian associated to any CRN is continuously differentiable in (x, p) and strictly convex in the second variable. According to Theorem 1.5.4, in order to verify that a non-negative real-valued function U is equal to the quasipotential on a neighbourhood of the steady state, it suffices to check

that $U(c) = 0$, $DU(c) = 0$, $U(x) > 0$ and $DU(x) \neq 0$ for $x \neq c$, and that U satisfies the HJB PDE $H(x, DU(x)) = 0$. Then $U(x)$ is equal to the quasipotential on $\{x : U(x) \leq \min_{y \in \partial D} U(y)\}$.

Let $D \subset \mathcal{S}_c$ be the domain of attraction of c. Note that $U(x)$ is a sum of strictly convex functions, hence strictly convex, and that $U(x)$ has its unique zero at $x = c$. The gradient is

$$DU(x) = \left(\log \frac{x_1}{c_1}, \ldots, \log \frac{x_d}{c_d} \right).$$

Each component of DU is equal to zero if and only if $x_i = c_i$ for all $i = 1, \ldots, d$. Thus, aside from satisfying the PDE, U satisfies all the prerequisite conditions for the theorem on D. We now verify that U in fact solves the HJB PDE $H(x, DU(x)) = 0$. We have,

$$\begin{aligned}
H(x, DU(x)) &= \sum_{(a,b) \in \mathcal{R}} \kappa_{a,b} x^a (e^{\langle DU(x), b-a \rangle} - 1) \\
&= \sum_{(a,b) \in \mathcal{R}} \kappa_{a,b} x^a \left(\frac{\left(\frac{x}{c}\right)^b}{\left(\frac{x}{c}\right)^a} - 1 \right) \\
&= \sum_{(a,b) \in \mathcal{R}} \kappa_{a,b} x^b \frac{c^a}{c^b} - \sum_{(a,b) \in \mathcal{R}} \kappa_{a,b} x^a. \\
&= \sum_{(a,b) \in \mathcal{R}} \left(\left[\sum_{z \in \mathcal{C}} 1_{\{b=z\}} \right] \kappa_{a,b} x^b \frac{c^a}{c^b} \right) - \sum_{(a,b) \in \mathcal{R}} \left(\left[\sum_{z \in \mathcal{C}} 1_{\{a=z\}} \right] \kappa_{a,b} x^a \right) \\
&= \sum_{z \in \mathcal{C}} \left[\sum_{a:(a,z) \in \mathcal{R}} \kappa_{a,z} x^z \frac{c^a}{c^z} - \sum_{b:(z,b) \in \mathcal{R}} \kappa_{z,b} x^z \right] \\
&= \sum_{z \in \mathcal{C}} \left[\frac{x^z}{c^z} \sum_{a:(a,z) \in \mathcal{R}} \kappa_{a,z} c^a - \sum_{b:(z,b) \in \mathcal{R}} \kappa_{z,b} x^z \right] \\
&= \sum_{z \in \mathcal{C}} \left[\frac{x^z}{c^z} \sum_{b:(z,b) \in \mathcal{R}} \kappa_{z,b} c^z - \sum_{b:(z,b) \in \mathcal{R}} \kappa_{z,b} x^z \right] \\
&= 0.
\end{aligned}$$

In the penultimate equality above we used the complex-balanced relation (3.2.1). □

We remark that the domain of attraction in Theorem 3.2.5 can in practice be taken to be all of $\mathcal{S}_c$, since one can verify directly that the complex balancing steady state is globally asymptotically stable. The general result for an arbitrary complex-balanced CRN is much more difficult and is the content of the *Global Attractor Conjecture*, proved recently in [11]. While the result seems accepted, it has not yet been published in any peer-reviewed journal, and as such we keep the condition on the domain of attraction in the statement of Theorem 3.2.5.

We also point out that in [30] the authors observe U as defined in (3.2.3) satisfies the HJB PDE. However, the authors there derive the HJB via a known approximation to the Kolmogorov forward equation and do not develop the connection to large deviations, nor do they ever mention that it is a Hamilton-Jacobi equation. Moreover, their result does not seem to be published or have made its way through the CRN community.

Finally, we point out that the "reverse" direction is also true: if for some $c \in \mathbb{R}^d_+$ (3.2.3) is a solution of the Hamilton-Jacobi-Bellman equation for any given CRN, then that CRN is complex-balanced and c is the unique equilibrium in its stoichiometric compatibility class. This result is not much of interest in the present context and so we do not state it formally.

3.2.2 Subsolutions for one-dimensional networks

This section studies one-dimensional CRNs. Though they are much simpler than their higher-dimensional counterparts, one-dimensional CRNs can exhibit most kind

of rare events we are interested in. The main goal of this section is to demonstrate how to build reaction networks with prescribed rare event behavior, and subsequently describe a method for constructing time-independent subsolutions for general one-dimensional CRNs. Taken together, these facts will allow us to construct asymptotically optimal splitting schemes for several different types of rare events.

We briefly outline the section. First, in Section 3.2.2.1 we gather facts about one-dimensional systems and show how every suitable one-dimensional polynomial corresponds to the steady state dynamics $\dot{x} = \Phi(x)$ of some CRN. In Section 3.2.2.2 we identify the quasipotential and discuss its use as a subsolution for the escape problem over increasing time intervals. We then derive an explicit expression for the gradient of the quasipotential in the case of birth-death systems in Section 3.2.2.3 and we apply it to the construction in Section 3.2.2.1.

3.2.2.1 From polynomials to CRNs.

This section discusses how to construct CRNs with certain desired properties. The class of one-dimensional CRNs includes well-known models. For instance, the $A \rightleftarrows \emptyset$ network corresponds to the $M/M/\infty$ queue in Kendall's notation, and CRNs with jump sizes ± 1 (i.e., $\mathcal{R} = \{\pm 1\}$) are birth-death processes.

It is natural to ask more generally what kind of dynamics a one-dimensional CRN can describe, in either a stochastic or deterministic sense. The following result is a step in this direction. It tells us that we can associate to a generic polynomial $q(x)$ with $q(0) \geq 0$ a CRN whose deterministic dynamics $\Phi(x) = D_p H(x, 0)$ are given by $q(x)$. The higher-dimensional case was solved by J. Tóth and V. Hárs in 1979 [42] and also provides a constructive proof. The following special case is our personal

construction and is included because it will be used throughout to build networks.

Lemma 3.2.6. *To any polynomial* $q(x) = \sum_{i=0}^{k} a_i x^i$ *with* $q(0) \geqslant 0$*, we can associate a single-species CRN whose mass-action dynamics are given by*

$$\dot{x}(t) = q(x(t)).$$

In particular, the equilibrium points of the CRN are the real roots of q*.*

Proof. Let $k = \deg(q)$ be the degree of the polynomial $q(x)$, and let $\{a_i\}_{i=0}^{k}$ be the coefficients of q. By assumption, $a_0 \geqslant 0$. We define a reaction network and check that it works. Let $\kappa_i = |a_i|$ and define the outflow and inflow subsets $N = \{i : a_i < 0\}$ and $P = \{1, \ldots, k\} \setminus N$. Let

$$\mathcal{C} = N \cup P \cup (N - 1) \cup (P + 1),$$

where $N - 1 = \{i - 1 : i \in N\}$ and $P + 1 = \{i + 1 : i \in P\}$. Given this set of complexes $\mathcal{C}$, let the reactions be

$$\mathcal{R} = \{(i, i-1) : i \in N\} \cup \{i, i+1 : i \in P\},$$

and choose the rate constants to be

$$\begin{aligned} \kappa_{i,i-1} &= |a_i|, \quad i \in N, \\ \kappa_{i,i+1} &= a_i, \quad i \in P. \end{aligned}$$

Note that $a_0 > 0$ implies $0 \in P$.

There is one reaction $\zeta_{i,i\pm1}$ for each coefficient a_i, and mass-action kinetics are

given by

$$\begin{aligned}\dot{x} &= \sum_{(a,b)\in\mathcal{R}} \lambda_{a,b}(x)\zeta_{a,b} \\ &= \sum_{(i,i+1)\in\mathcal{R}} \kappa_{i,i+1}x^i - \sum_{(i,i-1)\in\mathcal{R}} \kappa_{i,i-1}x^i \\ &= \sum_{i\in P} |a_i|x^i - \sum_{i\in N} |a_i|x^i.\end{aligned}$$

By definition of N and P, we can remove absolute values on the a_i and combine the sums on the last line to recover $q(x)$. □

Remark. This is not the only way of constructing a reaction network whose determinstic dynamics $\dot{x} = \Phi(x)$ satisfy $\Phi(x) = q(x)$. For instance, instead of associating the reaction $A \to 2A$ at rate $\kappa = 2$ to the monomial $2x$, one could to associate the reaction $A \to 3A$ at rate $\kappa = 1$. This leads to the same deterministic dynamics $\Phi(x)$, but the corresponding Hamiltonian terms are different: $2x(e^p - 1)$ and $x(e^{2p} - 1)$, respectively.

A consequence of Lemma 3.2.6 is that we can construct one-dimensional reaction networks with whatever admissible deterministic dynamics $\dot{x} = \Phi(x)$ we like. This is captured in Corollary 3.2.7.

Corollary 3.2.7. *Given K non-negative numbers $c_1, \ldots, c_K \geq 0$, not necessarily distinct, there is a CRN with steady state dynamics Φ which satisfies $\{x : \Phi(x) = 0\} = \{c_1, \ldots, c_K\}$.*

Proof. Let $Z = \{i : c_i = 0\}$ and recall that $|Z|$ denotes the cardinality of Z. Let

$$q(x) = (-1)^{K-|Z|} \prod_{i=1,\ldots,K, i\notin Z} (x - c_i). \tag{3.2.4}$$

If $|Z| \geq 1$ then $q(0) = 0$, otherwise a simple parity check shows $q(0) > 0$. By Lemma 3.2.6 there is a CRN corresponding to $q(x)$. □

It follows immediately from Corollary 3.2.7 that one-dimensional reaction networks have the capacity for an arbitrarily large but finite number of steady states. Of course, not all steady states can be stable, because whatever polynomial describes the steady state dynamics has to change sign between two stable steady states. We illustrate this issue in the next example.

Example 3.2.8. Suppose we wish to study a metastability problem, where the stochastic process X^n starts at a stable steady state c_2 and can go to one of two stable steady states c_1 or c_3 with $c_1 < c_2 < c_3$. In order for c_i to be stable for each $i = 1, 2, 3$, there must be neighbourhoods N_i of c_i such that $\Phi(x) > 0$ for $x \in N_i \cap \{x < c_i\}$ and $\Phi(x) < 0$ for $x \in N_i \cap \{x > c_i\}$. In particular, whatever polynomial $q(x)$ describes the deterministic dynamics, it will need to change sign between c_1 and c_2 and again between c_2 and c_3. If $q(x)$ is of the form (3.2.4), then there will necessarily be $c_{12} \in (c_1, c_2)$ and $c_{23} \in (c_2, c_3)$ which are unstable steady states. See Figure 3.2 for a simple one-dimensional phase diagram of $\Phi(x)$.

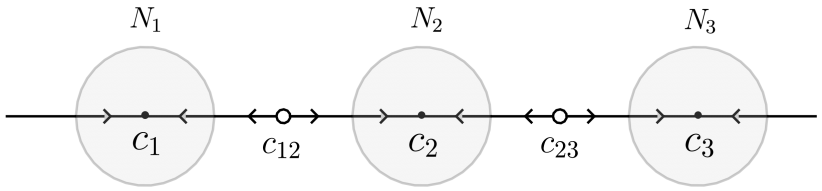

Figure 3.2: A one-dimensional phase diagram for the polynomial determined in Corollary 3.2.7 given the set $\{c_1, c_{12}, c_2, c_{23}, c_3\}$ of steady states. The states c_1, c_2, c_3 are stable, and c_{12}, c_{23} are unstable.

3.2.2.2 Subsolutions based on the quasipotential

This section identifies subsolutions for one-dimensional CRNs based on the quasipotential relative to a stable steady state. Recall from Definition 1.5.2 that the quasipotential relative to a stable steady state $c \in \mathbb{R}_+$ is defined as

$$Q(c,x) = \inf\left\{\int_0^T L(\varphi(s),\dot{\varphi}(s))ds : \varphi(0) = c, \varphi(T) = x, T < \infty\right\}. \tag{3.2.5}$$

Furthermore, in Section 2.1.4.1 we demonstrated how the quasipotential could be used to construct a subsolution for the problem of estimating $\mathbb{P}_{x_0}(X^n(t) \notin D \text{ for some } t \in [0,T])$ when x_0 is a rest point. When the problem is taken over increasing time intervals $[0,T(n)]$ with $T(n)$ being polynomial in n, Theorem 2.2.15 shows the corresponding splitting scheme is asymptotically optimal in the sense of Definition 2.1.2.

First, we explain how to compute the quasipotential, and then we explain how to "stitch" together the quasipotentials corresponding to two different stable steady states.

Computing the quasipotential. The way we identify the quasipotential relative to a stable steady state c is by solving the corresponding time-independent HJB,

$$H(x, DQ(x)) = 0, \quad Q(c) = 0.$$

As mentioned previously, it is a classical result of Freidlin and Wentzell's [35] that certain solutions of the time-independent HJB are equal to the quasipotential on a suitable neighourhood of the steady state c; see Theorem 1.5.4. A special case of this theorem tailored to the one-dimensional case is stated and proved in Corollary

1.5.6. We restate the Corollary here for convenience.

Corollary 1.5.6. *Consider a one-dimensional domain $D = (a, b) \subset \mathbb{R}$ which contains a single steady state $c \in D$ and assume c is stable. Let $p(x)$ be a continuous function with $p(x) < 0$ for $x < c$, $p(x) > 0$ for $x > c$ and which satisfies $H(x, p(x)) = 0$ for all $x \in D$. Then $p(x)$ is unique and*

$$U(x) = \int_c^x p(y)dy \tag{3.2.6}$$

satisfies the conditions of Theorem 1.5.4 and $U(x) = Q(c, x)$ for all $x \in D$.

In other words, in order to obtain the gradient of the quasipotential it suffices to find a function $p(x)$ such that $H(x, p(x)) = 0$, with $p(c) = 0$ and p increasing on a neighbourhood of c. It is then simply a question of numerical integration to determine the function $Q(c, x)$. Moreover, the existence of such a function is guaranteed by Lemma 1.5.7, which we restate here for convenience.

Lemma 1.5.7. *Let $c > 0$ be a stable steady state of a one-dimensional reaction network. There is a neighbourhood D of c and a continuous function $p : D \to \mathbb{R}$ which satisfies the conditions of Corollary 1.5.6.*

To illustrate Corollary 1.5.6, we identify the gradient of the quasipotential for Example 13 in [3]. Our procedure is much simpler than the one outlined in [3], which is based on computing the stationary distribution and taking an appropriate large deviations scaling; see Theorem 1.5.3.

Example 3.2.9. Consider the reaction network

$$A \to \emptyset \to 2A,$$

with associated Hamiltonian

$$H(x,p) = x(e^{-p} - 1) + e^{2p} - 1.$$

This is a cubic equation in $y(x) = e^{p(x)}$, and there are three real solutions in terms of $y(x)$:

$$y(x) = \pm\frac{1}{2}\left(\sqrt{1+4x} - 1\right), \quad y(x) = 1.$$

Since $y(x) = e^{p(x)}$ and $p(x) \in \mathbb{R}$, the solution $y(x) = -\left(\sqrt{1+4x}+1\right)/2$ is not admissible. The solution $y(x) = 1$ corresponds to the law of large numbers, and $p(x) = \log y(x) = 0$ does not satisfy the conditions of Corollary 1.5.6. The remaining root $y(x) = (\sqrt{1+4x}+1)/2$ satisfies $H(x, \log y(x)) = 0$, and hence $H(x, p(x)) = 0$. Moreover, it is easily verified that $p(x) = \log((\sqrt{1+4x}+1)/2)$ satisfies the conditions of Corollary 1.5.6. One can conclude that the gradient of the quasipotential Q satisfies

$$DQ(x) = \log\left(\sqrt{1+4x} - 1\right) - \log 2.$$

Indeed,

$$\begin{aligned}
&H(x, DQ)\\
&= x\left(\frac{1}{\frac{1}{2}\left(\sqrt{1+4x}-1\right)} - 1\right) + \left(\frac{1}{2}\right)^2\left(\sqrt{1+4x}-1\right)^2 - 1\\
&= x\left(\frac{\sqrt{1+4x}+1}{2x} - 1\right) + \left(\frac{1}{2}\right)^2\left(2 + 4x - 2\sqrt{1+4x}\right) - 1\\
&= 0.
\end{aligned}$$

Stitching together quasipotentials. The quasipotential is defined relative to an initial stable steady state and the cost of trajectories is measured relative to this steady state; see (3.2.5). It may be the case that there are multiple stable steady states in the domain to be escaped. In this case, some care may be required to "stitch" together the solutions $p(x)$ in Corollary 1.5.6 which yield the associated quasipotentials. Fortunately, this is straightforward enough in one-dimensional networks.

To illustrate, consider Example 3.2.8, which has steady states $c_1 < c_{12} < c_3 < c_{23} < c_3$ and consider the problem of estimating $\mathbb{P}_{c_1}(X^n(t) \notin D, \text{ for some } t \in [0,T])$ for $D = [0, c_{23})$. The domain D includes the steady states c_1 and c_2, and the quasipotential relative to c_1 gives the correct cost to go from c_1 to c_{23}; see Figure 3.3. Let $p_{c_i}, i = 1,2,3$ denote the solutions which satisfy the conditions of Corollary

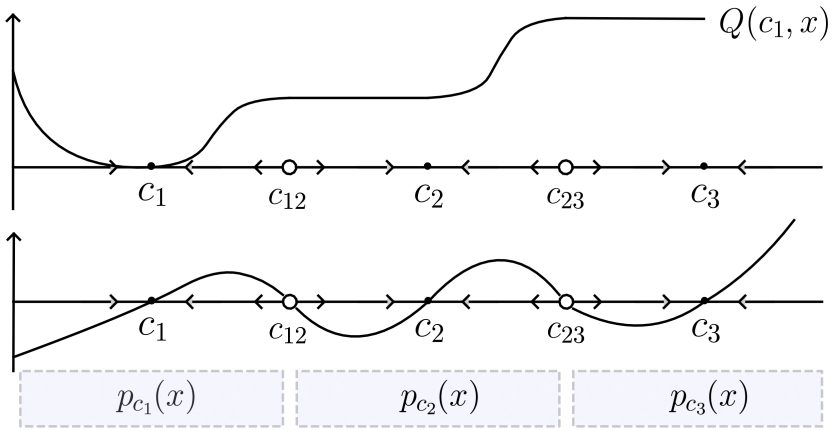

Figure 3.3: Top: sketch of the quasipotential with respect to the stable steady state c_1. The cost to go between c_{12} and c_2 is zero since every $x \in (c_{12}, c_2)$ is attracted to c_2 through the noiseless dynamics. Bottom: sketch of the function $p(x)$ which satisfies $H(x, p(x)) = 0$ and $p(x) \neq 0$ unless $\Phi(x) = 0$. The shaded boxes below the axis delineate the basins of attraction of the stable states and therefore the region where the expression (3.2.6) defines the quasipotential.

1.5.6 for the steady states $c_i, i = 1,2,3$. Each solution has an associated domain: for p_{c_1}, it is $[0, c_{12})$; for p_{c_2}, it is (c_{12}, c_{23}), and for p_{c_3}, it is (c_{23}, ∞). On their respective domains, the solutions p_{c_i} satisfy the conditions of Corollary 1.5.6, and their integral

as defined in (3.2.6) gives the quasipotential $Q(c_i, x)$.

In order to obtain $Q(c_1, x)$ by integrating the solutions p_{c_i}, we simply ignore the fact that p_{c_2} is negative on (c_{12}, c_2). It can be shown that

$$Q(c_1, x) = \int_{c_1}^{x} p_{c_1}(y)dy + \int_{c_2}^{x \wedge c_2} p_{c_2}(y)dy. \tag{3.2.7}$$

We will refer to this process as "stitching together" the quasipotential. In the next section we will derive an explicit expression for the solution $p(x)$ in a specific case, and the stitching process will simply be a question of taking the solution to be zero on some intervals.

Producing a subsolution. Recall from Definition 1.4.5 that a continuously differentiable time-independent subsolution $\bar{V}$ for the problem of escaping D over a time interval $[0, T]$ satisfies the HJB PDE

$$H(x, -D\bar{V}(x)) \leq 0, \quad x \in D, \tag{3.2.8}$$

and $\bar{V}(x) \leq 0$ for $x \notin D$. As discussed in Section 2.1.4.1, such a subsolution can be obtained by taking the negative of the quasipotential (so that it satisfies the PDE with equality) and shifting it by a constant (so that it satisfies the boundary condition). For instance, suppose D is a neighbourhood of a stable steady state c and every point in D is attracted to c. A subsolution for the problem of estimating $\mathbb{P}_c(X^n(t) \notin D \text{ for some } t \in [0, T])$ is given by

$$\bar{V}(x) = -Q(c, x) + k,$$

where $k = \inf_{x \in \partial D} Q(c, x)$. For further details consult Section 2.1.4.1.

3.2.2.3 Birth-death processes

The previous section gave an expression for the quasipotential in terms of a solution $p(x)$ to $H(x, p(x)) = 0$. We now derive an explicit expression for $p(x)$ in the case of birth-death processes, which are one-dimensional CRNs with jump sizes ± 1. Note that our method for associating CRNs to polynomials, outlined in Lemma 3.2.6, produces a birth-death process. Thus we will have a streamlined process for building a CRN with desired properties, computing the gradient of its quasipotential relative to any stable steady states, and generating the associated time-independent subsolutions.

The generic form of the Hamiltonian for a one-dimensional CRN with jump sizes ± 1 is

$$H(x,p) = \sum_{(n,n+1)\in\mathcal{R}} \kappa_{n,n+1} x^n (e^p - 1) + \sum_{(n,n-1)\in\mathcal{R}} \kappa_{n,n-1} x^n (e^{-p} - 1), \tag{3.2.9}$$

and we will show that the quasipotential around a stable steady state c is

$$Q(c,x) \doteq \int_c^x \log\left(\frac{\sum_{(n,n-1)\in\mathcal{R}} \kappa_{n,n-1} s^n}{\sum_{(n,n+1)\in\mathcal{R}} \kappa_{n,n+1} s^n}\right) \mathrm{ds}. \tag{3.2.10}$$

The expression (3.2.10) is well known but is typically obtained by writing down the stationary distribution of the process and taking a scaling limit. The stationary distribution for a general birth-death process is available; see for instance [36], p. 237 (7.1.13). Theorem 1.5.3 shows that by applying a large deviations style scaling to the stationary distribution and taking limits, one obtains the quasipotential relative to a stable steady state. In [3] the authors employ this approach to obtain a Lyapunov function around some stable steady state c.

Our approach is simpler than the one taken in [3].

Lemma 3.2.10. *Let c be a stable fixed point of the noiseless dynamics for a birth-death process. The gradient of the quasipotential relative to c is given for x in the basin of attraction of c by (3.2.10).*

Proof. Fix c and let D denote the basin of attraction of c. We find the roots of (3.2.9) in terms of e^p. The equation $H(x, p(x)) = 0$ can be made algebraic by using the substitution $y(x) \doteq e^{p(x)}$ and clearing the denominators from powers of y. For $x \in D$, consider

$$\begin{aligned} K(x,y) &\doteq\ y\ H(x, \log y) \\ &= \sum_{(n,n+1)\in\mathcal{R}} \kappa_{n,n+1} x^n (y^2 - y) + \sum_{(n,n-1)\in\mathcal{R}} \kappa_{n,n-1} x^n (1-y). \end{aligned}$$

Write $A = A(x) = \sum_{n\in P} \kappa_{n,n+1} x^n$ and $B = B(x) = \sum_{n\in N} \kappa_{n,n-1} x^n$. Then

$$K(x,y) = A(x)y^2 - (A(x) + B(x))y + B(x).$$

The solution to the quadratic is

$$y = \frac{1}{2A}\left[A + B \pm \sqrt{(A-B)^2}\right].$$

There are four cases and two possible solutions:

	$A > B$	$A \leqslant B$
$+$	1	B/A
$-$	B/A	1

.

Since $p(x) = \log y(x)$, we find that a solution is

$$p(x) = \log\left(\frac{B(x)}{A(x)}\right) = \log\left(\frac{\sum_{(n,n-1)\in\mathcal{R}} \kappa_{n,n-1}x^n}{\sum_{(n,n+1)\in\mathcal{R}} \kappa_{n,n+1}x^n}\right)$$

Observe that $\Phi(x) = 0$ if and only if $A(x) = B(x)$. If c is stable, then $\Phi(c) > 0$ for $x < c, x \in D$ and $\Phi(c) < 0$ for $x > c, x \in D$, which implies respectively that $A(x) < B(x)$ and $A(x) > B(x)$. Consequently, $p(x) < 0$ for $x < c, x \in D$ and $p(x) > 0$ for $x > c, x \in D$. In particular, p satisfies the conditions of Corollary 1.5.6 and (3.2.10) is equal to the quasipotential relative to c for $x \in D$. □

3.2.3 Subsolutions based on moderate deviations

In Section 2.3 we discussed the moderate deviations approximation for the estimation of terminal escape probabilities, that is, problems of the form

$$\mathbb{P}(X^n(T) - X^0(T) \notin D), \tag{3.2.11}$$

where X^n is a stochastic process and X^0 is its deterministic approximation, the law of large numbers limit of X^n. In particular, Theorem 2.3.7 established conditions on subsolutions for which the corresponding moderate deviations-based splitting schemes were asymptotically optimal. In this section, we detail the general construction of such subsolutions for CRNs.

We recall that the moderate deviations approximation can be thought of as a quadratic expansion of the cost around the law of large numbers dynamics. In contrast to the two previous techniques for finding subsolutions, which require either a complex-balancing structure or are limited to one-dimensional situations, sub-

solutions are generally available for the moderate-deviations approximation. The drawback is that one cannot use the approximation to design splitting schemes for events which fall out of the moderate deviations regime, which will be illustrated in the examples in Section 3.3.

The section is outlined as follows: first, in Section 3.2.3.1 we re-introduce the basic objects for moderate deviations now in the notation of CRNs; in Section 3.2.3.2, we describe a procedure to construct asymptotically optimal time-dependent subsolutions for terminal escape problems.

3.2.3.1 Basic objects

We first identify the objects required for a moderate deviations-based subsolution for a splitting scheme. We consider a general d-dimensional CRN G_κ with Hamiltonian

$$H(x,p) = \sum_{(a,b)\in\mathcal{R}} \kappa_{a,b} x^a (e^{\langle \zeta_{a,b}, p\rangle} - 1).$$

Expand around LLN. First we linearize around the law of large numbers, which corresponds to the noiseless dynamics

$$X^0(t) = x_0 + \int_0^t b(X^0(s))ds.$$

The drift $b(x)$ corresponds to the deterministic dynamics $\Phi(x) = D_p H(x,0)$, so throughout this section $b(x)$ will be defined by

$$b(x) = D_p H(x,0) = \sum_{(a,b)\in\mathcal{R}} \zeta_{a,b}\kappa_{a,b} x^a.$$

The drift $b(x)$ determines the state transition function Ψ from Definition 2.3.4, which must satisfy

$$\frac{d}{dt}\Psi(s,t) = Db(X^0(t))\Psi(s,t), s < t, \quad \Psi(s,s) = I,$$

where I denotes the $d \times d$ identity matrix. We remark for later use that that if x_0 is a steady state, i.e. $\Phi(x_0) = 0$, then

$$\Psi(s,t) = \exp(Db(x_0)(t-s)).$$

Quadratic Gaussian cost. The constructions in Section 2.3 are stated in terms of a centered cumulant generating function H_c, which is related to H via

$$H_c(x,p) = H(x,p) - \langle p, b(x)\rangle.$$

In particular, we must compute the Hessian of the centered cumulant generating function H_c from equation (2.3.2),

$$A(x) = D_p^2 H_c(x,0).$$

Observe that we obtain the same matrix $A(x)$ whether we take the Hessian of H or H_c. The Hessian of H evaluated at $(x,0)$ is

$$D_p^2 H(x,0) = \sum_{(a,b)\in\mathcal{R}} \kappa_{a,b} x^a \zeta_{a,b} \zeta_{a,b}^T,$$

where the superscript T denotes a transpose.

Centered process. The moderate deviations approximation is written in terms of the centered process, which in this section is defined to be

$$Y^n(t) \doteq X^n(t) - X^0(t), \quad Y^n(0) = 0.$$

We do not include any n-dependent scaling, because it was shown in Section 2.3.4 that the specific scaling is irrelevant for the design of the algorithm. The problem of interest (3.2.11) is therefore $\mathbb{P}(Y^n(T) \notin D)$. Note that $Y^n(0) = 0$ for all $n \geq 0$, regardless of the initial condition for X^n. Though the subsolutions will be constructed for the centered process Y^n, we will frequently write the rare event in terms of the original process X^n, as this is the object of interest and often provides a physical interpretation of the event. For this reason we reserve the y, z variables for the "centered" coordinates and x for the "original" ones.

Subsolution property and point-to-point cost. The subsolution property in the present context requires that

$$\partial_t \bar{V}(t,y) \geq \frac{1}{2}\|D_y \bar{V}(t,y)\|^2_{A(X^0(t))} - \langle D_y \bar{V}(t,y), D_x b(X^0(t))y\rangle \tag{3.2.12}$$

with $\bar{V}(T,y) \leq 0$ for $y \notin D$. We recall Theorem 2.3.5, which states that the cost to go from (t,y) to (T,z) for some fixed $t, T > 0$ and $y, z \in \mathbb{R}^d$ is

$$C(t,y,T,z) \doteq \frac{1}{2}\|z - \Psi(t,T)y\|^2_{M^{-1}_{t,T}},$$

where

$$M_{t,T} = \int_t^T \Psi(s,T)A(X^0(s))\Psi(s,T)^T ds.$$

3.2.3.2 Time-dependent subsolutions with affine terminal cost.

We now outline the construction of time-dependent subsolutions for terminal escape from a domain $D \subset \mathbb{R}^d$,

$$\mathbb{P}(X^n(T) - X^n(0) \notin D) = \mathbb{P}(Y^n(T) \notin D).$$

The goal is to construct $\bar{V}$ such that the resulting splitting scheme is asymptotically optimal in the sense of Definition 2.1.2. Theorem 2.3.7 from Section 2.3 establishes that $\bar{V}$ is asymptotically optimal if

$$\bar{V}(0,0) = \inf_{z \notin D} C(0,0,T,z). \tag{3.2.13}$$

To construct a $\bar{V}$ which achieves (3.2.13) for a general domain D, we proceed in five steps. First, we recall a two-parameter family of subsolutions and rewrite (3.2.13) for the case where D is half-space. Next, we solve the optimization problem determined by the right-hand side of (3.2.13) for this special case. We then identify the pair of parameters that yields an asymptotically optimal subsolution, and subsequently generalize the construction to the case where D is an intersection of half-spaces. Finally, we consider the approximation of a general domain D. A procedure which summarizes the numerical algorithm is described at the end of the section.

Step one: D as a half-space. Let $v \in \mathbb{R}^d$ and $\delta \in \mathbb{R}$, and define the half-space

$$\mathcal{H} \doteq \{z : \langle v, z \rangle < \delta\}.$$

To construct a subsolution $\bar{V}$ which satisfies (3.2.13) when $D = \mathcal{H}$, we will use Theorem 2.3.6, which implies that for given $\xi \in \mathbb{R}^d$ and $c \in \mathbb{R}$,

$$\bar{V}^{\xi,c}(t,y) \doteq \langle \xi, \Psi(t,T)y \rangle + c - \frac{1}{2}\|\xi\|^2_{M_{t,T}}$$

satisfies the PDE (3.2.12) in a classical sense with equality, but with terminal value

$$\bar{V}^{\xi,c}(T,y) = \langle \xi, y \rangle + c.$$

The function $\bar{V}^{\xi,c}$ can also be written as

$$\bar{V}^{\xi,c}(t,y) = \inf_{z \in \mathbb{R}^d} \{C(t,y,T,z) + \langle \xi, z \rangle + c\}.$$

Thus, to satisfy (3.2.13) when $D = \mathcal{H}$, the goal will be to choose ξ and c such that

$$\inf_{z \in \mathbb{R}^d} \{C(0,0,T,z) + \infty \cdot 1\{\langle v, z \rangle < \delta\}\} = \inf_{z \in \mathbb{R}^d} \{C(0,0,T,z) + \langle \xi, z \rangle + c\}. \quad (3.2.14)$$

Step two: Lagrange multipliers. We determine the value of the left-hand side of (3.2.14) by identifying the minimizer z explicitly. Let $M = M_{0,T}$ for notational convenience.

The infimization problem

$$\inf_{z \in \mathbb{R}^d} \{C(0,0,T,z) + \infty \cdot 1\{\langle v, z \rangle < \delta\}\}$$

has a unique minimizer which is given by

$$z_{v,\delta} \doteq \delta \frac{Mv}{\|v\|^2_M}.$$

We only sketch the proof. The problem is equivalent to minimizing $\|z\|^2_{M^{-1}}/2$ over z such that $\langle v, z\rangle \geqslant \delta$. We associate the Lagrangian

$$L(z, \lambda) = \frac{1}{2}\|z\|^2_{M^{-1}} + \lambda(\langle v, z\rangle - \delta),$$

with $\lambda \leqslant 0$ to ensure $\langle v, z\rangle \geqslant \delta$. Since M is symmetric, we find $D_z L(z, \lambda) = z^T M^{-1} + \lambda v^T = 0$, and multiplying by M on the right,

$$z = -\lambda M v.$$

Thus,

$$L(z(\lambda), \lambda) = \frac{1}{2}\| - \lambda M v\|^2_{M^{-1}} + \lambda(\langle v, -\lambda M v\rangle - \delta) = -\frac{\lambda^2}{2}\|v\|^2_M - \lambda\delta,$$

and setting the derivative with respect to λ of $L(z(\lambda), \lambda)$ equal to 0,

$$\lambda = -\frac{\delta}{\|v\|^2_M}, \quad z_{v,\delta} = \delta\frac{Mv}{\|v\|^2_M}.$$

The Karush-Kuhn-Tucker conditions guarantee the existence of a multiplier λ for the inequality constraint at an optimal z. Since the inequality constraint is convex and the function to be minimized is quadratic, the existence of the multiplier is sufficient to guarantee that $z_{v,\delta}$ is in fact optimal; see for instance Theorem 2.1.4. Chap VII.2. in [43].

Step three: asymptotic optimality on $\mathcal{H}$. Since M is symmetric, the gradient at a general point z is equal to

$$D_z C(0, 0, T, z) = D_z\frac{1}{2}\|z\|^2_{M^{-1}} = z^T M^{-1},$$

so the gradient at the minimizer $z_{v,\delta}$ is

$$\delta \frac{v^T}{\|v\|_M^2}.$$

We take ξ to be the negative of the gradient of the cost $C(0,0,T,z)$ at the minimizer and choose c to enforce the boundary condition $\bar{V}^{\xi,c}(T,z) = \langle \xi, z\rangle + c \leqslant 0$ for $z \notin \mathcal{H}$:

$$\xi = -\delta \frac{v}{\|v\|_M^2}, \quad c = \frac{\delta^2}{\|v\|_M^2}. \tag{3.2.15}$$

This determines the function $\bar{V}^{\xi,c}$ for all $t \in [0,T]$ and $y \in \mathbb{R}^d$. Moreover,

$$\bar{V}^{\xi,c}(0,0) = c - \frac{1}{2}\|\xi\|_M^2 = \frac{\delta^2}{\|v\|_M^2} - \frac{\delta^2}{2\|v\|_M^2} = \frac{1}{2}\|z_{v,d}\|_{M^{-1}}^2 = \inf_{z\notin\mathcal{H}} C(0,0,T,z), \tag{3.2.16}$$

which is equivalent to (3.2.13) for $D = \mathcal{H}$.

Step four: D as the intersection of half-spaces. Let $v_i, i = 1,\dots,H$ be vectors in $\mathbb{R}^d$ and $\delta_i \in \mathbb{R}$, and consider the case where D is given by the intersection of half-spaces,

$$D = \bigcap_{i=1}^{H} \mathcal{H}_i, \quad \mathcal{H}_i = \{\langle v_i, z\rangle < \delta_i\}.$$

For each i, if we choose ξ_i and c_i according to (3.2.15), then $\bar{V}^{\xi_i,c_i}$ achieves asymptotic optimality for terminal escape from $\mathcal{H}_i$, that is, the minimal cost to exit $\mathcal{H}_i$ starting at $(0,0)$ is $\bar{V}^{\xi_i,c_i}(0,0)$:

$$\bar{V}^{\xi_i,c_i}(0,0) = \inf_{z\notin\mathcal{H}_i} C(0,0,T,z).$$

Since D is the intersection of sets $\mathcal{H}_i$, the minimal cost to exit D must be the minimum over the cost to exit each $\mathcal{H}_i$:

$$\inf_{z\notin\cap_i\mathcal{H}_i} C(0,0,T,z) = \min_{i=1,\dots,H} \inf_{z\notin\mathcal{H}_i} C(0,0,T,z).$$

An asymptotically optimal subsolution for escaping D is therefore given by

$$\bar{V}(t,y) \doteq \min_{i=1}^{H} \bar{V}^{\xi_i,c_i}(t,y), \quad (t,y) \in [0,T] \times \mathbb{R}^d.$$

Step five: general set D. Recall that in the PDE characterization of a subsolution (3.2.12), the only dependence on D is in the boundary conditions. The boundary conditions require that the "cost to escape" at the terminal time be non-positive if escape has occurred, i.e. $V(T,y) \leqslant 0$ for $y \notin D$. It is therefore appropriate to *under*-estimate the domain D by an approximation $\hat{D} \subset D$. If $\bar{V}$ is a subsolution for the problem of escaping $\hat{D}$ at time T, then $\bar{V}(T,y) \leq y$ for $y \notin \hat{D}$ and in particular for all $y \in D^c \subset \hat{D}^c$. Thus $\bar{V}$ is also a subsolution for the problem of terminal escape from D.

Say we approximate D by an intersection of half-spaces, $\hat{D} = \cap_i \mathcal{H}_i$. The subsolution $\bar{V}$ will be asymptotically optimal for problem of terminal escape from D so long as $\hat{D}$ contains the most likely point of escape. Pre-asymptotically, however, the decay rate can take a hit if $\hat{D}$ does not approximate D "well enough". It is only in practice that one sees how well the scheme is performing. Another restriction on the

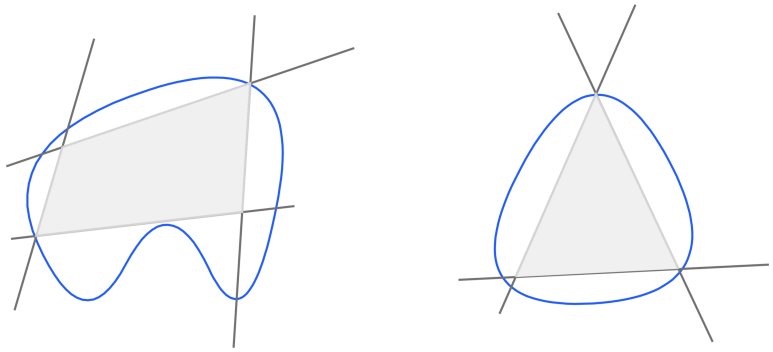

Figure 3.4: The approximation of two sets in $\mathbb{R}^2$.

domain is that we assume $D \subset \mathbb{R}^d$ is genuinely d-dimensional, in the sense that its

d-dimensional Euclidean volume is positive. If D is a subset of a plane or simplex, then one of the variables can be written as a linear combination of the others.

Procedure for constructing subsolutions. Suppose we are given a CRN G_κ and an arbitrary set $D \subset \mathbb{R}^d$. How does one go about numerically constructing a subsolution for the moderate deviations approximation? There are essentially two things to do: approximate the set D as the intersection of half-spaces, and compute $\Psi(t,T)$ and $M_{t,T}$, $t \in [0,T]$.

Approximating the domain D is the hardest part of the procedure, in the sense that it is the only step in which problem-specific constructions are required. If $D \subset \mathbb{R}^d$ but its d-dimensional volume is zero (for instance, it is supported on a hyperplane), it must be reparametrized in terms of fewer variables. This is not necessary, but does allow for a simpler expression of the optimal parameters (3.2.15).

Regardless of the construction, it will be necessary to compute the one-parameter family of matrices $\{M_{t,T} : t \in [0,T]\}$. Perhaps the easiest way to do this numerically is to integrate $\Psi(t,T)A(X^0(t))\Psi(t,T)^T$ from $t = 0$ to $t = T$:

$$M_{t,T} = \int_t^T \Psi(s,T)A(X^0(s))\Psi(s,T)^T ds.$$

Here the equality of matrices is understood componentwise, and we recall $\Psi(t,T)$ is the state transition function, which satisfies the matrix differential equation

$$\frac{d}{dt}\Phi(t,T) = Db(X^0(t))\Phi(t,T), \quad \Phi(t,t) = I_{d\times d}, \quad t \in [0,T].$$

In practice one first computes $\Psi(\cdot,T)$ and then $M_{\cdot,T}$. Note also that if $X^0(0)$ starts at a steady state of the law of large numbers dynamics, then $X^0(t) = X^0(0)$ for all

$t \geqslant 0$. In this case, $A(X^0(t))$ and $Db(X^0(t))$ are constant matrices, which simplifies the required computations.

3.3 Rare events in reaction networks

We now use the techniques developed in Section 3.2 to identify subsolutions for specific rare event problems arising in chemical reaction networks. We focus on three different kinds of rare events. First, in Section 3.3.1, we consider the problem of escape from rest points, i.e., the escape from a neighbourhood of a stable steady state. In Section 3.3.2 we use the techniques for one-dimensional reaction networks developed in Section 3.2.2 to study metastability issues, which is the passage of the stochastic process between different stable steady states. Finally, in Section 3.3.3, we explore rare events related to the behavior of the CRN near the boundary, in particular the event of *hitting* a boundary.

3.3.1 Escaping rest points

We first recall the notion of a rest point, and then present two examples which use two different techniques. A rest point c for a d-dimensional CRN is a stable steady state, that is, a point c which satisfies $\Phi(c) = 0$ and is locally stable in the sense of Definition 1.5.1. Started at the rest point, the law of large numbers will stay there forever, and the rare event of interest is the probability of escape, either at some future time or over a time interval. For some domain $D \subset \mathcal{S}_c$ containing c (see Definition 3.1.4) and contained in the domain of attraction of c (see Definition

1.5.1), we will consider problems of the form

$$\mathbb{P}_c(X^n(t) \notin D \text{ for some } t \in [0, T(n)]),$$

or

$$\mathbb{P}_c(X^n(T) \notin D).$$

For the first problem of escape over $[0, T(n)]$, the sequence of time $T(n)$ is expected to diverge to infinity and satisfy a subexponential growth condition; see Condition 2.2.6. In this case it is appropriate to use techniques from Section 3.2.1. For the second problem of terminal escape, it may be appropriate to use the moderate deviations approximation described in Section 3.2.3.

3.3.1.1 Using techniques from complex-balanced networks

Example 3.3.1. This example illustrates that certain networks which seem complicated may end up having a simple form for a subsolution, which makes the sampling of escape from the rest point a tractable problem using the splitting technique. We also demonstrate how the Deficiency Zero theorem (Theorem 3.2.4) can be applied to determine that a network is complex-balanced, and subsequently how the complex-balanced condition (3.2.1) can simplify the identification of the unique steady state.

CRN and rare event. Consider the reaction network G_κ defined by the graph

$$\begin{array}{ccc} A+2B & \to & 3C \\ \uparrow & & \downarrow \\ D & \rightleftarrows & C+E \end{array}$$

$$\begin{array}{ccc} C & \leftarrow & A \\ & \searrow \quad \nearrow & \\ & 2E & \end{array}$$

with all rate constants equal to one. Define the matrix M whose columns are the reaction vectors $\zeta_{a,b}$, $(a,b) \in \mathcal{R}$.

$$M = \begin{pmatrix} -1 & 1 & 0 & 0 & 0 & -1 & 0 & 1 \\ -2 & 2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 3 & 0 & 1 & -1 & -2 & 1 & -1 & 0 \\ 0 & -1 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 & 0 & 2 & -2 \end{pmatrix}.$$

This matrix has rank 5, there are 7 complexes, and there are 2 connected components of the graph, so the deficiency is $\delta = 7 - 5 - 2 = 0$. Since the network is weakly reversible, the deficiency zero theorem implies that there is a unique "complex balanced" equilibrium point c in each compatibility class. Since the rank of the matrix is 5 and there are 5 species, the state space of the process $X^n(t)$ is $n^{-1}\mathbb{Z}_+^5$ and there is only one compatibility class. The following system of equations must be satisfied for the unique complex-balanced equilibrium (inflow = outflow for each complex, see

(3.2.1)):

$$\begin{aligned} c_1c_2^2 &= c_4 \\ c_1c_2^2 &= c_3^3 \\ c_4 + c_3^3 &= c_3c_5 \\ c_3c_5 &= 2c_4 \\ c_1 &= c_5^2 \\ c_3 &= c_1 \\ c_3 &= c_5^2. \end{aligned}$$

We have $c_5 = \sqrt{c_1}$ and $c_3 = c_1$, so $c_4 = c_1^{3/2} - c_1^3$, $c_2 = c_1$ and $c_4 = c_1^3$. From the two relations for c_4 we find $2 = c_1^{3/2}$. Thus

$$\begin{aligned} c_1 &= 2^{-2/3} \\ c_2 &= 2^{-2/3} \\ c_3 &= 2^{-2/3} \\ c_4 &= 1/4 \\ c_5 &= 2^{-1/3}. \end{aligned}$$

We will consider a domain of the following product form:

$$D = \prod_{i=1}^{5} (c_i^-, c_i^+),$$

and the problem of estimating

$$\mathbb{P}_c(X^n(t) \notin D \text{ for some } t \in [0, T(n)]). \tag{3.3.1}$$

Subsolution. It follows from Theorem 3.2.5 that the quasipotential relative to this unique stable fixed point $c = (2^{-2/3}, 2^{-2/3}, 2^{-2/3}, 1/4, 2^{-1/3})$ is given by

$$Q(x) = \sum_{i=1}^{5} c_i \ell\left(\frac{x_i}{c_i}\right).$$

As discussed in Section 2.1.4.1, one can obtain a subsolution from the quasipotential by taking $\bar{V}(x) = -Q(x) + c_Q$, where $c_Q = \inf_{y \in \partial D} Q(y)$. Owing to the product form structure of the domain and the fact that Q is strictly increasing relative to c_i in each direction x_i, the minimizer will occur at one of the $2^5 = 32$ values $Q(\pm c_1, \ldots, \pm c_5)$. This then determines the subsolution for estimating (3.3.1).

3.3.1.2 Using the moderate deviations approximation

Example 3.3.2. In this example we consider the escape from a rest point when the quasipotential is not available.

CRN and rare event. Consider the network

$$\begin{aligned} B + C &\to 2A \\ A + C &\to 2B \\ A + B &\to 2C \end{aligned}$$

with Hamiltonian

$$H(x,p) = x_2 x_3 (e^{2p_1 - p_2 - p_3} - 1) + x_1 x_3 (e^{-p_1 + 2p_2 - p_3}) + x_1 x_2 (e^{-p_1 - p_2 + 2p_3} - 1).$$

and initial condition $x_0 = (1, 1, 1)$, which is a steady state. It can be seen immediately from the graphical structure that the network is not weakly reversible and hence cannot be complex balanced, so it would be difficult to identify the quasipotential explicitly. For the problem of estimating terminal escape probabilities from rest points, we can fall back on a moderate deviations approximation.

In order to work with this network, we will first need to rewrite one of the variables in terms of the other using the linear constraint $n \cdot (x - x_0) = 0$, where $n = (1, 1, 1)$. The vector n is normal to each of the reaction vectors. Setting

$$x_3 = 3 - x_2 - x_1,$$

and viewing the Markov process as a jump process with modified reaction rates and the same jump vectors, we find

$$H(x,p) = x_2(3-x_1-x_2)(e^{2p_1-p_2}-1)+x_1(3-x_1-x_2)(e^{-p_1+2p_1}-1)+x_1x_2(e^{-p_1-p_2}-1).$$

This is the Hamiltonian we will work with. The domain now becomes inherently two-dimensional.

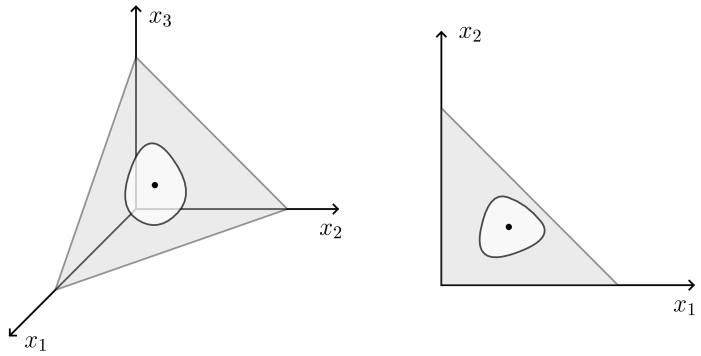

Figure 3.5: Left: the simplex $\{x : n \cdot x = 3\}$ containing the rest point and initial condition x_0, and a sample neighbourhood of x_0 with lighter shading. Right: the same state space containing the same rest point and neighoubrhood, but in (x_1, x_2) coordinates.

Precomputations. The law of large numbers dynamics are

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} -x_1x_2 - x_1(3 - x_1 - x_2) + 2x_2(3 - x_1 - x_2) \\ -x_1x_2 - x_2(3 - x_1 - x_2) + 2x_1(3 - x_1 - x_2) \end{pmatrix}, \quad x(0) = (1, 1),$$

and the associated matrices for computing $M_{t,T}$ are

$$\begin{aligned} Db(x) &= \begin{pmatrix} 2x_1 - 4x_2 - 3 & -2x_2 - 4x_1 + 6 \\ -2x_2 - 4x_1 + 6 & 2x_2 - 4x_1 - 3 \end{pmatrix} \\ A(x) &= \begin{pmatrix} 4 & -2 \\ -2 & 1 \end{pmatrix} x_2(3 - x_1 - x_2) + \begin{pmatrix} 1 & -2 \\ -2 & 4 \end{pmatrix} x_1(3 - x_1 - x_2) + \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} x_1x_2. \end{aligned}$$

Thus,

$$b(X_0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad Db(x_0) = \begin{pmatrix} -3 & 0 \\ 0 & -3 \end{pmatrix}, \quad A(x_0) = \begin{pmatrix} 6 & -3 \\ -3 & 6 \end{pmatrix},$$

and these last constant matrices are all that is required for the MDP approximation of the terminal escape problem.

Approximating the domain D. We will consider the problem of estimating the deviation from the rest point by more than $\delta > 0$ in any coordinate, i.e.

$$\mathbb{P}\left(\bigcup_{i=1}^{3}\{|X_i^n(T) - X_i^0(T)| \geqslant \delta\}\right).$$

Since we work with the centered process in two coordinates, it is straightforward to accomplish this for $i = 1, 2$ by taking

$$v_1 = (1, 0)^T, v_2 = -v_1, v_3 = (0, 1)^T, v_4 = -v_3,$$

and $\delta_j = \delta > 0$ for $j = 1, \ldots, 4$. For the third coordinate, $i = 3$, we note that

$$|z_1 + z_2| = |x_1 - 1 + x_2 - 1| = |x_3 - 1|,$$

and so we are looking for $|z_1 + z_2| \geqslant \delta$, which can be accomplished by taking $v_5 = (1, 1)$ and $v_6 = -v_5$, with $\delta_5 = \delta_6 = \delta > 0$.

Identifying the subsolution. Recall that, given the constraint $\langle v, z \rangle \geqslant \delta$, the gradient which yields $\bar{V}^{\xi,c}(0, 0) = V(0, 0)$ is given by

$$\xi = -\delta \frac{v}{\|v\|_M^2}, \quad c = \frac{\delta^2}{\|v\|_M^2}.$$

We take the minimum of six subsolutions, so that

$$\bar{V}(t, y) = \min_{i=1,\ldots,6} \bar{V}^{\xi_i, c_i}(t, y).$$

See Figure 3.12 for an illustration of the subsolutions at the initial and terminal times.

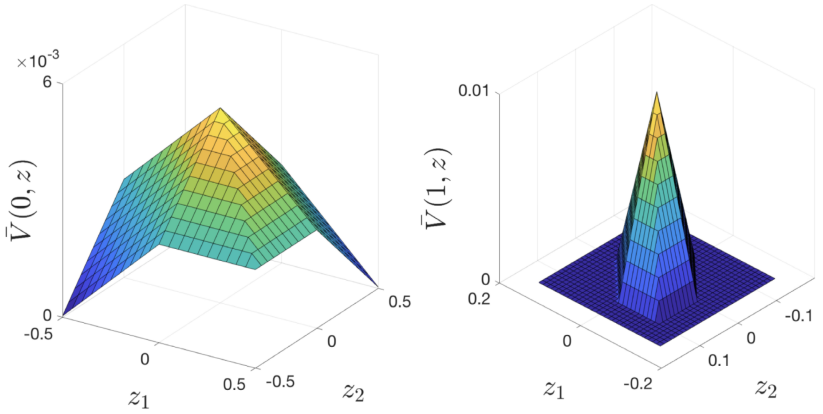

Figure 3.6: The subsolutions for the triangle example at times $t = 0$ (left) and $t = 1$ (right). Note that the scales are different, and that $V(t, 0)$ has increased by an order of magnitude.

3.3.2 Metastability

In this thesis, metastability is the phenomenon of a stochastic process jumping between steady states of the noiseless dynamics. In large deviations, such behavior is described by Freidlin-Wentzell theory and many rare event properties are captured in the quasipotential $Q(c, x)$ (3.2.5). In this section we will use the techniques for one-dimensional systems discussed in Section 3.2.2.2 to construct subsolutions for the problem of estimating

$$\mathbb{P}_{x_0}(X^n(t) \notin D \text{ for some } t \in [0, T(n)].$$

When the time interval $[0, T(n)]$ grows polynomially with the large deviations parameter n. Under appropriate conditions on the process, the subsolution based on the quasipotential is asymptotically optimal for the problem of escape from a rest point (see Theorem 2.2.15). If we take the domain to escape to be the domain of attraction of the stable steady state, then as soon as a trajectory escapes, it is in the domain of attraction of another stable steady state and will naturally tend towards it. This suggests that subsolutions based on the quasipotential will be appropriate for metastability problems over "long" time intervals. We note that the mean transition time between two steady states increases exponentially with n (see [13] or [35], Theorem 4.1 in Chapter 4).

In the examples we will take $T(n)$ to be fixed, though it will be "long" relative to the expected waiting time, and we will demonstrate in the next chapter through numerical simulations that subsolutions based on quasipotential are indeed viable. The first example considers the Schlögl model, which is a well-studied system exhibiting bistability. The second is a specific instance of Example 3.2.8, in which we will consider six different transition probabilities.

Example 3.3.3 (The Schlögl model)**.** The Schlögl model is perhaps the most famous CRN exhibiting metastable behavior and dates back to at least 1972 when Schlögl first introduced it [58]. Since then, it has been studied from various perspectives: [62] write down the quasipotential (though they do not call it that) and perform an eigenvector analysis, [27] uses a diffusion approximation and renormalisation approach, and [28] investigates the potential for bistability in bacteria and eukaryotic cells. Sargsyan [57] provides a writeup of his summer research project from 2006 which uses many of the large deviations techniques discussed here, though the project does not explore the use of HJB equations as we do. In [3], Example 10, the authors use the relationship in Theorem 1.5.3 to derive the quasipotential (3.2.10), and they write down the integrated form explicitly for a specific choice of rate constants – see equation (3.3.3).

CRN and rare event. The Schlögl model is the network

$$\emptyset \underset{\kappa_{10}}{\overset{\kappa_{01}}{\rightleftarrows}} A, \quad 2A \underset{\kappa_{32}}{\overset{\kappa_{23}}{\rightleftarrows}} 3A.$$

The Hamiltonian for the Schlögl model is

$$H(x,p) = \kappa_{01}(e^p - 1) + \kappa_{10}x(e^{-p} - 1) + \kappa_{2,3}x^2(e^p - 1) + \kappa_{3,2}x^3(e^{-p} - 1).$$

With the choice of $\kappa_{01} = 6$, $\kappa_{10} = 11$, $\kappa_{23} = 6$ and $\kappa_{32} = 1$, the deterministic dynamics can be neatly factored:

$$\Phi(x) = -x^3 + 6x^2 - 11x + 6 = -(x-1)(x-2)(x-3).$$

Thus there are three steady states, with two stable ones at $c_1 = 1$ and $c_2 = 3$ and an unstable one at $c_{12} = 2$. We will focus on estimating the following probability:

$$\mathbb{P}_{c_1}(X^n(t) \notin D, \text{ for some } t \in [0, T]), \tag{3.3.2}$$

where

$$D = [0, c_{12} + 0.5) = [0, 2.5).$$

Subsolution. From Lemma 3.2.10, the gradient of the quasipotential relative to one of the steady states is

$$Q(c, x) = \int_c^x \log\left(\frac{s^3 + 11s}{6s^2 + 6}\right) \mathrm{ds}$$

for x in $B(c)$, the domain of attraction of c (see Definition 1.5.1). The authors in [3] integrate the expression explicitly for $c = c_1$ and obtain

$$\begin{aligned} Q(c_1, x) &= x\left(\log\left(\frac{x(x^2+11)}{x^2+1}\right) - \log(6) - 1\right) + 2\sqrt{11}\arctan\left(\frac{x}{\sqrt{11}}\right) \\ &\quad -2\arctan(x) - 2\sqrt{11}\arctan\left(\frac{1}{\sqrt{11}}\right) + \frac{\pi}{2} + 1. \end{aligned} \tag{3.3.3}$$

In order to use the subsolution (3.3.3) for the problem of estimating (3.3.2), we need to make a modification similar to the one in equation (3.2.7). The reason is we do not need to encourage trajectories to split on the interval $(2, 2.5)$, since the law of large dynamics are drawn to the stable equilibrium 3 anyway. We therefore use the subsolution

$$\bar{V}(x) = \int_c^{x \wedge 2} \log\left(\frac{s^3 + 11s}{6s^2 + 6}\right) ds.$$

See Figure 3.7. We build and test a splitting scheme for this network in Section 4.2.2.

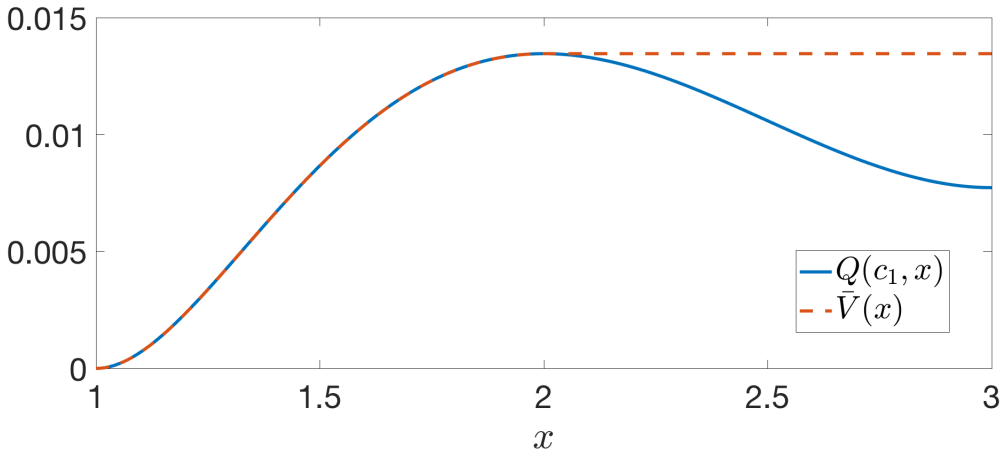

Figure 3.7: In solid blue is the analytic expression (3.3.3) derived in [3]. In dashed red is the numerically integrated subsolution $\bar{V}(x)$ for the problem of estimating $\mathbb{P}_1(X^n(t) > 2.5,$ for some $t \in [0, T(n)])$, where $T(n)$ is on the order of n^k for some $k \geq 1$.

Example 3.3.4. In this example we use Lemma 3.2.6 and Corollary 3.2.7 to construct a network which has three stable steady states c_1, c_2, c_3, as in Example 3.2.8. We consider the problem of estimating the probability of transitioning between pairs of distinct steady states (c_i, c_j), $i, j = 1, 2, 3$, $i \neq j$. We will construct a subsolution for all six ordered pairs of states which will allow us to compare the results and provide a check on the estimator.

CRN and rare events. We construct a CRN using Corollary 3.2.7. Let $c_1 = 1, c_2 = 12, c_3 = 50$ and $c_{12} = 5$, $c_{23} = 26$ (see Figure 3.2). The polynomial $q(x)$ is then

$$q(x) = -(x-1)(x-5)(x-12)(x-26)(x-50) = -x^5+94x^4-2745x^3+29312^2-104660+78000.$$

A reaction network G_κ with steady state dynamics given by $q(x)$ is

$$\begin{aligned} 5A &\underset{\kappa_{45}}{\overset{\kappa_{54}}{\rightleftarrows}} 4A \\ 3A &\underset{\kappa_{23}}{\overset{\kappa_{32}}{\rightleftarrows}} 2A \\ A &\underset{\kappa_{01}}{\overset{\kappa_{10}}{\rightleftarrows}} \emptyset, \end{aligned} \tag{3.3.4}$$

where

$$\kappa_{54} = 1, \kappa_{45} = 94, \kappa_{32} = 2745, \kappa_{23} = 29312, \kappa_{10} = 104660, \kappa_{01} = 78000.$$

For $\delta > 0$, let $B_\delta(c) = \{x : |x - c| < \delta\}$, and define

$$p_{ij}^n = p_{ij}^n(T, \delta) \doteq \mathbb{P}_{c_i}(X^n(t) \in B_\delta(c_j) \text{ for some } t \in [0, T]), \quad i, j \in \{1, 2, 3\}, i \neq j.$$

Since the underlying process is one-dimensional one can infer basic relationships between these probabilities, such as $p_{13}^n \leq p_{12}^n$. Moreover, if $\delta > 0$ is sufficiently small, one expects the probabilities to satisfy relationships such as

$$p_{13}^n(T + S, \delta) \geq p_{12}^n(T, \delta) p_{23}^n(S, \delta) \tag{3.3.5}$$

for $T, S > 0$. Estimating each of the p_{ij}^n's and considering these relationships provides a check as to whether the estimator is working correctly.

Subsolution. The CRN G_κ defined in (3.3.4) is a birth-death process. According to Lemma 3.2.10, the quasipotential relative to a stable steady state c on a

neighbourhood where c is attracting is given by

$$Q(c,x) = \int_c^x \log\left(\frac{s^5 + 2745s^3 + 104660s}{94s^4 + 29312s^2 + 78000}\right) ds.$$

As discussed in Section 3.2.2.2, in order to estimate the probabilities $p_{13}^n(\delta)$ and $p_{31}^n(\delta)$, one will have to stitch together the quasipotentials with a formula analogous to (3.2.7); see Figure 3.3. The subsolution is then given by taking the negative of the quasipotential and shifting it by an appropriate constant so as to satisfy the boundary condition of the underlying PDE; see equation (3.2.8) in Section 3.2.2.2.

It is insightful to consider the "potential landscape" of the network, i.e., the function

$$U(x) = \int_0^x p(y)dy + c$$

where p satisfies the conditions of Corollary 1.5.6, as from this one can infer each of the six subsolutions constructed; see Figure 3.8. The six subsolutions corresponding to the problem of estimating $p_{ij}^n(\delta)$ are also given in Table 3.1 and can be seen in Figure 3.9.

$\bar{V}_{12}(x) = -Q(c_1, (x \vee c_1) \wedge c_{12}) + Q(c_1, c_{12})$	$\bar{V}_{21}(x) = -Q(c_2, (x \wedge c_2) \vee c_{12}) + Q(c_2, c_1)$
$\bar{V}_{23}(x) = -Q(c_2, (x \vee c_2) \wedge c_{23}) + Q(c_2, c_{23})$	$\bar{V}_{32}(x) = -Q(c_3, (x \wedge c_2) \vee c_{23}) + Q(c_3, c_{23})$
$\bar{V}_{13}(x) = \bar{V}_{12}(x) + \bar{V}_{23}(x)$	$\bar{V}_{31}(x) = \bar{V}_{32}(x) + \bar{V}_{21}(x).$

Table 3.1: Six subsolutions for the network constructed in Example 3.3.4.

3.3.3 Boundary behavior

In this section we consider the boundary behavior of a CRN and construct subsolutions for estimating probabilities related to hitting and leaving the boundary. In order to study rare events in this context, it is necessary to first understand what

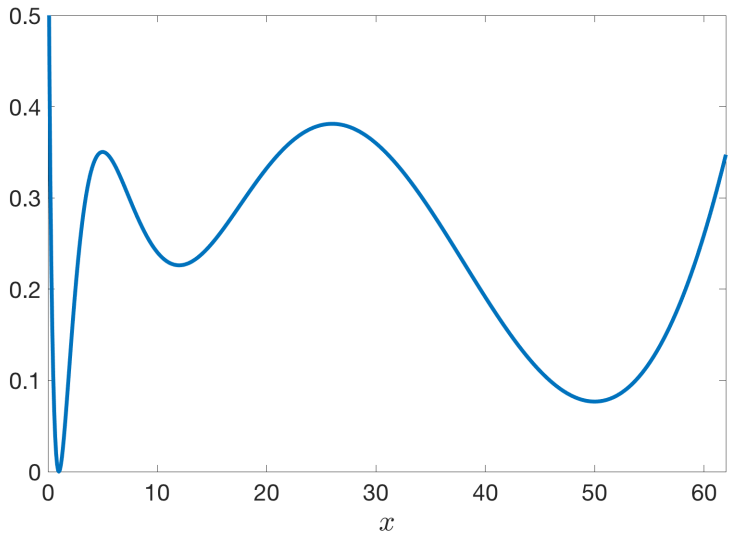

Figure 3.8: The "potential" $U(x) = \int_0^x p(y)dy + c$, where $p(x) \neq 0$ unless x is a steady state and $H(x, p(x)) = 0$ for all $x \in [0, \infty)$, with the constant c chosen so that $U(x) \geq 0$ for all x and $U(1) = 0$.

typical events (as dictated by the law of large numbers) might look like.

To this end, in Section 3.3.3.1 we introduce a few definitions which allow us to describe the deterministic dynamics which start on the boundary. These are older results due to Vol'pert ([64] or [65], Section 11) which provide qualitative and quantitative information on the trajectories using only the graphical structure of the network. We then use this information in Section 3.3.3.2 to identify the law of large numbers dynamics and construct subsolutions for a rare event problem when leaving a corner. Finally, in Section 3.3.3.2 we explore rare events problems for hitting the boundary.

Recall that a CRN G_κ is defined as a tuple $(\mathcal{C}, \mathcal{R}, \kappa)$, where $\mathcal{C}$ is the set of complexes a, each of which is identified with a vector in $\mathbb{Z}_+^d$, $\mathcal{R} \subset \mathcal{C} \times \mathcal{C}$ is the set of reactions (or edges of the graph), and $\kappa = \{\kappa_{a,b} : (a, b) \in \mathcal{R}\}$ is a collection of positive scalars called rate constants. It will be useful throughout this section to order the species, so that for a d-dimensional CRN the index $i = 1, \ldots, d$ determines

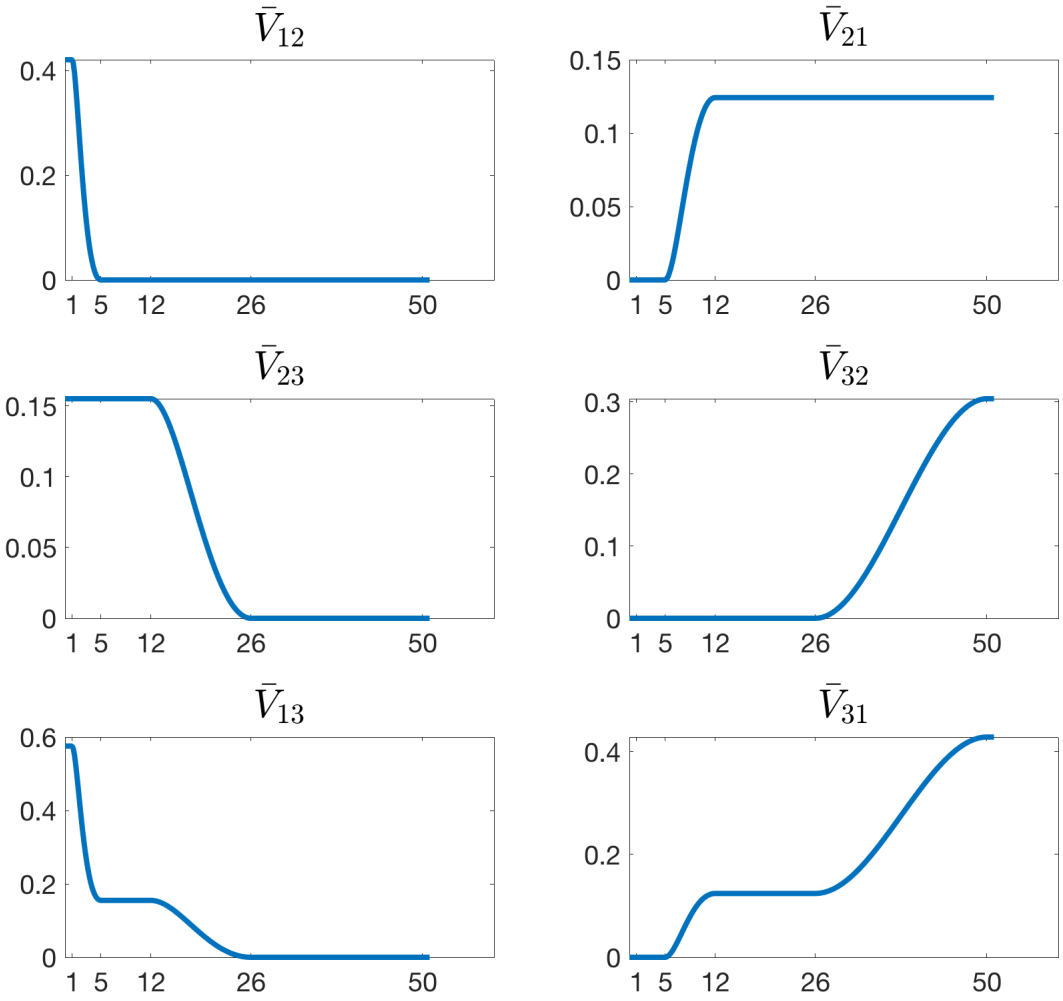

Figure 3.9: The six subsolutions corresponding to the problem of estimating $p_{ij}^n(\delta)$ for $i, j \in \{1, 2, 3\}, i \neq j$. The subsolution is independent of $\delta > 0$ so long as $\delta < 1$.

the dimension unambiguously, so that each dimension corresponds to a chemical *species*.

3.3.3.1 Law of large numbers for leaving boundary

In this section we focus on the following question. Let G_κ be a d-dimensional CRN with domain $\mathcal{S}$ (see Definition 3.1.4) and suppose $x(t) = (x_1(t), \ldots, x_d(t))$ solves $\dot{x} = \Phi(x)$ with $x(0) \in \partial\mathcal{S}$. Being on the boundary implies that $x_i(0) = 0$ for some

$i \in \{1, \ldots, d\}$. Do we have $x_i(t) > 0$ for all $i \in \{1, \ldots, d\}$ and $t > 0$, and if so, can one estimate $x_i(t)/x_j(t)$ for $i \neq j$, $t > 0$?

It turns out one can answer these questions strictly from the structure of the graph underlying the CRN. The statements below closely follow the exposition and work of Gnacadja (see pp. 9-10, [39]). We will use his definitions to state a result of Vol'pert's ([64, 65]).

Definition 3.3.5. $\mathcal{I} \subset \{1, \ldots, d\}$ be a subset of indices, let $a \in \mathcal{C}$, and let a_i denote the i^{th} coordinate of the vector in $\mathbb{Z}_+^d$ with which a is naturally identified (see Remark 3.1.2). We write $a \subset \mathcal{I}$ if, for any i such that $a_i > 0$, we have $i \in \mathcal{I}$. We write $\text{supp}(a)$ for the support of a, i.e. set of indices where a is positive:

$$\text{supp}(a) = \{i = 1, \ldots, d : a_i > 0\}.$$

Define the sets

$$\text{React}(\mathcal{I}) \doteq \{(a, b) \in \mathcal{R} : a \subset \mathcal{I}\},$$

$$\text{Prod}(\mathcal{I}) \doteq \{\text{supp}(b) : (a, b) \in \text{React}(\mathcal{I})\} \setminus \mathcal{I}.$$

$\text{React}(\mathcal{I})$ is the set of reactions that would occur if only the species in $\mathcal{I}$ were present. $\text{Prod}(\mathcal{I})$ is the set of species which would be created once these reactions occurred. Using these definitions we recursively define the r^{th} reachable set for $r \geqslant 0$ by

$$\begin{aligned} \text{Reach}_0(\mathcal{I}) &= \mathcal{I} \\ \text{Reach}_r(\mathcal{I}) &= \text{Prod}\left(\bigcup_{\rho=0}^{r-1} \text{Reach}_\rho(\mathcal{I})\right), \quad r \geqslant 1. \end{aligned}$$

Note that the sets $\text{Reach}_r(\mathcal{I})$ are pairwise disjoint. If $\text{Reach}_{r_0}(\mathcal{I}) = \emptyset$ for some $r_0 \geqslant 0$, then $\text{Reach}_r(\mathcal{I}) = \emptyset$ for all $r \geqslant r_0$, and since the dimension d of any CRN is assumed to be finite, such r_0 must exist and be finite. The species in $\text{Reach}_r(\mathcal{I})$ are said to

have *reachability index* r.

Finally, we define the reachable and non-reachable sets,

$$\mathrm{Reach}(\mathcal{I}) = \bigcup_{r=0}^{\infty} \mathrm{Reach}_r(\mathcal{I}), \quad \mathrm{NonReach}(\mathcal{I}) = \{1, \ldots, d\} \setminus \mathrm{Reach}(\mathcal{I}).$$

Below is the promised Theorem of Vol'pert, restated in the notation of this thesis. We will use $x^{(r)}(t) = (x_1^{(r)}(t), \ldots, x_d^{(r)}(t))$ to denote the r^{th} derivative of the trajectory $x(t)$,

$$x^{(r)}(t) \doteq \left.\frac{d^r}{dt^r}\right|_{t=0} x_i(t).$$

Theorem 3.3.6 (Vol'pert)**.** *Let $G = (\mathcal{C}, \mathcal{R})$ be a reaction network with $0 \notin \mathcal{C}$. Let $x(t; x(0))$ denote a solution to the system*

$$\dot{x}(t) = \Phi(x(t)), \quad x(0) = x_0,$$

where we recall that $\Phi(x) = \nabla_p H(x, 0)$. Set $\mathcal{I} = \mathrm{supp}(x_0)$ and let $i \in \{1, \ldots, d\}$ be any species. If $i \in \mathrm{NonReach}(\mathcal{I})$ then $x_i(t) = 0$ for all $t \geqslant 0$. If $i \in \mathrm{Reach}(\mathcal{I})$, let r denote the corresponding reachability index. We have

$$\begin{aligned} x_i^{(\rho)}(t) &= 0, \quad \rho = 0, \ldots, r-1, \\ x_i(t) &= \frac{1}{r!} x_i^{(r)}(0) t^r + O(t^{r+1}), \quad t \geqslant 0. \end{aligned}$$

Note that the statement does not indicate whether or not $x_i^{(r)}(t) > 0$, only that the zero at $t = 0$ of x_i is at least of order r, where r is the reachability index of species i. It may be greater than r. As pointed out in [39], if $i, i' \in \{1, \ldots, d\}$ are two distinct species with reachability indices r, r' and $0 < r < r'$, then there is $t_0 > 0$

such that

$$x_{i'}(t) < x_i(t), \quad t \in (0, t_0).$$

In other words, the higher the reachability index, the slower the initial production of the species. A proof of Vol'pert's Theorem can be found in the Appendix of [39] (Theorem 7.1). Vol'pert also shows that in the case of *linear* reaction rates, the reachability index determines the order of the zero exactly (see [64], Theorem 7).

Example 3.3.7. Consider the CRN

$$\emptyset \to A \to B.$$

Since the reaction rates $r_1(x) = 1$ and $r_2(x) = x_1$ are linear, the reachability index determines the order of the zero. The reachability indices for A and B are 0 and 1, respectively. Theorem 3.3.6 tells us that starting from the initial configuration $x(0) = (0, 0)$, we can expect

$$x_1(t) = O(1)$$
$$x_2(t) = O(t).$$

Qualitatively speaking, this means that the curve in dashed red in Figure 3.10 will never occur under the deterministic dynamics and is therefore a rare event.

3.3.3.2 Subsolutions based on moderate deviations to leave boundary

Example 3.3.8. Consider the network

$$\emptyset \xrightarrow{\kappa_1} A \xrightarrow{\kappa_2} B,$$

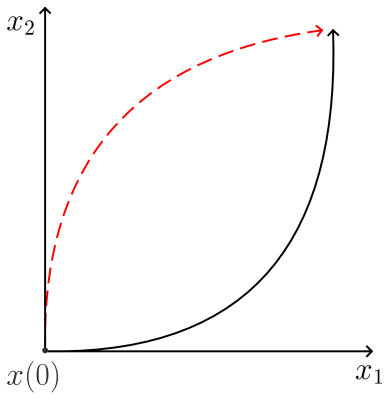

Figure 3.10: Solid black curve: a sketch of the trajectory $x(t) = (x_1(t), x_2(t))$ for the deterministic dynamics of the network $\emptyset \to A \to B$ starting at the initial condition $x(0) = (0, 0)$. Dashed red curve: an unlikely way of leaving the boundary for the stochastic dynamics.

with Hamiltonian

$$H(x, p) = \kappa_1(e^{p_1} - 1) + \kappa_2(e^{-p_1+p_2} - 1)$$

and initial condition $x_0 = (0, 0)$. We first determine the boundary behavior of the network. Note that $\text{Reach}_0(\emptyset) = \{A\}$ and $\text{Reach}_1(\emptyset) = \{B\}$. Thus, we expect the curve $X^0(t) = (X_1^0(t), X_2^0(t))$ to be tangent to the x_1 axis at $t = 0$.

There are various rare events one could consider, but the moderate deviations approximation will only be appropriate for some of them. We will consider the problem of estimating

$$\mathbb{P}(X_1^n(T) - X_1^0(T) \leqslant -\delta),$$

for some fixed $T > 0$ and various $\delta > 0$. The moderate deviations approximation would not be ideally suited to estimate this probability when $\delta = X_1^0(T)$, since this would require an increasingly dense spacing of thresholds near the boundary as n increases; ssee for instance Figure 3.13 in the next section on hitting the boundary. The reason such an event occurs in a way not captured by the quadratic approximation around the law of large numbers X^0; see Figure 3.11.

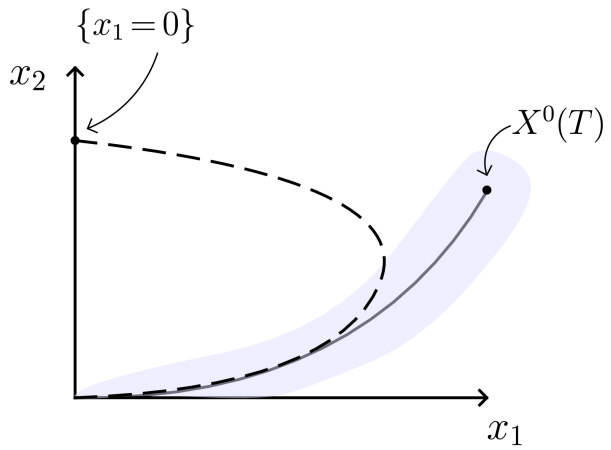

Figure 3.11: Solid curve: the LLN trajectory X^0, with its location at time T marked. Dashed line: deviation from the LLN trajectory which hits the shaded gray region at time T. The deviation from the LLN begins much earlier than time T, and hitting the boundary requires densely spaced thresholds which do not occur with the moderate deviations approximation; see Figure 3.12.

Precomputations. The law of large numbers dynamics are

$$\begin{pmatrix} \dot{X}_1^0 \\ \dot{X}_2^0 \end{pmatrix} = \begin{pmatrix} \kappa_1 - \kappa_2 X_1^0 \\ \kappa_2 X_1^0 \end{pmatrix}, \quad X^0(0) = x_0,$$

and the associated matrices for computing $M_{t,T}$ are

$$Db(x) = \begin{pmatrix} -\kappa_2 & 0 \\ \kappa_2 & 0 \end{pmatrix}, \quad A(x) = \begin{pmatrix} \kappa_1 x_1 + \kappa_2 & -\kappa_1 x_1 \\ -\kappa_1 x_1 & \kappa_1 x_1 \end{pmatrix}.$$

Approximating the domain D. We are interested in

$$\mathbb{P}(X_1^n(T) - X_1^0(T) \leqslant -\delta).$$

In this case, the shifted domain to be $\{z_1 \leqslant -\delta\}$, which corresponds to choosing $v = (-1, 0)^T$ and $\delta > 0$ as given. Then, if $z = x - X^0(T)$,

$$\{\langle v, z\rangle \geqslant \delta\} = \{-(x_1 - X_1^0(T)) \geqslant \delta\} = \{x_1 - X_1^0(T) \leqslant -\delta\},$$

as required.

Identifying the subsolution. Recall that, given the constraint $\langle v, z\rangle \geqslant \delta$, the gradient which yields $\bar{V}^{\xi,c}(0,0) = V(0,0)$ is given by

$$\xi = -\delta \frac{v}{\|v\|_M^2}, \quad c = \frac{\delta^2}{\|v\|_M^2}.$$

Since we have only one constraint for the domain with $v = (-1, 0)$, we have only a single subsolution, $\bar{V}^{\xi,c}(t, y) = \langle \xi, \Psi(t, T)y\rangle + c - \|\xi\|^2_{M_{t,T}}/2$.

For each fixed $t \in [0, T]$ the subsolutions are affine in y. It turns out that in this case $\Psi(s, t)$ is lower-triangular, and since the second coordinate of ξ is 0, the subsolutions depend only on the first coordinate. As t increases to T, the slope of the affine function $\bar{V}^{\xi,c}(t, y)$ increases. Figure 3.12 illustrates the possible thresholds at some fixed time t.

3.3.3.3 Subsolutions for hitting the boundary of complex-balanced systems

In this section we discuss stochastic trajectories of complex-balanced systems which hit a boundary of the domain (see Definition 3.1.4) and the cost of the associated

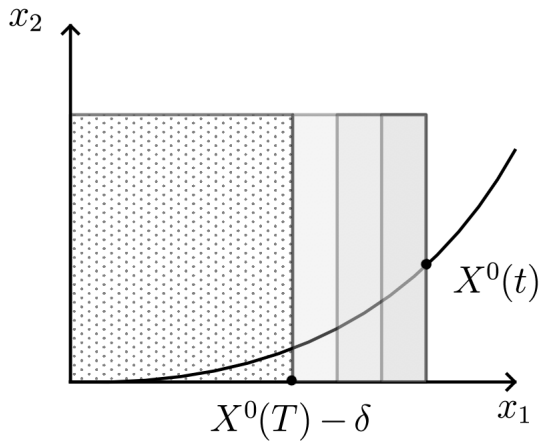

Figure 3.12: A snapshot of the thresholds determined by $\bar{V}^{\xi,c}$ with $v = (-1, 0)$ for some fixed t, translated back into x coordinates. The function $\bar{V}^{\xi,c}(t, \cdot)$ is affine and depends only on the first coordinate, so the thresholds are equally spaced.

event. This will then be used to estimate probabilities of the form

$$\mathbb{P}_{x_0}(X_i^n(t) = 0 \text{ for some } t \in [0, T(n)]),$$

where $i \in \mathcal{I} \subset \{1, \ldots, d\}$ for a d-dimensional CRN. We will call $\mathcal{I}$ the *vanishing indices* or the *vanishing species*.

The law of large numbers (LLN) dynamics can never hit the boundary in finite time, so there is no analogue of Theorem 3.3.6 theorem which describes the "typical behavior" of doing so. Instead, we will use the complex-balanced structure to compute the quasipotential for various networks. Then, since the quasipotential $Q(c, x)$ gives the optimal cost of moving between a stable steady state c and a point x in its domain of attraction over an arbitrarly long but finite amount of time, by comparing the quasipotentials we have a rudimentary way of comparing the "rarity" of different events.

We present four simple examples. For each one, we will consider the problem

of starting at the unique stable steady states of the complex-balanced system and hitting one or more boundaries. We do so by identifying the quasipotential using Theorem 3.2.5 and comparing the costs of the different events occurring. We recall that in the context of splitting for time-independent problems, the quasipotential will ultimately determine the subsolution, and its spatial gradient determines how closely spaced together the thresholds are. The greater the directional derivative, the more closely spaced the thresholds will be in that direction.

Example 3.3.9. This first example is perhaps the simplest one-dimensional reaction network with a unique stable steady state for which the event of hitting the boundary can occur. Around the boundary the behavior of the local cost $L(x, \beta)$ is singular and we use this example to illustrate some of the issues that come up. It will subsequently be used as a reference point and a basis of comparison in later examples.

CRN and rare event.

$$A \rightleftarrows \emptyset,$$

with associated Hamiltonian

$$H(x, p) = x(e^{-p} - 1) + e^{p} - 1.$$

The associated pair of Hamilton's equations is

$$\begin{aligned} \dot{x} &= -xe^{-p} + e^{p} \\ \dot{p} &= 1 - e^{-p}, \end{aligned} \tag{3.3.6}$$

with the LLN dynamics being $\Phi(x) = D_p H(x, 0) = 1 - x$. Setting $\Phi(x) = 0$, we find the unique stable steady state is at $c = 1$. We will want to estimate the rare event

probability

$$\mathbb{P}_1(X^n(t) = 0 \text{ for some } t \in [0, T(n)]), \tag{3.3.7}$$

where $T(n)$ is a sequence of increasing time intervals with polynomial n-dependence.

We note that it is not a priori clear a trajectory can hit the boundary in finite time with finite cost. First, observe that if $\dot{x}$ is to be negative and x is to hit 0 at time T, then from (3.3.6) we require $p(t) \to -\infty$ as $t \to T$. The local cost $L(x, \beta)$, however, has a superlinear factor in terms of e^{-p}:

$$L(x, \beta) = \inf_{p \in \mathbb{R}} \{x\ell(e^{-p}) + \ell(e^{p}) : \beta = -xe^{-p} + e^{p}, \}$$

where we recall that $\ell(x) = x \log x - x + 1$. It turns out that the factor of x in front of $\ell(e^{-p})$ is sufficient to ensure the integrated cost is finite for some trajectories which hit the boundary. For example, one can check the trajectory $\varphi(t) = 1 - t$ starts at 1, hits 0 at time 1, and satisfies $I_1(\varphi) < 3$. What is not possible is *staying at the boundary*: no trajectory φ can satisfy $I_T(\varphi) < \infty$ and $\varphi(t) = 0$ for $t \in (t_0, t_1)$, $t_0 < t_1$.

Subsolution. The network is weakly reversible (Definition 3.1.6) and has a deficiency of zero (Definition 3.2.2), so by the Deficiency Zero theorem (Theorem 3.2.4) it is complex-balanced and has a unique stable steady state which we know is $c = 1$. From Theorem 3.2.5 in Section 3.2.1, the quasipotential $Q(x)$ relative to c is given by

$$Q_{A \rightleftarrows \emptyset}(x) = x \log x - x + 1.$$

The cost to hit 0 starting at 1 while allowing as much time as one likes is $Q(0) = 1$. The subsolution for the problem of estimating (3.3.7) is given by

$$\bar{V}(x) = 1 - Q_{A \rightleftarrows \emptyset}(x) = x - x \log x.$$

We remark that the majority of the cost is coming from "fighting" the reaction $\emptyset \to A$. Moreover, the derivative of the quasipotential is equal to

$$DQ_{A \rightleftarrows \emptyset}(x) = \log x,$$

and the derivative diverges to $-\infty$ at a logarithmic rate, which implies the thresholds will become increasingly closely spaced near the boundary.

Example 3.3.10. This next example shows that the "order" of the reaction doesn't affect the long-term cost of hitting the boundary.

CRN and rare event. We consider the CRN

$$2A \rightleftarrows \emptyset,$$

with LLN dynamics $\Phi(x) = 1 - x^2$. The unique steady state is at $x = 1$, and we will seek to estimate

$$\mathbb{P}_1(X^n(t) = 0 \text{ for some } t \in [0, T(n)])$$

over a sequence of increasing time intervals, just as in (3.3.7).

Subsolution. For the same reasons as for the previous example, this network is complex-balanced. Despite the different "order" ($2A$ instead of A), the steady state

is the same and so the quasipotentials are as well:

$$Q_{2A\rightleftarrows\emptyset}(x) = Q_{A\rightleftarrows\emptyset}(x).$$

Since the two quasipotentials are equal, we expect the sequence of probabilities of hitting the boundary over increasing time intervals to eventually be on the same order of magnitude for both $A \rightleftarrows \emptyset$ and $2A \rightleftarrows \emptyset$. As before, the time-independent subsolution for the problem of estimating the hitting probability over time $[0, T(n)]$ starting at 1 is

$$\bar{V}(x) = x - x\log x.$$

Example 3.3.11. The third example is meant to convey that components of the spatial gradient of Q which correspond to the vanishing dimensions is more sensitive to the vanishing rates when near the boundary. This confirms what natural intuition might suggest, that thresholds should be closely spaced for the vanishing dimensions as one approaches the boundary.

CRN and rare event. Consider the CRN

$$A \rightleftarrows B.$$

This CRN is complex-balanced by the Deficiency Zero theorem (Theorem 3.2.4), so in each stoichiometric compatibility class there is a unique steady state $c = (c_1, c_2)$. Under the assumption that all rate constants are equal to one, we have $c_1 = c_2$. In this case, the stoichiometric compatibility classes can be parametrized in terms of the steady state as $\mathcal{S}_c = \{x_1 + x_2 = 2c_1\}$, and the initial condition $x(0)$ determines the compatibility class. In this sense, the network is one-dimensional. We will assume that $x_1(0) + x_2(0) = 1$, and hence $c = (1/2, 1/2)$, and consider the probability of

hitting the facet $\{x_1 = 0\}$,

$$\mathbb{P}_{(1/2,1/2)}(X_1^n(t) = 0 \text{ for some } t \in [0, T(n)]).$$

Subsolution. By Theorem 3.2.5,

$$Q_{A\rightleftarrows B}(x) = c_1\ell\left(\frac{x_1}{c_1}\right) + c_2\ell\left(\frac{x_2}{c_2}\right).$$

Note that the quasipotential can be considered as a function of x_1 or x_2 alone, since $x_1 = 1 - x_2$; see Figure 3.13. The subsolution is obtained by taking the negative of the quasipotential and ensuring the boundary condition $\bar{V}(x) \leq 0$ holds for $x_1 = 0$:

$$\bar{V}(x) = \frac{1}{2} + \frac{1}{2}\ell(2) - \frac{1}{2}\left(\ell(2x_1) + \ell(2x_2)\right) = \log(2) - x_1\log(2x_1) - x_2\log(2x_1),$$

where for the last equality we used the definition of $\ell(\cdot)$ and the fact that $x_1 + x_2 = 1$. The spatial gradient is

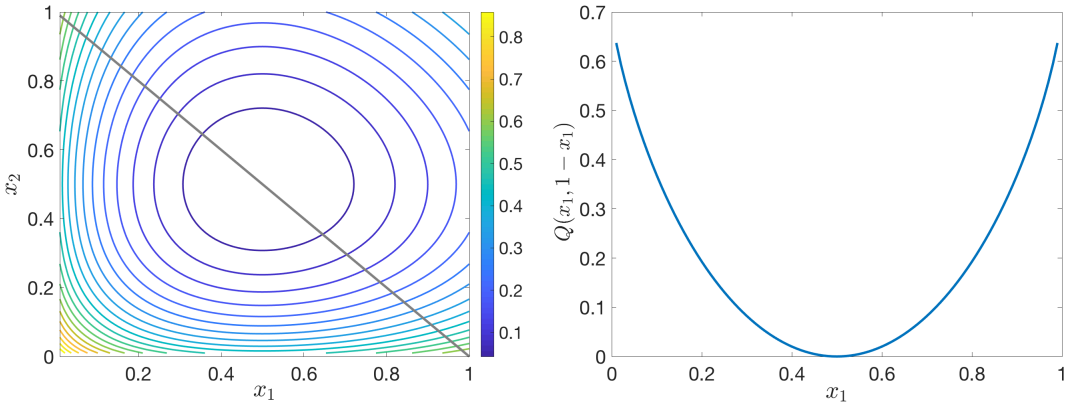

Figure 3.13: Left: contours of level sets of the quasipotential $Q_{A\rightleftarrows B}(x)$ in the (x_1, x_2) plane. Notice that the spacing increases as x_1, x_2 get close the the boundary. The dynamics of the CRN confine the trajectories to a simplex which was chosen to be $\{x_1 + x_2 = 1\} \cap \mathbb{R}^2_+$. Right: the same quasipotential restricted to the simplex $\{x_1 + x_2 = 1\} \cap \mathbb{R}^2_+$.

$$DQ_{A\rightleftarrows B}(x) = \left(\log\left(\frac{x_1}{c_1}\right),\ \log\left(\frac{x_2}{c_2}\right)\right).$$

As $x_1 \to 0$, the directional derivative in the direction $(1,0)$ diverges to $-\infty$ at rate $\log x_1$ (up to an additive constant), the same as for $DQ_{A\rightleftarrows\emptyset}(x) = \log(x)$.

Example 3.3.12. This final example illustrates some of the issues which arise when the domain is legitimately multidimensional. It also demonstrates some of the limitations of the approach, and rare events which cannot be estimated.

CRN and rare events. We take the network

$$B \rightleftarrows A \rightleftarrows \emptyset$$

This CRN is complex-balanced, with unique steady state $(c_1, c_2) = (1,1)$. We consider the family of rare events

$$p^n_{\mathcal{I}} = \mathbb{P}_{(1,1)}(X^n_i(t) = 0 \text{ for some } t \in [0, T(n)] \text{ and all } i \in \mathcal{I}),$$

where $\mathcal{I} = \{1\}$, $\{2\}$ or $\{1,2\}$. We will be able to estimate each of these probabilities. However, at least for the case of hitting a *corner* such as $(x_1, x_2) = (0,0)$, there is the additional question of "how" one might expect to hit the boundary. Do the B species vanish first, then the A species, the other way around, or at equal rates? In order to determine this, we would have to track the trajectories over a time interval, or at least check the derivatives at the hitting time. These situations are not covered by the splitting techniques we have discussed, although splitting can be used to learn the conditionally most likely way things happen.

Subsolution. From Theorem 3.2.5, the quasipotential is equal to

$$Q_{B\rightleftarrows A\rightleftarrows\emptyset}(x) = \ell(x_1) + \ell(x_2).$$

The smallest cost to hit either $\{x_1 = 0\}$ or $\{x_2 = 0\}$ is equal to 1 in both cases, with the quasipotential being minimized at $(0, 1)$ or $(1, 0)$, respectively. However, over any finite time interval $[0, T(n)]$, the probability of hitting 0 will be greater for species B than species A, since the latter must overcome both the reactions $B \to A$ and $\emptyset \to A$, which we have argued are the main contributors to the cost. Thus, we expect $p^n_{\{2\}} > p^n_{\{1\}}$. Moreover, the cost to hit $(0, 0)$ is twice as large as the cost to hit either $\{x_1 = 0\}$ or $\{x_2 = 0\}$, and we expect $p^n_{\{1\}} > p^n_{\{1,2\}}$. The subsolutions $\bar{V}_{\mathcal{I}}$ are

$$\bar{V}_{\{1\}}(x) = \bar{V}_{\{2\}}(x) = 1 - \ell(x_1) - \ell(x_2),$$
$$\bar{V}_{\{1,2\}}(x) = \bar{V}_{\{1,2\}}(x) = 2 - \ell(x_1) - \ell(x_2).$$

Chapter Four

Numerical examples

4.1 Introduction

This section takes the subsolutions constructed in Section 3.3 and uses them to construct importance functions for the RESTART splitting scheme introduced and analyzed in Chapter 2. Before proceeding to the examples themselves, in Section 4.1.1 we briefly discuss the implementation of the algorithm, and in Section 4.1.2 we define the statistics recorded for each example.

4.1.1 Implementation

As with any numerical method, the machine architecture and quality of the code can have a huge effect on overall performance. This section serves as a disclaimer, since no effort was made to make the algorithm fast, and one of the goals of performing the numerics is to verify that the algorithms do work and produce reasonable results with exponentially fewer samples than would be necessary with standard Monte Carlo. In addition, we mention a small approximation that is made in the case of terminal time problems. Finally, we make a note on the use of importance functions which are not subsolutions.

Efficient memory management of the particles. Much of the computational effort of the algorithm is spent on memory management. The way the RESTART algorithm is described in Algorithm 1 represents a "parallel" version of the method, in which all the particles are evolved simultaneously: those which are killed are discarded, those which are branched are added to the dynamic list, and the rest are simply stored. The dynamic list which stores the active particles must also store the killing threshold.

The issue with such a dynamic list is that, in theory, it resizes every step. If we ignore the cost of generating random variables, the computational effort of updating the times and positions of the particles is negligible in comparison with the cost of retrieving information about the particles and fetching additional memory when the number of particles is increased. This issue is central to good performance of the algorithm.

Terminal time problems. Terminal time escape problems are ones of the form

$$\mathbb{P}_{x_0}(X^n(T) \notin D).$$

For any given particle in the RESTART simulation, two values are stored: the time t, and the position $X^n(t)$. Recall that X^n is a piecewise constant, right continuous process, and since it is a Markov process, the next time increment τ is exponentially distributed with parameter equal to the total jump intensity at state $X^n(t)$. The problem is that the event of interest occurs when $t < T$ and $t + \tau \geq T$, and this requires the storage of the state of a particle at two time steps. This is not a huge computational expense, but this extra conditional check would have to be executed at every time step, for every particle.

To make the distinction more precise, let $\{\eta_j^n\}_{j\geq 1}$ denote the increasing sequence of jump times of the process X^n, and let $\theta = \{j : \eta_j^n \geq T\}$. Since X^n is right-continuous and the jump intensities are bounded on compact sets, for any random variable h such that $h \in (0, \eta_j^n - \eta_{j-1}^n)$ almost surely, we have $X^n(\eta_j^n) = X^n(\eta_j^n + h_j)$ a.s. . The difference is between estimating $\mathbb{P}_{x_0}(X^n(T) \notin D) = \mathbb{P}_{x_0}(X^n(\eta_{\theta-1}^n) \notin D)$ and $\mathbb{P}(X^n(\eta_\theta^n) \notin D)$.

Since the jump intensities are continuous functions of the position, and the jump

sizes are on the order of $1/n$, we expect the difference $\eta_\theta^n - \eta_{\theta-1}^n$ to be on order of $1/n$. We do not deem it necessary for our purposes to store the extra state information, since we expect the error due to the slight time difference to be comparable or smaller than the statistical error. In all the simulations where a terminal time escape problem is discussed, we will actually estimate $\mathbb{P}_{x_0}(X^n(\eta_\theta^n) \notin D)$.

Using importance functions which are not subsolutions. We advocate the use of (appropriately scaled) subsolutions as importance functions because in many of the cases we study, they are guaranteed to produce a splitting scheme which is asymptotically optimal in the sense of maximizing the decay rate (see Definition 2.1.2. A related reason is that a splitting scheme based on a subsolution is guaranteed to be stable, i.e. one does subexponential work. However, we wish to make it clear that any admissible importance function will produce an unbiased splitting scheme.

4.1.2 Statistics recorded

We introduce some terminology used throughout the section to refer unambiguously to different computational tasks, and we define the statistics recorded in the numerical simulations.

We will denote a single sample of the RESTART estimator by γ^n, where n is the large deviations scaling parameter. We note that the simulation of a single sample γ^n requires the simulation of multiple trajectories started at points where a threshold is crossed. For a given estimation problem, we generate multiple samples and average the result to obtain an estimate. We refer to the collection of all samples and the resulting estimate as a single *trial*.

Remark. We make a remark on scientific notation. We will use the exponential notation e-k to denote $\times 10^{-k}$, for instance, 3.02e-2 = 0.0302. This will save space in the tables.

Suppose γ_k^n, $k = 1, \ldots, K$ are samples of the estimator. Together, the samples $\{\gamma_k^n\}_{k=1}^K$ comprise a single trial. We record the following quantities:

- the estimate
$$\bar{\gamma}^n = \frac{1}{K} \sum_{k=1}^{K} \gamma_k^n,$$
- the formal confidence interval
$$[\bar{\gamma}^n - 1.96\, \sigma^n/\sqrt{K}, \bar{\gamma}^n + 1.96\, \sigma^n/\sqrt{K}],$$
where
$$\sigma^n = \sqrt{\frac{1}{K-1} \sum_{k=1}^{K} (\gamma_k^n - \bar{\gamma}^n)^2},$$
- the relative error
$$\frac{\sigma^n}{\gamma^n \sqrt{K}},$$
- the empirical decay rate
$$\frac{-\log\left(\frac{1}{K}\sum_{k=1}^{K}(\gamma_k^n)^2\right)}{-\log \bar{\gamma}^n},$$
- we denote the number of particles for the k^{th} sample γ_k^n at the i^{th} step of the algorithm by $N_{i,k}^n$, and we record the average maximum number of particles at

any given time i in the simulation,

$$\bar{N}^n = \frac{1}{K}\sum_{k=1}^{K} \max_i N_{i,k}^n,$$

and we also record the maximum over all trials,

$$\max_{i,k} N_{i,k}^n,$$

- the average work

$$\frac{1}{K}\sum_{k=1}^{K} w_k^n,$$

where the work w_k^n for the sample γ_k^n is the total number of individual time-steps computed,

$$w_k^n \doteq \sum_{i=0}^{\infty} N_{i,k}^n,$$

- the number of thresholds $J^n \doteq \lceil (\bar{V}(0) - \inf_{x\in D} \bar{V}(x))/\log R \rceil$.

4.2 Numerical examples

4.2.1 Escaping rest points

4.2.1.1 Using techniques from complex-balanced networks

Example 3.3.1. This example illustrates the utility of the splitting method around rest points over long time intervals. The subsolution based on the quasipotential will prove to be an effective one for the purpose of sampling near a rest point, as predicted by Theorem 2.2.15. We note that the question of what constitutes a long

simulation of a reaction network is heavily dependent on the average time-step, as will be discussed in the analysis section.

Problem data. The CRN is

$$\begin{array}{ccc} A+2B & \to & 3C \\ \uparrow & & \downarrow \\ D & \rightleftarrows & C+E \end{array}$$

$$\begin{array}{ccc} C & \leftarrow & A \\ & \searrow \quad \nearrow & \\ & 2E & \end{array}$$

with all rate constants equal to 1. The network is complex balanced and the initial condition is the unique complex-balanced equilibrium,

$$c = (2^{-2/3}, 2^{-2/3}, 2^{-2/3}, 1/4, 2^{-1/3}.$$

We take the domain to escape to be

$$D = \{|x_1 - c_1| < \delta\} \cup \{|x_2 - c_2| < \delta\} \cup \{|x_3 - c_3| < \delta\} \cup \{|x_4 - c_4| < \delta/2\}.$$

The subsolution is the one based on the quasipotential,

$$\bar{V}(x) = c_Q - \sum_{i=1}^{5} c_i \ell \left(\frac{x_i}{c_i} \right),$$

where $\ell(x) = xlogx - x + 1$ and c_Q is the minimum cost to escape $\bar{V}(x)$ (though we remind the reader that this constant does not actually affect the design of the splitting scheme).

Parameters. The parameters for the splitting simulation are: the large deviations scaling parameter n, the n-dependent time $T(n)$, the number of samples K, the deviation $\delta > 0$ in the set D, the splitting rate R, and the number of samples per trial K. We take the following values:

$$\begin{aligned} n &= 30, 40, 50 \\ T(n) &= n/6 \\ \delta &= 0.34 \\ R &= 2 \\ K &= 100. \end{aligned}$$

As will be discussed in the analysis section, the time scaling $T(n) = n/6$ was chosen to be long relative to the average timestep of the process. The relatively small choice of $K = 100$ samples per trial is due to the computational cost of the simulation.

Results. The results are presented in Table 4.1.

Analysis. The domain to escape requires that the first four coordinates must all deviate from their equilibrium value simultaneously, and this constraint can be satisfied because the domain of the network is all of $\mathbb{R}^5_+$. The domain is unbounded, so the number of thresholds could potentially be taken as large as one likes, since escape could at a location with arbitrarily large cost relative to the rest point. We

n	30	40	50
Average	7.53e-06	1.91e-06	4.30e-08
Relative error	1.56e-01	1.47e-01	1.31e-01
Confidence interval	[5.22e-06, 9.83e-06]	[1.36e-06, 2.46e-06]	[3.20e-08, 5.41e-08]
Decay	1.90	1.91	1.94
Avg. work	73084	262672	545808
Avg. max no. particles	662	1342	1845
Max no. particles	1863	3627	8692
No. thresholds	29	38	47

Table 4.1: Results for escape from D over the time interval $[0, T(n)]$ using the time-independent subsolution based on the quasipotential.

take a point with significantly greater cost than the minimal cost to exit c_Q and note the number of thresholds. Note that near the rest point, the average timestep is on the order of 5×10^{-3}.

What is most promising about the results is the fairly constant relative error and decay rate. This would certainly not occur with importance sampling, as trajectories which linger near the rest point would inevitably occur, and these would corrupt the performance of the scheme for reasons discussed in Section 2.2.3. In this sense, the combined results of asymptotic optimality and the solution of the HJB PDE for complex-balanced systems allow us to approximate previously inaccessible probabilities.

It is interesting to note that splitting occurs due to the variation in the fifth coordinate as well, even though this coordinate has no influence on whether or not escape occurs. It would be reasonable to attempt the same numerical example using an importance function of the form $-\sum_{i=1}^{4} c_i \ell(x_i/c_i)$, shifted by some constant so as to make it positive.

4.2.1.2 Using the moderate deviations approximation

Example 3.3.2 . This example illustrates the utility of the moderate deviations approximation when the rare event of interest is not "too rare", in the sense that the cost of trajectories is well approximated by a quadratic expansion around the law of large numbers. The subsolution proposed for Example 3.3.2 returns values which are quite small, and so the large deviations scaling parameter n must be large enough to ensure that there is more than one threshold (we recall that the number of thresholds is roughly $\bar{V}(0,0)/(n \log R)$, where $\bar{V}$ is the time-dependent subsolution and R is the splitting rate).

Problem data. The CRN is

$$\begin{aligned} B + C &\to 2A \\ A + C &\to 2B \\ A + B &\to 2C \end{aligned}$$

with all rate constants equal to one. The initial condition for the uncentered process is at the stable steady state $X^n(0) = (1,1,1)$. In this stoichiometric compatibility class, each coordinate X_i^n, $i = 1,2,3$ of the uncentered process X^n takes values in $[0,3]$. The problem of interest is to estimate

$$\mathbb{P}_{(1,1,1)}(X^n(T) \notin D),$$

where

$$D_\delta \doteq \bigcap_{i=1}^{3} \{x : x_i \in (1-\delta, 1+\delta)\}.$$

The subsolution is a piecewise minimum of affine functions,

$$\bar{V}(t,y) = \min_{i=1,\dots,6} \bar{V}^{\xi_i,c_i}(t,y).$$

Parameters. The parameters for the splitting simulation are: the terminal time T, the large deviations scaling parameter n, the deviation $\delta > 0$ which determines the terminal escape set D_δ, the number of samples K, and the splitting rate R. We will take the following parameter values:

$$\delta = 0.2, 0.25, 0.3$$

$$n = 200, 250, 300$$

$$R = 2$$

$$T = 1$$

$$K = 2000.$$

For $\delta = 0.2, 0.25, 0.3$, the values of the subsolution at the origin are respectively 2.01×10^{-2}, 3.13×10^{-2} and 4.51×10^{-2}. Due to the small cost of escape, we need to take n large and a small splitting rate to ensure that we have at least a few thresholds. For instance, with $n = 200, \delta = 0.2$ and a splitting rate of $R = 2$, we have 6 thresholds.

Results. We present results for $\delta = 0.2, 0.25, 0.3$ in Tables 4.2, 4.3 and 4.4.

Analysis. We expect the moderate deviations to behave well in this case for a few reasons. First, the network structure is symmetric, and so all three dimensions in which the trajectories must deviate carry equal cost, and we expect a quadratic

n	200	250	300
Average	1.26e-02	4.86e-03	1.49e-03
Relative error	4.67e-02	4.91e-02	5.22e-02
Confidence interval	[1.15e-02, 1.38e-02]	[4.39e-03, 5.32e-03]	[1.34e-03, 1.64e-03]
Decay	1.62	1.67	1.71
Avg. work	990	1339	1709
Avg. max no. particles	6	8	8
Max no. particles	25	35	45
No. thresholds	6	8	9

Table 4.2: Results for terminal escape from D_δ with $\delta = 0.2$ using the moderate deviations approximation.

n	200	250	300
Average	1.21e-03	2.00e-04	3.61e-05
Relative error	5.75e-02	6.16e-02	7.04e-02
Confidence interval	[1.07e-03, 1.35e-03]	[1.76e-04, 2.24e-04]	[3.11e-05, 4.11e-05]
Decay	1.70	1.75	1.77
Avg. work	1156	1529	1931
Avg. max no. particles	8	9	10
Max no. particles	53	61	88
No. thresholds	10	12	14

Table 4.3: Results for terminal escape from D_δ with $\delta = 0.25$ using the moderate deviations approximation.

approximation to be suitable. If one made the network reversible (adding the reactions (b, a) for all $(a, b) \in \mathcal{R}$), it would be complex balanced, and the quasipotential would accurately describe the cost over long time intervals. The quasipotential at a rest point for a complex-balanced system is also well-approximated by a quadratic. The domain to escape is also symmetric and an intersection of half-spaces, so the piecewise affine subsolution is expected to be a good approximation.

The results confirm the expectations. The relative error is less than 10% across all values of n and δ tested. While this is no guarantee of accuracy, it does suggest that the true variance of the estimator is relatively small compared to the probabilities of interest. The decay also increases steadily, going from 1.62 for $(n, \delta) = (200, 0.2)$ to 1.82 for $(n, \delta) = (300, 0.3)$. The probabilities are on the order of 10^{-7}, which is

n	200	250	300
Average	7.46e-05	8.38e-06	6.35e-07
Relative error	6.41e-02	7.10e-02	7.76e-02
Confidence interval	[6.53e-05, 8.40e-05]	[7.22e-06, 9.55e-06]	[5.39e-07, 7.32e-07]
Decay	1.77	1.79	1.82
Avg. work	1316	1727	2236
Avg. max no. particles	10	11	12
Max no. particles	56	75	91
No. thresholds	14	17	20

Table 4.4: Results for terminal escape from D_δ with $\delta = 0.3$ using the moderate deviations approximation.

well within the range of the moderate deviation-based splitting scheme. It would be much more expensive to estimate this probability with standard Monte Carlo.

4.2.2 Metastability

4.2.2.1 Using techniques for one-dimensional networks

Example 3.3.3. The Schlögl model is a classical example in the CRN literature and a natural candidate for the study of rare events associated to metastability. The quasipotential is available in analytic form from [3], Example 10, but it is just as easily found by numerical quadrature.

Problem data. The CRN is

$$3A \underset{\kappa_{23}}{\overset{\kappa_{32}}{\rightleftarrows}} 2A, \quad A \underset{\kappa_{01}}{\overset{\kappa_{10}}{\rightleftarrows}} \emptyset,$$

where the rate constants are

$$\kappa_{32} = 1, \kappa_{23} = 6, \kappa_{10} = 11, \kappa_{01} = 6.$$

There are two stable steady states $c_1 = 1$ and $c_2 = 3$, and an unstable one, $c_{12} = 2$. We will consider the problem of estimating

$$\mathbb{P}_1(X^n(t) \geq 2.5 \text{ for some } t \in [0, T]).$$

Parameters. The parameters in this simulation are: the large deviations scaling parameter n, the time T, the number of samples K, and the splitting rate R. We take

$$\begin{aligned} n &= 400 \\ T(n) &= 5 \\ R &= 2 \\ K &= 200. \end{aligned}$$

Results. The results for the Schlögl model are shown in Table 4.5. We run three separate trials for the same value of n.

Analysis. Despite the relatively few number of samples, each trial is on the same order of magnitude and demonstrates similar performance. The third trial is a bit of an outlier in terms of estimate, though this is not detected in the decay.

While the decay is significantly better than what it would have been for ordinary

Trial #	1	2	3
Average	1.95e-03	1.91e-03	1.64e-03
Relative error	2.02e-01	2.10e-01	2.15e-01
Confidence interval	[1.18e-03, 2.73e-03]	[1.13e-03, 2.70e-03]	[9.49e-04, 2.33e-03]
Decay	1.65	1.64	1.64
Avg. work	120143	113965	108138
Avg. max no. particles	10	10	10
Max no. particles	29	30	32
No. thresholds	8	8	8

Table 4.5: Three trials for the transition probability of the Schlögl model.

Monte Carlo, it is not close to 2. There are at least two reasons for this. First, the large deviations scaling parameter n might simply not be large enough, so that the large deviations principle does not describe the rare event behavior accurately enough. Indeed, the estimates are on the order of 10^{-3}, which is not terribly small compared to the probabilities estimated in other examples which exhibit good performance. A related reason is that the splitting scheme does not encourage particles to continue past the unstable steady state at $c_{12} = 2$, and some could cross back into the region $(0, 2)$.

Example 3.3.4. This example explores a network with five steady states, three of which are stable, and the probabilities of transitioning between the three domains of attraction. The cost (as determined by the quasipotential) to transition between the various states is on the order of 10^{-3}, and the large deviations scaling parameter will need to be quite large in order for the transitions to be "rare". In turn, the expected difference between successive jump times will be small, on the order of 10^{-7} for $n = 10^3$.

Problem data. The CRN is

$$\begin{aligned} 5A &\underset{\kappa_{45}}{\overset{\kappa_{54}}{\rightleftarrows}} 4A \\ 3A &\underset{\kappa_{23}}{\overset{\kappa_{32}}{\rightleftarrows}} 2A \\ A &\underset{\kappa_{01}}{\overset{\kappa_{10}}{\rightleftarrows}} \emptyset, \end{aligned}$$

where

$$\kappa_{54} = 1, \kappa_{45} = 94, \kappa_{32} = 2745, \kappa_{23} = 29312, \kappa_{10} = 104660, \kappa_{01} = 78000.$$

There are three stable steady states: $c_1 = 1, c_2 = 12, c_3 = 50$, and two unstable ones, $c_{12} = 5$ and $c_{23} = 26$. The initial condition $X^n(0)$ will vary depending on the problem, but all initial conditions are in the only stoichiometric compatibility class $\mathcal{S} = [0, \infty)$. The problem is to estimate

$$p_{ij}^n(\delta) = \mathbb{P}_{c_i}(X^n(t) \in B_{\delta_{ij}}(c_j) \text{ for some } t \in [0, T]), \quad i, j \in \{1, 2, 3\}, i \neq j,$$

where we recall that $B_\delta(c) = \{x : |x - c| < \delta\}$.

Parameters. There are a few parameters for the splitting simulation: the large deviations scaling parameter n, the number of samples K, the splitting rate R, the time T, and the distances from the neighbourhoods $\delta_{ij} > 0$. Moreover, there are six different probabilities to estimate.

Due to the large rate constants, the average time-step is very small, and consequently the time intervals will be a lot smaller than for the other examples. Nevertheless, they will be long for what might be considered the "natural timescale"

of the process, which for our purposes is simply the expectation of the difference between two subsequent jump times (the holding time). We focus on comparing the probabilities for the various pairs (i, j). The average time-step can change by an order of magnitude depending on the pair, and so we will use different values of T for the different pairs. We denote these times T_{ij}. We use the following parameter values: $n = 30$, $K = 100$, $R = 2$, and

Pair	δ_{ij}	T_{ij}	Pair	δ_{ij}	T_{ij}
(1, 2)	4	1e-3	(2, 1)	2	5e-4
(2, 3)	20	2e-3	(3, 2)	4	5e-5
(1, 3)	20	3e-3	(3, 1)	2	5.5e-4

We use large values of δ_{ij} because we wish to capture the probability of transitioning to another domain of attraction. Once in the new domain, the splitting scheme does not encourage particles to continue towards the new stable steady state, and there is always the possibility that the trajectory doubles back to the domain of attraction of the initial steady state. In some cases this "recrossing" event is important, and in this case the δ's would be taken smaller.

The time intervals chosen are simply our guesses obtained through trial and error. The time intervals cannot be too long, because for fixed n the probability of any transition occurring goes to 1 as the time interval increases. The time intervals cannot be too short, because one still expects the mean transition time from i to j to scale like $e^{n\bar{V}_{ij}(c_i)}$.

Results. The results for the transition probabilities associated to the pairs (1, 2), (2, 3) and (1, 3) can be seen in Table 4.6. The results for the pairs (3, 2), (2, 1) and (3, 1) can be seen in Table 4.7.

Pair	(1, 2)	(2, 3)	(1, 3)
Average	3.97e-05	5.80e-02	6.63e-06
Relative error	1.49e-01	9.12e-02	1.05e-01
Confidence interval	[2.81e-05, 5.13e-05]	[4.76e-02, 6.83e-02]	[5.26e-06, 8.00e-06]
Decay	1.89	1.79	1.94
Avg. work	31637	1071893	4376529
Avg. max no. particles	21	17	37
Max no. particles	49	30	72
No. thresholds	16	7	22

Table 4.6: Results for the three transition probabilities for the pairs (1, 2), (2, 3) and (1, 3).

Pair	(2, 1)	(3, 2)	(3, 1)
Average	7.53e-02	8.13e-04	9.19e-05
Relative error	1.17e-01	8.48e-02	4.96e-02
Confidence interval	[5.80e-02, 9.26e-02]	[6.78e-04, 9.48e-04]	[8.30e-05, 1.01e-04]
Decay	1.67	1.92	1.98
Avg. work	429995	4622456	53937776
Avg. max no. particles	12	34	56
Max no. particles	22	62	92
No. thresholds	6	14	19

Table 4.7: Results for the three transition probabilities for the pairs (2, 1), (3, 2) and (3, 1).

Analysis. The values of the six subsolutions at the respective initial steady states represents the total "cost" of respective transitions. By evaluating $\bar{V}_{ij}(c_i)$ for different pairs (i, j), one can order the costs from greatest to smallest: (1, 3), (3, 1), (1, 2), (3, 2), (2, 3), (2, 1). The transition costs for (2, 1) and (2, 3) are also roughly on the same order of magnitude, as are those of (1, 3) and (3, 1).

The relative costs give some hint as to the order of the transition probabilities p_{ij}^n. For this to be the case, the time intervals must have been chosen appropriately, and as mentioned previously the chosen times T_{ij} were simply educated guesses, with the stipulation that $T_{13} = T_{12} + T_{23}$ and $T_{31} = T_{21} + T_{32}$. Nevertheless, the order of the costs correlates well with the estimated probabilities. For instance, p_{13}^n is the smallest transition probability, and p_{23}^n and p_{21}^n are on the same order of magnitude.

Moreover, the heuristic relationships such as

$$p_{13}^n \geq p_{12}^n p_{23}^n$$

seem to hold; see (3.3.5) in the previous chapter.

It is worth noting that some of estimates required a tremendous amount of computational work, with the (3, 1) pair being an extreme case: 5.36×10^7 individiual time steps were computed for the simulation of at most 56 particles over a time interval of length 5.5×10^{-4}.

The decay rate seems smaller for the larger probabilities but is close to 2 for probabilities on the order of 10^{-4} and smaller. This confirms the belief that the quaispotential will be appropriate for these relatively simple one-dimensional metastability problems over long time intervals. The relative error is also promising, since it is less than 15% for every trial.

4.2.3 Boundary behavior

4.2.3.1 Subsolutions based on moderate deviations to leave boundary

Example 3.3.8. This example illustrates the flexibility of moderate deviations-based subsolutions: one can explicitly center around the law of large numbers dynamics and describe deviations away from them at the terminal time. This is not possible with a time-independent subsolution (such as one based on the quasipotential), while a time-dependent subsolution can be very difficult to identify when no approximation is made. As in Example 3.3.2, an issue with the affine function used

for the subsolution is that it takes values which are very small, and so a large n is required to see even a few thresholds.

Problem data. The CRN is

$$\emptyset \to A \to B.$$

with all rate constants equal to one. The initial condition is $X^n(0) = (0, 0)$ for all $n \geq 0$, and X^n takes values in the only stationary compatibility class $\mathcal{S} = \mathbb{R}^2_+$. The problem is to estimate

$$\mathbb{P}(X_1^n(T) - X_1^0(T) \leqslant -\delta),$$

The subsolution $\bar{V}(t, y)$ is time-dependent in such a way that for each fixed $t \in [0, 1]$, $\bar{V}(t, \cdot)$ is affine.

Parameters. The parameters for the splitting simulation are: the large deviations scaling parameter n, the number of samples K, and the splitting rate R. We take the following parameter values:

$$\delta = 0.15, 0.20, 0.25$$

$$n = 300$$

$$R = 3$$

$$T = 1$$

$$K = 8000.$$

Results. The results are presented in Table 4.8.

δ	0.15	0.20	0.25
Average	1.21e-04	1.35e-06	1.10e-09
Relative error	1.13e-01	1.45e-01	3.66e-01
Confidence interval	[9.45e-05, 1.48e-04]	[9.68e-07, 1.74e-06]	[3.11e-10, 1.88e-09]
Decay	1.49	1.62	1.66
Avg. work	468	488	469
Avg. max no. particles	2	3	2
Max no. particles	25	40	37
No. thresholds	5	9	14

Table 4.8: Results of the simulation for the network $\emptyset \to A \to B$. The simulations occur until the terminal escape time $T = 1$.

Analysis. As is typical of the moderate deviations approximation, the maximum number of particles is not exceedingly large. The estimated decay rate is far from 2, but significantly better than 1, as one would obtain with standard Monte Carlo. Indeed, with only $K = 8000$ samples of the RESTART estimator, probabilities on the order of 10^{-9} can be estimated. We note, however, that the relative error is quite large, especially for the $\delta = 0.25$ trial.

4.2.3.2 Subsolutions for hitting the boundary of complex-balanced systems

As discussed in Section 3.3.3.3 of Chapter 3, it is possible to estimate the probability of hitting the boundary over increasing time intervals using a time-independent subsolution based on the quasipotential. Four networks were considered: $A \rightleftarrows \emptyset$, $2A \rightleftarrows \emptyset$, $A \rightleftarrows B$, and $A \rightleftarrows B \rightleftarrows \emptyset$, all with rate constants taken equal to 1. In this section we investigate this problem numerically.

Example 3.3.9.

Problem data. The CRN is

$$A \rightleftarrows \emptyset.$$

with all rate constants equal to one. The initial condition is $X^n(0) = 1$ for all $n \geq 0$, and X^n takes values in the only stationary compatibility class $\mathcal{S} = [0, \infty)$. The problem of interest is

$$\mathbb{P}_1(X^n(t) = 0 \text{ for some } t \in [0, T(n)]),$$

where $T(n)$ is linear in n. This can be formulated as an escape problem by taking $D = [0, \infty)$ and demanding that $X^n(t) \notin D$ instead. The subsolution $\bar{V}$ is The subsolution is a piecewise minimum of affine functions,

$$\bar{V}(x) = x - x \log x,$$

and we note that $\bar{V}(x) \leq 0$ for $x = 0$.

Parameters. The parameters for the splitting simulation are: the increasing time intervals $T(n)$, the large deviations scaling parameter n, the number of samples K, and the splitting rate R. We take the following parameter values:

$$n = 20, 40, 60$$
$$R = 3$$
$$T = n/2$$
$$K = 2000.$$

Results. The results are presented in Table 4.9

n	20	40	60
Average	3.41e-07	3.28e-15	1.41e-23
Relative error	4.49e-02	3.92e-02	3.92e-02
Confidence interval	[3.11e-07, 3.71e-07]	[3.03e-15, 3.53e-15]	[1.31e-23, 1.52e-23]
Decay	1.89	1.96	1.97
Avg. work	866	4313	10665
Avg. max no. particles	81	234	431
Max no. particles	711	2093	3603
No. thresholds	19	37	55

Table 4.9: Results of the simulation for the network $A \rightleftarrows \emptyset$. The simulations occur over time intervals $[0, T(n)]$ with $T(n) = n/2$.

Analysis. The results show excellent performance. The empirical decay approaches 2 as n increases, and the relative error remains at around 4%, despite the fact that we are using only $K = 2000$ samples of the RESTART estimator for probabilities on the order of 10^{-7}. One reason the splitting scheme performs exceptionally well is that the process is one-dimensional, and there is no additional variance incurred due to escape at different locations; the weight associated to each escaped particle is always the same.

Example 3.3.10.

Problem data. The CRN is

$$2A \rightleftarrows \emptyset.$$

with all rate constants equal to one. The initial condition is $X^n(0) = 1$ for all $n \geq 0$, and X^n takes values in the only stationary compatibility class $\mathcal{S} = [0, \infty)$. The problem of interest is

$$\mathbb{P}_1(X^n(t) = 0 \text{ for some } t \in [0, T(n)]),$$

where $T(n)$ is linear in n. The subsolution $\bar{V}$ is the same as for the network $A \rightleftarrows \emptyset$, namely,

$$\bar{V}(x) = x - x \log x,$$

and we note that $\bar{V}(x) \leq 0$ for $x = 0$.

Parameters. The parameters for the splitting simulation are the same as for $A \rightleftarrows \emptyset$: the increasing time intervals $T(n)$, the large deviations scaling parameter n, the number of samples K, and the splitting rate R. We use the same parameter values as before:

$$n = 20, 40, 60$$

$$R = 3$$

$$T = n/2$$

$$K = 2000.$$

Results. The results are presented in Table 4.10.

n	20	40	60
Average	4.48e-06	2.36e-14	9.80e-23
Relative error	2.76e-02	3.71e-02	4.49e-02
Confidence interval	[4.24e-06, 4.73e-06]	[2.19e-14, 2.54e-14]	[8.94e-23, 1.07e-22]
Decay	1.92	1.96	1.97
Avg. work	2689	6358	11266
Avg. max no. particles	1080	2530	4364
Max no. particles	6567	28446	60037
No. thresholds	19	37	55

Table 4.10: Results of the simulation for the network $2A \rightleftarrows \emptyset$. The simulations occur over time intervals $[0, T(n)]$ with $T(n) = n/2$.

Analysis. We compare the results of this network to those of the previous one. We note that the relative error behaves just as well, and the decay gets close to 2 just as before. Since the subsolution used is the same, the number of thresholds does not change. However, owing to the increased jump size of $2/n$, the hitting probability has noticeably increased by an order of magnitude for $n = 20, 40$. This is to be expected, since in the pre-limit it is easier for a single jump of size $2/n$ to occur, as opposed to two successive jumps of size $1/n$. As n increases, the jump sizes shrink and the time interval gets larger, and this difference disappears. This can be seen when $n = 60$.

Example 3.3.11.

Problem data. The CRN is

$$A \rightleftarrows B.$$

with both rate constants equal to one. The initial condition is $X^n(0) = (1/2, 1/2)$ for all $n \geq 0$, and X^n takes values in the stoichiometric compatibility class $\mathcal{S}_{(1/2,1/2)} = \{x : x_1 + x_2 = 1\} \cap \mathbb{R}^2_+$. The problem of interest is

$$\mathbb{P}_1(X^n_1(t) = 0 \text{ for some } t \in [0, T(n)]),$$

where $T(n)$ is linear in n. The subsolution $\bar{V}$ is

$$\bar{V}(x) = \log(2) - x_1 \log(2x_1) - x_2 \log(2x_2),$$

and we note that $\bar{V}(x) \leq 0$ for $x = 0$.

Parameters. The parameters for the splitting simulation are: the increasing time intervals $T(n)$, the large deviations scaling parameter n, the number of samples K, and the splitting rate R. We use the same parameter values as the two previous examples, so as to have a basis for comparison:

$$n = 20, 40, 60$$

$$R = 3$$

$$T = n/2$$

$$K = 2000.$$

Results. The results are presented in Table 4.11.

n	20	40	60
Average	1.73e-04	3.20e-10	4.81e-16
Relative error	3.39e-02	4.82e-02	4.90e-02
Confidence interval	[1.61e-04, 1.84e-04]	[2.90e-10, 3.51e-10]	[4.35e-16, 5.28e-16]
Decay	1.86	1.92	1.95
Avg. work	730	2128	3797
Avg. max no. particles	99	335	277
Max no. particles	424	2047	1765
No. thresholds	13	26	38

Table 4.11: Results for the problem of hitting the boundary $\{x_1 = 0\}$.

Analysis. As discussed in Section 3.3.3.3 on the behavior of stochastic trajectories which hit the boundary, the nature of the singularity that the gradient of the quasipotential has is similar for the networks $A \rightleftarrows \emptyset, 2A \rightleftarrows \emptyset$ and $A \rightleftarrows B$. In this last case, on the domain $\mathcal{S}_{(1/2,1/2)}$ the concentration of B goes to 1 as the concentration of A goes to 0. The "push" exerted by the $B \to A$ becomes equal to the push $\emptyset \to A$ at the boundary $\{x_1 = 0\}$.

Nevertheless, since the rate of $B \to A$ is $x_2 \leq 1$, the cost to hit the boundary in

this case should be smaller. This can be seen by directly evaluating the quasipotential: $Q_{A\rightleftarrows B}((0,1)) = 1/2 + \log(2)/2 \approx 0.85$. It is also reflected in the results.

Example 3.3.12. In this example we compare the probability of hitting a facet $\{x_i = 0\}$ with the probability of hitting a corner. The cost (in the sense of the quasipotential) to hit a corner is twice the cost to hit a facet, and consequently the number of thresholds is doubled. When the process is near the initial condition, the average time step is roughly half of what it was for the previous examples, which further increases the computational cost of this example. We will therefore use very few samples for each trial.

Problem data. The CRN is

$$A \rightleftarrows B \rightleftarrows \emptyset.$$

with both rate constants equal to one. The initial condition is $X^n(0) = (1,1)$ for all $n \geq 0$, and X^n takes values in the only stoichiometric compatibility class $\mathcal{S} = \mathbb{R}^2_+$. The problem of interest is

$$\mathbb{P}_1(X_1^n(t) = 0 \text{ for some } t \in [0, T(n)]),$$

where $T(n)$ is linear in n. The subsolution $\bar{V}$ is

$$\bar{V}(x) = \log(2) - x_1 \log(2x_1) - x_2 \log(2x_2),$$

and we note that $\bar{V}(x) \leq 0$ for $x = 0$.

Parameters. The parameters for the splitting simulation are: the increasing time intervals $T(n)$, the large deviations scaling parameter n, the number of samples K, and the splitting rate R. We use the same parameter values for R, n and $T(n)$ as the two previous examples so as to have a basis for comparison, but we use significantly fewer samples due to the greater computational cost.

$$n = 20, 40, 60$$
$$R = 3$$
$$T = n/2$$
$$K = 20.$$

We study the problem of hitting the facets $\{x_1 = 0\}$, $\{x_2 = 0\}$, and the corner $\{x_1 = 0, x_2 = 0\}$.

Results. The results for hitting $\{x_1 = 0\}$ and $\{x_2 = 0\}$ are presented in Tables 4.12 and 4.13. The results for hitting the corner $\{x_1 = 0, x_2 = 0\}$ are presented in Table 4.14.

n	20	40	60
Average	1.55e-05	1.47e-13	6.76e-22
Relative error	9.85e-02	8.14e-02	7.73e-02
Confidence interval	[1.25e-05, 1.85e-05]	[1.23e-13, 1.70e-13]	[5.73e-22, 7.78e-22]
Decay	1.98	2.00	2.00
Avg. work	244203	1971107	6102630
Avg. max no. particles	520	1499	2735
Max no. particles	755	2174	3790
No. thresholds	19	37	55

Table 4.12: Results for the problem of hitting the facet $\{x_1 = 0\}$.

Analysis. As discussed in Section 3.3.3.3 on the behavior of stochastic trajectories which hit the boundary, the nature of the singularity in the gradient of the quasipo-

n	20	40	60
Average	7.95e-06	6.83e-14	2.52e-22
Relative error	1.28e-01	9.55e-02	1.07e-01
Confidence interval	[5.96e-06, 9.95e-06]	[5.55e-14, 8.11e-14]	[1.99e-22, 3.05e-22]
Decay	1.98	1.99	2.00
Avg. work	234686	1796300	6321918
Avg. max no. particles	519	1498	2412
Max no. particles	916	2347	3713
No. thresholds	19	37	55

Table 4.13: Results for the problem of hitting the facet $\{x_2 = 0\}$.

n	20	40	60
Average	1.55e-14	2.36e-31	2.47e-48
Relative error	1.46e-01	1.45e-01	2.07e-01
Confidence interval	[1.11e-14, 2.00e-14]	[1.69e-31, 3.03e-31]	[1.47e-48, 3.48e-48]
Decay	1.98	1.97	1.99
Avg. work	458124	3434189	13603039
Avg. max no. particles	1149	3433	7534
Max no. particles	1574	5573	14262
No. thresholds	37	73	110

Table 4.14: Results for the problem of hitting the corner.

tential is similar for the networks $A \rightleftarrows \emptyset, 2A \rightleftarrows \emptyset$ and $A \rightleftarrows B$. Moreover, the overall cost (in terms of the quasipotential) to hit a facet is the same as for the networks $A \rightleftarrows \emptyset$ and $2A \rightleftarrows \emptyset$. Comparing the estimated averages of Tables 4.12 and 4.13 with those in Tables 4.9 and 4.10, we see that the probabilities are indeed on the same order of magnitude.

It is worth noting, however, that in this network the average work is orders of magnitude larger. This is in part due to splitting occurring in both the x_1 and x_2 coordinates, even when the problem is to hit a facet $\{x_i = 0\}$. In addition, the average timestep is smaller than in the two previous examples, and so reaching a time of $n/2$ becomes more expensive.

Bibliography

[1] Andrea Agazzi, Amir Dembo, and Jean-Pierre Eckmann. Large deviations theory for markov jump models of chemical reaction networks. *The Annals of Applied Probability*, 28(3):1821–1855, 2018.

[2] Andrea Agazzi, Amir Dembo, and Jean-Pierre Eckmann. On the geometry of chemical reaction networks: Lyapunov function and large deviations. *Arxiv preprint*, 2018.

[3] D Anderson, G Craciun, M Gopalkrishnan, and C Wiuf. Lyapunov functions, stationary distributions, and non-equilibrium potential for reaction networks. *Bulletin of Mathematical Biology*, 77(9):1744–1767, 2015.

[4] David Anderson, Gheorghe Craciun, and Thomas Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010.

[5] Anthony Bartholomay. Stochastic models for chemical reactions. i. theory of the unimolecular reaction process. *Bull. Math. Biophys*, 20:175–190, 1958.

[6] Anthony Bartholomay. Stochastic models for chemical reactions. ii. the unimolecular rate constant. *Bull. Math. Biophys*, 21(363-373), 1959.

[7] Patrick Billingsley. *Convergence of Probability Measures*. Wiley, 2nd edition, 1999.

[8] Amarjit Budhiraja, Paul Dupuis, Markus Fischer, and Kavita Ramanan. Limits of relative entropies associated with weakly interacting particle systsems. *Arxiv Preprint*, 2015.

[9] Terence Chan. Large deviations and quasi-stationarity for density-dependent birth-death processes. *J. Australian. Math. Soc. Ser. B*, 40:238–256, 1998.

[10] Kai Lai Chung. *A Course in Probability Theory*. Academic Press, third edition, 2001.

[11] G. Craciun. Toric differential inclusions and a proof of the global attractor conjecture. *Arxiv preprint*, 2015.

[12] Michael G. Crandall and Pierre-Louis Lions. Viscosity solutions of hamilton-jacobi equations. *Trans. Amer. Math. Soc.*, 277:1–42, 1983.

[13] M. Day. On the exponential exit law in the small parameter exit problem. *Stochastics*, 8(4):297–333, 1983.

[14] P. T. de Boer. Some observations on importance sampling and restart. *Proc. of the 6th International Workshop on Rare Event Simulation*, 2006.

[15] Max Delbrück. Statistical fluctuations in autocatalytic reactions. *Journal of Chemical Physics*, 8(1):120–124, 1940.

[16] P. Dupuis, K. Spiliopoulos, and X. Zhou. Escaping from an attractor: importance sampling and rest points i. *The Annals of Applied Probability*, 25(5): 2909–2598, 2015.

[17] Paul Dupuis and Amarjit Budhiraja. *Representations and Weak Convergence Methods for the Analysis and Approximation of Rare Events*. Springer-Verlag, 2019.

[18] Paul Dupuis and Thomas Dean. Splitting for rare event simulation: a large deviation approach to design and analysis. *Stochastic Processes and their Applications*, 119(2):562–587, 2009.

[19] Paul Dupuis and Thomas Dean. The design and analysis of a generalized restart/dpr algorithm for rare event simulation. *Annals of Operations Research*, 189:63–102, 2011.

[20] Paul Dupuis and Richard Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley, 1997.

[21] Paul Dupuis and Hui Wang. Subsolutions of an isaacs equation and efficient schemes for importance sampling. *Mathematics of Operations Research*, 32(3): 723–757, 2007.

[22] Paul Dupuis, Hui Wang, and Ali Devin-Sezer. Dynamic importance sampling for queuing networks. *The Annals of Applied Probability*, 17(4):1306–1346, 2007.

[23] Paul Dupuis, Amarjit Budhiraja, Markus Fischer, and Kavita Ramanan. Local stability of kolmogorov forward equations for finite state nonlinear markov processes. *Arxiv Preprint*, 2015.

[24] Paul Dupuis, Anne Buijsrogge, and Michael Snarski. Splitting algorithms for rare event simulation over long time intervals. *Arxiv preprint*, 2018.

[25] Paul Dupuis, Guo-Jhen Wu, and Michael Snarski. Infinite swapping using i.i.d. samples. *(Arxiv preprint)*, 2018.

[26] M. I. Dykman, E. Mori, J. Ross, and P.M. Hunt. Large fluctuations and optimal paths in chemical kinetics. *J. Chem. Phys.*, 100(8), 1994.

[27] D Elderfield and D D Vvedensky. Non-equilibrium scaling in the schlogl model. *Journal of Physics A*, 18:2591–2601, 1985.

[28] Robert Endres. Bistability: Requirements on cell-volume, protein diffusion and thermodynamics. *PLoS One*, 10(4), 2015.

[29] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, 1986.

[30] Zhou Fang and Chuanhou Gao. Lyapunov function partial differential equations for chemical reaction networks. *Arxiv Preprint*, 2016.

[31] Martin Feinberg. Complex balancing in general kinetic systems. *Arch. Rational Mech. Anal*, 49:172–186, 1972.

[32] Martin Feinberg. Lectures on chemical reaction networks, 1979. URL https://crnt.osu.edu/LecturesOnReactionNetworks.

[33] Martin Feinberg. Chemical reaction network structure and the stability of complex isothermal reac- tors - i. the deficiency zero and deficiency one theorems. *Chem. Eng. Sci*, 42:2229–2268, 1987.

[34] Martin Feinberg. The existence and uniqueness of steady states for a class of chemical reaction networks. *Archive for Rational Mechanics and Analysis*, 132: 311–370, 1995.

[35] Mark Freidlin and Alexander Wentzell. *Random Perturbations of Dynamical Systems*. Springer, 3rd edition, 2012.

[36] Crispin Gardiner. *Handbook of Stochastic Methods*. Springer, 1997.

[37] M. J. J. Garvels. *The splitting method in rare event simulation*. PhD thesis, University of Twente, 2000.

[38] P. W. Glynn and W. Whitt. The asymptotic efficiency of simulation estimators. *Operations Research*, 40(3):505–520, 1992.

[39] Gilles Gnacadja. Reachabaility, persistence, and constructive chemical reaction networks (part i). *Journal of Mathematical Chemistry*, 2011.

[40] Jeremy Gunawardena. Chemical reaction network theory for in-silico biologists, 2003. URL http://vcp.med.harvard.edu/papers/crnt.pdf.

[41] Jack Hale. *Ordinary differential equations*. Robert E. Krieger Publishing Company, second edition, 1980.

[42] V Hars and J Toth. On the inverse problem of reaction kinetics. *Colloquia Mathematica Societatis Janos Bolyai 30*, 1979.

[43] Jean-Baptiste Hiriart-Urruty and Claude Lemarechal. *Convex Analysis and Minimization Algorithms*. Springer, 1993.

[44] F. Horn. Necessary and sufficient conditions for complex balancing in chemical kinetics. *Arch. Rational Mech. Anal., 49, 172-186*, 1972.

[45] R. Jackson and F. Horn. General mass action kinetics. *Archive for Rational Mechanics and Analysis*, 47(2), 1972.

[46] A. Joffe and M. Metivier. Weak convergence of sequences of semimartingales with applications to multitype branching processes. *Adv. Appl. Prob.*, 18:20–65, 1986.

[47] Dane Johnson and Paul Dupuis. Moderate deviations for recursive stochastic algorithms. *Stochastic Systems*, 5(1):87–119, 2015.

[48] Alex Kamenev. *Field Theory of Non-Equilibrium Systems*. Cambridge University Press, 2011.

[49] Ioannis Karatzas and Steve Shreve. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, 1991.

[50] Min Ke, Zhou Fang, and Chuanhou Gao. Generalized gibbs' free energy for stability analysis of chemical reaction networks, 2017.

[51] Thomas Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *Journal of Chemical Physics*, 57(7):2976–2978, 1972.

[52] Thomas Kurtz. Strong approximation theorems for density dependent markov chains. *Stochastic Processes Applications*, 6:223–240, 1977/78.

[53] Thomas Liggett. *Continuous Time Markov Processes*, volume 113. American Mathematical Society, 2010.

[54] Robert Patterson and Michiel Renger. Large deviations of reaction fluxes. *Arxiv Preprint*, 2018.

[55] Riccardo Rao and Massimilano Esposito. Nonequilibrium thermodynamics of chemical reaction networks: Wisdom from stochastic thermodynamics. *Physical Review X*, 6, 2016.

[56] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.

[57] Khachik Sargsyan. Fluctuations in chemical reactions in a large volume, 2006.

[58] F Schlogl. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift Fur Physik A*, 253:147–161, 1972.

[59] Adam Schwartz and Alan Weiss. *Large deviations for performance analysis*. Chapman and Hall, 1995.

[60] Anne Shiu. Algebraic methods for biochemical reaction network theory. *PhD Thesis*, 2010.

[61] B. Sturmfels, G. Craciun, A. Dickenstein, and A. Shiu. Toric dynamical systems. *Journal of Symbolic Computation*, 44, 2009.

[62] Melissa Vellela and Hong Qian. Stochastic dynamics and non-equilibrium thermodynamics of a bistable chemical system: the schlögl model revisited. *Journal of the Royal Society Interface*, 6(39):925–940, 2009.

[63] M. Villen-Altamirano and J. Villen-Altamirano. Restart: A method for accelerating rare event simulations. *Proc. of the 13th International Teletraffic Congress, Queuing, Performance and Control in ATM*, pages 71–76, 1991.

[64] A. I. Vol'pert. Differential equations on graphs. *Mat. Sb.*, 88(4):578–588, 1972.

[65] A. I. Vol'pert and S. V Hudjaev. *Analysis in Classes of Discontinuous Functions and Equations of Mathematical Physics*, volume 8. Springer, 1985.

Index