



ON A STRONGLY CONVEX APPROXIMATION
OF A STOCHASTIC OPTIMAL CONTROL PROBLEM
FOR IMPORTANCE SAMPLING OF METASTABLE DIFFUSIONS

Dissertation
zur Erlangung des Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.)

am Fachbereich Mathematik und Informatik
der Freien Universität Berlin

vorgelegt von
Han Cheng LIE

Berlin, 2015

Copyright © 2015 Han Cheng Lie

Erstgutachter: Prof. Dr. Christof Schütte

Zweitgutachter: Prof. Dr. Grigorios Pavliotis

Tag der Disputation: 21 März 2016

Acknowledgements

I thank my supervisor, Christof Schütte for his time, for his support in dealing with administrative matters, for offering me a six-month contract with the SFB 1114 at the end of my doctoral studies, and for giving me the freedom to fail and succeed on my own. I thank Carsten Hartmann for sharing his time and ideas, for inviting me to workshops in Edinburgh and Champéry, and especially for his feedback on the first draft of this thesis. I thank both Christof and Carsten for their encouragement.

I thank my mentor, Günter Ziegler, for his time, for the conversations we had, for candidly answering questions I had about the profession, and for his support. I am extremely grateful to have had a mentor who took mentoring seriously and mentored well.

I thank the people who, more than three years ago, decided to grant me a Ph.D. scholarship from the International Max Planck Research School (IMPRS) for Computational Biology and Scientific Computing. While several people may have had a part to play, I wish to thank Martin Vingron in particular, for his interest in and support of my work.

I thank the many people who helped me with many administrative matters. At the FU Berlin, these include Dorothé Auth and Christian Wendt of the Biocomputing Group; Bettina Junkes, Stephanie Neuberg-Schönborn and Stephanie Auerbach of the Geschäftsstelle Verbundprojekte; Dr. Nina Fabjančič and Dr. Kristine al-Zoukra, of the SFB 1114; Maria Berschadski of the Personalstelle, and Angelika Pasanec of the Promotionsbüro. I also thank Annerose Steinke and Katrin Fedtke in the administration department of the Zuse Institut Berlin (ZIB), and Fabian Feutlinske and Kirsten Kelleher of the IMPRS.

Most of the last three years were spent in the mathematics department at the Humboldt-Universität zu Berlin. I wish to thank Nicolas Perkowski and Peter Imkeller, for their encouragement at the optimal transport seminar; Todor, Victor, and Klébert for answering some questions I had on stochastic analysis; Daniele, Pedro, Todor, Gregor, Ignacio, Juan, Patrick, Sidy, Marcel (from Spain), Marcel (from Berlin), Victor, Gabriel, Marta, Miguel, Irfan, Klebert, Antareep, Emre, Shirley, and Maryna, all of whom were at one point in time in the BMS lounge and provided (to quote Robin Hartshorne) ‘an enriched human context for my life’. I thank Irfan in particular for his hospitality in July and August 2015. Among many other things, I am grateful to Mara for her time, for introducing me to the music of Iron Maiden, Oscar Peterson, and Rush, and for taking me to the Berlin Philharmonic.

In the past five months that I have spent writing, I enjoyed good times with many people: Borong, Farnaz, and Adam of the ZIB; Esam, Felix, Benjamin, Guillermo, Nuria, Sebastian, Kasia, Marjan, Maureen, Kaveh, Sulav, Max, Evgenia, Péter, Vikram, Ralf, Ana, Marco, Victor, Iliusi, Nada, Sunil, Ling, Alexandra, Rudolf, Iurii, and Stefan of Arnimallee 6; Lara, Animesh, Patrick, Stefan, Luigi, Wei, Omar, and Jannes of Arnimallee 9; Ágnes, Julie, Arne, Atul, Adrián, Ximena, Giovanni, Anna, Alberto, Giuseppe, Sara, Massimo, Eugenio, Matti, Katerina, Aditya, Antareep, Marco, Sandra, Mats, Stan, Mark, José, Jakub, Tatiana, Albert, Lucia, Chris, and Wayne. I thank all these people. I thank Jannes for reading and giving feedback on my thesis, and Julie for correcting my German in the first iteration of my Zusammenfassung.

I thank my family for supporting me in this endeavour, and for trying to understand what it is precisely that I am doing and how the profession works. In particular, I thank my parents. The fact that I can appreciate beauty in mathematics is largely due to the education that they made possible for me to enjoy. I dedicate this thesis to them.

Contents

Table of Contents	v
1 Introduction	1
2 Optimal control of diffusions	5
2.1 Requisite theory of continuous-time stochastic processes	5
2.1.1 Girsanov's formula and change of measure	9
2.1.2 The Dirichlet problem	17
2.2 Stochastic optimal control of diffusions	20
2.2.1 Path functionals	21
2.2.2 Optimal control and importance sampling	24
3 Strongly convex approximations	34
3.1 Variational analysis of the control functional	36
3.1.1 The first variation and its first variation	36
3.2 Relations deriving from Itô's formula	44
3.2.1 Martingales constructed from bounded basis functions	45
3.2.2 Pairwise independence via non-overlap condition on supports	47
3.3 Strong convexity	52
3.3.1 Admissible classes spanned by linearly independent basis functions	52
3.3.2 Basis functions with non-overlapping supports	55
4 Solution via a gradient flow	62
4.1 The flow associated to a gradient descent algorithm	62
4.1.1 Exponential convergence to a unique equilibrium	62
4.1.2 The unique equilibrium and optimal projection	66
4.2 Variance reduction	69
5 Conclusion	75
6 Zusammenfassung	78
Bibliography	79

Chapter 1

Introduction

Random dynamical systems with continuous paths are ubiquitous in science and engineering. Such systems are often referred to as ‘diffusive systems’, or simply ‘diffusions’. Diffusions are particularly important whenever a large collection of interacting particles is involved, since one can often infer the statistical properties of the entire collection from the behaviour of individual particles. For this reason, diffusions are used to model many statistical-physical phenomena. Indeed, much of the current research in computational biophysics involves studying the motion of a large molecule immersed in a noisy environment consisting of many smaller molecules. Diffusions are attractive systems for computational study because remarkable increases in computing power and methods have made it possible to perform quantitative studies of complex systems with unprecedented accuracy. In this context, mathematics plays an important role in the analysis and development of algorithms that efficiently use both computing power and data.

This aim of this thesis is to analyse and develop some aspects of a Monte Carlo-based gradient descent algorithm for computing statistical properties of metastable diffusions. Roughly speaking, metastability refers to the tendency for the diffusion to spend long periods of time in certain areas. If the walk of a drunkard provides a reasonable analogy for the path of a diffusion, then the tendency for the drunkard to spend long times in certain locations - such as bars - provides a reasonable analogy for the property of metastability. Metastability is important in biophysics, not only because many biologically relevant molecules appear to be metastable, but because their metastability appears to be essential to their function. From an algorithmic point of view, metastability is important because it renders the standard Monte Carlo method for estimating statistical properties too time-consuming to be efficient.

The optimal control problem that we study in this thesis attempts to circumvent the problem of inefficiency of the standard Monte Carlo method by finding another statistical estimation problem in which the standard Monte Carlo method is more efficient. To each (admissible) control function, one associates a random variable and measure. One can therefore also associate to the admissible control function the mean of the alternative random variable with respect to the alternative measure. This procedure defines a continuous function from the set of admissible control functions to the real numbers. We shall refer to this function as the *control functional* of the optimal control problem. The desired unknown value that we wish to estimate is equal to the global minimum of the control functional. The gradient descent algorithm seeks the nearest local minimum to the point

on the graph of the control functional at which the algorithm is initialised, by following the gradient ‘downwards’ on the graph of the control functional.

A desirable feature of any algorithm is that the algorithm provides the correct answer to a given question. For the optimal control problem described earlier, one can derive a closed-form solution for the correct answer, i.e. the *unique* optimal control, and the corresponding value, i.e. the global minimum of the control functional. Since the gradient descent algorithm identifies only the ‘nearest’ attracting local minimum, it is possible that the algorithm yields an estimate that is systematically incorrect, in the sense that the difference between the estimate and the true value of the desired statistical property cannot be explained by random fluctuations alone. This outcome is possible if the graph of the control functional has multiple local minima with different corresponding values.

The main result of this thesis is the identification of conditions which suffice to guarantee, given a certain finite-dimensional approximating subset of the infinite-dimensional set of admissible control functions, that the restriction of the control functional to the approximating subset is strongly convex. Strong convexity is important because it guarantees that the restricted control functional has one local (and hence global) minimum. The result of strong convexity holds independently of both the dimension of the state space of the diffusion and the dimension of the approximating subset. An important immediate consequence of strong convexity is the proof that the gradient descent algorithm produces iterates that converge at an exponential rate to the global minimum. To the best of our knowledge, the uniqueness result and the characterisation of the convergence to the global minimum are new. Given the specific nonlinear dependence of the control functional on the admissible control function, the strong convexity result may be considered to be somewhat surprising.

The second main result of this thesis is the correspondence between the global minimum described earlier and the best approximation of the solution to a nonlinear elliptic boundary value problem. Since the existence of the unique global minimum holds independently of the dimension of the state space, this result suggests a method for solving elliptic boundary value problems that does not suffer from the curse of dimensionality. Moreover, if one can construct a gradient descent algorithm which has smaller lower computational cost than methods for solving partial differential equations, then the theoretical result thus described may potentially be used to construct efficient solvers for boundary value problems defined over high-dimensional domains.

The third result we obtain is to show that a certain random variable - namely, a martingale - that is associated in a natural way to the alternative random variable, can be used to reduce the variance of the alternative random variable, even when the corresponding control is not optimal. This result is relevant in the context of Monte Carlo methods, where smaller variances generally lead to better Monte Carlo algorithms in the sense that the algorithm produces estimates of the value of the control functional that converge more rapidly.

Thus, two of the main results of this thesis - the strong convexity of the restricted control functional, and the correspondence between the global minimum and best approximation - are analytical in nature. These results provide the theoretical justification for the gradient descent algorithm. The third result concerns the development of the variance-reduction component of the gradient descent algorithm.

Since the statistical estimation of properties of rare events is a common problem in

computational physics, it is not surprising that many methods for rare event estimation have been proposed, especially in the case of free energy computations. Well-known methods include constraint-based techniques, in which one applies so-called holonomic (i.e. coordinate-dependent but momentum-independent) constraints; one example of such a constraint method is the Blue Moon Ensemble method [13,14]. Other methods consist of driving a system by the application of a force, e.g. the adaptive biasing force method [16] or hyperdynamics [64,65], or by changing the energy landscape on which the system lives, e.g. metadynamics [11,42], and action-based methods such as the Passerone-Parrinello method [50], in which one applies a least action principle to obtain a variational problem. Other methods for simulating rare events involve studying transition paths, such as the transition path sampling method [19], or sampling only those paths which dominate the calculation of rare event quantities, such as the forward flux sampling method [1,2]. Some reviews which may help one navigate the zoo of methods are [18,20,27].

The approach described earlier, of formulating an importance sampling problem as a stochastic optimal control problem, appears in the work [22] of Dupuis, Sezer, and Wang, who sought to estimate probabilities of rare events taking place on telecommunications networks. Dupuis and Wang also observed in [24] that the formulation of importance sampling (equivalently, of variance reduction) problems in terms of stochastic optimal control could be viewed as two-player stochastic differential games. Fleming further studied the connection between stochastic control and differential games in [29] by introducing the feature of risk-sensitivity into the problem. The connection between stochastic control and dynamic games was already known to Dai Pra, Meneghini and Runggaldier, who also considered the duality relationship between free-energy like quantities and relative entropy in [15]. The idea of optimising over a set of measures - a central idea in this thesis - was applied to Markov decision processes in Borkar's work [8]. Some recent work on importance sampling and rare event simulation for multiscale diffusions includes the work of Dupuis, Spiliopoulos and Wang [23], and Vanden-Eijnden and Weare [61]. The method of cross-entropy minimisation introduced by Kullback [39] was also applied to importance sampling for rare event estimation by Rubenstein in the work [57], and recently by Asmussen, Kroese, and Rubenstein in [4]. These ideas have also been applied to rare event simulation in the context of molecular dynamics; see, e.g. the articles by Hartmann and Schütte [35], Hartmann et. al. [34], and Zhang et. al. [66].

We now outline the content of the thesis. In Chapter 2, we review the theory and the results from the literature which underpin the problem considered here. We will take most of the theory for granted, in the sense that we will mostly state results without proving them. We do however point the reader to the texts of Durrett [25], Karatzas and Shreve [37], Øksendal [47], and Revuz and Yor [55] for diffusions, and Gilbarg and Trudinger [33] for elliptic partial differential equations. Given that the dynamical systems of interest are diffusions, we deal first with the theory of stochastic processes that are continuous in both space and time in Section §2.1. The Cameron-Martin-Girsanov theorem on change of measure for diffusions in Euclidean space (Theorem 2.1.17) is essential to the reformulation of the importance sampling problem as a stochastic optimal control problem, because it provides a correspondence between a change of drift (i.e. a control function) with an alternative measure. In particular, we can view the stochastic optimal control problem as a problem of optimising over a set of measures, precisely because of the Cameron-Martin-Girsanov theorem. We shall use Itô's formula (Theorem 2.1.4), and

weaker versions (Lemma 2.1.6) of the Burkholder-Davis-Gundy martingale inequalities in the derivation of our main result concerning strong convexity of the restricted control functional. The Feynman-Kac theorem (Theorem 2.1.20) establishes an important connection between solutions of certain parabolic partial differential equations and functionals of diffusions, and is central to our third main result concerning the link between the global minimum of the restricted control functional and the best approximation of the solution to an elliptic boundary value problem. In Section §2.2, we define the stochastic optimal control problem and show the existence and uniqueness of the optimal control (Theorem 2.2.8), following the treatment given in [35].

Chapter 3 concerns the main result of strong convexity of the restricted control functional. In Section §3.1, we compute the first- and second-order variations of the control functional using the Cameron-Martin-Girsanov change of measure theorem. The first- and second-order variations will be useful in order to show that the restricted control functional, when viewed as a function over a finite-dimensional parameter space, is twice continuously differentiable, and therefore has a Hessian. In Section §3.2, we apply Itô's formula to obtain relations concerning continuous local martingale terms that appear in the first- and second-order variations. We state sufficient conditions for these local martingales to be pairwise independent. Independence is a powerful tool that we shall use to simplify the expressions for the first- and second-order variations. The simplified expressions will be central to our proof that the Hessian of a certain function is uniformly positive definite. In Section §3.3, we apply the results from Sections §3.1 and §3.2 in order to construct a function over a finite-dimensional parameter space. The function corresponds to the restriction of the control functional to a subset of the set of admissible controls, the subset consists of linear combinations of finitely many basis functions, and the parameter space corresponds to the space of expansion coefficients for the linear combinations. In the main result of this thesis, Theorem 3.3.9, we show how two important conditions - a non-overlap condition on the supports of the basis functions, and a uniform lower bound on the path functional corresponding to the statistical property of interest - may be used to show that the Hessian of the function is uniformly positive definite. This implies that the function is strongly convex.

In Chapter 4, we apply the result of strong convexity to study the gradient descent algorithm proposed by Hartmann and Schütte for solving the finite-dimensional approximation of the optimal control problem. In Section §4.1.1, we use the strong convexity result to show that the gradient descent iterates converge at an exponential rate to a unique global minimum. In Section §4.1.2 we show that, under the assumption that the basis functions and the value function of the optimal control problem belong to the same Hilbert space, the unique global minimum corresponds to the best approximation of the value function of the optimal control problem in the finite-dimensional subset of the set of admissible controls. In Section §4.2, we return to the central concern of Monte Carlo methods, namely variance reduction. We show that a martingale term that arises from the Cameron-Martin-Girsanov change of measure theorem is a suitable control variate and state a method for scaling the control variate at every point along a solution of the gradient descent algorithm. We show that this method yields the maximum variance reduction attainable for the martingale control variate.

In Chapter 5, we conclude with a review of the results and the methods by which we obtained them. We critique the methods employed and discuss ideas for future work.

Chapter 2

Optimal control of diffusions

In this section, we examine the optimal control problem and show how it arises from importance sampling for metastable diffusions, following the approach of [35]. We first recall the requisite theory of diffusions in Section §2.1. Section §2.1.1 covers most of the theory of continuous stochastic processes that we will need, with the fundamental results being the preservation of the continuous local martingale property under the Itô integral (Theorem 2.1.1), Itô's formula (Theorem 2.1.4), and the Cameron-Martin-Girsanov theorem for change of measure via change of drift (Theorem 2.1.17). The key objects are continuous local martingales, their quadratic variation processes, and the Doléans exponential martingale. In Section §2.1.2, we recall some fundamental results relating diffusions and elliptic boundary value problems, with the main result being the Feynman-Kac representation (Theorem 2.1.20) of solutions to linear elliptic boundary value problems as expected values of suitably defined path functionals. In Section §2.2, we use the theory presented in Section §2.1 in order to formulate the problem (given a path functional) of finding an optimal importance sampling measure in terms of an optimal control problem. We provide some examples of statistical properties of interest in Section §2.2.1 and then formulate the optimal control problem in Section §2.2.2. Theorem 2.2.8 proves the existence and uniqueness of the optimal control to the problem, while Theorem 2.2.6 provides sufficient conditions for the value function of the optimal control problem to be a classical solution to a nonlinear elliptic boundary value problem involving the Hamilton-Jacobi-Bellman equation.

Canonical filtered probability space: Let $d \in \mathbb{N}$ be fixed, and let the sample space be the set of continuous paths in \mathbb{R}^d , defined on the nonnegative half-line:

$$\Omega := \left\{ \omega : [0, \infty) \rightarrow \mathbb{R}^d \mid \omega \text{ is continuous} \right\} = C([0, \infty); \mathbb{R}^d).$$

Let P be the Wiener measure, so that $B = \omega$ denotes the standard d -dimensional standard Brownian motion with respect to P , and $(\mathcal{F}_t)_{t \geq 0}$ denote the Brownian filtration.

2.1 Requisite theory of continuous-time stochastic processes

The presentation of the theory of continuous-time stochastic processes in this section was adapted mainly from [25, Chapters 2,5].

For arbitrary $x \in \mathbb{R}^d$, let P^x denote the probability measure concentrated on the set of all paths satisfying $\omega_0 = x$,

$$P^x(A) = P(\omega \in A \mid \omega_0 = x), \quad \forall A \in \mathcal{F}.$$

Recall that a continuous martingale $M = (M_t)_t$ on the canonical filtered probability space is a L^1 , adapted process that satisfies the martingale property. That is, for all $t \geq 0$, $M_t \in L^1(P^x)$, is \mathcal{F}_t -measurable, and

$$E_{P^x}[M_t \mid \mathcal{F}_s] = M_s, \quad \forall s \in [0, t]. \quad (2.1)$$

The martingale property is preserved under affine transformations.

A stopping time with respect to a given filtration $(\mathcal{F}_t)_t$ is a random variable $T : \Omega \rightarrow [0, \infty)$ such that

$$\{T \leq t\} \in \mathcal{F}_t.$$

A continuous stochastic process M is said to be a continuous local martingale if there exists a sequence of stopping times $(T_n)_{n \in \mathbb{N}}$ such that $T_n \rightarrow \infty$ and such that for every $n \in \mathbb{N}$, $(M_t^{T_n})_{n \in \mathbb{N}}$ is a martingale with respect to the filtration $(\mathcal{F}_{t \wedge T_n})_{t \geq 0}$, where

$$M_t^{T_n} := \begin{cases} M_{T_n \wedge t} & \{0 < T_n\} \\ 0 & \{0 = T_n\} \end{cases} \quad (2.2)$$

and where we use the notation

$$s \wedge t := \min\{s, t\}.$$

We shall define the class \mathcal{M}^2 to be the set of all continuous local martingales M such that $M_t \in L^2(P^x)$ for all $t \geq 0$, in a sense made precise below in (2.4). For an arbitrary continuous local martingale $M \in \mathcal{M}^2$, one can define, by the Doob-Meyer decomposition, the quadratic variation $\langle M \rangle$ to be the P^x -almost surely unique, increasing process such that $\langle M \rangle_0 = 0$ and $(M_t^2 - \langle M \rangle_t)_{t \geq 0}$ defines a continuous local martingale. In particular, if $M_0 = 0$, then the quadratic variation process satisfies

$$E_{P^x}[M_t^2] = E_{P^x}[\langle M \rangle_t].$$

The quadratic variation process is a continuous-time stochastic process that is locally of bounded variation. The quadratic variation of any process that is locally of bounded variation is zero for all time. Finally, given two continuous local martingales M and N , one can define their covariance process by the polarisation identity

$$\langle M, N \rangle_t = \frac{1}{4} (\langle M + N \rangle_t - \langle M - N \rangle_t).$$

The covariance process is the P -almost surely unique, continuous process that is locally of bounded variation; it satisfies $\langle M, N \rangle_0 = 0$ and makes $(M_t N_t - \langle M, N \rangle_t)_t$ a continuous local martingale, but it is not necessarily increasing. An important property of the quadratic variation (and hence of the covariance process) is that, for any stopping time T and any two continuous local martingales M and N ,

$$\langle M^T, N^T \rangle = \langle M, N \rangle^T. \quad (2.3)$$

Given the filtration $(\mathcal{F}_t)_{t \geq 0}$, a stochastic process H is said to be previsible if H_t is measurable with respect to $\mathcal{F}_{t-} := \lim_{s \uparrow t} \mathcal{F}_s$. Previsible processes are important, because the Itô integral of a previsible process with respect to a continuous local martingale over $[0, t]$ is defined as the limit (taken over increasingly finer subdivisions of $[0, t]$) of the Lebesgue-Stieltjes integral in which the integrand is evaluated at the left endpoints of subintervals. For a continuous local martingale M , we define the norm

$$\|M\|_2 := \left(\sup_{t \geq 0} E_{P^x}[M_t^2] \right)^{1/2}$$

and

$$\mathcal{M}^2 := \{ \text{martingales adapted to } (\mathcal{F}_t)_{t \geq 0} \mid \|M\|_2 < \infty \}. \quad (2.4)$$

The space \mathcal{M}^2 is complete and isomorphic to the Hilbert space $L^2(\Omega, \mathcal{F}_\infty, P^x)$. We define

$$\Pi_3(M) := \left\{ H \text{ previsible} \mid \int_0^t H_s^2 d\langle M \rangle_s < \infty \text{ } P^x - \text{a.s. } \forall t \geq 0 \right\}, \quad (2.5)$$

where the integral $\int_0^t H_s^2 d\langle M \rangle_s$ should be understood as a Lebesgue-Stieltjes integral. The class $\Pi_3(M)$ denotes the largest set of integrand processes H such that the Itô integral of H with respect to the continuous local martingale M , denoted by

$$(H \cdot X)_\cdot = \int_0^\cdot H_s dX_s,$$

exists [25, Section 2.6]. In particular, we have the following result [25, Chapter 2, Theorem 6.3]:

Theorem 2.1.1. *If X is a continuous local martingale and $H \in \Pi_3(X)$ (not necessarily continuous), then the process defined by the Itô integral*

$$\int_0^t H_s dX_s$$

is a continuous local martingale.

A continuous semimartingale is any process that admits a decomposition into the sum of a continuous local martingale and a continuous stochastic process that is locally of bounded variation. In particular, if $X = M + A$ and $Y = M' + A'$ are two semimartingales, where M and M' are continuous local martingales and A and A' are locally of bounded variation, then

$$\langle X, Y \rangle = \langle M, M' \rangle,$$

since the covariance process is uniformly zero whenever one of the processes is locally of bounded variation, i.e. since

$$\langle X, A \rangle \equiv 0$$

almost surely for any continuous processes X and A , if A is locally of bounded variation. The Kunita-Watanabe inequality, which we shall use later, bounds the absolute value of the covariance of two continuous local martingales in terms of their quadratic variations,

$$|\langle X, Y \rangle_\infty| \leq (\langle X \rangle_\infty \langle Y \rangle_\infty)^{1/2}. \quad (2.6)$$

We shall use the Kunita-Watanabe inequality later in the variational analysis of a certain control functional.

We define the set of locally bounded, previsible processes

$$\ell b\Pi := \{ \exists (T_n)_{n \in \mathbb{N}} \mid T_n \rightarrow \infty, s \leq T_n \Rightarrow |H(s, \omega)| \leq n \}. \quad (2.7)$$

Analogously to Theorem 2.1.1, we have the following result [25, Chapter 2, Theorem 8.3]

Theorem 2.1.2. *If X is a continuous semimartingale and $H \in \ell b\Pi$ (not necessarily continuous), then the Itô integral $(H \cdot X)$ is again a continuous semimartingale.*

An important property of continuous semimartingales that we shall use is expressed in the following formula for the covariance of Itô integrals [25, Chapter 2, Theorem 8.7]:

Proposition 2.1.3. *Let $(X^{(i)})_{1 \leq i \leq n}$ and $(Y^{(i)})_{1 \leq i \leq n}$ be finite collections of continuous semimartingales, $(H^{(i)})_{1 \leq i \leq n}$ and $(K^{(i)})_{1 \leq i \leq n}$ be subsets of $\ell b\Pi$, and let $X := \sum_i H^{(i)} \cdot X^{(i)}$ and $Y := \sum_i K^{(i)} \cdot Y^{(i)}$ denote finite sums of Itô integrals. Then the covariance of X and Y satisfies*

$$\langle X, Y \rangle_t = \sum_{i,j} \int_0^t H_s^{(i)} K_s^{(j)} d\langle X^{(i)}, Y^{(j)} \rangle_s. \quad (2.8)$$

Given two continuous semimartingales X and Y , we have the integration by parts formula

$$X_t Y_t = X_0 Y_0 + \int_0^t X_s dY_s + \int_0^t Y_s dX_s + \langle X, Y \rangle_t. \quad (2.9)$$

The integration by parts formula is a special case of Itô's formula, which we state below (see [25, Section 2.10, Theorem 10.2]):

Theorem 2.1.4 (Itô's formula). *Let $(X^{(i)})_{1 \leq i \leq d}$ be a finite collection of \mathbb{R} -valued semimartingales, and $f \in C^2(\mathbb{R}^d; \mathbb{R})$. Then*

$$f(X_t) - f(X_0) = \sum_{i=1}^d \int_0^t \frac{\partial}{\partial x_i} f(X_s) dX_s^{(i)} + \frac{1}{2} \sum_{i,j=1}^d \int_0^t \frac{\partial^2}{\partial x_i \partial x_j} f(X_s) d\langle X^{(i)}, X^{(j)} \rangle_s. \quad (2.10)$$

In some cases it will be useful to bound expected values of powers of martingales by expected values of their quadratic variation. Recall the following inequalities [25, Section 3.5]:

Theorem 2.1.5 (Burkholder-Davis-Gundy inequalities). *Let X be a continuous local martingale with $X_0 = 0$. For any $p \in (0, \infty)$, and for any stopping time τ , there exist constants $c, C \in (0, \infty)$ so that*

$$cE[\langle X \rangle_\tau^{p/2}] \leq E\left[\left(\sup_{s \leq \tau} |X_s|\right)^p\right] \leq CE[\langle X \rangle_\tau^{p/2}]. \quad (2.11)$$

In the next chapter, we will need to use a weaker version of the Burkholder-Davis-Gundy inequalities [17, Equations (1.1) and (1.2)]:

Lemma 2.1.6. For $p \in (0, \infty)$, there exists a constant A_p such that

$$E[|M_\tau|^p] \leq A_p E[\langle M \rangle_\tau^{p/2}], \quad (2.12)$$

and for $p \in (1, \infty)$, there exists a constant a_p such that

$$a_p E[\langle M \rangle_\tau^{p/2}] \leq E[|M_\tau|^p]. \quad (2.13)$$

In particular, we shall need to use a result [46] due to Novikov for the case that $p = 4$, in which the best possible values for a_4 and A_4 are the smallest and largest positive roots of the Hermite polynomial H_4 . Given that

$$H_4(x) = 16x^4 - 48x^2 + 12,$$

the two positive roots of H_4 are given by

$$a_4 := \rho^{(1)} = \sqrt{\frac{3}{2} - \sqrt{\frac{3}{2}}} \approx 0.5246 \quad (2.14)$$

$$A_4 := \rho^{(2)} = \sqrt{\frac{3}{2} + \sqrt{\frac{3}{2}}} \approx 1.6507. \quad (2.15)$$

Remark 1. Novikov's remarkable result on the identification of sharp constants for the martingale inequalities (2.13) and (2.12) inspired much research, especially concerning the best values of the constants for other values of p . The recurring theme of the research in this direction is that the extremal roots of corresponding Hermite polynomials provide the sharp constants, see, e.g. [17, 53].

2.1.1 Girsanov's formula and change of measure

An important result that we will use later is the change of measure theorem for diffusions, which builds upon Girsanov's formula. Recall that two measures Q and P on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0})$ are said to be *locally equivalent* if, for each $t \geq 0$, the probability measures $Q|_{\mathcal{F}_t}$ and $P|_{\mathcal{F}_t}$ are mutually absolutely continuous, i.e.

$$\forall A \in \mathcal{F}_t, \quad Q(A) = 0 \Leftrightarrow P(A) = 0.$$

If $Q|_{\mathcal{F}_t}$ and $P|_{\mathcal{F}_t}$ are mutually absolutely continuous, we may define the Radon-Nikodym derivative

$$\alpha_t = \frac{dQ|_{\mathcal{F}_t}}{dP|_{\mathcal{F}_t}},$$

and vice versa. Note that $\alpha_t > 0$, by mutual absolute continuity. Given two locally equivalent measures on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0})$, we obtain a stochastic process $(\alpha_t)_{t \geq 0}$. The process α is important because a process X is a (local) martingale with respect to Q if and only if αX is a (local) martingale with respect to P [25, Section 2.12, Lemma 12.1]. In particular, since constants are martingales, it follows that the process α is itself a martingale with respect to P [25, Section 2.12, Corollary 12.2].

Lemma 2.1.7. *Given a strictly positive process α that is a martingale with respect to P , there is a unique, locally equivalent probability measure Q so that*

$$\frac{dQ}{dP} = \alpha.$$

The next result is important because it provides a way to construct a process A from the process α and a P -local martingale X , such that A contains that part of X which is not a local martingale with respect to Q . The result is useful because, in many cases (for example, the case of optimal control), one wishes to change a stochastic process by adding another process, instead of by multiplying with another process. In particular, one wishes to know what to add to a process that is a continuous local martingale with respect to some measure P such that the modified process is a continuous local martingale with respect to another measure Q .

Theorem 2.1.8 (Girsanov's formula). *Suppose X is a continuous local martingale with respect to P , and define a process A by*

$$A_t = \int_0^t \alpha_s^{-1} d\langle \alpha, X \rangle_s.$$

Then $X - A$ is a local martingale with respect to Q .

Theorem 2.1.8 is important because it implies that, if an arbitrary process is a semimartingale with respect to P , then it is a semimartingale with respect to Q . This implication follows from the observation that A is a process that is locally of bounded variation. We will need the following result [25, Section 2.12, Theorem 12.6] in the proof of the change of measure theorem:

Theorem 2.1.9. *The quadratic variation is given by the same random variable under locally equivalent changes of measure.*

Now recall that B is a d -dimensional Brownian motion on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, P)$. Given functions $b : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, a solution to the stochastic differential equation

$$dX_t = b(X_t)dt + \sigma(X_t)dB_t \tag{2.16}$$

is a triple $((X_t)_t, (B_t)_t, (\mathcal{F}_t)_t)$, where B is a Brownian motion adapted to the filtration $(\mathcal{F}_t)_t$, and the random variable $X : \Omega \rightarrow \Omega$ satisfies the integral equation

$$X_t = X_0 + \int_0^t b(X_s)ds + \int_0^t \sigma(X_s)dB_s. \tag{2.17}$$

The function b is sometimes referred to as the ‘infinitesimal drift’ (or more commonly the ‘drift’) and the function σ as the ‘infinitesimal standard deviation’ term, while the function $a : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, defined by $a := \sigma\sigma^\top$, is often called the ‘diffusion matrix’ or ‘infinitesimal covariance’. This is because, at $t = 0$, the rate of change of the expected value $E[X_t^i]$ of the i -th coordinate process equals $b_i(X_0)$, and similarly the rate of change of $\text{cov}[X_t^i, X_t^j]$ equals $a_{ij}(x)$.

Given the data (b, σ) of the stochastic differential equation (2.16), the corresponding martingale problem $MP(b, a)$, where $a = \sigma\sigma^\top$, is defined as follows: find $X : \Omega \rightarrow \Omega$ such that, for $1 \leq i, j \leq n$,

$$\begin{aligned} M_t^i &:= X_t^i - \int_0^t b_i(X_s) ds \\ M_t^{ij} &:= M_t^i M_t^j - \int_0^t a_{ij}(X_s) ds \end{aligned}$$

are continuous local martingales (see [55, Chapter VII, Definition (2.3) and Proposition (2.4)]). The random variable X is said to solve $MP(b, a)$. Now suppose that $(X, B, (\mathcal{F}_t)_t)$ solves the stochastic differential equation (2.16). Note that, by (2.17), we have

$$M_t^i = X_0 + \int_0^t \sum_{\ell} \sigma_{i\ell}(X_s) dB_s^\ell.$$

If $X_0 \in \mathbb{R}^d$ is constant, then the random component of M consists entirely of an Itô integral term, which by Theorem 2.1.1 is a continuous local martingale, as required. We also have from Proposition 2.1.3 that

$$\begin{aligned} \langle M^i, M^j \rangle_t &= \sum_{k, \ell} \int_0^t \sigma_{i\ell}(X_s) \sigma_{kj}(X_s) d\langle B^\ell, B^k \rangle_s \\ &= \sum_{k, \ell} \int_0^t \sigma_{i\ell}(X_s) \sigma_{kj}(X_s) \delta_{k\ell} ds \\ &= \int_0^t a_{ij}(X_s) ds. \end{aligned}$$

Since the covariance process $\langle M^i, M^j \rangle$ makes $M^i M^j - \langle M^i, M^j \rangle$ a continuous local martingale, it follows that M^{ij} is a continuous local martingale, as desired.

Since $X : \Omega \rightarrow \Omega$, it follows that the law of X with respect to P , i.e. $P \circ X^{-1}$, is again a probability distribution on (Ω, \mathcal{F}) . Therefore, one may consider a solution to the martingale problem $MP(b, a)$ to be either a random variable $X : \Omega \rightarrow \Omega$, or a distribution on the canonical filtered probability space. This is important, because it implies the following connection between solutions of (2.16) and solutions of the martingale problem (see [25, Section 5.4, Theorem 4.5]):

Theorem 2.1.10. *The martingale problem $MP(b, a)$ has a unique solution if and only if the stochastic differential equation defined by the data b and σ has a solution that is unique in distribution.*

Another link between solutions of the stochastic differential equation (2.16) and solutions of the martingale problem may given by the following theorem that we obtain from [25, Theorem 4.5]:

Theorem 2.1.11. *If X is a solution of $MP(b, a)$ and σ is a measurable square root of a , then there exists a Brownian motion B , so that (X, B, \mathcal{F}^X) solves (2.16).*

In particular, for the special case that a is uniformly elliptic in the sense that

$$z \cdot a(x)z \geq \kappa|z|^2$$

for some constant $\kappa > 0$ not depending on x , then there exists a measurable square root of a , and one obtains a bijection between the solutions of the stochastic differential equation (2.16) and the martingale problem $MP(b, a)$.

Formulating questions about solutions of stochastic differential equations in terms of martingale problems yields some advantages. One advantage is that we can use the theory of continuous local martingales to study solutions of stochastic differential equations. In particular, we will use Girsanov's formula (Theorem 2.1.8), which is stated in terms of local martingales, in order to prove the change of measure theorem (Theorem 2.1.17). Another advantage is that the martingale problem provides an easy way to check whether the solution to a stochastic differential equation has the Markov property (see [25, Section 5.4, Theorem 4.6]):

Theorem 2.1.12. *If a and b are locally bounded, and $MP(b, a)$ has a unique solution, then the strong Markov property holds for X .*

We shall use the strong Markov property in Section §3.2.2 to show that a certain collection of continuous local martingales is pairwise independent. The following definition is drawn from [25, Section 1.3, Definition 3.7]:

Definition 2.1.1. Let $Y : [0, \infty) \times \Omega \rightarrow \mathbb{R}$ be bounded, and for any stopping time S , let $\theta_S : \Omega \rightarrow \Omega$ be the shift operator defined by $\theta_S((\omega_t)_{t \geq 0}) = (\omega_{S+t})_{t \geq 0}$. Then the random variable X satisfies the strong Markov property if

$$E_{P^x}[Y(S, \theta_S(X)) | \mathcal{F}_S^X] = E_{P^{X_S}}[Y(S, X)] \quad \text{on } \{S < \infty\}. \quad (2.18)$$

In other words, given the history of the process X up to and including the stopping time S , the expected value of the functional Y on the path of X from S onwards depends only on the value of the stopping time S and the state X_S of the process. The rough interpretation of (2.18) is that, at any random time, the only relevant information for predictions of future states of strongly Markovian processes is contained entirely in the current state of the process. Note that both sides of the equation (2.18) are random variables.

We now proceed towards the denouement of this section on the stochastic analysis of continuous diffusion processes. Suppose that, for a measure P and some given data (b, σ) of the stochastic differential equation (2.16), one perturbs the drift term by some suitable function β . The next main result of this section (Theorem 2.1.17), which is known as the Cameron-Martin formula [55, Chapter IX, §1, Theorem 1.10] or the change of measure theorem [25, Section 5.5], specifies conditions on the perturbing function β such that the random variable X is a solution of the martingale problem $MP(b, a)$ under P and a solution of the martingale problem $MP(b + \beta, a)$ under Q . The idea of the change of measure theorem is to use the definition of the solution to a martingale problem, and to use Girsanov's formula for finding locally equivalent changes of measure that preserve the (local) martingale property. By the hypotheses,

$$M_t^i := X_t^i - \int_0^t b_i(X_s) ds$$

is a continuous local martingale with respect to P . We shall first show that

$$N_t^i := X_t^i - \int_0^t b_i(X_s) ds - \int_0^t \beta_i(X_s) ds$$

is a continuous local martingale with respect to Q , and then proceed to show that

$$N_t^{ij} := N_t^i N_t^j - \int_0^t a_{ij}(X_s) ds$$

is a continuous local martingale with respect to Q . Before we do this, we first consider the following results [25, Section 2.2, Theorem 2.5 and Corollary 2.6]:

Theorem 2.1.13. *If X is a continuous local submartingale, and*

$$E \left[\sup_{s \leq t} |X_s| \right] < \infty \quad (2.19)$$

for each $t \geq 0$, then X is a submartingale.

Corollary 2.1.14. *If X is a continuous local martingale and (2.19) holds, then X is a martingale.*

We use Corollary 2.1.14 in the next Theorem [25, Section 3.3, Theorem 3.7]:

Theorem 2.1.15. *Let X be a continuous local martingale such that $X_0 = 0$ and*

$$\langle X \rangle_t \leq Mt \quad (2.20)$$

for all $t \geq 0$. Then the process $\mathcal{E}(X)$ defined by

$$\mathcal{E}(X)_t = \exp \left(X_t - \frac{1}{2} \langle X \rangle_t \right) \quad (2.21)$$

is a continuous martingale that satisfies $E^x [\mathcal{E}(X)_t] = 1$ for all $t \geq 0$. Furthermore, $\mathcal{E}(X)_t \in L^2(P^x)$ for all $t \geq 0$.

We reproduce the proof given in [25].

Proof. By Itô's formula, $\mathcal{E}(X)$ is a continuous local martingale, since if we define

$$f(X_t, \langle X \rangle_t) := \exp \left(X_t - \frac{1}{2} \langle X \rangle_t \right)$$

then applying Itô's formula yields

$$f(X_t, \langle X \rangle_t) = \int_0^t f(X_s, \langle X \rangle_s) dX_s - \frac{1}{2} \int_0^t f(X_s, \langle X \rangle_s) d\langle X \rangle_s + \frac{1}{2} \int_0^t f(X_s, \langle X \rangle_s) d\langle X \rangle_s$$

and the local martingale property is preserved under the Itô integral, by Theorem 2.1.1. By the definition of $\mathcal{E}(X)$ for any continuous local martingale X , we have

$$(\mathcal{E}(X)_t)^2 = \exp(2X_t - \langle X \rangle_t) = \mathcal{E}(2X)_t \exp(\langle X \rangle_t). \quad (2.22)$$

Let $(T_n)_n$ be a sequence of stopping times that reduces $\mathcal{E}(2X)$. By Doob's maximal L^2 inequality to $\mathcal{E}(X)_{t \wedge T_n}$, (2.22), the sublinear growth condition (2.20) on the quadratic variation of X , and the martingale property, we obtain

$$\begin{aligned} E^x \left[\sup_{s \leq t} \mathcal{E}(X)_{s \wedge T_n}^2 \right] &\leq 4E \left[\mathcal{E}(X)_{t \wedge T_n}^2 \right] \\ &= 4E^x \left[\mathcal{E}(2X)_{t \wedge T_n} \exp(\langle X \rangle_{t \wedge T_n}) \right] \\ &\leq 4E^x \left[\mathcal{E}(2X)_{t \wedge T_n} \right] \exp(Mt) \\ &= 4 \exp(Mt). \end{aligned} \tag{2.23}$$

By Jensen's inequality, it therefore holds that

$$\left(E \left[\sup_{s \leq t} \mathcal{E}(X)_{s \wedge T_n} \right] \right)^2 \leq E \left[\sup_{s \leq t} \mathcal{E}(X)_{s \wedge T_n}^2 \right].$$

Since $T_n \rightarrow \infty$ as $n \rightarrow \infty$, applying the monotone convergence theorem yields

$$\left(E \left[\sup_{s \leq t} |\mathcal{E}(X)_s| \right] \right)^2 \leq 4 \exp(Mt).$$

The inequalities above imply that $\mathcal{E}(X)$ is a martingale, by Corollary 2.1.14. \square

Corollary 2.1.16. *If the stopping time τ is P^x -almost surely bounded, then $\mathcal{E}(X)_\tau \in L^2(P^x)$.*

The significance of Corollary 2.1.16 will become apparent later in Section §3.1.1, when we need the square integrability of an exponential martingale at a stopping time in order to prove the existence of the first variation of a certain control functional. We note that the assumption of τ being P^x -almost surely bounded may be restrictive in some contexts. For example, in the context of the exit of a diffusion from a bounded Lipschitz domain, sufficient conditions are known for which the first exit time of the diffusion from the domain is almost surely finite (see Lemma 2.1.19 in Section §2.1.2), but conditions for which the first exit time are bounded do not appear to be as well known. Therefore it would be useful to specify conditions for which the conclusion of Corollary 2.1.16 holds under less restrictive conditions on the stopping time. By examining the proof of Theorem 2.1.15, one way to relax the boundedness condition on the stopping time τ would be to impose a stronger growth condition on the quadratic variation, e.g. by replacing (2.20) with

$$\langle M \rangle_t \leq \gamma(t) \tag{2.24}$$

where $\gamma : [0, \infty) \rightarrow [0, \infty)$ is any increasing function that is asymptotically bounded in the sense that

$$\lim_{t \rightarrow \infty} \gamma(t) \leq C < +\infty$$

P^x -almost surely, for some constant C . Then in (2.23), the deterministic exponential term $\exp(Mt)$ can be replaced by $\exp(\gamma(t))$, and by taking the limit, one establishes that the exponential martingale $\mathcal{E}(X)$ is in fact L^2 -bounded. Unfortunately, it is not clear how one could guarantee that a bound of the form (2.24) holds.

We now use Theorem 2.1.15 to prove the next

Theorem 2.1.17 (Change of measure). *Let X be a solution of the martingale problem $MP(\beta, a)$, defined on the canonical filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, P)$ of continuous paths in \mathbb{R}^d . Suppose $a(x)$ is invertible for all $x \in \mathbb{R}^d$, that $b : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is measurable, and that*

$$|b(s, x) \cdot a^{-1}(x)b(s, x)| \leq M \quad (2.25)$$

for some constant M . Then there exists a probability measure Q on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t)$ that is locally equivalent to P , such that X solves $MP(b + \beta, a)$ with respect to Q .

Remark 2. In the proof of Theorem 2.1.15, we used condition (2.20) to show that the continuous local martingale $\mathcal{E}(X)$ was in fact a martingale. We shall use condition (2.25) to show a related exponential continuous local martingale is in fact a martingale, and obtain a change of measure by Girsanov's theorem. Requiring that the change of drift be bounded with respect to some norm is a strong sufficient condition for the existence of a change of measure. Another well-known sufficient condition for the change of drift to give rise to a change of measure is Novikov's condition

$$E^x \left[\exp \left(\frac{1}{2} \int_0^\infty |b(X_s)|^2 ds \right) \right] < \infty,$$

which is a stronger requirement than Kazamaki's condition

$$\sup \left\{ E^x \left[\exp \left(\frac{1}{2} \int_0^\tau b(X_s) dB_s \right) \right] \mid E^x \left[\exp \left(\frac{1}{2} \int_0^\tau b(X_s) dB_s \right) \right] < \infty, \tau \text{ bounded} \right\} < \infty$$

in the sense that Novikov's condition implies Kazamaki's condition. Note that Kazamaki's condition on the change of drift can also be expressed as the requirement that $\exp \left(\frac{1}{2} \int_0^\tau b(X_s) dB_s \right)$ be a uniformly integrable submartingale, see [55, Chapter VIII, §1, Proposition 1.14]. Both Kazamaki's and Novikov's conditions imply that the exponential martingale is in fact a uniformly integrable martingale, and hence are stronger than the boundedness condition (2.25) we specified above. In the context of numerical methods, condition (2.25) has the advantage of being easier to verify than Novikov's condition. From a conceptual point of view, it is also natural to approximate functions by bounded functions.

We now prove Theorem 2.1.17. The key idea of the proof is to establish that the associated exponential local martingale $\mathcal{E}(Y)$ is in fact a true martingale, which implies that $E^x[\mathcal{E}(Y(X))_t] = 1$ for all $t \geq 0$. Once this has been established, the existence of a locally equivalent measure is guaranteed by Lemma 2.1.7.

Proof. We follow the proof given in [25, Section 5.5, Theorem 5.1]. Define

$$\bar{X}_t := X_t - \int_0^t \beta(X_s) ds \quad (2.26)$$

$$Y_t := \sum_{i,j} \int_0^t (a^{-1})_{ij}(X_s) b_j(s, X_s) d\bar{X}_s^i. \quad (2.27)$$

Note that \bar{X} is a d -dimensional continuous local martingale, and Y is a 1-dimensional continuous local martingale. By (2.8), and by the condition (2.25), it holds that

$$\langle Y \rangle_t = \int_0^t b(s, X_s) \cdot a^{-1}(X_s) b(s, X_s) ds \leq Mt$$

so that the quadratic variation of Y satisfies the growth condition (2.20). By Theorem 2.1.15, $\mathcal{E}(Y)$ is a strictly positive martingale. By Lemma 2.1.7, $\mathcal{E}(Y)$ defines a locally equivalent measure Q . By Itô's formula and the associative law,

$$\mathcal{E}(Y)_t - 1 = \int_0^t \mathcal{E}(Y)_s dY_s = \sum_{i,j} \int_0^t \mathcal{E}(Y)_s (a^{-1})_{ij}(X_s) b_j(s, X_s) d\bar{X}_s^i.$$

By the formula (2.8) for the covariance of Itô integrals with respect to semimartingales, we have

$$\begin{aligned} \langle \mathcal{E}(Y), \bar{X}^k \rangle_t &= \sum_{i,j} \int_0^t \mathcal{E}(Y)_s (a^{-1})_{ij}(X_s) b_j(s, X_s) d\langle \bar{X}^i, \bar{X}^k \rangle_s \\ &= \int_0^t \mathcal{E}(Y)_s b_k(s, X_s) ds, \end{aligned}$$

where we have used that

$$\langle \bar{X}^i, \bar{X}^k \rangle_t = \int_0^t a_{ik}(X_s) ds$$

by definition of X being a solution to $MP(\beta, a)$. We define a d -dimensional process A by

$$A_t^i = \int_0^t \mathcal{E}(Y)_s^{-1} d\langle \mathcal{E}(Y), \bar{X}^i \rangle_s = \int_0^t b_i(s, X_s) ds.$$

By Girsanov's formula (Theorem 2.1.8), it holds that

$$\bar{X}_t^i - \int_0^t b_i(s, X_s) ds = X_t^i - \int_0^t \beta_i(X_s) + b_i(s, X_s) ds$$

is a continuous local martingale with respect to Q . This satisfies the first requirement that a solution of $MP(\beta + b, a)$ must fulfill. The second and last requirement is satisfied, since

$$\left\langle X^i - \int_0^\cdot \beta_i(X_s) + b_i(s, X_s) ds, X^j - \int_0^\cdot \beta_j(X_s) + b_j(s, X_s) ds \right\rangle_t = \int_0^t a_{ij}(X_s) ds,$$

and since Theorem 2.1.9 implies that the covariance process is given by the same random variable under locally equivalent changes of measure. \square

One consequence of Theorem 2.1.17 that we shall exploit hereafter is the following property. Let X denote the solution of the stochastic differential equation

$$dX_t = \beta(X_t)dt + \sigma(X_t)dB_t, \quad (2.28)$$

and let X^b denote the solution of

$$dX_t^b = [b(X_t) + \beta(X_t)]dt + \sigma(X_t)dB_t, \quad (2.29)$$

where B is a P -Brownian motion. Let $a = \sigma\sigma^\top$. By Theorem 2.1.11, it follows that $P \circ X^{-1}$ solves $MP(\beta, a)$ and $P \circ (X^b)^{-1}$ solves $MP(b + \beta, a)$. Theorem 2.1.17 states that $Q \circ X^{-1}$ also solves $MP(b + \beta, a)$. If we assume that pathwise uniqueness holds for X and X^b , then by Theorem 2.1.10, the solution of $MP(b + \beta, a)$ is unique, and therefore

$$P \circ (X^b)^{-1} = Q \circ X^{-1}. \quad (2.30)$$

In particular, since

$$\frac{dQ \circ X^{-1}}{dP \circ X^{-1}} = \mathcal{E}(Y), \quad (2.31)$$

it holds that for any random variable $Z \in L^1(P \circ X^{-1})$,

$$E_{P \circ X^{-1}}[Z] = E_{P \circ (X^b)^{-1}} \left[Z \frac{dP \circ X^{-1}}{dP \circ (X^b)^{-1}} \right] = E_{Q \circ X^{-1}} \left[Z \frac{dP \circ X^{-1}}{dQ \circ X^{-1}} \right],$$

which we rewrite as

$$E_P[Z(X)] = E_P \left[Z(X^b) \mathcal{E}(Y(X^b)) \right]. \quad (2.32)$$

The relevance of (2.32) to Monte Carlo algorithms may be described as follows. In Monte Carlo algorithms, it is often desirable to be able to sample from alternative probability distributions which have the same mean but different variance; indeed, this is essentially the principal idea behind importance sampling. In the case of random variables expressed in terms of diffusions, (2.32) means that one can estimate the mean of the probability distribution $P \circ (Z(X))^{-1}$ by searching in the collection of probability distributions $P \circ (Z(X^b) \mathcal{E}(Y(X^b)))^{-1}$, since all probability distributions in the collection have the same mean. The goal is then to find the probability distribution in the collection that has the smallest variance.

2.1.2 The Dirichlet problem

In this section, we recall the connection between stochastic differential equations and solutions to elliptic Dirichlet boundary value problems. The material in this section concerning the connection between diffusions and elliptic boundary value problems is drawn from [37, 47]. Perhaps one of the best-known references on the theory of elliptic partial differential equations is the treatise [33] by Gilbarg and Trudinger, from which we also draw some material.

Let $D \subset \mathbb{R}^d$ be a bounded domain (i.e. a bounded, open, connected set) with sufficiently smooth boundary, and denote by $\omega^x \in \Omega$ any continuous function $\omega : [0, \infty) \rightarrow \mathbb{R}^d$ that satisfies $\omega_0 := \omega(0) = x$. Then we denote the first exit time for ω^x from D by

$$\tau_D(\omega^x) := \inf \{ t > 0 \mid \omega_t^x \notin D \}. \quad (2.33)$$

For the stochastic differential equation

$$dX_t = b(X_t)dt + \sigma(X_t)dB_t, \quad (2.34)$$

where $b \in C(\mathbb{R}^d; \mathbb{R}^d)$ and $\sigma_{ij} \in C(\mathbb{R}^d \rightarrow \mathbb{R})$ satisfy the linear growth condition

$$|b(x)|^2 + |\sigma(x)|^2 \leq K^2(1 + |x|^2), \quad (2.35)$$

and B is a d -dimensional Brownian motion, the infinitesimal generator is a second-order partial differential operator defined by

$$L\psi = \frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} \psi + \sum_{i=1}^d b_i \frac{\partial}{\partial x_i} \psi \quad (2.36)$$

on a subset of $C^2(\mathbb{R}^d; \mathbb{R})$, where $a = \sigma\sigma^\top$ is the diffusion matrix. The operator L is said to be *elliptic (uniformly elliptic)* in D if $a(x)$ is (uniformly) positive definite for every $x \in D$. The Dirichlet problem is to find, given functions $q \in C(\bar{D}; \mathbb{R})$, $f \in C(\bar{D}; [0, \infty))$ and $k \in C(\partial D; \mathbb{R})$, a solution $\psi \in C^2(D) \cap C(\bar{D})$ such that

$$L\psi(x) - f(x)\psi(x) = -q(x) \quad x \in D \quad (2.37)$$

$$\psi(y) = k(y) \quad y \in \partial D. \quad (2.38)$$

The next result (see [37, Section 5.7, Proposition 7.2]) establishes a key probabilistic representation of the Dirichlet problem:

Proposition 2.1.18. *Suppose that (2.35) holds, and that the solution to the martingale problem corresponding to (2.34) exists and is unique. If ψ solves (2.37)–(2.38) in the bounded domain D , and if*

$$E^x [\tau_D(X)] < \infty \quad \forall x \in D, \quad (2.39)$$

then it holds that the solution admits a representation of the form

$$\psi(x) = E^x \left[k(X_{\tau_D(X)}) \exp \left(- \int_0^{\tau_D(X)} f(X_s) ds \right) + \int_0^{\tau_D(X)} q(X_t) \exp \left(- \int_0^t f(X_s) ds \right) dt \right] \quad (2.40)$$

for every $x \in \bar{D}$.

The representation (2.40) is sometimes known as the ‘Feynman-Kac representation’. The next result (see [37, Section 5.7, Lemma 7.4]) provides a sufficient condition for (2.39) to hold, i.e. for the first exit time of X from D to have finite first moment:

Lemma 2.1.19. *Suppose that for some $i \in \{1, \dots, d\}$, it holds that*

$$\min_{x \in \bar{D}} a_{ii}(x) > 0. \quad (2.41)$$

Then (2.39) holds.

The next result (see, e.g. [37, Section 5.7, Remark 7.5]) of this short section will be relevant to the situations in which we are interested.

Theorem 2.1.20. *Assume the following:*

- (i) *the infinitesimal generator L is uniformly elliptic in D ,*
- (ii) *the functions $(a_{ij})_{i,j}$, $(b_i)_i$, f , q , and k are Hölder continuous on their respective domains, and*
- (iii) *the exterior sphere property holds for every point $y \in \partial D$, i.e. there exists a ball B such that*

$$\bar{B} \cap D = \emptyset \text{ and } \bar{B} \cap \partial D = \{y\}. \quad (2.42)$$

Then there exists a function $\psi \in C^{2,\alpha}(D) \cap C(\bar{D})$ that uniquely solves (2.37)–(2.38). Furthermore, ψ admits the representation (2.40).

The exterior sphere property (2.42) is interpreted as requiring that, at every point y on the boundary ∂D , one can place a closed ball \bar{B} that touches the closure \bar{D} of the domain precisely at y . The exterior sphere property ensures that the domain does not have a cusp singularity anywhere on its boundary. The theorem above is based on Itô's formula, and the application of a result concerning linear elliptic boundary value problems (see [33, Chapter 6, Theorem 6.13]). Recall that the class $C^{2,\alpha}(D)$ refers to the set of functions whose second partial derivatives are uniformly Hölder continuous of exponent $\alpha \in (0, 1)$ in D , and that $C^{2,\alpha}(D)$ is a subset of the set $C^2(D)$.

The Feynman-Kac representation is a result of fundamental importance in stochastic analysis, because it provides another connection between diffusions and partial differential equations, in addition to the forward and backward Kolmogorov equations. For the sake of comparison, we provide another statement, adapted from [47, Chapter 9, Exercise 9.12], that specifies slightly different conditions under which the Feynman-Kac representation (2.40) holds:

Theorem 2.1.21. *Let X be a solution of the stochastic differential equation (2.34), with the infinitesimal generator (2.36), and let D be a domain (not necessarily bounded). Assume the following:*

- (i) $k \in C(\partial D; \mathbb{R})$ is bounded,
- (ii) $E^x [\int_0^{\tau_D} |q(X_t)| dt] < \infty$,
- (iii) the first exit time $\tau_D(X)$ is P^x -almost surely finite for all $x \in D$.

If $\psi \in C^2(D; \mathbb{R}) \cap C_b(\bar{D}; \mathbb{R})$ solves the boundary value problem

$$L\psi(x) - f(x)\psi(x) = -q(x) \quad x \in D \tag{2.43}$$

$$\lim_{x \rightarrow y} \psi(x) = k(y) \quad y \in \partial D, \tag{2.44}$$

then the solution ψ admits the Feynman-Kac representation (2.40).

The key difference between Proposition 2.1.18 and Theorem 2.1.21 regards whether the domain D is bounded or not. In particular, if D is not bounded, then the Feynman-Kac representation holds under the condition that the solution ψ to (2.43)–(2.44) is bounded. To determine when the ψ is bounded, we need the next result, which we have adapted from [33, Theorem 3.5, Section 3.2]:

Theorem 2.1.22 (Hopf's strong maximum principle). *Let L be a uniformly elliptic partial differential operator of the form (2.36), with $L\psi \geq 0$ (≤ 0) in a domain D where D is not necessarily bounded. If ψ achieves its maximum (minimum) in the interior of D , then ψ is constant. If the function f in the Dirichlet problem (2.37) is such that f/λ is bounded, where λ is the smallest eigenvalue of a , then ψ cannot achieve a non-negative maximum (non-positive minimum) in the interior of D unless ψ is constant. The conclusion remains valid if L is locally uniformly elliptic and the functions f and q are locally bounded.*

We conclude this section on the Dirichlet problem by considering what additional conditions are needed in order to extend the regularity of ψ to the closure of the bounded domain D . The motivation for doing so is that, in the next section, we will formulate a

stochastic optimal control problem that is related to the Dirichlet problem. We then study an approximation of the control functional in Chapter 3 in terms of a strongly convex function. This approximation is given in terms of an approximating subset of the set of feedback control functions. If the approximating subspace is spanned by bounded controls, it is of interest to determine when the optimal control is bounded.

Recall the following definition (see [33, Section 6.2]):

Definition 2.1.2. Let D be a bounded open subset of \mathbb{R}^d . Then D is said to be a *bounded domain of class $C^{k,\alpha}$* for $0 \leq \alpha \leq 1$ and $k \in \mathbb{N} \cup \{0\}$ if, for all $x_0 \in \partial D$, there exists a ball B centred at x_0 and a bijective mapping ψ of B onto D such that

- (i) $\psi(B \cap D) \subset \mathbb{R}_+^d = \{x \in \mathbb{R}^d \mid x_n > 0\}$,
- (ii) $\psi(B \cap \partial D) \subset \partial \mathbb{R}_+^d$, and
- (iii) $\psi \in C^{k,\alpha}(B)$, $\psi^{-1} \in C^{k,\alpha}(D)$.

Roughly speaking, a domain D is of class $C^{k,\alpha}$ if its boundary is locally the graph of a $C^{k,\alpha}$ function of all but one of the coordinates (x_1, \dots, x_n) . One can show that a bounded domain of class $C^{2,\alpha}$ satisfies the exterior sphere property; this fact is used in the proof [33, Section 6.3, Theorem 6.14], for example. Now recall that $C^k(\overline{D})$ is the subset of $C^k(D)$ functions whose partial derivatives of order less than or equal to k may be continuously extended to \overline{D} . The space $C^{k,\alpha}(\overline{D})$ is then a subset of $C^k(\overline{D})$ consisting of functions whose k -th order partial derivatives are uniformly Hölder continuous with exponent α , where $0 < \alpha \leq 1$. The next result, which is a slightly modified version of [33, Section 6.3, Theorem 6.14], concerns sufficient conditions for which $\psi \in C^{2,\alpha}(\overline{D})$.

Theorem 2.1.23. *Assume the following:*

- (i) *the infinitesimal generator L is uniformly elliptic in D ,*
- (ii) *the functions $(a_{ij})_{i,j}$, $(b_i)_i$, f , and q are uniformly Hölder continuous of exponent α on \overline{D} (i.e. they are elements of $C^\alpha(\overline{D})$),*
- (iii) *$k \in C^{2,\alpha}(\overline{D})$, and*
- (iv) *D is a bounded domain of class $C^{2,\alpha}$.*

Then the Dirichlet problem (2.37)–(2.38) has a unique solution $\psi \in C^{2,\alpha}(\overline{D})$.

Thus, if the hypotheses of Theorem 2.1.23 hold, then the gradient $\nabla_x \psi(\sigma, x)$ admits a continuous extension to \overline{D} , and hence is bounded on \overline{D} . We shall use this fact to show that, if it exists, then the optimal control for a given stochastic optimal control problem may be continuously extended to a bounded function on the bounded domain D .

2.2 Stochastic optimal control of diffusions

In this section, we formulate and study a stochastic optimal control problem. We assume that conditions (i)–(iv) of Theorem 2.1.20 hold for a bounded, open, connected set $D \subset \mathbb{R}^d$, for the drift vector b , and for the diffusion matrix a . We first consider three examples of path

functionals corresponding to statistical properties of diffusions that appear in the physics literature in Section §2.2.1. Section §2.2.2 concerns the connection between importance sampling and stochastic optimal control of diffusions, as well as sufficient conditions for the existence and uniqueness of a solution to the stochastic optimal control problem.

2.2.1 Path functionals

Let $f \in C(\overline{D}; [0, \infty))$ and $g \in C(\partial D; \mathbb{R})$ for a bounded domain D . Define the path functional

$$W(\omega) := \int_0^{\tau_D(\omega)} f(\omega_t) dt + g(\omega_{\tau_D(\omega)}). \quad (2.45)$$

We shall refer to W as the ‘work path functional’, or simply as the ‘work’; we will relate this choice to free energy-like quantities later. An immediate consequence of the conditions imposed so far is

Lemma 2.2.1. *There exists a constant $-\infty < C < +\infty$ such that*

$$C \leq W, \quad (2.46)$$

i.e. W is bounded from below.

Proof. Since g is continuous on a closed and bounded set in \mathbb{R}^d , it is bounded. Combined with the nonnegativity of f , it holds that

$$W \geq -\|g\|_{\partial D, \infty}. \quad \square$$

In the literature on controlled Markov processes, the functions f and g are sometimes referred to as the ‘running cost’ and the ‘terminal cost’ respectively. We shall be interested in computing statistical properties of W with respect to the measure

$$\mu^{0,x} := P^x \circ X^{-1}, \quad (2.47)$$

where $\mu^{0,x}$ solves the martingale problem $MP(b, a)$. We have the following

Lemma 2.2.2. *If the diffusion matrix a is uniformly elliptic on the domain D , then $W \in L^1(\mu^{0,x})$.*

Proof. By the conditions imposed on D , f and g , it holds that

$$E^{0,x}[|W|] \leq \|f\|_{\overline{D}, \infty} E^{0,x}[\tau_D] + \|g\|_{\partial D, \infty}. \quad (2.48)$$

Since a is uniformly elliptic, the right-hand side is finite, by Lemma 2.1.19. \square

The functional W and its moments are of interest in the analysis of metastable diffusions. For example, by setting

$$g \equiv 0, \quad f \equiv 1 \quad (2.49)$$

we can obtain that $W = \tau_D$, and in this case $E^x[W(X)] = E^x[\tau_D(X)]$ is known as the ‘mean first passage time’ or the ‘mean first exit time’ of the diffusion X from the bounded domain D . The mean first passage time is important because it may be used to obtain a

meaningful definition of metastability: A metastable diffusion is precisely a diffusion whose state space contains two or more bounded domains $(D_i)_i$, such that the mean first passage time $E^x[\tau_{D_i}(X)]$ is large for $x \in D_i$ relative to the mean first passage time $E^x[\tau_{D^c}(X)]$, where $D^c = (\cup_i D_i)^c$. In other words, the metastable diffusion spends long (short) times inside (outside) the metastable sets, on average. The metastable sets of the diffusion are then defined to be the domains $(D_i)_i$. In computational biophysics, conformations of molecules are modeled as metastable sets, and chemical reactions are modeled as transitions between them.

A second example is as follows. Suppose that the bounded domain D is such that ∂D admits a decomposition into the disjoint union of two or more sets:

$$\partial D = \cup_{i=1}^n B_i, \quad i \neq j \Rightarrow B_i \cap B_j = \emptyset.$$

For example, in the one-dimensional setting, if $D = (a, b)$ for $a < b$, then $\partial D = \{a\} \cup \{b\}$. For an arbitrary i , let

$$g \equiv 1_{B_i}, \quad f \equiv 0, \tag{2.50}$$

and let τ_D denote the first exit time from the domain D as before. Then the expected value of the resulting work function W with respect to $\mu^{0,x}$ gives the probability that, when the system X exits the domain D , it does so along the boundary set B_i :

$$E_{\mu^{0,x}}[W] = P^x(X_{\tau_D(X)} \in B_i).$$

This choice of running and terminal cost functions appears in [34]. In computational chemistry and biophysics, the function $q_{B_i}(x) := P^x(X_{\tau_D(X)} \in B_i)$ is known as the ‘committor function’, and has special relevance in the identification of transition states in chemical reactions (see, e.g. [21, 32]); Onsager’s study [48] into the dissociation of electron pairs is often cited as being the first to define and use such commitment probabilities. The states x for which the probability that the system X commits to exit via B_i equals the probability that the system does *not* exit via B_i are sometimes referred as ‘transition states’.

The third and final example that we wish to consider here makes use of Itô’s formula and yields a quantity that is similar to the thermodynamic free energy difference between two states of a statistical mechanical system. Recall that, since many systems of interest are complicated and high-dimensional, a common approach is to find a small number of collective variables that can describe most of the effective behaviour of the system. The idea is that, given a good choice of collective variables, one can group together states of the system in the full state space, according to their corresponding collective variables values. Given a probability density p on the collective variable space, one notion of the thermodynamic free energy may be obtained by taking the logarithm of the probability density, i.e.

$$G(x) = -\log p(x). \tag{2.51}$$

The motivation of this definition is that there is a relatively high (low) probability to find the statistical-mechanical system in regions in the collective variable space with low (high) free energy. Indeed, the overdamped Langevin equation

$$dX_t = -\nabla V(X_t)dt + \sqrt{2\varepsilon}dB_t \tag{2.52}$$

in collective variable space may be understood as describing the path taken by a statistical-mechanical system in equilibrium with a heat bath that seeks to minimise the value of some energy function V . The distribution of the system X in the collective variable space \mathcal{X} that corresponds to the dynamics above has the density

$$p(x) = \exp(-V(x))/Z, \quad Z = \int_{\mathcal{X}} \exp(-V(x))dx. \quad (2.53)$$

In the ideal case, the collective variable space is such that the probability distribution $p(x)$ defined in (2.53) is multi-modal, or equivalently, that there exist multiple local minima on the free energy landscape (i.e. the graph of G over \mathcal{X}). To each local minima there corresponds a metastable set D_i . The metastable sets are separated from one other by free energy barriers that the system X must cross in order to transition from one metastable set to another. To quantitatively determine which metastable sets are more stable than other, the differences between pairs of local minima - which may be thought of as *free energy differences* in this path setting - may be useful. Suppose that we wish to compute the free energy difference of one local minimum x with respect to another local minimum y . Fix $0 < \delta \ll 1$, and suppose that we can define a bounded domain $D = D(\delta)$ so that

$$\tau_D(X) := \inf \{t > 0 \mid |X_t - y| < \delta\}.$$

Then, for some energy function $V \in C^2(\mathcal{X}; \mathbb{R})$ that is bounded from below,

$$\begin{aligned} V(X_{\tau_D(X)}) - V(x) &= \int_0^{\tau_D(X)} \nabla V(X_s) dX_s + \frac{1}{2} \int_0^{\tau_D(X)} \sum_{i,j} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} V(X_s) d\langle X^i, X^j \rangle_s \\ &= - \int_0^{\tau_D(X)} |\nabla V(X_s)|^2 + \varepsilon \Delta V(X_s) ds + \sqrt{2\varepsilon} \int_0^{\tau_D(X)} \nabla V(X_s) dB_s \end{aligned}$$

by Itô's formula, (2.52), the formula (2.8) for the covariance of stochastic integrals, and the associative law. Taking expectations with respect to the conditioned measure P^x , we obtain

$$E^x \left[V(X_{\tau_D(X)}) \right] - V(x) = E^x \left[\int_0^{\tau_D(X)} -|\nabla V(X_s)|^2 + \varepsilon \Delta V(X_s) ds \right].$$

If one can interchange limits with expectations, then letting δ decrease to zero above would yield the difference in the 'energy' of the system at the minima at x and y . Thus, by setting

$$f(x) = -|\nabla V(x)|^2 + \varepsilon \Delta V(x), \quad g \equiv 0$$

in the work path functional W , we could compute the free energy difference between the minima at x and y . The difficulty in applying this idea to computing thermodynamic free energy differences is that one does not in general have precise knowledge of the free energy function or its derivatives. Furthermore, since it is not clear if the function f defined above is nonnegative, we cannot apply either Proposition 2.1.18 or Theorem 2.1.21. A mathematically-oriented survey of the problem of free energy computations may be found in [43].

2.2.2 Optimal control and importance sampling

In this section, we study the stochastic optimal control problem that was presented in [35] for performing importance sampling on metastable diffusions. Throughout this section, we shall consider $\sigma > 0$ to be a fixed, strictly positive real number.

Consider the stochastic differential equation

$$dX_t^c = [c - \nabla V](X_t^c)dt + \sqrt{2\varepsilon}dB_t, \quad (2.54)$$

where we assume the growth condition

$$|c(x)|^2 + |\nabla V(x)|^2 \leq C(1 + |x|^2)$$

for some constant C , analogous to (2.35), as well as the continuity and local Lipschitz conditions required for the existence and uniqueness of a strong solution X^c to (2.54). Since X^c is a random variable taking values in Ω , it induces a probability measure

$$\mu^{c,x} := P^x \circ (X^c)^{-1} \quad (2.55)$$

on (Ω, \mathcal{F}) . We shall write

$$E^{c,x}[\phi] = E[\phi(X^c)|X_0^c = x] \quad (2.56)$$

to denote the expectation of a random variable $\phi : \Omega \rightarrow \mathbb{R}$ with respect to $\mu^{c,x}$.

By varying the control function c , we obtain a family $(X^c)_{c \in \mathcal{A}}$ of random variables, parametrised by controls c belonging to an admissible class \mathcal{A} of control functions. Since each X^c is a Ω -valued random variable, we therefore obtain a family $(\mu^{c,x})_{c \in \mathcal{A}}$ of probability measures on (Ω, \mathcal{F}) . The reference measure in which we are interested is $\mu^{0,x}$, the law of the random variable X . We shall refer to $\mu^{0,x}$ as the ‘equilibrium measure’ and X as the ‘equilibrium system’ respectively. Accordingly, any measure $\mu^{c,x}$ for $c \neq 0$ shall be a ‘nonequilibrium measure’, and likewise X^c shall be a ‘nonequilibrium system’ for $c \neq 0$. In Section §2.2.1, a path functional W of the form

$$W(X) = \int_0^\tau f(X_t)dt + g(X_\tau)$$

was defined for a suitably defined stopping time τ and cost functions $f \in C(\bar{D}; [0, \infty))$ and $g \in C(\partial D; \mathbb{R})$, where D was a bounded domain satisfying an exterior sphere condition. The significance of the expected value

$$E^{0,x}[W] = E[W(X^0)|X_0^0 = x]$$

was explained by way of some examples, such as the mean first passage time and committor probabilities. The optimal control problem in which we are interested involves the real-valued functions ψ and F defined on $(0, \infty) \times \mathbb{R}^d$, defined by

$$\psi(\sigma, x) := E^{0,x}[\exp(-\sigma W)] \quad (2.57)$$

$$F(\sigma, x) := -\sigma^{-1} \log E^{0,x}[\exp(-\sigma W)]. \quad (2.58)$$

Remark 3. One of the most important results in statistical mechanics is the equation known as Jarzynski's equality [36], which establishes the following identity for the thermodynamic free energy difference ΔF_{AB} between two states A and B of a system in equilibrium with a heat bath:

$$\Delta F_{AB} = -\varepsilon \log \int_{\Omega} \exp(-W/\varepsilon)(\omega) d\omega.$$

In the equation above, the integral is interpreted as being the limit over infinitely many repetitions of a 'switching process', in which the system is driven out of equilibrium by the application of an external forcing or 'control protocol' that brings the system from state A to B . The remarkable nature of Jarzynski's equality is that it relates an equilibrium quantity (the thermodynamic free energy difference between two states of the system when it is in equilibrium) to nonequilibrium measurements. Prior to Jarzynski's result, one understood thermodynamic free energy differences in terms of ideal, 'infinitely slow' switching processes, in which the system was brought from state A to B in such a way as to keep the system in equilibrium with the heat bath at every point during the switching process. Since the ideal of an infinitely slow switching process was not experimentally realisable, measurements of thermodynamic free energy differences exhibited bias, due to the dissipation accumulated over the course of the switching processes. Jarzynski's equality showed that equilibrium free energy differences could in principle be attained by experimentally realisable switching processes. Note however that Jarzynski's equality does not remove the problem of bias, due to the nonlinearity of the logarithm function.

We now establish some useful estimates on the functions ψ and F .

Lemma 2.2.3. *It holds that ψ is strictly positive on its domain, and*

$$\psi(\sigma, x) \leq \exp(\sigma \|g\|_{\partial D, \infty}), \quad \forall (\sigma, x) \in (0, \infty) \times \mathbb{R}^d. \quad (2.59)$$

Proof. The second statement (2.59) follows immediately from the definition (2.57) of ψ and Lemma 2.2.1, since

$$-\sigma W \leq -\sigma(-\|g\|_{\partial D, \infty}).$$

To show that ψ is strictly positive, it suffices to show that W is finite $\mu^{0,x}$ -almost surely, since Jensen's inequality yields that $\psi(\sigma, x)$ is bounded from below according to

$$\exp(E^{0,x}[-\sigma W]) \leq E^{0,x}[\exp(-\sigma W)] = \psi(\sigma, x).$$

Since W was shown to be integrable with respect to $\mu^{0,x}$ in Lemma 2.2.2, almost sure finiteness follows, and hence so does the strict positivity of ψ . \square

Corollary 2.2.4. *The function F is finite and bounded from below according to*

$$-\|g\|_{\partial D, \infty} \leq F(\sigma, x), \quad \forall (\sigma, x) \in (0, \infty) \times \mathbb{R}^d. \quad (2.60)$$

The significance of the function F is that, for fixed x and $m \in \mathbb{N}$, the m -th moment of W with respect to $\mu^{0,x}$ may be obtained by differentiation:

$$\lim_{\sigma \downarrow 0} \frac{d^m}{d\sigma^m} F(\sigma, x) = E^{0,x}[W^m].$$

In other words, one can potentially obtain more information about W than just the expected value of W .

Recall that we are interested in computing statistical properties of metastable diffusions. The metastability of the diffusions implies that, in order to estimate statistical properties of path functionals (e.g. those defined in Section §2.2.1), one must wait a long time on average; this is because the first mean passage time $E^{0,x}[\tau_D]$ from a metastable set D is large, relative to a smaller time scale that needs to be preserved. In the context of molecular dynamics and biophysics, one must preserve motions that take place on very short time-scale (such as bond vibrations, that occur on the femtosecond time scale) in order to simulate events, such as protein folding events, that take place on longer time scales. Many events of interest involve transition events between metastable sets for which the mean first passage times are on the order of milliseconds or even larger time scales. Thus, the time scale of interest is many orders of magnitude larger than the shortest time scale that must be preserved in simulations, and it is reasonable to say that the events in which one is interested are *rare* with respect to the shortest time scale of the simulation.

The rareness of the events of interest implies that the probability distributions of observables (i.e. of path functionals) have tails that do not decay rapidly at more extreme values. For such probability distributions, standard Monte Carlo methods are inefficient, because more sample values are needed to bring the relative error (the standard deviation of the estimator divided by its mean) to within a preset tolerance level as the event becomes rarer. This problem is often circumvented by *importance sampling*, which consists of finding an alternative distribution such that the alternative distribution agrees with the original distribution for a given statistical property (e.g. the mean) but has smaller variance. In our case, the original distribution is the equilibrium distribution $\mu^{0,x}$, and the alternative distribution may be chosen from the family of nonequilibrium distributions $(\mu^{c,x})_{c \in \mathcal{A}}$ for the (as yet undefined) admissible class \mathcal{A} of control functions c .

As mentioned earlier, the quantity of interest is the value of

$$F(\sigma, x) = -\sigma^{-1} \log E^{0,x}[\exp(-\sigma W)],$$

for fixed x and σ . By the reweighting formula (2.32) seen earlier, we may rewrite the mean of $\exp(-\sigma W)$ with respect to $\mu^{0,x}$ as the mean (with respect to $\mu^{c,x}$) of $\exp(-\sigma W)$, reweighted by the appropriate Radon-Nikodym derivative:

$$E^{0,x}[\exp(-\sigma W)] = E^{c,x} \left[\exp(-\sigma W) \frac{d\mu^{0,x}}{d\mu^{c,x}} \right], \quad (2.61)$$

provided that the Radon-Nikodym derivative exists. For example, we could require that the control function c satisfies the condition (2.25) needed for the existence of the Radon-Nikodym derivative of $\mu^{0,x}$ with respect to $\mu^{c,x}$, i.e. we could require that

$$|c(x)|^2 \leq M \quad \forall x \in D \quad (2.62)$$

for some constant M .

Proposition 2.2.5. *Let $c : D \rightarrow \mathbb{R}^d$ be bounded and measurable. Then the Radon-Nikodym derivative of $\mu^{0,x}$ with respect to $\mu^{c,x}$ admits the following closed-form expression in terms*

of the control function c :

$$\left. \frac{d\mu^{0,x}}{d\mu^{c,x}} \right|_{\mathcal{F}_t} = \exp \left(-\frac{1}{\sqrt{2\varepsilon}} \int_0^t c(X_s^c) dB_s - \frac{1}{4\varepsilon} \int_0^t |c(X_s^c)|^2 ds \right). \quad (2.63)$$

Proof. The statement follows from Theorem 2.1.17. By the boundedness condition (2.62), the hypothesis (2.25) of Theorem 2.1.17 holds. Hence, there exists a probability measure Q that is locally equivalent to P such that $Q \circ (X^c)^{-1}$ solves $MP(-\nabla V, 2\varepsilon I_d)$, where I_d denotes the $d \times d$ identity matrix, and the Radon-Nikodym derivative of $Q \circ (X^c)^{-1}$ with respect to $P \circ (X^c)^{-1}$ is given by

$$\left. \frac{dQ \circ (X^c)^{-1}}{dP \circ (X^c)^{-1}} \right|_{\mathcal{F}_t} = \mathcal{E}(Y)_t.$$

By (2.54) and (2.55), it holds that for $c \neq 0$, $\mu^{c,x}$ solves $MP(c-\nabla, 2\varepsilon I_d)$. Setting $\beta = c-\nabla V$ and $b + \beta = -\nabla V$ in Theorem 2.1.17 yields that $b = -c$. Thus, the processes \bar{X} and Y defined in (2.26)–(2.27) are equal to

$$\begin{aligned} \bar{X}_t &= \sqrt{2\varepsilon} B_t \\ Y_t &= -\frac{1}{\sqrt{2\varepsilon}} \int_0^t c(X_s) dB_s, \end{aligned}$$

and hence

$$\mathcal{E}(Y)_t = \exp \left(-\frac{1}{\sqrt{2\varepsilon}} \int_0^t c(X_s^c) dB_s - \frac{1}{4\varepsilon} \int_0^t |c(X_s^c)|^2 ds \right).$$

By the uniqueness of solutions to the martingale problem $MP(-\nabla V, 2\varepsilon I_d)$, it holds that

$$Q \circ (X^c)^{-1} = P \circ (X^0)^{-1} = \mu^{0,x}.$$

Therefore,

$$\mathcal{E}(Y) = \frac{dQ \circ (X^c)^{-1}}{dP \circ (X^c)^{-1}} = \frac{d\mu^{0,x}}{d\mu^{c,x}},$$

as desired. \square

Notation. Recall that, for a martingale M in the class \mathcal{M}^2 of square-integrable martingales, the set $\Pi_3(M)$ defined in (2.5) denotes the largest class of previsible (not necessarily continuous) integrand processes H such that the Itô integral of H with respect to the square integrable martingale M is well-defined. By Theorem 2.1.1, the Itô integral of any process in $\Pi_3(M)$ with respect to M is again a continuous local martingale. For arbitrary $h, h' \in \Pi_3(B)$, for arbitrary $\omega \in \Omega$, and for arbitrary $t \geq 0$, we shall write

$$M_t^h := \frac{1}{\sqrt{2\varepsilon}} \int_0^t h(\omega_s) dB_s, \quad (2.64)$$

$$\langle M^h \rangle_t(\omega) := \frac{1}{2\varepsilon} \int_0^t |h(\omega_s)|^2 ds. \quad (2.65)$$

$$\langle M^h, M^{h'} \rangle_t(\omega) := \frac{1}{2\varepsilon} \int_0^t h \cdot h'(\omega_s) ds. \quad (2.66)$$

Then we may rewrite the Radon-Nikodym derivative of $\mu^{0,x}$ with respect to $\mu^{c,x}$ by

$$\frac{d\mu^{0,x}}{d\mu^{c,x}} = \exp\left(M^{-c} - \frac{1}{2}\langle M^{-c} \rangle\right) = \mathcal{E}(M^{-c}).$$

Rewriting (2.63), substituting the resulting expression into (2.61), taking the logarithm and dividing by $-\sigma$ yields

$$F(\sigma, x) = -\sigma^{-1} \log E^{c,x} \left[\exp\left(-\sigma \left(W + \frac{M_\tau^c}{\sigma} + \frac{\langle M^c \rangle_\tau}{2\sigma}\right)\right) \right] \quad (2.67)$$

where we have stopped M^c and $\langle M^c \rangle$ at $\tau(X^c)$ since the smallest sigma-algebra with respect to which the path functional W is measurable is $\mathcal{F}_{\tau(X^c)}$. By Jensen's inequality, we obtain the key inequality

$$F(\sigma, x) \leq E^{c,x} \left[W + \frac{M_\tau^c}{\sigma} + \frac{\langle M^c \rangle_\tau}{2\sigma} \right], \quad (2.68)$$

which holds for all $c \in \mathcal{A}$. The inequality (2.68) is important because it defines the eponymous stochastic optimal control problem in which we are interested. Define for a Boolean variable α the functional

$$K^{\sigma,c,\alpha} := W + (2\sigma)^{-1} \langle M^c \rangle_\tau + \alpha \sigma^{-1} M_\tau^c. \quad (2.69)$$

We define a control functional $\bar{\phi}^{\sigma,x} : \mathcal{A} \rightarrow [F(\sigma, x), \infty)$ by

$$\bar{\phi}^{\sigma,x}(c) := E^{c,x} [K^{\sigma,c,0}], \quad (2.70)$$

where we justify the choice of $\alpha = 0$ in (2.70) by the observation that the value of the parameter α does not change the value of $\bar{\phi}^{\sigma,x}$, since M^c is a continuous local martingale with $M_0^c = 0$.

The stochastic optimal control problem is:

$$\min_c \bar{\phi}^{\sigma,x}(c) \quad (2.71)$$

$$\text{s.t. } dX_t^c = [c(X_t^c) - \nabla V(X_t^c)]dt + \sqrt{2\varepsilon}dB_t, \quad (2.72)$$

where the optimisation problem (2.71) is performed over the set of feedback control functions. In the theory of stochastic optimal control, one often performs the optimisation problem over the set of all previsible control processes. Since a control process need not be generated by a feedback control, one can modify the problem (2.71)–(2.72) to be defined in terms of a control process $u = (u_t)_{t \geq 0}$. However, in Theorem 2.2.8, we shall show that the control functional that we defined attains its global minimum on the set of feedback control functions, so nothing is lost by defining the optimal control problem as we have done. Recall that for $x \in D$,

$$\mu^{0,x} := P^x \circ X^{-1}$$

is the law of the random variable X , where X is the pathwise unique strong solution to

$$dX_t = -\nabla V(X_t)dt + \sqrt{2\varepsilon}dB_t.$$

The associated infinitesimal generator of the above stochastic differential equation is defined by

$$L\psi(x) := \varepsilon\Delta\psi(x) - \nabla V(x) \cdot \nabla\psi(x).$$

The infinitesimal generator above is uniformly elliptic, since the diffusion matrix a is a strictly positive multiple of the identity matrix. Recall that D was assumed to be a bounded domain that satisfies the exterior sphere condition (2.42). Suppose that the functions f and g in (2.37)–(2.38) are Hölder continuous on their respective domains. By Theorem 2.1.20, it follows that there exists a unique classical solution ψ , i.e. some function $\psi \in C^2(D; \mathbb{R}) \cap C(\bar{D}; \mathbb{R})$, that solves

$$L\psi(x) - \sigma f(x)\psi(x) = 0 \quad x \in D \tag{2.73}$$

$$\psi(y) = \exp(-\sigma g(y)) \quad y \in \partial D, \tag{2.74}$$

and that ψ is given by (2.40) for $k(y) := \exp(-\sigma g(y))$ and $q \equiv 0$. In particular, it holds that

$$\psi(x) = E^{0,x} [\exp(-\sigma W)],$$

which equals the function $\psi(\sigma, x)$ defined in (2.57). By Lemma 2.2.3, $\psi(\sigma, x)$ is strictly positive on D . In order to show the existence and uniqueness of the control that solves the optimal control problem (2.71)–(2.72), we need to establish that the value function associated to the optimal control problem is a classical solution to a nonlinear elliptic problem. The inequality (2.68) suggests that we would like the value function to be given by the function $F(\sigma, x)$. By substituting $\psi(\sigma, x) = \exp(-\sigma F(\sigma, x))$ into the linear elliptic boundary value problem above and by applying the chain rule, we obtain the desired nonlinear boundary value problem

$$f(x) + LF(\sigma, x) - \varepsilon\sigma|\nabla_x F(\sigma, x)|^2 = 0 \quad x \in D, \tag{2.75}$$

$$F(\sigma, y) = g(y) \quad y \in \partial D. \tag{2.76}$$

Note that the nonlinearity in the boundary value problem (2.75)–(2.76) is entirely due to the quadratic term in (2.75). In the optimal control literature, (2.75) would be referred to as the ‘Hamilton-Jacobi-Bellman’ or ‘dynamic programming’ equation of the optimal control problem (2.71)–(2.72), and the solution $F(\sigma, x)$ (for fixed $\sigma > 0$) is known as the value function. In general, the Hamilton-Jacobi-Bellman equation is a parabolic partial differential equation in time t and x , and the value function $v(t, x)$ describes the optimal cost of solving the optimal control problem, given that the controlled diffusion is at x at time t . For the stochastic optimal control problem considered here, $t = 0$ is held fixed, and hence the boundary value problem (2.75)–(2.76) exhibits no time dependence.

For a general stochastic optimal control problem, there are no classical solutions to (2.75)–(2.76), due to the difficulties inherent to working in the nonlinear case, and one must search for a *viscosity solution*. Although viscosity solutions are interesting objects in their own right, we shall restrict ourselves to classical solutions. In [30, IV.5], the existence of a classical solution to (2.75)–(2.76) for the infinite-time horizon case is proven by appealing to a result [33, Theorem 17.17] concerning the existence of a classical solution to nonlinear elliptic boundary value problems. A detailed proof would lead too far into the theory of nonlinear elliptic partial differential equations, so we shall only sketch the

main idea of the result here. Let $\Gamma := D \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d \times d}$, $\gamma = (x, z, p, r) \in \Gamma$ denote an arbitrary element of Γ , and $\mathcal{H} : \Gamma \rightarrow \mathbb{R}$ be defined by

$$\mathcal{H}(x, z, p, r) = f(x) - \nabla V(x) \cdot p - \varepsilon \sigma |p|^2 + \varepsilon \left(\sum_i r_{ii} \right). \quad (2.77)$$

From (2.77), it follows that \mathcal{H} is constant with respect to z , concave with respect to p , and linear with respect to r . Note that we can rewrite (2.75) as

$$\mathcal{H}(x, F(\sigma, x), \nabla_x F(\sigma, x), \nabla_x^2 F(\sigma, x)) = 0 \quad x \in D.$$

We write

$$\mathcal{H}_{ij}(x, z, p, r) := \frac{\partial \mathcal{H}}{\partial r_{ij}}(x, z, p, r).$$

Given (2.77), it holds that

$$\mathcal{H}_{ij}(x, z, p, r) = \delta_{ij} \varepsilon \quad (2.78)$$

where δ_{ij} is the Kronecker delta. Define the so-called ‘structural conditions’

$$0 < \lambda |\xi|^2 \leq \sum_{i,j} \mathcal{H}_{ij} \xi_i \xi_j \leq \Lambda |\xi|^2 \quad (2.79a)$$

$$|\mathcal{H}_p|, |\mathcal{H}_z|, |\mathcal{H}_{rx}|, |\mathcal{H}_{px}|, |\mathcal{H}_{zx}| \leq \mu \lambda, \quad (2.79b)$$

$$|\mathcal{H}_x|, |\mathcal{H}_{xx}| \leq \mu \lambda (1 + |p| + |r|), \quad (2.79c)$$

where $\xi \in \mathbb{R}^d$ is nonzero, $(x, z, p, r) \in \Gamma$ is arbitrary, λ is a nonincreasing function of $|z|$, and Λ and μ are nondecreasing functions of $|z|$. These structural conditions are used to obtain so-called *interior estimates* on the norm of the function $F(\sigma, \cdot)$ and its derivatives.

With these preparations in mind, we adapt [33, Theorem 17.17] below:

Theorem 2.2.6. *Let D be a bounded domain in \mathbb{R}^d satisfying an exterior sphere condition at each boundary point and suppose the function $\mathcal{H} \in C^2(\Gamma; \mathbb{R})$ is concave (or convex) with respect to z , p , r , is nonincreasing with respect to z , and satisfies the structural conditions (2.79). Then the classical Dirichlet problem,*

$$\begin{aligned} \mathcal{H}(x, F(\sigma, x), \nabla_x F(\sigma, x), \nabla_x^2 F(\sigma, x)) &= 0 \quad x \in D \\ F(\sigma, y) &= g(y) \quad y \in \partial D \end{aligned}$$

is uniquely solvable in $C^2(D; \mathbb{R}) \cap C(\bar{D}; \mathbb{R})$ for any $g \in C(\partial D; \mathbb{R})$.

The exterior sphere condition corresponds to condition (iv) in Theorem 2.1.20. For $\mathcal{H} \in C^2(\Gamma; \mathbb{R})$, we must make the following

Assumption 2.2.7. *The drift term and running cost function are twice continuously differentiable functions of the spatial variable x .*

Given that \mathcal{H} is constant with respect to z , concave with respect to p , and linear with respect to r , it remains to verify that the structural conditions hold. From (2.78), we can set λ and Λ to be constant functions equal to ε . Thus we can verify that λ and Λ are

nonincreasing and nondecreasing functions of $|z|$ respectively, and (2.79a) holds. For the remaining conditions, observe that by (2.77), we have

$$|\mathcal{H}_z| = |\mathcal{H}_{rz}| = |\mathcal{H}_{zx}| = 0.$$

Since

$$\frac{\partial \mathcal{H}}{\partial p_i} = -\frac{\partial V}{\partial x_i} - 2\varepsilon \sigma p_i \quad (2.80a)$$

$$\frac{\partial^2 \mathcal{H}}{\partial p_i \partial x} = -\frac{\partial^2 V}{\partial x_i^2} \quad (2.80b)$$

$$\frac{\partial \mathcal{H}}{\partial x_i} = -\frac{\partial}{\partial x_i} (f(x) - \nabla V(x) \cdot p) \quad (2.80c)$$

$$\frac{\partial^2 \mathcal{H}}{\partial x_i \partial x_j} = -\frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} (f(x) - \nabla V(x) \cdot p), \quad (2.80d)$$

and since the right-hand sides of (2.80) are constant with respect to $|z|$, the remaining structural conditions (2.79b)–(2.79c) are satisfied. Thus, we have shown that the sufficient conditions for the function $\psi(\sigma, \cdot)$ to be a classical solution of the linear elliptic boundary value problem (2.37)–(2.38) do *not* suffice for $F^\sigma(\cdot) := -\sigma^{-1} \log \psi(\sigma, \cdot)$ to be a classical solution of the nonlinear elliptic boundary value problem (2.75)–(2.76). In particular, the first-order coefficient vector ∇V of the operator L must be twice differentiable with respect to x in order for nonlinear Dirichlet problem to have a classical solution, whereas uniform Hölder continuity suffices for the linear Dirichlet problem. This observation exemplifies the general principle that nonlinear equations are generally more complicated than their linear counterparts.

We now adapt the main result in [35], namely, the existence and uniqueness of the optimal control that solves (2.71)–(2.72):

Theorem 2.2.8. *Let D be a bounded domain satisfying the exterior sphere condition (2.42), f and g be Hölder continuous on their respective domains, τ be the first time of exit from D , $\bar{\phi}^{\sigma,x}$ be given as in (2.70), and assume that Assumption 2.2.7 holds. Then the optimal control c_{opt}^σ that satisfies*

$$\bar{\phi}^{\sigma,x}(c_{opt}^\sigma) = \min_c \bar{\phi}^{\sigma,x}(c), \quad (2.81)$$

where the optimisation problem on the right-hand side is done over the set of feedback control functions, is a feedback control of the form

$$c_{opt}^\sigma(x) := -2\varepsilon \sigma \nabla_x F(\sigma, x), \quad (2.82)$$

where $F(\sigma, \cdot) \in C^2(D; \mathbb{R}) \cap C(\bar{D}; \mathbb{R})$. Furthermore, the optimal control has the property that

$$F(\sigma, x) = \bar{\phi}^{\sigma,x}(c_{opt}^\sigma). \quad (2.83)$$

Proof. The proof follows a common strategy in infinite time-horizon optimal control problems where the control is linear in the dynamics and quadratic in the objective. The

nonlinearity in the Hamilton-Jacobi-Bellman equation (2.75) that arises due to the quadratic cost of control is particularly useful, because it holds that

$$-\varepsilon\sigma|\nabla_x F(\sigma, x)|^2 = \min_{v \in \mathbb{R}^d} \left\{ v \cdot \nabla_x F(\sigma, x) + \frac{1}{4\varepsilon\sigma} |v|^2 \right\}. \quad (2.84)$$

Therefore, the Hamilton-Jacobi-Bellman equation is equivalent to

$$f(x) + LF(\sigma, x) + \min_{v \in \mathbb{R}^d} \left\{ v \cdot \nabla_x F(\sigma, x) + \frac{1}{4\varepsilon\sigma} |v|^2 \right\} = 0, \quad x \in D. \quad (2.85)$$

The unique optimal control which solves this problem is the argument that minimises the quantity in the curly braces in (2.84). By substituting $c_{\text{opt}}^\sigma(x)$ defined by (2.82) for v in (2.84), the conclusion follows. \square

In the next chapter, we will construct an approximation scheme using bounded basis functions. Although one often approximates objects that are not a priori bounded by their bounded counterparts in mathematics (e.g. nonnegative Lebesgue integrable functions by simple functions), it is nevertheless of interest to consider conditions under which the optimal control c_{opt}^σ is bounded on the domain \bar{D} . Observe that, by the formula (2.82) for the optimal control and by the fact that $F(\sigma, x) = -\sigma^{-1} \log \psi(\sigma, x)$, it holds that the optimal control c_{opt}^σ may be rewritten as

$$c_{\text{opt}}^\sigma(x) = -2\varepsilon\sigma \nabla_x F(\sigma, x) = 2\varepsilon \frac{\nabla_x \psi(\sigma, x)}{\psi(\sigma, x)}.$$

Since D is bounded, it follows from the Heine-Borel theorem that \bar{D} is compact. By Theorem 2.1.20, $\psi(\sigma, \cdot) \in C(\bar{D}; \mathbb{R})$, and hence $\psi(\sigma, \cdot)$ attains its minimum on \bar{D} . By the strong maximum principle (Theorem 2.1.22), $\psi(\sigma, \cdot)$ attains its minimum in D . In Lemma 2.2.3, we showed that $\psi(\sigma, x)$ is strictly positive for all $x \in D$. Hence $\psi(\sigma, \cdot)$ attains a strictly positive minimum in D . Thus, to show that the optimal control c_{opt}^σ is bounded on the domain \bar{D} , it suffices to show that $\nabla_x \psi(\sigma, x)$ can be continuously extended to a bounded, continuous function on \bar{D} . This can be done by imposing more stringent regularity conditions on the data f and g of the linear elliptic boundary value problem (2.37)–(2.38), and on the domain D . In particular, we must assume that the functions f and g admit uniformly Hölder continuous extensions to the closure of D , and that the bounded domain is of class $C^{2,\alpha}$ for $0 \leq \alpha \leq 1$; see Theorem 2.1.23.

Recall the formula (2.69) for the path functional $K^{\sigma,c,\alpha}$.

Corollary 2.2.9. *The estimator $K^{\sigma,c,1}$ is $\mu^{c,x}$ -almost surely constant if and only if the control c is optimal.*

Proof. Recall that $\mu^{c,x} = P^x \circ (X^c)^{-1}$, where X^c is the controlled diffusion that solves (2.72). The inequality (2.68)

$$F(\sigma, x) \leq E^{c,x} \left[K^{\sigma,c,1} \right] \quad (2.86)$$

that defines the stochastic optimal control problem is a consequence of the reweighting formula (2.67), which we rewrite in terms of $K^{\sigma,c,1}$ as

$$E^{0,x} [\exp(-\sigma W)] = E^{c,x} \left[\exp \left(-\sigma K^{\sigma,c,1} \right) \right],$$

and Jensen's inequality. Since the logarithm is a nonlinear, strictly concave function, equality holds in Jensen's inequality if and only if the random variable $K^{\sigma,c,1}$ is $\mu^{c,x}$ -almost surely constant (see, e.g. [44, Theorem 5]). By equation (2.83) in Theorem 2.2.8, c_{opt}^σ is the unique control for which equality holds in the inequality above, and hence the conclusion follows. \square

Remark 4. It does not follow from Corollary 2.2.9 that $K^{\sigma,c,0}$ is $\mu^{c,x}$ -almost surely constant when $c = c_{\text{opt}}^\sigma$. In fact, since the variance of $K^{\sigma,c,0}$ is given by

$$\text{var}^{c,x} \left[K^{\sigma,c,0} \right] = \text{var}^{c,x} \left[K^{\sigma,c,1} \right] + \text{var}^{c,x} \left[-\frac{M_\tau^c}{\sigma} \right] - 2\text{cov}^{c,x} \left[K^{\sigma,c,1}, \frac{M_\tau^c}{\sigma} \right],$$

it follows that, if $c = c_{\text{opt}}^\sigma$, then the variance of $K^{\sigma,c,0}$ is proportional to the variance of M_τ^c , which is strictly positive provided that c is nonzero and is supported in the domain D .

We end this section by putting into context the results that we have presented thus far. Recall that the relative entropy or Kullback-Leibler divergence [40] of $\mu^{0,x}$ with respect to $\mu^{c,x}$ on the probability space $(\Omega, \mathcal{F}_T, \mu^{c,x})$ for a stopping time $T \leq \infty$ is defined by (compare with [15, Equation (2)])

$$KL(\mu^{c,x} | \mu^{0,x})|_{\mathcal{F}_T} := \begin{cases} E^{c,x} \left[\log \frac{d\mu^{c,x}}{d\mu^{0,x}} \Big|_{\mathcal{F}_T} \right] & \mu^{c,x} \ll \mu^{0,x}, \frac{d\mu^{c,x}}{d\mu^{0,x}} \in L^1(\mu^{c,x}) \\ \infty & \text{otherwise.} \end{cases} \quad (2.87)$$

The inequality (2.68) provides an example of Legendre-type duality relationships between free energy and relative entropy

$$F(\sigma; x) = \inf \left\{ E^{c,x} [W] + \sigma^{-1} KL(\mu^{c,x} | \mu^{0,x})|_{\mathcal{F}_\tau} \mid \mu^{c,x} \ll \mu^{0,x} \text{ on } \mathcal{F}_\tau \right\}, \quad (2.88)$$

(see, e.g. [9, 15, 34]). The optimal control c_{opt}^σ given in (2.82) that solves the stochastic optimal control problem (2.71) corresponds to an optimal measure $\mu^{c_{\text{opt}}^\sigma, x}$ that solves the problem (2.88). By Corollary 2.2.9, it holds that

$$\mu^{c_{\text{opt}}^\sigma, x} \left(\sigma W - \sigma F(\sigma, x) = -M_\tau^{c_{\text{opt}}^\sigma} - 2^{-1} \langle M^{c_{\text{opt}}^\sigma} \rangle_\tau \right) = 1. \quad (2.89)$$

By Proposition 2.2.5, the importance sampling measure $\mu^{c,x}$ corresponding to an arbitrary suitable control is defined by $\frac{d\mu^{0,x}}{d\mu^{c,x}} = \exp \left(-M^c - \frac{1}{2} \langle M^c \rangle \right)$. Hence, it follows from (2.89) that

$$\frac{d\mu^{0,x}}{d\mu^{c_{\text{opt}}^\sigma, x}} \Big|_{\mathcal{F}_\tau} = \exp(\sigma W - \sigma F(\sigma, x)), \quad (2.90)$$

and since $F(\sigma, x) = -\sigma^{-1} \log E^{0,x}[\exp(-\sigma W)]$, we have

$$\frac{d\mu^{c_{\text{opt}}^\sigma, x}}{d\mu^{0,x}} \Big|_{\mathcal{F}_\tau} = \frac{\exp(-\sigma W)}{E^{0,x}[\exp(-\sigma W)]}. \quad (2.91)$$

Substituting the above into the optimisation problem (2.88) yields equality, as desired. The formula (2.91) for the optimal importance sampling measure is characteristic of all optimal importance sampling measures, in that the optimal importance sampling measure is defined in terms of the quantity that one wishes to estimate (see, e.g. [49]).

Chapter 3

Strongly convex approximations

In this chapter, we present the first of the three main results of this thesis, namely the construction of a strongly convex approximation to the control functional defined in Section §2.2.2. In Section §3.1, we state conditions that suffice to guarantee the existence of the first and second variation of $\bar{\phi}^{\sigma,x}$ in the direction of some suitable perturbing function. An important observation in Section §3.1 is that the expressions for the first and second variations involve stopped martingales. In Section §3.2, we build upon this observation, and apply fundamental results from stochastic analysis concerning continuous local martingales, in order to derive representations of the first and second variation that shall prove useful in Section §3.3. We show in Section §3.2.2 that imposing a non-overlap condition on the supports of a finite collection of perturbing functions gives rise to an associated collection of *independent* martingales. We then apply Itô's formula and certain martingale inequalities to these independent martingales in Section §3.3.2 in order to construct the strongly convex approximation of the control functional.

Recall that $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t>0})$ denotes the canonical filtered probability space, where Ω is the set of continuous functions $\omega : [0, \infty) \rightarrow \mathbb{R}^d$. Recall also that P denotes the Wiener measure, that P^x indicates conditioning on the set of all paths satisfying $\omega_0 = x$, and that $B = \omega$ denotes the standard Brownian motion or Wiener process. Let X^c denote the pathwise unique strong solution to

$$dX_t^c = [c(X_t^c) - \nabla V(X_t^c)]dt + \sqrt{2\varepsilon}dB_t. \quad (3.1)$$

In the previous chapter, we considered the following problem: given a bounded domain $D \subset \mathbb{R}^d$ whose boundary satisfies certain regularity conditions (e.g. an exterior sphere condition), $f \in C(\bar{D}; [0, \infty))$, $g \in C_b(\partial D; \mathbb{R})$, and $\tau(\omega)$ being the first exit time of ω from D for $\omega_0 = x \in D$, we wish to estimate statistical properties of the functional

$$W(\omega) := \int_0^\tau f(\omega_s)ds + g(\omega_\tau) \quad (3.2)$$

with respect to the measure $\mu^{0,x} = P^x \circ (X^0)^{-1}$ (i.e. $\mu^{0,x} \equiv \mu^{c,x}$ for $c \equiv 0$), so that $E_{\mu^{0,x}}[W] = E_{P^x}[W(X)]$. We shall write $E^{0,x}[W] = E_{\mu^{0,x}}[W]$. Our goal is to estimate the value of $F(\sigma, x) = -\sigma^{-1} \log E^{0,x}[\exp(-\sigma W)]$ for $\sigma > 0$ and $x \in D$. The value of $F(\sigma, x)$ is the value of the optimal control problem

$$\min_c \bar{\phi}^{\sigma,x}(c) \quad (3.3)$$

where

$$\bar{\phi}^{\sigma,x}(c) := E^{c,x} [K^{\sigma,c,0}] \quad (3.4)$$

denotes the expectation with respect to $\mu^{c,x} = P^x \circ (X^c)^{-1}$ of the path functional $K^{\sigma,c,0}$, where

$$K^{\sigma,c,\alpha} := W + (2\sigma)^{-1} \langle M^c \rangle_\tau + \alpha \sigma^{-1} M_\tau^c, \quad \alpha \in \{0, 1\}, \quad (3.5)$$

and where the optimisation problem (3.3) is defined over the set of feedback control functions.

Since the diffusion matrix corresponding to (3.1) is a positive scalar multiple of the identity matrix, it follows that if the feedback control c is bounded (i.e. $c \in L^\infty(\bar{D})$), then the assumptions of the change of measure theorem (Theorem 2.1.17) are satisfied. Hence μ^0 and μ^c are locally equivalent, as are their conditioned versions $\mu^{0,x}$ and $\mu^{c,x}$, and the associated Radon-Nikodym derivatives exist and are exponential martingales. Note that Proposition 2.2.5 does not require continuity of c .

Notation: For $\varphi, \zeta \in L^\infty(\bar{D})$, recall that we write

$$\begin{aligned} M_t^\varphi(\omega) &= \frac{1}{\sqrt{2\varepsilon}} \int_0^t \varphi(\omega_s) dB_s(\omega) \\ \langle M^\varphi, M^\zeta \rangle_t(\omega) &= \frac{1}{2\varepsilon} \int_0^t \varphi(\omega_s) \cdot \zeta(\omega_s) ds. \end{aligned} \quad (3.6)$$

Remark 5. Recall that, given a continuous local martingale M , the set $\Pi_3(M)$ defined in (2.5) is the largest class of integrand processes H for which the Itô integral with respect to M , denoted by $H \cdot M$, is a well-defined, continuous local martingale whose quadratic variation process is almost-surely finite at all deterministic times $t \geq 0$. This implies that, for any $\varphi \in L^\infty(\bar{D})$, the integrand process H defined by $H_t := \varphi(X_t^c)$ is an element of $\Pi_3(B)$, since

$$\langle H \cdot B \rangle_t = \int_0^t H_s^2 d\langle B \rangle_s = \int_0^t |\varphi(X_s^c)|^2 ds \leq \|\varphi\|_\infty^2 t < \infty \quad \forall t \geq 0.$$

For $a, b \in \mathbb{R}$, we have

$$\begin{aligned} M_t^{a\varphi+b\zeta} &= aM_t^\varphi + bM_t^\zeta \\ \langle M^{a\varphi+b\zeta} \rangle_t &= a^2 \langle M^\varphi \rangle_t + b^2 \langle M^\zeta \rangle_t + 2ab \langle M^\varphi, M^\zeta \rangle_t \end{aligned} \quad (3.7)$$

The covariance process $\langle M^\varphi, M^\zeta \rangle$ has the property that

$$E^x [M_t^\varphi M_t^\zeta] = E^x [\langle M^\varphi, M^\zeta \rangle_t]. \quad (3.8)$$

Recall that the Radon-Nikodym derivative of $\mu^c := P \circ (X^c)^{-1}$ with respect to $\mu^{c+\varphi}$ is given by a Doleans exponential martingale, where comparing (2.21) and (2.63) yield

$$\left. \frac{d\mu^c}{d\mu^{c+\varphi}} \right|_{\mathcal{F}_t} = \exp \left[M_t^{-\varphi} - 2^{-1} \langle M^{-\varphi} \rangle_t \right] = \mathcal{E}(M^{-\varphi})_t \quad \forall t \geq 0. \quad (3.9)$$

Observe also that the relation above may be used to show that

$$\left. \frac{d\mu^{c+\varphi}}{d\mu^c} \right|_{\mathcal{F}_t} = \exp \left[M_t^\varphi - 2^{-1} \langle M^\varphi \rangle_t \right] = \mathcal{E}(M^\varphi)_t \quad \forall t \geq 0.$$

We shall write

$$D_t^\varphi := \left. \frac{d\mu^{c+\varphi}}{d\mu^c} \right|_{\mathcal{F}_t} = \mathcal{E}(M^\varphi)_t \quad (3.10)$$

to denote the value of the Radon-Nikodym derivative of $\mu^{c+\varphi}$ with respect to μ^c at time t . We shall also use the fact that, for a functional $F \in L^1(\mu^{c+\varphi, x})$ that is measurable with respect to the sigma-algebra \mathcal{F}_T (i.e. the sigma algebra corresponding to the filtration up to some stopping time $T > 0$), then

$$E^{c+\varphi, x}[F] = E^{c, x}[FD_T^\varphi], \quad (3.11)$$

provided that the Radon-Nikodym derivative exists.

3.1 Variational analysis of the control functional

In this section, we obtain expressions for the first- and second-order variations of the control functional $\bar{\phi}^{\sigma, x}$ that was defined in (3.4), in the direction of suitable perturbing functions. Using the change of measure theorem and the formula for the Doleans exponential martingale $\mathcal{E}(M)$, we compute the first variation by taking limits and showing L^1 convergence of the approximating functional to the desired expression. In what follows, we shall use the inequality of the arithmetic and geometric means,

$$|ab| \leq 2^{-1}(a^2 + b^2) \quad (3.12)$$

for $a, b \in \mathbb{R}$, where equality holds if and only if $a = b$.

3.1.1 The first variation and its first variation

Recall that, given a functional F defined on some domain X , the first variation of F in an admissible direction y evaluated at some point $x \in X$ is defined by the limit

$$\lim_{n \rightarrow \infty} \frac{F(x + \delta_n y) - F(x)}{\delta_n}$$

for any sequence $(\delta_n)_n$ satisfying $\delta_n \rightarrow 0$, provided that the limit exists and is finite. By the definition (3.10), and by the property (3.7), it holds that

$$D_\tau^{\delta\varphi} = \exp \left[M_\tau^{\delta\varphi} - \frac{1}{2} \langle M^{\delta\varphi} \rangle_\tau \right] = \exp \left[\delta \left(M_\tau^\varphi - \frac{\delta}{2} \langle M^\varphi \rangle_\tau \right) \right]$$

for a scalar $\delta \in \mathbb{R}$, a Borel-measurable function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and a stopping time τ . The series expansion of the exponential thus yields

$$D_\tau^{\delta\varphi} = 1 + \sum_{m=1}^{\infty} \frac{\delta^m (M_\tau^\varphi - \delta 2^{-1} \langle M^\varphi \rangle_\tau)^m}{m!} \quad (3.13)$$

and hence

$$\frac{D_\tau^{\delta\varphi} - 1}{\delta} - M_\tau^\varphi = -\frac{\delta}{2} \langle M^\varphi \rangle_\tau + \sum_{m=2}^{\infty} \frac{\delta^{m-1} (M_\tau^\varphi - \delta 2^{-1} \langle M^\varphi \rangle_\tau)^m}{m!}. \quad (3.14)$$

The definition of the first variation, combined with the identities (3.14) and (3.11), suggest that the martingale M^φ will appear in the first variation of a functional on Ω . For a \mathcal{F}_T -measurable functional $F \in L^1(\mu^{c,x})$ that does not depend on the function c , for any $\varphi \in L^\infty(\bar{D})$, and for any nonzero sequence $(\delta_n)_n$ decreasing to zero,

$$\lim_{n \rightarrow \infty} \frac{E^{c+\delta_n \varphi, x}[F] - E^{c,x}[F]}{\delta_n} = \lim_{n \rightarrow \infty} E^{c,x} \left[F \frac{D_T^{\delta_n \varphi} - 1}{\delta_n} \right] = E^x[F(X^c)M_T^\varphi(X^c)],$$

In order to justify the second equation using the dominated convergence theorem, we shall need the following result.

Lemma 3.1.1. *For any $\mu^{c,x}$ -almost surely bounded stopping time τ , for $\varphi \in L^\infty(\bar{D})$, and for any nonzero sequence $(\delta_n)_n \subset (-1, 1)$ satisfying $\delta_n \rightarrow 0$ as $n \rightarrow \infty$, it holds that*

$$\lim_{n \rightarrow \infty} E^{c,x} \left[\left(D_\tau^{\delta_n \varphi} - 1 \right)^2 \right] = 0 = \lim_{n \rightarrow \infty} E^{c,x} \left[\left(\frac{D_\tau^{\delta_n \varphi} - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right].$$

Proof. Let T be such that $\tau \leq T$ $\mu^{c,x}$ -almost surely. To prove the first limit, observe that

$$\begin{aligned} E^{c,x} \left[\left(D_\tau^{\delta_n \varphi} - 1 \right)^2 \right] &\leq 2E^{c,x} \left[\left(D_\tau^{\delta_n \varphi} \right)^2 + 1 \right] \\ &\leq 2E^{c,x} \left[\sup_{s \leq T} \left(D_s^{\delta_n \varphi} \right)^2 + 1 \right] \\ &\leq 8 \exp \left(\delta_n^2 \|\varphi\|_\infty^2 T \right) + 2, \end{aligned}$$

where the first inequality follows from (3.12), the second follows from the fact that $\tau \leq T$ $\mu^{c,x}$ -almost surely, and the third follows from the inequality (2.23). Since $(\delta_n)_n \subset (-1, 1)$, it follows that the sequence $((D_\tau^{\delta_n \varphi} - 1)^2)_n$ is bounded in $L^1(\mu^{c,x})$ by $8 \exp(\|\varphi\|_\infty^2 T) + 2$. Therefore,

$$\lim_{n \rightarrow \infty} E^{c,x} \left[\left(D_\tau^{\delta_n \varphi} - 1 \right)^2 \right] = E^{c,x} \left[\lim_{n \rightarrow \infty} \left(D_\tau^{\delta_n \varphi} - 1 \right)^2 \right] = 0,$$

by Lebesgue's dominated convergence theorem and by the identity (3.13).

To prove the second limit, it follows from applying (3.13) with $\delta = \delta_n$ that, for any n ,

$$\left(D_\tau^{\delta_n \varphi} - 1 \right)^2 = \left(D_\tau^{\delta_n \varphi} \right)^2 - 2D_\tau^{\delta_n \varphi} + 1 = \left(M_\tau^{\delta_n \varphi} \right)^2 + R_n, \quad (3.15)$$

where we have omitted the dependence on ω , and $R_n = O(\delta_n^3)$. Then

$$E^{c,x} \left[\left(\frac{D_\tau^{\delta_n \varphi} - 1}{\delta_n} \right)^2 \right] = E^{c,x} \left[\langle M^\varphi \rangle_\tau + \delta_n^{-2} R_n \right] \leq \|\varphi\|_\infty^2 T + E^{c,x} \left[\delta_n^{-2} R_n \right],$$

using the Itô isometry for the equation, and using the fact that φ is bounded and that $\tau \leq T$ $\mu^{c,x}$ -almost surely for the inequality. Since $(\delta_n)_n \subset (-1, 1)$ and since $E^{c,x}[\delta_n^{-2} R_n] = O(\delta_n)$, we may define the constant

$$C = \sup_n E^{c,x}[\delta_n^{-2} R_n]$$

that depends on the sequence $(\delta_n)_n$ but not on n . Thus, the sequence $(\delta_n^{-2}(D^{\delta_n\varphi} - 1)^2)_n$ is bounded in $L^1(\mu^{c,x})$ by $\|\varphi\|_\infty^2 T + C$. Since (3.12) implies that

$$E^{c,x} \left[\left(\frac{D_\tau^{\delta_n\varphi} - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] \leq 2E^{c,x} \left[\left(\frac{D_\tau^{\delta_n\varphi} - 1}{\delta_n} \right)^2 + (M_\tau^\varphi)^2 \right],$$

it follows that the sequence $((\delta_n^{-2}(D^{\delta_n\varphi} - 1) - M_\tau^\varphi)^2)_n$ is bounded in $L^1(\mu^{c,x})$. Therefore,

$$\lim_{n \rightarrow \infty} E^{c,x} \left[\left(\frac{D_\tau^{\delta_n\varphi} - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] = E^{c,x} \left[\lim_{n \rightarrow \infty} \left(\frac{D_\tau^{\delta_n\varphi} - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] = 0,$$

by Lebesgue's dominated convergence theorem and by the identity (3.14). \square

We now proceed to obtaining expressions for the first variation and mixed second variation (i.e. the first variation of the first variation) of $\bar{\phi}^{\sigma,x}$.

Lemma 3.1.2. *Let $\delta \in \mathbb{R}$ and $\varphi \in L^\infty(\bar{D})$ be such that*

$$\bar{\phi}^{\sigma,x}(c + \delta\varphi) = E^{c+\delta\varphi,x} \left[\left(W + (2\sigma)^{-1} \langle M^{c+\delta\varphi} \rangle_\tau \right) \right] \quad (3.16)$$

exists and is finite. Then

$$\begin{aligned} & \bar{\phi}^{\sigma,x}(c + \delta\varphi) - \bar{\phi}^{\sigma,x}(c) \\ &= E^{c,x} \left[K^{\sigma,c,0} \left(D_\tau^{\delta\varphi} - 1 \right) + \frac{\delta}{\sigma} \left(\frac{\delta}{2} \langle M^\varphi \rangle_\tau + \langle M^c, M^\varphi \rangle_\tau \right) D_\tau^{\delta\varphi} \right]. \end{aligned} \quad (3.17)$$

Proof. Since $\varphi \in L^\infty(\bar{D})$, it follows that $\mu^{c+\delta\varphi}$ and μ^c are locally equivalent, by the change of measure theorem, and hence the relevant Radon-Nikodym derivatives exist and are exponential martingales. We rewrite (3.16) as an expectation with respect to $\mu^{c,x}$. By (3.7),

$$\langle M^{c+\delta\varphi} \rangle_\tau = \langle M^c \rangle_\tau + \delta^2 \langle M^\varphi \rangle_\tau + 2\delta \langle M^c, M^\varphi \rangle_\tau. \quad (3.18)$$

Substituting (3.18) in (3.16), and applying (3.11), we obtain

$$\bar{\phi}^{\sigma,x}(c + \delta\varphi) = E^{c,x} \left[\left(W + \frac{\langle M^c \rangle_\tau}{2\sigma} + \frac{\delta^2}{2\sigma} \langle M^\varphi \rangle_\tau + \frac{\delta}{\sigma} \langle M^c, M^\varphi \rangle_\tau \right) D_\tau^{\delta\varphi} \right].$$

Taking differences yields (3.17). \square

As before, let $(\delta_n)_n \subset (-1, 1)$ be an arbitrary, nonzero sequence such that $\delta_n \rightarrow 0$. Define

$$h_n := \delta_n^{-1} K^{\sigma,c,0} (D_\tau^{\delta_n\varphi} - 1) + \sigma^{-1} \left(\langle M^c, M^\varphi \rangle_\tau + \delta_n 2^{-1} \langle M^\varphi \rangle_\tau \right) D_\tau^{\delta_n\varphi} \quad (3.19)$$

$$h := K^{\sigma,c,0} M_\tau^\varphi + \sigma^{-1} \langle M^c, M^\varphi \rangle_\tau, \quad (3.20)$$

Note that h_n is defined so that

$$E^{c,x}[h_n] = \frac{\bar{\phi}^{c,x}(c + \delta_n\varphi) - \bar{\phi}^{\sigma,x}(c)}{\delta_n}.$$

We shall see that $E^{c,x}[h]$ coincides with the first variation of $\bar{\phi}^{\sigma,x}(c)$ in the direction φ . First, we need to establish convergence of h_n to h in $L^1(\mu^{c,x})$, using Lemma 3.1.1.

Proposition 3.1.3. *Let $c, \varphi \in L^\infty(\bar{D})$, $T > 0$ be a fixed number such that $\tau \leq T$ $\mu^{c,x}$ -almost surely, and $W \in L^2(\mu^{c,x})$. Then the sequence h_n converges to h in $L^1(\mu^{c,x})$.*

Proof. The assumptions on c , φ and τ imply that $\langle M^c \rangle_\tau$ and $\langle M^\varphi \rangle_\tau$ are bounded $\mu^{c,x}$ -almost surely. Furthermore, the almost-sure boundedness of τ guarantees that $D_\tau^{\delta_n \varphi} \in L^2(\mu^{c,x})$ for any $\delta \in \mathbb{R}$, by Corollary 2.1.16. The proof proceeds in two steps: showing that $\{h\} \cup \{h_n; n \in \mathbb{N}\} \subset L^1(\mu^{c,x})$, and then showing convergence of h_n to h .

Integrability of h_n and h . By the triangle inequality,

$$\begin{aligned} E^{c,x} [|h_n|] &\leq |\delta_n|^{-1} E^{c,x} \left[\left| K^{\sigma,c,0} D_\tau^{\delta_n \varphi} \right| + \left| K^{\sigma,c,0} \right| \right] \\ &\quad + \sigma^{-1} E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau \right| D_\tau^{\delta_n \varphi} \right] + |\delta_n| (2\sigma)^{-1} E^{c,x} \left[\langle M^\varphi \rangle_\tau D_\tau^{\delta_n \varphi} \right] \\ E^{c,x} [|h|] &\leq E^{c,x} \left[\left| K^{\sigma,c,0} M_\tau^\varphi \right| \right] + \sigma^{-1} E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau \right| \right]. \end{aligned}$$

By the inequality (3.12) of arithmetic and geometric means and the Kunita-Watanabe inequality (2.6) for the covariance process,

$$\begin{aligned} E^{c,x} \left[\left| K^{\sigma,c,0} \right|^2 \right] &\leq 2E^{c,x} \left[W^2 + (2\sigma)^{-2} \langle M^c \rangle_\tau^2 \right] \\ E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau \right|^2 \right] &\leq E^{c,x} \left[\langle M^c \rangle_\tau^2 + \langle M^\varphi \rangle_\tau^2 \right] \end{aligned} \quad (3.21)$$

$$\begin{aligned} E^{c,x} \left[\left| K^{\sigma,c,0} D_\tau^{\delta_n \varphi} \right| \right] &\leq E^{c,x} \left[(K^{\sigma,c,0})^2 + (D_\tau^{\delta_n \varphi})^2 \right] \\ E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau D_\tau^{\delta_n \varphi} \right| \right] &\leq E^{c,x} \left[\langle M^c, M^\varphi \rangle_\tau^2 + (D_\tau^{\delta_n \varphi})^2 \right] \\ E^{c,x} \left[\langle M^\varphi \rangle_\tau D_\tau^{\delta_n \varphi} \right] &\leq E^{c,x} \left[\langle M^\varphi \rangle_\tau^2 + (D_\tau^{\delta_n \varphi})^2 \right] \\ E^{c,x} \left[\left| K^{\sigma,c,0} M_\tau^\varphi \right| \right] &\leq E^{c,x} \left[(K^{\sigma,c,0})^2 + \langle M^\varphi \rangle_\tau \right]. \end{aligned} \quad (3.22)$$

By the hypotheses of the proposition, the right-hand sides of the inequalities above are finite. Hence, we have established that $\{h\} \cup \{h_n; n \in \mathbb{N}\} \subset L^1(\mu^{c,x})$. Since $L^1(\mu^{c,x})$ is a Banach space with the norm $\|X\| := E^{c,x}[|X|]$, it follows that $h_n - h \in L^1(\mu^{c,x})$ for all $n \in \mathbb{N}$, and the limit of any convergent sequence in $L^1(\mu^{c,x})$ lies in $L^1(\mu^{c,x})$.

Convergence of h_n to h . To show convergence, we apply the triangle inequality to obtain

$$\begin{aligned} E^{c,x} [|h_n - h|] &\leq E^{c,x} \left[\left| K^{\sigma,c,0} \left(\frac{D_\tau^{\delta_n \varphi} - 1}{\delta_n} - M_\tau^\varphi \right) \right| \right] \\ &\quad + \frac{1}{\sigma} E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau (D_\tau^{\delta_n \varphi} - 1) \right| \right] + \frac{|\delta_n|}{2\sigma} E^{c,x} \left[\langle M^\varphi \rangle_\tau D_\tau^{\delta_n \varphi} \right]. \end{aligned} \quad (3.23)$$

By (3.22), the third term in the inequality (3.23) converges to zero. By the Cauchy-Schwarz inequality,

$$\begin{aligned} E^{c,x} \left[\left| K^{\sigma,c,0} \frac{D_\tau^{\delta_n \varphi} - 1}{\delta_n} - M_\tau^\varphi \right|^2 \right] &\leq E^{c,x} \left[(K^{\sigma,c,0})^2 \right] E^{c,x} \left[\left(\frac{D_\tau^{\delta_n \varphi} - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] \\ E^{c,x} \left[\left| \langle M^c, M^\varphi \rangle_\tau (D_\tau^{\delta_n \varphi} - 1) \right|^2 \right] &\leq E^{c,x} \left[\langle M^c, M^\varphi \rangle_\tau^2 \right] E^{c,x} \left[(D_\tau^{\delta_n \varphi} - 1)^2 \right]. \end{aligned}$$

By Lemma 3.1.1, the first and second terms in the inequality (3.23) converge to zero. This completes the proof. \square

Theorem 3.1.4. *Suppose that the hypotheses of Proposition 3.1.3 hold. Then the first variation of $\bar{\phi}^{\sigma,x}(c)$ in the direction of φ is given by the functional*

$$\Phi^{\sigma,x}(c; \varphi) = E^{c,x} \left[K^{\sigma,c,0} M_\tau^\varphi + \sigma^{-1} \langle M^c, M^\varphi \rangle_\tau \right]. \quad (3.24)$$

Proof. The random variables h_n and h defined in (3.19) and (3.20) satisfy

$$E^{c,x}[h_n] = \frac{\bar{\phi}^{\sigma,x}(c + \delta_n \varphi) - \bar{\phi}^{\sigma,x}(c)}{\delta_n}$$

and $E^{c,x}[h] = \Phi^{\sigma,x}(c, \varphi)$. Thus, if we can show

$$\lim_{n \rightarrow \infty} E^{c,x}[h_n - h] = 0,$$

then we will have proved the theorem. Since

$$|E^{c,x}[h - h_n]| \leq E^{c,x}[|h - h_n|], \quad (3.25)$$

the conclusion follows from Proposition 3.1.3. \square

Remark 6. On the left-hand side of (3.24), the term $\Phi^{\sigma,x}(c; \varphi)$ is the directional derivative of the first variation of $\phi^{\sigma,x}(c)$ in the direction φ , and one can show (see Proposition 3.2.1 below) that the right-hand side can be recast as the expectation of $K^{\sigma,c,1}$ with M_τ^φ . Thus the formula (3.24) is reminiscent of the duality relationship or integration-by-parts formula in the Malliavin calculus (see, e.g. [31, Proposition 3.1])

$$E[\langle DX, h \rangle] = E[XW(h)],$$

where $\langle DX, h \rangle$ denotes the inner product in a Hilbert space of the Malliavin derivative of a random variable X satisfying certain conditions, h is an arbitrary element of a real, separable Hilbert space H of L^2 processes defined on some parameter set, and $W(h)$ denotes the Wiener integral of h ,

$$W(h) = \int_0^\infty h_s dB_s.$$

Note that $W(h)$ may also be written in terms of the divergence operator or Skorokhod integral δ associated to the Malliavin derivative operator D , as $W(h) = \delta(h)$ (Property P3 and P4 on [31, pp. 395-396]). In particular, both the left-hand sides of (3.24) and the integration by parts formula above involve directional derivatives. However, some observations show that the relation (3.24) does not have the duality relationship interpretation, at least not in the Malliavin calculus setting. One observation is that the left-hand side of (3.24) involves the derivative with respect to an infinitesimal change in the drift, while in general the Malliavin derivative is taken with respect to the random path ω . Another observation is that, in the Malliavin calculus setting, the random variable X must be *smooth* in the sense that there exists a $n \in \mathbb{N}$, a collection $(h_i)_{1 \leq i \leq n} \subset H$ and a function $f \in C_p^\infty(\mathbb{R}^n)$ such that

$$X = f(W(h_1), \dots, W(h_n)).$$

In (3.24), the random variable $K^{\sigma,c,0}$ is not smooth: since neither the running cost function f nor the terminal cost function g are assumed to be smooth, the work path functional W is not smooth. The final observation follows from Proposition 3.2.1 below, since the right-hand side is equal to the expectation of $K^{\sigma,c,1}$ with M_τ^φ , and since $K^{\sigma,c,0} = K^{\sigma,c,0}$ almost surely if and only if $c \equiv 0$.

We shall now consider the first variation $\Phi^{\sigma,x}(c; \varphi)$ defined in Theorem 3.1.4 as a functional that is parametrised by a function φ and has the control function c as its argument. We shall compute the first variation of $\Phi(c; \varphi)$ in the direction of a perturbing function ζ .

Lemma 3.1.5. *Let $\delta \in \mathbb{R}$ and $\zeta \in L^\infty(\overline{D})$ be such that*

$$\Phi^{\sigma,x}(c + \delta\zeta; \varphi) = E^{c+\delta\zeta,x} \left[K^{\sigma,c+\delta\zeta,0} M_\tau^\varphi + \sigma^{-1} \langle M^{c+\delta\zeta}, M^\varphi \rangle_\tau \right] \quad (3.26)$$

exists and is finite. Then

$$\begin{aligned} \Phi^{\sigma,x}(c + \delta\zeta; \varphi) - \Phi^{\sigma,x}(c; \varphi) &= E^{c,x} \left[\left(K^{\sigma,c,0} M_\tau^\varphi + \frac{1}{\sigma} \langle M^c, M^\varphi \rangle_\tau \right) (D^{\delta\zeta} - 1) \right] \\ &\quad + E^{c,x} \left[\frac{\delta}{\sigma} \left(\langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi + \langle M^\varphi, M^\zeta \rangle_\tau + \frac{\delta}{2} \langle M^\zeta \rangle_\tau \right) D_\tau^{\delta\zeta} \right]. \end{aligned} \quad (3.27)$$

The proof of the statement proceeds exactly as the proof of Lemma 3.1.2, relying on the reweighting formula (3.11) and the expansion (3.18) of the quadratic variation of the perturbed, controlled process.

Proof. By (3.18), we obtain an expression analogous to (3.17):

$$E^{c+\delta\zeta,x} \left[K^{\sigma,c+\delta\zeta,0} M_\tau^\varphi \right] = E^{c+\delta\zeta,x} \left[\left(K^{\sigma,c,0} + \frac{\delta}{\sigma} \left(\frac{\delta}{2} \langle M^\zeta \rangle_\tau + \langle M^c, M^\zeta \rangle_\tau \right) \right) M_\tau^\varphi \right].$$

Since $\zeta \in L^\infty(\overline{D})$, it follows that $\mu^{c+\delta\zeta}$ and μ^c are locally equivalent, by the change of measure theorem for $\delta \in \mathbb{R}$, and hence the relevant Radon-Nikodym derivatives exist and are exponential martingales. Using the reweighting formula (3.11), the bilinearity (3.7) of the quadratic variation, and the expression (3.24) for the first variation, we obtain (3.27). \square

Let $(\delta_n)_n \subset (-1, 1)$ be an arbitrary nonzero sequence such that $\delta_n \rightarrow 0$.

$$\begin{aligned} h_n &:= \left(K^{\sigma,c,0} M_\tau^\varphi + \frac{1}{\sigma} \langle M^c, M^\varphi \rangle_\tau \right) \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} \\ &\quad + \frac{1}{\sigma} \left(\langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi + \langle M^\varphi, M^\zeta \rangle_\tau + \frac{\delta_n}{2} \langle M^\zeta \rangle_\tau \right) D_\tau^{\delta_n \zeta} \end{aligned} \quad (3.28)$$

$$h := K^{\sigma,c,0} M_\tau^\varphi M_\tau^\zeta + \frac{1}{\sigma} \left(\langle M^\varphi, M^\zeta \rangle_\tau + M_\tau^\zeta \langle M^c, M^\varphi \rangle_\tau + M_\tau^\varphi \langle M^c, M^\zeta \rangle_\tau \right). \quad (3.29)$$

As before, h_n has been defined so that its expectation with respect to $\mu^{c,x}$ coincides with the difference $\delta_n^{-1}(\Phi^{\sigma,x}(c + \delta\zeta; \varphi) - \Phi^{\sigma,x}(c; \varphi))$, and h has been defined so that its expectation with respect to $\mu^{c,x}$ coincides with the desired first variation of $\Phi^{\sigma,x}(c; \varphi)$ in the direction of ζ . We now prove convergence in $L^1(\mu^{c,x})$ of h_n to h .

Proposition 3.1.6. *Let $c, \varphi, \zeta \in L^\infty(\overline{D})$, $T > 0$ be a fixed number such that $\tau \leq T$ $\mu^{c,x}$ -almost surely, and $W^2 \in L^p(\mu^{c,x})$ for some $p > 1$. Then, the sequence of random variables h_n defined in (3.28) converges to the random variable h defined in (3.29) in $L^1(\mu^{c,x})$.*

Note that in the assumptions of Proposition 3.1.6, which we shall use to obtain the mixed second-order variation, we require that W is more than square-integrable, whereas in Proposition 3.1.3, which we used to obtain the first variation, we only required that W be square integrable.

Proof. The assumptions on c , φ , ζ , and τ imply that $\langle M^c \rangle_\tau$, $\langle M^\varphi \rangle_\tau$, and $\langle M^\zeta \rangle_\tau$ are bounded $\mu^{c,x}$ -almost surely. Furthermore, the almost-sure boundedness of τ guarantees that $D_\tau^{\delta_n \zeta} \in L^2(\mu^{c,x})$, by Corollary 2.1.16. By Lemma 2.1.6, we have

$$E^{c,x} \left[(M_\tau^\varphi)^4 \right] \leq \sqrt{\frac{3}{2} + \sqrt{\frac{3}{2}}} E^{c,x} \left[\langle M^\varphi \rangle_\tau^2 \right],$$

and so $M_\tau^\varphi \in L^4(\mu^{c,x})$. Similarly, $M_\tau^\zeta \in L^4(\mu^{c,x})$.

The proof follows the same ideas as the proof of Proposition 3.1.3: we first show that $\{h\} \cup \{h_n; n \in \mathbb{N}\} \subset L^1(\mu^{c,x})$, using the triangle inequality, the inequality (3.12) of arithmetic and geometric means, and the Kunita-Watanabe inequality (2.6), and then show that $E^{c,x}[|h_n - h|]$ is bounded by a quantity that decreases to zero as $n \rightarrow \infty$.

Integrability of h_n and h . We begin first with h_n . Observe that

$$\begin{aligned} & E^{c,x} \left[\left(K^{\sigma,c,0} M_\tau^\varphi + \frac{1}{\sigma} \langle M^c, M^\varphi \rangle_\tau \right) \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} \right] \\ & \leq E^{c,x} \left[\left| K^{c,0} M_\tau^\varphi \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} \right| + \sigma^{-1} \langle M^c, M^\varphi \rangle_\tau \left| \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} \right| \right] \\ & \leq E^{c,x} \left[|K^{\sigma,c,0} M_\tau^\varphi|^2 + 2 \left| \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} \right|^2 + \sigma^{-2} \langle M^c, M^\varphi \rangle_\tau^2 \right]. \end{aligned}$$

We showed that $|\delta_n^{-1}(D_\tau^{\delta_n \zeta} - 1)|^2 \in L^1(\mu^{c,x})$ in Lemma 3.1.1. We need to show that the other terms on the right-hand side of the last inequality above are integrable. The integrability of $\langle M^c, M^\varphi \rangle_\tau^2$ follows from (3.21). To show that $K^{\sigma,c,0} M_\tau^\varphi \in L^2(\mu^{c,x})$, observe that by Holder's inequality (with $q := (1 - p^{-1})^{-1}$), we have

$$\begin{aligned} E^{c,x} \left[|K^{\sigma,c,0} M_\tau^\varphi|^2 \right] & \leq E^{c,x} \left[\left(|W M_\tau^\varphi| + (2\sigma)^{-1} |\langle M^c \rangle_\tau M_\tau^\varphi| \right)^2 \right] \\ & \leq 2E^{c,x} \left[(W M_\tau^\varphi)^2 + (2\sigma)^{-2} \langle M^c \rangle_\tau^2 (M_\tau^\varphi)^2 \right] \\ & \leq 2 \left(E^{c,x} [W^{2p}]^{1/p} E^{c,x} \left[(M_\tau^\varphi)^{2q} \right]^{1/q} + (2\sigma)^{-2} E^{c,x} \left[\langle M^c \rangle_\tau^4 + (M_\tau^\varphi)^4 \right] \right). \end{aligned}$$

Since we assumed that $W^2 \in L^p(\mu^{c,x})$ for some $p > 1$, and since $M_\tau^\varphi \in L^{2q}(\mu^{c,x})$ follows from the application of the martingale inequality (2.12) in Lemma 2.1.6, it follows that $K^{\sigma,c,0} M_\tau^\varphi \in L^2(\mu^{c,x})$. Proceeding with the other terms in h_n , we observe that the arithmetic-geometric inequality (3.12) yields

$$\begin{aligned} & E^{c,x} \left[\left| \langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi \right| D_\tau^{\delta_n \zeta} + \left| \langle M^\varphi, M^\zeta \rangle_\tau \right| D_\tau^{\delta_n \zeta} + \frac{|\delta_n|}{2} \langle M^\zeta \rangle_\tau D_\tau^{\delta_n \zeta} \right] \\ & \leq E^{c,x} \left[\langle M^c, M^\zeta \rangle_\tau^2 (M_\tau^\varphi)^2 + \langle M^\varphi, M^\zeta \rangle_\tau^2 + \frac{\delta_n^2}{4} \langle M^\zeta \rangle_\tau^2 + 3(D_\tau^{\delta_n \zeta})^2 \right] \end{aligned} \quad (3.30)$$

Since we assumed that $\langle M^\zeta \rangle_\tau$ is $\mu^{c,x}$ -almost surely bounded, and since $\langle M^\varphi, M^\zeta \rangle_\tau^2$ is in-

tegrable by (3.21), we need to show that $\langle M^c, M^\zeta \rangle_\tau^2 (M_\tau^\varphi)^2$ is integrable. This follows from

$$\begin{aligned} E^{c,x} \left[\langle M^c, M^\zeta \rangle_\tau^2 (M_\tau^\varphi)^2 \right] &\leq E^{c,x} \left[\langle M^c, M^\zeta \rangle_\tau^4 + (M_\tau^\varphi)^4 \right] \\ &\leq E^{c,x} \left[\langle M^c \rangle_\tau^2 \langle M^\zeta \rangle_\tau^2 + (M_\tau^\varphi)^4 \right] \\ &\leq E^{c,x} \left[\langle M^c \rangle_\tau^4 + \langle M^\zeta \rangle_\tau^4 + (M_\tau^\varphi)^4 \right], \end{aligned} \quad (3.31)$$

where the first inequality follows from the arithmetic-geometric inequality (3.12), the second inequality follows from the Kunita-Watanabe inequality (2.6), and the third follows from the arithmetic-geometric inequality (3.12). This proves that $h_n \in L^1(\mu^{c,x})$.

The proof that $h \in L^1(\mu^{c,x})$ consists of two parts. The first consists of showing that the terms inside the parentheses in (3.29) are integrable with respect to $\mu^{c,x}$, and this can be shown using the inequality (3.21) with the fact that $M_\tau^\zeta, M_\tau^\varphi \in L^2(\mu^{c,x})$. The second part consists of showing that $K^{\sigma,c,0} M_\tau^\varphi M_\tau^\zeta \in L^1(\mu^{c,x})$, but this follows from the arithmetic-geometric inequality and the fact that $K^{\sigma,c,0} M_\tau^\varphi \in L^2(\mu^{c,x})$ shown earlier.

Convergence of h_n to h . By grouping together similar terms and applying the triangle inequality, we obtain

$$\begin{aligned} E^{c,x} [|h_n - h|] &\leq E^{c,x} \left[\left(|K^{\sigma,c,0} M_\tau^\varphi| + \sigma^{-1} |\langle M^c, M^\varphi \rangle_\tau| \right) \left| \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} - M_\tau^\zeta \right| \right] \\ &\quad + \frac{1}{\sigma} E^{c,x} \left[\left(|\langle M^\varphi, M^\zeta \rangle_\tau| + |\langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi| \right) |D_\tau^{\delta_n \zeta} - 1| \right] \\ &\quad + \frac{|\delta_n|}{2\sigma} E^{c,x} \left[\langle M^\zeta \rangle_\tau D_\tau^{\delta_n \zeta} \right]. \end{aligned} \quad (3.32)$$

In order to use the convergence results of Lemma 3.1.1, we need to verify that the terms on the right-hand side of (3.32) converge to zero, after using the Cauchy-Schwarz inequality, and after checking the relevant terms are sufficiently integrable. For example, for the first row, we obtain the inequality

$$\begin{aligned} &E^{c,x} \left[\left(|K^{\sigma,c,0} M_\tau^\varphi| + \sigma^{-1} |\langle M^c, M^\varphi \rangle_\tau| \right) \left| \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} - M_\tau^\zeta \right| \right] \\ &\leq E^{c,x} \left[\left(|K^{\sigma,c,0} M_\tau^\varphi| + \sigma^{-1} |\langle M^c, M^\varphi \rangle_\tau| \right)^2 \right]^{1/2} E^{c,x} \left[\left| \frac{D_\tau^{\delta_n \zeta} - 1}{\delta_n} - M_\tau^\zeta \right|^2 \right]. \end{aligned}$$

We showed that $K^{\sigma,c,0} M_\tau^\varphi \in L^2(\mu^{c,x})$ above. The other terms on the right-hand side of the inequality above have been shown to exhibit the required degree of integrability in the proof of Lemma 3.1.1 and in (3.21). The integrability and convergence of the remaining terms on the right-hand side of (3.32) follow from the integrability statements derived earlier and Lemma 3.1.1 respectively. \square

Theorem 3.1.7. *Suppose that the hypotheses of Proposition 3.1.6 hold. Then the first variation of $\bar{\phi}^{\sigma,x}(c)$ in the direction of φ exists and is equal to $\Phi^{\sigma,x}(c; \varphi)$ as defined in (3.24), and the first variation of $\Phi^{\sigma,x}(c; \varphi)$ in the direction of ζ is given by the functional*

$$\begin{aligned} \Phi^{\sigma,x}(c; \varphi, \zeta) &:= E^{c,x} \left[K^{\sigma,c,0} M_\tau^\varphi M_\tau^\zeta \right] \\ &\quad + E^{c,x} \left[\sigma^{-1} \left(\langle M^c, M^\varphi \rangle_\tau M_\tau^\zeta + \langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi + \langle M^\varphi, M^\zeta \rangle_\tau \right) \right]. \end{aligned} \quad (3.33)$$

Note that the functional $\Phi^{\sigma,x}(c; \cdot, \cdot)$ is symmetric in the parameters, i.e. $\Phi^{\sigma,x}(c; \varphi, \zeta) = \Phi^{\sigma,x}(c; \zeta, \varphi)$.

Proof. The random variables h_n and h defined in (3.28) and (3.29) satisfy

$$E^{c,x}[h_n] = \frac{\Phi^{\sigma,x}(c + \delta_n \zeta; \varphi) - \Phi^{\sigma,x}(c; \varphi)}{\delta_n}$$

and $E^{c,x}[h] = \Phi^{\sigma,x}(c; \varphi, \zeta)$. By (3.25) and by Proposition 3.1.6, it holds that $E^{c,x}[h_n - h] \rightarrow 0$, and hence the theorem is proved. \square

By setting $\zeta = \varphi$ in Theorem 3.1.7, we obtain the following expression for the second variation of $\bar{\phi}^{\sigma,x}(c)$ in the direction of φ :

$$\Phi^{\sigma,x}(c; \varphi, \varphi) = E^{c,x} \left[K^{\sigma,c,0}(M_\tau^\varphi)^2 + \sigma^{-1} (\langle M^\varphi \rangle_\tau + 2\langle M^c, M^\varphi \rangle_\tau M_\tau^\varphi) \right]. \quad (3.34)$$

Recall that the second variation of a functional at some point in its domain possesses the strong positive definiteness property over some set A if there exists some $m > 0$ such that

$$\Phi^{\sigma,x}(c; \varphi, \varphi) \geq m \|\varphi\|^2, \quad \forall \varphi \in A, \quad (3.35)$$

where m does not depend on φ . Note that the property of being strongly positive definite must be specified with respect to some norm on the set of admissible perturbations. A suitable norm might be

$$\|\varphi\|_{\mu^{c,x},2}^2 := E^x \left[\int_0^\tau |\varphi(X_s^c)|^2 ds \right] = E^{c,x}[\langle M^\varphi \rangle_\tau].$$

However, without specifying additional conditions on $K^{\sigma,c,0}$, c and φ , it is not immediately obvious whether the second variation of $\bar{\phi}^{\sigma,x}$ is a strongly positive definite operator on a given set of admissible perturbations. In the next section, we state some conditions for which strong positive definiteness holds for the second variation.

To conclude this section, we briefly discuss the integrability assumptions of Propositions 3.1.3 and 3.1.6. In order for the mixed second variation to be well-defined, we assumed that the control c , the perturbing functions φ and ζ , and the stopping time τ are bounded. These strong assumptions simplified matters greatly, since we could then use Corollary 2.1.16 to assert that the stopped Radon-Nikodym derivative terms $D_\tau^{\delta_n \varphi}$ were square-integrable. We could also use that the stopped quadratic variation terms were also bounded, which implied integrability statements on powers of stopped martingale terms via the martingale inequalities in Lemma 2.1.6. If one were to weaken the almost sure boundedness assumption on τ , the first task would be to ensure that the stopped Radon-Nikodym derivative terms $D_\tau^{\delta_n \varphi}$ are sufficiently integrable in order to obtain convergence statements of the kind described in Lemma 3.1.1.

3.2 Relations deriving from Itô's formula

In this section, we use Itô's formula (2.10) in order to derive useful expressions involving the martingale terms that appeared in the first variation functionals defined in (3.24) and (3.24). The most important results of this section are contained in §3.2.2, in which we use pairwise independence to obtain expressions with which to simplify the mixed second-order variations of the control functional $\bar{\phi}^{\sigma,x}$.

3.2.1 Martingales constructed from bounded basis functions

In this section, we derive some expressions involving products of the values of continuous local martingales at some stopping time τ , where the local martingales are Itô integrals of arbitrary bounded integrand processes with respect to the Brownian motion B . We use these expressions to reformulate the expression for the mixed second variation of the control functional $\bar{\phi}^{\sigma,x}$.

Proposition 3.2.1. *The first variation $\Phi^{\sigma,x}(c; \varphi)$ defined in (3.24) satisfies*

$$\Phi^{\sigma,x}(c; \varphi) = E^{c,x} [K^{c,1} M_\tau^\varphi]. \quad (3.36)$$

Proof. By the definition (3.5) of the random variable $K^{c,\alpha}$, and by the property (3.8) of the covariance process,

$$E^{c,x} [K^{\sigma,c,0} M_\tau^\varphi + \sigma^{-1} \langle M^c, M^\varphi \rangle_\tau] = E^{c,x} [K^{\sigma,c,0} M_\tau^\varphi + \sigma^{-1} M_\tau^c M_\tau^\varphi] = E^{c,x} [K^{c,1} M_\tau^\varphi] \quad \square$$

Recall that the optimal control c^* has the property that $K^{c^*,1}$ is $\mu^{c^*,x}$ -almost surely constant. By Lemma 3.2.1, and the mean-zero property of M_τ^φ , this implies that c^* is a critical point of the functional $\bar{\phi}^{\sigma,x}$, as expected.

Lemma 3.2.2. *Suppose that $\varphi_i \in L^\infty(\bar{D})$ for $i = 1, 2, 3$. Then*

$$E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} M_\tau^{\varphi_3}] = E^{c,x} [M_\tau^{\varphi_1} \langle M^{\varphi_2}, M^{\varphi_3} \rangle_\tau + M_\tau^{\varphi_2} \langle M^{\varphi_3}, M^{\varphi_1} \rangle_\tau + M_\tau^{\varphi_3} \langle M^{\varphi_1}, M^{\varphi_2} \rangle_\tau]. \quad (3.37)$$

Proof. As was observed in Remark 5, the assumption that $\varphi_i \in L^\infty(\bar{D})$ implies that M^{φ_i} is a well-defined continuous local martingale with almost-surely finite quadratic variation for all $t \geq 0$. By Itô's formula, it holds that

$$\begin{aligned} E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} M_\tau^{\varphi_3}] &= E^{c,x} \left[\int_0^\tau M_s^{\varphi_1} d\langle M^{\varphi_2}, M^{\varphi_3} \rangle_s \right. \\ &\quad \left. + \int_0^\tau M_s^{\varphi_2} d\langle M^{\varphi_3}, M^{\varphi_1} \rangle_s + \int_0^\tau M_s^{\varphi_3} d\langle M^{\varphi_1}, M^{\varphi_2} \rangle_s \right]. \end{aligned}$$

On the other hand, using the integration by parts formula (2.9), it holds that

$$E^{c,x} [M_\tau^{\varphi_i} \langle M^{\varphi_j}, M^{\varphi_k} \rangle_\tau] = E^{c,x} \left[\int_0^\tau M_s^{\varphi_i} d\langle M^{\varphi_j}, M^{\varphi_k} \rangle_s \right]. \quad \square$$

Corollary 3.2.3. *For $\varphi \in L^\infty(\bar{D})$,*

$$E^{c,x} [(M_\tau^\varphi)^3] = 3E^{c,x} [M_\tau^\varphi \langle M^\varphi \rangle_\tau].$$

Recall that the covariance process of two continuous local martingales on a filtered probability space is defined precisely so that

$$\bar{M}_t^{\varphi_1, \varphi_2} := M_t^{\varphi_1} M_t^{\varphi_2} - \langle M^{\varphi_1}, M^{\varphi_2} \rangle_t, \quad (3.38)$$

defines a continuous local martingale on the same filtered probability space. Rearranging (3.37) and applying the definition (3.38) yields

Corollary 3.2.4. *It holds that*

$$E^{c,x} \left[\langle \overline{M}^{\varphi_1, \varphi_2}, M^{\varphi_3} \rangle_\tau \right] = E^{c,x} \left[M_\tau^{\varphi_1} \langle M^{\varphi_3}, M^{\varphi_2} \rangle_\tau + M_\tau^{\varphi_2} \langle M^{\varphi_3}, M^{\varphi_1} \rangle_\tau \right]. \quad (3.39)$$

The significance of Corollary 3.2.4 is that (3.39) provides an alternative characterisation of the martingale $\overline{M}^{\varphi_1, \varphi_2}$ in addition to the definition (3.38). Furthermore, since one in general has better control over processes that are locally of bounded variation, it is useful to know that the stochastic process inside the expectation on the right-hand side of (3.39) - which is a sum of products of continuous local martingales with processes of locally bounded variation - can be treated as a process that is locally of bounded variation. Corollary 3.2.4 also yields the following

Proposition 3.2.5. *The mixed second variation $\Phi^{\sigma,x}(c; \varphi, \zeta)$ described in (3.24) satisfies*

$$\Phi^{\sigma,x}(c; \varphi, \zeta) = E^{c,x} \left[K^{\sigma,c,1} M_\tau^\varphi M_\tau^\zeta + \sigma^{-1} (1 - M_\tau^c) \langle M^\varphi, M^\zeta \rangle_\tau \right]. \quad (3.40)$$

Proof. Setting $\varphi_1 = \varphi$, $\varphi_2 = \zeta$, and $\varphi_3 = c$ in (3.39), and using (3.38), we obtain

$$\begin{aligned} & E^{c,x} \left[\langle M^c, M^\varphi \rangle_\tau M_\tau^\zeta + \langle M^c, M^\zeta \rangle_\tau M_\tau^\varphi + \langle M^\varphi, M^\zeta \rangle_\tau \right] \\ &= E^{c,x} \left[\langle M^c, \overline{M}^{\varphi, \zeta} \rangle_\tau + \langle M^\varphi, M^\zeta \rangle_\tau \right] \\ &= E^{c,x} \left[M_\tau^c \overline{M}_\tau^{\varphi, \zeta} + \langle M^\varphi, M^\zeta \rangle_\tau \right] \\ &= E^{c,x} \left[M_\tau^c (M_\tau^\varphi M_\tau^\zeta - \langle M^\varphi, M^\zeta \rangle_\tau) + \langle M^\varphi, M^\zeta \rangle_\tau \right] \\ &= E^{c,x} \left[M_\tau^c M_\tau^\varphi M_\tau^\zeta + (1 - M_\tau^c) \langle M^\varphi, M^\zeta \rangle_\tau \right] \end{aligned} \quad (3.41)$$

The conclusion (3.40) follows by substituting (3.41) into the definition (3.33) of $\Phi^{\sigma,x}(c; \varphi, \zeta)$. \square

Corollary 3.2.6. *It holds that the second variation of $\overline{\phi}^{\sigma,x}$ in the direction of φ is given by*

$$\Phi^{\sigma,x}(c; \varphi, \varphi) = E^{c,x} \left[K^{\sigma,c,1} (M_\tau^\varphi)^2 + \sigma^{-1} (1 - M_\tau^c) \langle M^\varphi \rangle_\tau \right].$$

Because $\Phi^{\sigma,x}(c; \varphi, \varphi)$ depends on c through the path measure $\mu^{c,x}$, the quadratic variation term $\langle M^c \rangle_\tau$ in $K^{\sigma,c,0}$, and the martingale M_τ^c , it is not clear whether the second variation of $\overline{\phi}^{\sigma,x}$ is positive definite. This suggests that additional conditions must be imposed on c and φ in order to guarantee that $\Phi^{\sigma,x}(c; \cdot, \cdot)$ is a positive definite operator.

The final result in this subsection concerns products of four martingale terms. The need for this result is due to the presence of such products in the second variation (e.g. $\varphi = c$).

Lemma 3.2.7. *Let $\varphi_i \in L^\infty(\overline{D})$ for $1 \leq i \leq 4$. Then*

$$E^{c,x} \left[M_\tau^{\varphi_1} M_\tau^{\varphi_2} M_\tau^{\varphi_3} M_\tau^{\varphi_4} \right] = \sum_{i=1}^3 \sum_{i < j} E^{c,x} \left[\int_0^\tau \prod_{k \neq i, j} M_s^{\varphi_k} d \langle M^{\varphi_i}, M^{\varphi_j} \rangle_s \right]. \quad (3.42)$$

and

$$\begin{aligned} E^{c,x}[M_\tau^{\varphi_1} M_\tau^{\varphi_2} \langle M^{\varphi_3}, M^{\varphi_4} \rangle_\tau] &= E^{c,x} \left[\int_0^\tau M_s^{\varphi_1} M_s^{\varphi_2} d\langle M^{\varphi_3}, M^{\varphi_4} \rangle_s \right] \\ &+ E^{c,x} \left[\int_0^\tau \langle M^{\varphi_3}, M^{\varphi_4} \rangle_s d\langle M^{\varphi_1}, M^{\varphi_2} \rangle_s \right]. \end{aligned} \quad (3.43)$$

Proof. By Itô's formula,

$$\prod_{i=1}^4 M_\tau^{\varphi_i} = \sum_{i=1}^4 \left(\int_0^\tau \prod_{j \neq i} M_s^{\varphi_j} dM_s^{\varphi_i} + \sum_{j \neq i} \frac{1}{2} \int_0^\tau M_s^{\varphi_k} M_s^{\varphi_\ell} d\langle M^{\varphi_i}, M^{\varphi_j} \rangle_s \right)$$

where $\{k, \ell\}$ and $\{i, j\}$ constitute a disjoint partition of $\{1, 2, 3, 4\}$. Taking expectations of the equation above, using that Itô integrals with respect to continuous local martingales are again continuous local martingales, and taking into account that certain summands appear twice, we obtain (3.42). By the same reasoning, we obtain (3.43). \square

Although one could use Itô's formula to derive equations for expected values of five or more random variables, we shall only consider expected values of products of at most four random variables. This is because the expression (3.33) contains products of at most four martingale terms, if one associates quadratic variation terms with squares of martingale terms.

3.2.2 Pairwise independence via non-overlap condition on supports

In the previous section, we derived some identities, e.g. Corollary 3.2.3 and Lemma 3.2.7, involving the martingale terms that appear in the expressions (3.36) and (3.40) for the first variation $\Phi^{\sigma,x}(c; \varphi)$ and the mixed second variation $\Phi^{\sigma,x}(c; \varphi, \zeta)$. In this section, we construct finite collections of these martingale terms, and use an important assumption (Assumption 3.2.8) to establish certain independence relationships between the martingales. These independence relationships will prove crucial in the final section §3.3.1 of this chapter, in which we specify an approximating subset \mathcal{A} of feedback control functions such that the control functional $\bar{\phi}^{\sigma,x}$ is strictly convex over \mathcal{A} .

Recall that $\{\varphi_i\}_{1 \leq i \leq n} \subset L^\infty(\bar{D})$. Denote the support of φ_i by

$$S_i := \text{supp}(\varphi_i) = \{x \mid \varphi_i(x) \neq 0\} \subset \mathbb{R}^d. \quad (3.44)$$

Let \bar{S}_i denote the closure of S_i . The results we shall show next derive from the following

Assumption 3.2.8. *The supports $(S_i)_{i \leq n}$ are connected, open subsets of the domain D with strictly positive Lebesgue measure, whose closures form a partition of the closure of D , i.e. $\bar{D} = \cup_i \bar{S}_i$, and that satisfy the non-overlap condition*

$$\bar{S}_i \cap \bar{S}_j \subset \partial S_i \cap \partial S_j, \quad i \neq j. \quad (3.45)$$

Since the closures of the supports intersect at most in the intersection of the boundaries of the supports, and since the supports are open, we have

$$i \neq j \implies \lambda \left\{ x \in \mathbb{R}^d \mid \varphi_i(x) \cdot \varphi_j(x) \neq 0 \right\} = 0, \quad (3.46)$$

where λ denotes Lebesgue measure. The first result of this section is

Lemma 3.2.9. *If $i \neq j$, then $\mu^{c,x}$ -almost surely, $\langle M^{\varphi_i}, M^{\varphi_j} \rangle_t = 0$ for all $t \geq 0$.*

Proof. The result follows from the definition (3.6) of the covariance process, and the fact expressed in (3.46) that the integrand in the covariance process is nonzero only on a set of Lebesgue measure zero. \square

By (3.8), Lemma 3.2.9 yields

Corollary 3.2.10. *For $i \neq j$, the expectation $E^{c,x}[M_\tau^{\varphi_i} M_\tau^{\varphi_j}]$ equals zero.*

A stronger result than Corollary 3.2.10 is that $M_\tau^{\varphi_i}$ and $M_\tau^{\varphi_j}$ are independent. This is the content of Theorem 3.2.16, which we shall state later. To prove the independence of $M_\tau^{\varphi_i}$ and $M_\tau^{\varphi_j}$, we must show that the increments of M^{φ_i} are independent of the increments of M^{φ_j} . We will accomplish this by using the strong Markov property, in conjunction with Assumption 3.2.8, which is essential here. The idea is that, since M^{φ_i} remains constant whenever X^c is not in the support of φ_i , and since the intersections of the supports are of Lebesgue measure zero, the process M^{φ_i} is nonconstant over a given random time interval (σ, ρ) if and only if M^{φ_j} is constant over the same time interval, for all $j \neq i$. Since a nonconstant stochastic process and a constant stochastic process are independent, the independence of the increments of M^{φ_i} and M^{φ_j} follows. For a rigorous proof, we need to construct a family of random time intervals and then show that the increments of the martingales $(M^{\varphi_i})_i$ and the increments of quadratic variation processes $(\langle M^{\varphi_i} \rangle)_i$ form two collections of mutually and pairwise independent random variables.

Fix $i \in \{1, \dots, n\}$. Let $\sigma_{i,0} := 0$, and for $m \in \mathbb{N}$, define the time $\rho_{i,m}$ of the m -th entry into S_i and the time $\sigma_{i,m}$ of the m -th exit from S_i respectively by

$$\rho_{i,m}(\omega) := \inf \{t > \sigma_{i,m-1} \mid \omega_t \in S_i\} \quad (3.47)$$

$$\sigma_{i,m}(\omega) := \inf \{t > \rho_{i,m} \mid \omega_t \notin S_i\}. \quad (3.48)$$

The definitions (3.47) and (3.48) imply that

$$0 = \sigma_{i,0} \leq \rho_{i,1} \leq \sigma_{i,1} \leq \rho_{i,2} \leq \sigma_{i,2} \leq \dots \quad (3.49)$$

defines a sequence $(T_{i,m})_{m \in \mathbb{N}}$ of stopping times increasing to infinity, via

$$T_{i,m} := \begin{cases} \rho_{i,(m+1)/2} & m \text{ odd} \\ \sigma_{i,m/2} & m \text{ even.} \end{cases} \quad (3.50)$$

The definition yields the correspondence

$$\begin{array}{c} \rho_{i,1} \leq \sigma_{i,1} \leq \rho_{i,2} \leq \sigma_{i,2} \leq \dots \\ \updownarrow \\ T_{i,1} \leq T_{i,2} \leq T_{i,3} \leq T_{i,4} \leq \dots \end{array}$$

Lemma 3.2.11. *If Assumption 3.2.8 holds, then the countable collection*

$$\{T_{i,m} \mid 1 \leq i \leq n, m \in \mathbb{N}\}$$

of stopping times may be ordered such that for any $i, j \in \{1, \dots, n\}$ and for any $m \in \mathbb{N}$, there exists a $\ell \in \mathbb{N}$ such that

$$T_{j,2\ell+1} \leq T_{j,2\ell+2} \leq T_{i,2m+1} \leq T_{i,2m+2} \leq T_{j,2(\ell+1)+1} \leq T_{j,2(\ell+1)+2}. \quad (3.51)$$

Proof. Assumption 3.2.8 is equivalent to the condition that, almost surely, a random path ω cannot be in more than one support at any given time, except on a set of Lebesgue measure zero. Therefore, almost surely, if a path enters some support S_i , it must leave S_i before entering any other support S_j for $j \neq i$; this is the content of (3.51). \square

Since $\varphi_i \in L^\infty(\overline{D})$, it follows that, for the stopping time τ in (3.2), $(M_{t \wedge \tau}^{\varphi_i})_{t \geq 0}$ is a continuous local martingale. Hence, it follows that $(M_{\tau \wedge T_{i,n}}^{\varphi_i})_{n \in \mathbb{N}}$ is a martingale, and the increments $(\xi_{i,m})_{m \in \mathbb{N}}$ defined by

$$\xi_{i,m} := M_{\tau \wedge T_{i,m}}^{\varphi_i} - M_{\tau \wedge T_{i,m-1}}^{\varphi_i} \quad (3.52)$$

are independent. For the quadratic variation of M^{φ_i} , we define the increments by

$$\theta_{i,m} := \langle M^{\varphi_i} \rangle_{\tau \wedge T_{i,m}} - \langle M^{\varphi_i} \rangle_{\tau \wedge T_{i,m-1}}. \quad (3.53)$$

By the integral form (3.6) of the quadratic variation and covariance processes, the increments $(\theta_{i,m})_{m \in \mathbb{N}}$ are also independent.

Lemma 3.2.12. *For given i and m , the random variables $\xi_{i,m}$ and $\theta_{i,m}$ are measurable with respect to $\mathcal{F}_{\tau \wedge T_{i,m}}$ and independent of the sigma-algebra $\mathcal{F}_{\tau \wedge T_{i,m-1}}$.*

Proof. The measurability follows from the fact that $\xi_{i,m}$ depends on the history of the process X^c up to, and including, the time $\tau \wedge T_{i,m}$. The independence follows from the strong Markov property of X^c , which is guaranteed by Theorem 2.1.12. \square

By (3.49) and (3.50), we have

Lemma 3.2.13. *Let $\varphi_i \in L^\infty(\overline{D})$, and let $(T_{i,m})_{m \in \mathbb{N}}$ be defined as in (3.50). Then the sequence of increments $(\xi_{i,m})_{m \in \mathbb{N}}$ defined by (3.52) satisfies*

$$\xi_{i,m} = \begin{cases} 0 & m \text{ odd} \\ M_{\tau \wedge \sigma_{i,m/2}}^{\varphi_i} - M_{\tau \wedge \rho_{i,m/2}}^{\varphi_i} & m \text{ even} \end{cases} \quad (3.54)$$

and the sequence of increments $(\theta_{i,m})_{m \in \mathbb{N}}$ defined by (3.53) satisfies

$$\theta_{i,m} = \begin{cases} 0 & m \text{ odd} \\ \langle M^{\varphi_i} \rangle_{\tau \wedge \sigma_{i,m/2}} - \langle M^{\varphi_i} \rangle_{\tau \wedge \rho_{i,m/2}} & m \text{ even}. \end{cases} \quad (3.55)$$

Proof. The assertions follow from the definition (3.50) of $T_{i,m}$. In particular, when m is odd, then the increments are computed by taking the difference of the respective processes at the endpoints of the interval $(\tau \wedge \sigma_{i,(m-1)/2}, \tau \wedge \rho_{i,(m+1)/2})$. At every point s in the interval $(\tau \wedge \sigma_{i,(m-1)/2}, \tau \wedge \rho_{i,(m+1)/2})$, it holds that $X_s^c \notin S_i$, by the definitions of the stopping times $\rho_{i,\cdot}$ and $\sigma_{i,\cdot}$ and by the assumption that S_i is open. Thus M^{φ_i} and $\langle M^{\varphi_i} \rangle$ are constant over this interval, and the corresponding increments are zero. \square

In the following proposition, $\Delta_{i,m}$ and $\overline{\Delta}_{i,m}$ should be considered as placeholders for the increments $\xi_{i,m}$ and $\theta_{i,m}$.

Proposition 3.2.14. *Let $\varphi_i, \varphi_j \in L^\infty(\bar{D})$ be distinct functions satisfying Assumption 3.2.8, let $(T_{i,m})_{m \in \mathbb{N}}$ and $(T_{j,\ell})_{\ell \in \mathbb{N}}$ be the corresponding sequences of stopping times defined by (3.50), and let $(\Delta_{i,m})_{m \in \mathbb{N}}$ and $(\bar{\Delta}_{j,\ell})_{\ell \in \mathbb{N}}$ be two sequences of random variables, such that $\Delta_{i,m}$ is measurable with respect to $\mathcal{F}_{\tau \wedge T_{i,m}}$ and independent of $\mathcal{F}_{\tau \wedge T_{i,m-1}}$, and $\bar{\Delta}_{j,\ell}$ (not necessarily equal to $\Delta_{j,\ell}$) is measurable with respect to $\mathcal{F}_{\tau \wedge T_{j,\ell}}$ and independent of $\mathcal{F}_{\tau \wedge T_{j,\ell-1}}$. Then for any $m, \ell \in \mathbb{N}$, the increments $\Delta_{i,m}$ and $\bar{\Delta}_{j,\ell}$ are independent, and in particular*

$$E^{c,x}[\Delta_{i,m} \bar{\Delta}_{j,\ell}] = E^{c,x}[\Delta_{i,m}] E^{c,x}[\bar{\Delta}_{j,\ell}]. \quad (3.56)$$

Proof. We may assume that $m, \ell \in 2\mathbb{N}$, since otherwise the conclusion (3.56) holds trivially, by Lemma 3.2.13. Then, since Assumption 3.2.8 holds, then so does Lemma 3.2.11, and we may assume (after switching indices, if necessary) that

$$T_{j,\ell} \leq T_{i,m-1} \leq T_{i,m}. \quad (3.57)$$

Consequently,

$$\mathcal{F}_{\tau \wedge T_{j,\ell}} \subseteq \mathcal{F}_{\tau \wedge T_{i,m-1}} \subseteq \mathcal{F}_{\tau \wedge T_{i,m}}, \quad (3.58)$$

which means that the diffusion X^c exited S_j before entering S_i . This yields the independence of $\Delta_{i,m}$ and $\bar{\Delta}_{j,\ell}$, by the strong Markov property. To prove (3.56), we observe that

$$E^{c,x}[\Delta_{i,m} \bar{\Delta}_{j,\ell}] = E^{c,x}\left[E^{c,x}\left[\Delta_{i,m} \bar{\Delta}_{j,\ell} \mid \mathcal{F}_{\tau \wedge T_{j,\ell}}\right]\right] \quad (3.59)$$

$$= E^{c,x}\left[\bar{\Delta}_{j,\ell} E^{c,x}\left[\Delta_{i,m} \mid \mathcal{F}_{\tau \wedge T_{j,\ell}}\right]\right] \quad (3.60)$$

$$= E^{c,x}\left[\bar{\Delta}_{j,\ell} E^{c,x}[\Delta_{i,m}]\right]. \quad (3.61)$$

The first equation follows from the tower property of conditional expectation. The second equation follows from the assumption that $\bar{\Delta}_{j,\ell}$ is measurable with respect to $\mathcal{F}_{\tau \wedge T_{j,\ell}}$. The third equation follows from the assumption that $\Delta_{i,m}$ is independent of $\mathcal{F}_{\tau \wedge T_{i,m-1}}$ by Lemma 3.2.12, and hence independent of $\mathcal{F}_{\tau \wedge T_{j,\ell}}$, by (3.58). \square

Corollary 3.2.15. *Suppose that the sums $\sum_m \Delta_{i,m}$ and $\sum_\ell \bar{\Delta}_{j,\ell}$ have finite expectations, and that the product of the sums has finite expectation. Then*

$$E^{c,x}\left[\left(\sum_m \Delta_{i,m}\right) \left(\sum_\ell \bar{\Delta}_{j,\ell}\right)\right] = E^{c,x}\left[\sum_m \Delta_{i,m}\right] E^{c,x}\left[\sum_\ell \bar{\Delta}_{j,\ell}\right]. \quad (3.62)$$

The preceding two results yield the following result, in which we show that the martingales and the associated quadratic variation processes are mutually and pairwise independent:

Theorem 3.2.16. *Let $\varphi_i, \varphi_j \in L^\infty(\bar{D})$ be distinct functions satisfying Assumption 3.2.8. Then for $i \neq j$, $M_\tau^{\varphi_i}$ and $\langle M^{\varphi_i} \rangle_\tau$ are independent of $M_\tau^{\varphi_j}$ and $\langle M^{\varphi_j} \rangle_\tau$, and in particular*

$$E^{c,x}[M_\tau^{\varphi_i} M_\tau^{\varphi_j}] = E^{c,x}[M_\tau^{\varphi_i}] E^{c,x}[M_\tau^{\varphi_j}] \quad (3.63)$$

$$E^{c,x}[M_\tau^{\varphi_i} \langle M^{\varphi_j} \rangle_\tau] = E^{c,x}[M_\tau^{\varphi_i}] E^{c,x}[\langle M^{\varphi_j} \rangle_\tau] \quad (3.64)$$

$$E^{c,x}[\langle M^{\varphi_i} \rangle_\tau \langle M^{\varphi_j} \rangle_\tau] = E^{c,x}[\langle M^{\varphi_i} \rangle_\tau] E^{c,x}[\langle M^{\varphi_j} \rangle_\tau]. \quad (3.65)$$

Proof. For every $1 \leq i \leq n$, we have the decompositions

$$M_\tau^{\varphi_i} = \sum_{m=1} \xi_{i,m}$$

$$\langle M^{\varphi_i} \rangle_\tau = \sum_{m=1} \theta_{i,m},$$

for $\xi_{i,m}$ and $\theta_{i,m}$ as given in (3.52) and (3.53) respectively. By Lemma 3.2.12, both $\xi_{i,m}$ and $\theta_{i,m}$ satisfy the hypotheses placed on the random variables $\Delta_{i,m}$ and $\bar{\Delta}_{j,\ell}$ in Proposition 3.2.14. Therefore, the equations (3.63)–(3.65) are proved by choosing $\Delta_{i,m}$ and $\bar{\Delta}_{j,\ell}$ appropriately in Corollary 3.2.15 from $\{\xi_{i,m}, \theta_{i,m}\}$ and $\{\xi_{j,\ell}, \theta_{j,\ell}\}$ respectively. \square

To conclude this section, we revisit the relations presented in §3.2.1 in light of the independence result given in Theorem 3.2.16 above. We shall apply these relations in §3.3, where we show that the control functional $\bar{\phi}^{\sigma,x}$ exhibits the property of strong convexity when restricted to classes of control functions spanned by finitely many basis functions with non-overlapping supports.

Lemma 3.2.17. *Let $\varphi_1, \varphi_2, \varphi_3 \in L^\infty(\bar{D})$ satisfy Assumption 3.2.8. Then*

$$E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} M_\tau^{\varphi_3}] = \begin{cases} 3E^{c,x} [M_\tau^{\varphi_1} \langle M^{\varphi_1} \rangle_\tau] & \varphi_1 = \varphi_2 = \varphi_3 \\ 0 & \text{otherwise.} \end{cases} \quad (3.66)$$

Proof. By Corollary 3.2.3, it suffices to prove that if $\varphi_1 = \varphi_2 = \varphi_3$ does *not* hold, then the expression on the right-hand side of (3.37) in Lemma 3.2.2 vanishes. Without loss of generality, suppose that $\varphi_1 \neq \varphi_2 = \varphi_3$. Then by Lemma 3.2.2 and Theorem 3.2.16,

$$E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} M_\tau^{\varphi_3}] = E^{c,x} [M_\tau^{\varphi_1}] E^{c,x} [\langle M^{\varphi_2}, M^{\varphi_3} \rangle_\tau] \\ + E^{c,x} [M_\tau^{\varphi_2} \langle M^{\varphi_3}, M^{\varphi_1} \rangle_\tau + M_\tau^{\varphi_3} \langle M^{\varphi_1}, M^{\varphi_2} \rangle_\tau].$$

By Lemma 3.2.9, the terms $\langle M^{\varphi_3}, M^{\varphi_1} \rangle_\tau$ and $\langle M^{\varphi_1}, M^{\varphi_2} \rangle_\tau$ equal zero, and by the martingale property of M^φ , $E^{c,x} [M_\tau^{\varphi_1}] = 0$. \square

Lemma 3.2.18. *Let $\varphi_1, \varphi_2, \varphi_3 \in L^\infty(\bar{D})$ satisfy Assumption 3.2.8. Then*

$$E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} \langle M^{\varphi_3} \rangle_\tau] = \begin{cases} E^{c,x} [6^{-1} (M_\tau^{\varphi_1})^4 + 2^{-1} \langle M^{\varphi_1} \rangle_\tau^2] & \varphi_1 = \varphi_2 = \varphi_3 \\ E^{c,x} [(M_\tau^{\varphi_1})^2 (M_\tau^{\varphi_3})^2] & \varphi_1 = \varphi_2 \neq \varphi_3 \\ 0 & \text{otherwise.} \end{cases} \quad (3.67)$$

Proof. If $\varphi_1 \neq \varphi_2$ and $\varphi_2 \neq \varphi_3$, then by Theorem 3.2.16,

$$E^{c,x} [M_\tau^{\varphi_1} M_\tau^{\varphi_2} \langle M^{\varphi_3} \rangle_\tau] = \begin{cases} E^{c,x} [M_\tau^{\varphi_2}] E^{c,x} [M_\tau^{\varphi_1} \langle M^{\varphi_1} \rangle_\tau] & \varphi_1 = \varphi_3 \neq \varphi_2 \\ E^{c,x} [M_\tau^{\varphi_1}] E^{c,x} [M_\tau^{\varphi_2}] E^{c,x} [\langle M^{\varphi_3} \rangle_\tau] & \varphi_1 \neq \varphi_3 \neq \varphi_2 \\ E^{c,x} [M_\tau^{\varphi_1}] E^{c,x} [M_\tau^{\varphi_3} \langle M^{\varphi_3} \rangle_\tau] & \varphi_1 \neq \varphi_3 = \varphi_2, \end{cases}$$

and in all cases the right-hand sides vanish, by the martingale property. If $\varphi_1 = \varphi_2 \neq \varphi_3$, we have by Theorem 3.2.16 and by (3.8) that

$$E^{c,x} [(M_\tau^{\varphi_1})^2 \langle M^{\varphi_3} \rangle_\tau] = E^{c,x} [(M_\tau^{\varphi_1})^2] E^{c,x} [\langle M^{\varphi_3} \rangle_\tau] \\ = E^{c,x} [(M_\tau^{\varphi_1})^2] E^{c,x} [(M_\tau^{\varphi_3})^2] = E^{c,x} [(M_\tau^{\varphi_1})^2 (M_\tau^{\varphi_3})^2].$$

Setting $\varphi_1 = \varphi_2 = \varphi_3 = \varphi_4$ in (3.43) yields, by the integration by parts formula (2.9) and the martingale property, that

$$E^{c,x} \left[(M_\tau^{\varphi_1})^2 \langle M^{\varphi_1} \rangle_\tau \right] = E^{c,x} \left[\int_0^\tau (M_s^{\varphi_1})^2 d\langle M^{\varphi_1} \rangle_s + \int_0^\tau \langle M^{\varphi_1} \rangle_s d\langle M^{\varphi_1} \rangle_s \right]. \quad (3.68)$$

On the other hand, since the sum on the right-hand side of (3.42) consists of six terms, it follows that setting $\varphi_1 = \varphi_2 = \varphi_3 = \varphi_4$ in (3.42) yields

$$E^{c,x} \left[(M_\tau^{\varphi_1})^4 \right] = 6E^{c,x} \left[\int_0^\tau (M_s^{\varphi_1})^2 d\langle M^{\varphi_1} \rangle_s \right]. \quad (3.69)$$

The integration by parts formula (2.9) yields

$$E^{c,x} \left[\langle M^{\varphi_1} \rangle_\tau^2 \right] = 2E^{c,x} \left[\int_0^\tau \langle M^{\varphi_1} \rangle_s d\langle M^{\varphi_1} \rangle_s \right]. \quad (3.70)$$

Rearranging (3.69) and (3.70), and substituting the resulting expressions into (3.68), yields the remaining case on the right-hand side of (3.67). \square

We shall use Lemma 3.2.17 and Lemma 3.2.18 in order to determine conditions that imply the existence and uniqueness of solutions of the restriction of the optimal control problem (2.71) to a finite-dimensional subset of the set \mathcal{U} of admissible controls.

3.3 Strong convexity

In this section, we construct an approximating subset \mathcal{A} of controls such that the control functional $\bar{\phi}^{\sigma,x}$ is strongly convex over \mathcal{A} . The strong convexity result is the main result of this chapter, and one of the main results of this thesis. The strong convexity result is important because it guarantees that, under certain conditions, the gradient descent algorithm that we describe in Chapter 4 produces a unique solution in an approximating subset to the optimal control problem described in §2.2.

3.3.1 Admissible classes spanned by linearly independent basis functions

Let $\{\varphi_i\}_{1 \leq i \leq n} \subset L^\infty(\bar{D})$ be linearly independent. Define the associated approximating subset

$$\mathcal{A} := \left\{ c \in L^\infty(\bar{D}) \mid \exists a \in \mathbb{R}^n \text{ s.t. } c(x) = \sum_{i=1}^n a_i \varphi_i(x) \ \forall x \in \mathbb{R}^d \right\}. \quad (3.71)$$

Every element of the approximating subset \mathcal{A} is in one-to-one correspondence with an element $a \in \mathbb{R}^n$, by linear independence of the φ_i . Therefore, we shall denote every control function $c \in \mathcal{A}$ by

$$c^a(x) := \sum_{i=1}^n a_i \varphi_i. \quad (3.72)$$

We shall approximate the original optimisation problem (3.3), that was defined over the set of all feedback controls, by restricting the optimisation problem to the set \mathcal{A} . We shall refer to the set \mathcal{A} as the ‘approximating subset’ of the set of feedback controls.

Notation: In this section, we shall reserve the symbols a and z (and variants thereof) for vectors in \mathbb{R}^n . If either a or z appears as a superscript, then that indicates the dependence upon a parametrised control of the form (3.72) belonging to \mathcal{A} . For example, we shall denote by X^a the solution of the stochastic differential equation

$$dX_t^a = [c^a - \nabla V](X_t^a)dt + \sqrt{2\varepsilon}dB_t, \quad (3.73)$$

and we shall denote by $\mu^{a,x} = P^x \circ (X^a)^{-1}$ the distribution of X^a with respect to P^x , i.e.

$$\forall A \in \mathcal{F}, \quad E^{a,x}[1_A] = \mu^{a,x}(A) = P^x(X^a \in A). \quad (3.74)$$

We also denote the parametrised functionals

$$M_t^a(\omega) := \frac{1}{\sqrt{2\varepsilon}} \int_0^t c^a(\omega_s)dB_s \quad (3.75)$$

$$\langle M^a \rangle_t(\omega) := \frac{1}{2\varepsilon} \int_0^t |c^a(\omega_s)|^2 ds \quad (3.76)$$

$$K^{\sigma,a,\alpha}(\omega) := W(\omega) + (2\sigma)^{-1} \langle M^a \rangle_\tau(\omega) + \alpha \sigma^{-1} M_\tau^a(\omega), \quad \alpha \in \{0, 1\}. \quad (3.77)$$

Unless otherwise stated, any other superscripts shall refer to functions, so that

$$M_t^{\varphi_i}(\omega) = \frac{1}{\sqrt{2\varepsilon}} \int_0^t \varphi_i(\omega_s)dB_s,$$

$$\langle M^{\varphi_i} \rangle_t(\omega) = \frac{1}{2\varepsilon} \int_0^t |\varphi_i(\omega_s)|^2 ds.$$

In particular, we have the following simple, but useful relations:

$$M_t^a = \sum_i a_i M_t^{\varphi_i} \quad (3.78)$$

$$\langle M^a \rangle_t = \sum_{i,j} a_i a_j \langle M^{\varphi_i}, M^{\varphi_j} \rangle_t. \quad (3.79)$$

To conclude this digression into notation, we define the objective function $\phi^{\sigma,x}$ by

$$\begin{aligned} \phi^{\sigma,x} : \mathbb{R}^n &\rightarrow \mathbb{R} \\ a &\mapsto \bar{\phi}^{\sigma,x}(c^a). \end{aligned} \quad (3.80)$$

Thus, $\phi^{\sigma,x}$ is the restriction of $\bar{\phi}^{\sigma,x}$ to the approximating subset \mathcal{A} defined in (3.71). We shall refer to $\phi^{\sigma,x}$ as the ‘approximating function’ or simply the ‘approximation’ of $\bar{\phi}^{\sigma,x}$ given the approximating subset \mathcal{A} .

We now express some key results from §3.1 and §3.2 using the new notation introduced above. Let $(e_i)_{i=1}^n$ denote the canonical orthonormal basis of \mathbb{R}^n .

Lemma 3.3.1. *Let $c^a \in \mathcal{A}$ and φ_i be one of the basis elements of the approximating subset \mathcal{A} . If $W \in L^2(\mu^{a,x})$ and $T > 0$ is such that $\tau \leq T$ $\mu^{a,x}$ -almost surely, then the i -th partial derivative of the function $\phi^{\sigma,x}$ defined in (3.80) exists, and is given by*

$$\frac{\partial}{\partial a_i} \phi^{\sigma,x}(a) = E^{a,x} \left[K^{\sigma,a,1} M_\tau^{\varphi_i} \right]. \quad (3.81)$$

Proof. The existence of the first variation of $\bar{\phi}^{\sigma,x}$ follows from Theorem 3.1.4. Thus

$$\frac{\partial}{\partial a_i} \phi^{\sigma,x}(a) = \lim_{\delta \rightarrow 0} \frac{\phi^{\sigma,x}(a + \delta e_i) - \phi^{\sigma,x}(a)}{\delta} = \lim_{\delta \rightarrow 0} \frac{\bar{\phi}^{\sigma,x}(c^a + \delta \varphi_i) - \bar{\phi}^{\sigma,x}(c^a)}{\delta} = \Phi^{\sigma,x}(c^a; \varphi_i).$$

By Proposition 3.2.1, the conclusion follows. \square

Lemma 3.3.2. *Let $c^a \in \mathcal{A}$ and φ_i, φ_j be any two basis elements of the approximating subset \mathcal{A} . If $W^2 \in L^p(\mu^{a,x})$ for some $p > 1$ and $T > 0$ is such that $\tau \leq T$ $\mu^{a,x}$ -almost surely, then the second partial derivative of $\phi^{\sigma,x}$ with respect to a_i and a_j exists, and is given by*

$$\frac{\partial}{\partial a_j} \frac{\partial}{\partial a_i} \phi^{\sigma,x}(a) = E^{a,x} \left[K^{\sigma,a,1} M_\tau^{\varphi_i} M_\tau^{\varphi_j} + \sigma^{-1} (1 - M_\tau^a) \langle M^{\varphi_i}, M^{\varphi_j} \rangle_\tau \right]. \quad (3.82)$$

Proof. The existence of the mixed second variation of $\bar{\phi}^{\sigma,x}$ follows from Theorem 3.1.7. We have by (3.81) that

$$\begin{aligned} \frac{\partial}{\partial a_j} \frac{\partial}{\partial a_i} \phi^{\sigma,x}(a) &= \frac{\partial}{\partial a_j} \Phi^{\sigma,x}(c^a; \varphi_i) \\ &= \lim_{\delta \rightarrow 0} \frac{\Phi^{\sigma,x}(c^a + \delta \varphi_j; \varphi_i) - \Phi^{\sigma,x}(c^a; \varphi_i)}{\delta} = \Phi^{\sigma,x}(c^a; \varphi_i, \varphi_j). \end{aligned}$$

By Proposition 3.2.5, the conclusion follows. \square

Remark 7. Since the measure $\mu^{a,x}$ and the random variables $K^{\sigma,a,1}$ and M_τ^a depend continuously on a , the mixed second partial derivative in (3.82) is a continuous function of a . Therefore, over the subset of \mathbb{R}^n such that the right-hand side of (3.82) exists and is finite, it holds by Clairaut's theorem that the matrix of second-order partial derivatives (i.e. the Hessian) of $\phi^{\sigma,x}$,

$$\nabla^2 \phi^{\sigma,x}(a) := \left(\frac{\partial}{\partial a_i} \frac{\partial}{\partial a_j} \phi^{\sigma,x} \right)_{1 \leq i, j \leq n} (a), \quad (3.83)$$

is a symmetric matrix.

Lemma 3.3.3. *Let $z \in \mathbb{R}^n$ be arbitrary. Then for $M^z := \sum_i z_i M^{\varphi_i}$,*

$$z^\top \nabla^2 \phi^{\sigma,x}(a) z = E^{a,x} \left[K^{\sigma,a,1} (M_\tau^z)^2 + \sigma^{-1} (1 - M_\tau^a) \langle M^z \rangle_\tau \right]. \quad (3.84)$$

Proof. The assertion follows directly from (3.82) and the constancy of z . \square

The expression (3.84) suggests that proving convexity of $\phi^{\sigma,x}$ may be difficult, primarily because M_τ^a assumes negative values with positive probability. In particular the right-hand side of (3.84) will be nonnegative for a in some convex set and arbitrary z if and only if

$$E^{a,x} \left[K^{\sigma,a,0} (M_\tau^z)^2 + \sigma^{-1} \langle M^z \rangle_\tau \right] \geq E^{a,x} \left[\sigma^{-1} M_\tau^a \left(\langle M^z \rangle_\tau - (M_\tau^z)^2 \right) \right].$$

Even if we assume that $K^{\sigma,a,0}$ is nonnegative, it is not clear that the relation above will hold, especially if z is arbitrary. In order to guarantee that the Hessian is positive semidefinite, we shall impose additional structure on the vector space \mathcal{A} defined in (3.71), by imposing the non-overlap condition on the supports stated in Assumption 3.2.8.

3.3.2 Basis functions with non-overlapping supports

In this section, we apply the identities from the previous section, in conjunction with the martingale inequality (2.12), in order to derive a useful lower bound on the quantity (3.84). We then use the lower bound to show that the function $\phi^{\sigma,x}$ is strongly convex on \mathbb{R}^n , provided that the basis elements $\{\varphi_i\}_{1 \leq i \leq n}$ satisfy Assumption 3.2.8. Define the nonnegative functions $(C_k)_{k \leq n}$, where for each k , $C_k : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by

$$C_k(a) := \sum_{i \neq k} a_i^2 E^{a,x} [\langle M^{\varphi_i} \rangle_\tau]. \quad (3.85)$$

We have the following

Lemma 3.3.4. *Suppose the collection $\{\varphi_i\}_{1 \leq i \leq n}$ of basis functions of the approximating subset \mathcal{A} satisfies Assumption 3.2.8. Then the following identities hold for $a, z \in \mathbb{R}^n$:*

$$E^{a,x} \left[M_\tau^a \left((M_\tau^z)^2 - \langle M^z \rangle_\tau \right) \right] = \frac{2}{3} \sum_k a_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^3 \right] \quad (3.86)$$

$$E^{a,x} \left[\langle M^a \rangle_\tau (M_\tau^z)^2 \right] = \sum_{k=1}^n z_k^2 E^{a,x} \left[a_k^2 \left(\frac{(M_\tau^{\varphi_k})^4}{6} + \frac{\langle M^{\varphi_k} \rangle_\tau^2}{2} \right) + (M_\tau^{\varphi_k})^2 C_k(a) \right]. \quad (3.87)$$

Proof. Since Assumption 3.2.8 holds, we may apply the independence relation (3.63) to obtain

$$E^{a,x} \left[M_\tau^a (M_\tau^z)^2 \right] = \sum_{i,k,\ell} a_k z_i z_\ell E^{a,x} \left[M_\tau^{\varphi_k} M_\tau^{\varphi_i} M_\tau^{\varphi_\ell} \right] = \sum_k a_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^3 \right]. \quad (3.88)$$

Expanding M^a and M^z , applying Lemma 3.2.9, the independence relation (3.64), and applying Lemma 3.2.17, we obtain

$$\begin{aligned} E^{a,x} \left[M_\tau^a \langle M^z \rangle_\tau \right] &= \sum_{i,k,\ell} a_k z_i z_\ell E^{a,x} \left[M_\tau^{\varphi_k} \langle M^{\varphi_i}, M^{\varphi_\ell} \rangle_\tau \right] \\ &= \sum_{k,\ell} a_k z_\ell^2 E^{a,x} \left[M_\tau^{\varphi_k} \langle M^{\varphi_\ell} \rangle_\tau \right] \\ &= \sum_k a_k z_k^2 E^{a,x} \left[M_\tau^{\varphi_k} \langle M^{\varphi_k} \rangle_\tau \right] \\ &= 3^{-1} \sum_k a_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^3 \right]. \end{aligned} \quad (3.89)$$

Taking the difference of (3.88) and (3.89) yields (3.86).

Expanding M^a and M^z , and applying Lemma 3.2.9, and Lemma 3.2.18, we obtain

$$\begin{aligned} E^{a,x} \left[\langle M^a \rangle_\tau (M_\tau^z)^2 \right] &= \sum_{i,j,k,\ell} a_j a_k z_i z_\ell E^{a,x} \left[\langle M^{\varphi_j}, M^{\varphi_k} \rangle_\tau M_\tau^{\varphi_i} M_\tau^{\varphi_\ell} \right] \\ &= \sum_{i,k,\ell} a_k^2 z_i z_\ell E^{a,x} \left[\langle M^{\varphi_k} \rangle_\tau M_\tau^{\varphi_i} M_\tau^{\varphi_\ell} \right] \\ &= \sum_k a_k^2 z_k^2 E^{a,x} \left[6^{-1} (M_\tau^{\varphi_k})^4 + 2^{-1} \langle M^{\varphi_k} \rangle_\tau^2 \right] \\ &\quad + \sum_k \sum_{i \neq k} a_i^2 z_k^2 E^{a,x} \left[(M_\tau^{\varphi_i})^2 (M_\tau^{\varphi_k})^2 \right]. \end{aligned}$$

By Theorem 3.2.16, $M_\tau^{\varphi_i}$ and $M_\tau^{\varphi_k}$ are independent. Hence, by the definition (3.85),

$$\sum_{i \neq k} a_i^2 E^{a,x} [(M_\tau^{\varphi_i})^2 (M_\tau^{\varphi_k})^2] = \sum_{i \neq k} a_i^2 E^{a,x} [(M_\tau^{\varphi_i})^2] E^{a,x} [(M_\tau^{\varphi_k})^2] = C_k(a) E^{a,x} [(M_\tau^{\varphi_k})^2],$$

where we have used the Itô isometry and the definition (3.85) of $C_k(a)$ in the last equation. \square

We now state another assumption for proving strong convexity of the function $\phi^{\sigma,x}$.

Assumption 3.3.5. *The supports $(S_i)_{i \leq n}$ are such that*

$$P^x(\forall 1 \leq i \leq n, \exists u_i < v_i \text{ s.t. } r \in (u_i, v_i) \Rightarrow X_r^0 \in S_i) > 0. \quad (3.90)$$

Assumption 3.2.8 means that the diffusion X^0 that solves (3.1) when $c \equiv 0$ spends a strictly positive amount of time in each support S_i with positive P^x -probability.

Lemma 3.3.6. *Suppose that Assumption 3.3.5 holds, and that $c^a \in \mathcal{A}$. Then*

$$\mu^{a,x}(\forall 1 \leq i \leq n, M_\tau^{\varphi_i} \neq 0) > 0. \quad (3.91)$$

Proof. Since $c^a \in \mathcal{A}$, the change of measure theorem (Theorem 2.1.17) holds, and hence $\mu^{a,x}$ and $\mu^{0,x}$ are locally equivalent. Since (3.90) holds, it follows that

$$\mu^{0,x}(\forall 1 \leq i \leq n, M_\tau^{\varphi_i} \neq 0) > 0 \quad (3.92)$$

holds, and (3.91) follows by local equivalence of $\mu^{a,x}$ and $\mu^{0,x}$. \square

In preparation for the proof of strong convexity, recall that $\rho^{(2)} = \sqrt{3/2 + \sqrt{3/2}}$ is the largest root of the fourth-order Hermite polynomial (see Lemma 2.1.6 in §2.1). Define the constant

$$C := \frac{4}{9} \left(\frac{1}{3} + \frac{1}{\rho^{(2)}} \right)^{-1} \approx 0.4732. \quad (3.93)$$

We also define a finite collection $(p_k)_{k \leq n}$ of functions $p_k : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$, where for all $a \in \mathbb{R}^n$, each function $p_k(\cdot, \cdot; a)$ is a fourth-degree polynomial in the first argument and a second-degree polynomial in the second argument:

$$p_k(x, y; a) := \left(\frac{1}{12} + \frac{1}{4\rho^{(2)}} \right) x^4 y^2 + \frac{2}{3} x^3 y + \left(C + \frac{1}{2} C_k(a) \right) x^2. \quad (3.94)$$

Using the relations obtained in Section §3.2.2, we shall reformulate the number $z^\top \nabla^2 \phi^{\sigma,x} z(a)$ as the sum of expected values of the polynomials p_k , evaluated at the pair $(M_\tau^{\varphi_k}, a_k)$. Therefore it is of interest to characterise these polynomials. A crucial characterisation is provided in the next

Lemma 3.3.7. *For all $k \in \{1, \dots, n\}$ and for all $a \in \mathbb{R}^n$, $p_k(\cdot, \cdot; a)$ is nonnegative.*

Proof. Since $p_k(x, y; a)$ factorises according to

$$p_k(x, y; a) = x^2 q_k(x, y; a) \quad (3.95)$$

for

$$q_k(x, y; a) := \left(\frac{1}{12} + \frac{1}{4\rho^{(2)}} \right) x^2 y^2 + \frac{2}{3} xy + \left(C + \frac{1}{2} C_k(a) \right), \quad (3.96)$$

it follows that $p_k(\cdot, \cdot; a)$ is nonnegative if and only if $q_k(\cdot, \cdot; a)$ is nonnegative on $\mathbb{R} \times \mathbb{R}$. Let $k \in \{1, \dots, n\}$ be arbitrary, and fix y . If $y = 0$, then

$$q_k(x, 0; a) = C + \frac{1}{2} C_k(a).$$

By (3.93) and (3.85), both C and $C_k(a)$ are nonnegative. Thus, it follows that $q_k(x, 0; a)$ is a nonnegative constant. If $y \neq 0$, then by the discriminant test for quadratic polynomials, $q_k(x, y; a)$ is a nonnegative function of x if and only if

$$\left(\frac{2}{3} y \right)^2 - 4y^2 \left(\frac{1}{12} + \frac{1}{4\rho^{(2)}} \right) \left(C + \frac{1}{2} C_k(a) \right) \leq 0. \quad (3.97)$$

Since $y \neq 0$, we may divide both sides of (3.97) by y^2 , with the result that (3.97) holds if and only if

$$\frac{4}{9} \left(\frac{1}{3} + \frac{1}{\rho^{(2)}} \right)^{-1} \leq \left(C + \frac{1}{2} C_k(a) \right). \quad (3.98)$$

By (3.93), it follows that (3.98) is satisfied, regardless of the value of $C_k(a)$, and from the latter it follows that (3.97) is satisfied. Hence q_k is a nonnegative function, and so is p_k . \square

Corollary 3.3.8. *For all $k \in \{1, \dots, n\}$, $y \in \mathbb{R}$, and $a \in \mathbb{R}^n$, there is at most one real, nonzero root of $p_k(\cdot, y; a) = 0$ in \mathbb{R} .*

Recall that $\phi^{\sigma, x} \in C^2(\mathbb{R}^n; \mathbb{R})$ is said to be *strongly convex* if there exists a number $m > 0$ such that

$$\forall z \neq 0, \quad z^\top \nabla^2 \phi^{\sigma, x}(a) z \geq m |z|^2, \quad (3.99)$$

where m does not depend of a . Note that (3.99) implies that the spectrum of the Hessian is contained in the set $[m, \infty)$.

Let $U \subset \mathbb{R}^n$ be an open, convex set, and denote its closure by \bar{U} . Define

$$m^x(\bar{U}) := \min_{a \in \bar{U}} \min_k E^{a, x} [p_k(M_\tau^{\varphi_k}, a_k; a)]. \quad (3.100)$$

The following theorem is the main result of this chapter and one of the main results of the thesis.

Theorem 3.3.9. *Assume that the following hold:*

- (i) the basis functions $\{\varphi_i\}_{1 \leq i \leq n}$ of the approximating subset \mathcal{A} belong to $L^\infty(\bar{D})$,
- (ii) the supports of $(S_i)_{1 \leq i \leq n}$ satisfy Assumptions 3.2.8 and 3.3.5,
- (iii) the stopping time τ is $\mu^{0, x}$ -almost surely bounded for all $x \in D$, and

(iv) the terminal cost function $g \in C(\partial D; \mathbb{R})$ satisfies the uniform lower bound

$$-\sigma \|g\|_\infty \geq (C - 1) \quad (3.101)$$

for the constant C defined in (3.93).

Then for any open, convex set $U \subseteq \mathbb{R}^n$, the restriction of $\phi^{\sigma,x}$ to U is strongly convex, and (3.99) holds with $m = m^x(\bar{U})$ as given in (3.100).

Proof. Since we stipulated at the beginning of this chapter that the running cost function f is an element of $C(\bar{D}; [0, \infty))$, it follows that f is bounded on \bar{D} ; since the terminal cost function g is an element of $C(\partial D; \mathbb{R})$, and since D is a bounded domain, condition (iii) implies that the random variable W is $\mu^{0,x}$ -almost surely bounded. On the other hand, condition (i) implies that for all $c^a \in \mathcal{A}$, the corresponding path measure $\mu^{a,x}$ is locally equivalent to $\mu^{0,x}$, by the change of measure theorem, so W is $\mu^{a,x}$ -almost surely bounded for any $a \in U$. Therefore $W^2 \in L^p(\mu^{a,x})$ for some $p > 1$, and thus the assumptions of Lemma 3.3.2 are satisfied, which guarantees that the Hessian of $\phi^{\sigma,x}$ is well-defined on U . We will use conditions (ii) and (iv) to show that the Hessian $\nabla^2 \phi^{\sigma,x}(a)$ is positive definite for all $a \in \mathbb{R}^n$, and then strengthen this result to show uniform positive definiteness in the sense of (3.99). By the theory of convex functions, uniform positive definiteness of the Hessian of $\phi^{\sigma,x}$ implies strong convexity of $\phi^{\sigma,x}$.

Substituting the expression (3.77) for $K^{\sigma,a,1}$ into the expression (3.84) and pulling the factor σ^{-1} out of the expectation yields

$$z^\top \nabla^2 \phi^{\sigma,x}(a) z = \sigma^{-1} E^{a,x} \left[\left(\sigma W + 2^{-1} \langle M^a \rangle_\tau + M_\tau^a \right) (M_\tau^z)^2 + (1 - M_\tau^a) \langle M^z \rangle_\tau \right].$$

By the Itô isometry (3.8),

$$z^\top \nabla^2 \phi^{\sigma,x}(a) z = \sigma^{-1} E^{a,x} \left[\left(\sigma W + 1 + 2^{-1} \langle M^a \rangle_\tau + M_\tau^a \right) (M_\tau^z)^2 - M_\tau^a \langle M^z \rangle_\tau \right].$$

Since $f \geq 0$, it follows that $\sigma W \geq -\sigma \|g\|_\infty$, and (3.101) in condition (iv) implies that $\sigma W \geq C - 1$. Applying this uniform lower bound on W to the equation above and rearranging terms yields

$$\begin{aligned} z^\top \nabla^2 \phi^{\sigma,x}(a) z &\geq \sigma^{-1} E^{a,x} \left[\left(C + 2^{-1} \langle M^a \rangle_\tau + M_\tau^a \right) (M_\tau^z)^2 - M_\tau^a \langle M^z \rangle_\tau \right] \\ &= \sigma^{-1} E^{a,x} \left[2^{-1} \langle M^a \rangle_\tau (M_\tau^z)^2 + M_\tau^a \left((M_\tau^z)^2 - \langle M^z \rangle_\tau \right) + C (M_\tau^z)^2 \right]. \end{aligned}$$

By using the expression (3.87) for the expectation of $\langle M^a \rangle_\tau (M_\tau^z)^2$, the expression (3.86) for the expectation of $M_\tau^a \left((M_\tau^z)^2 - \langle M^z \rangle_\tau \right)$, and the independence relation (3.63) (which

holds, by condition (ii), we have

$$\begin{aligned}
& z^\top \nabla^2 \phi^{\sigma,x}(a) z \\
& \geq (2\sigma)^{-1} \sum_k z_k^2 E^{a,x} \left[a_k^2 \left(6^{-1} (M_\tau^{\varphi_k})^4 + 2^{-1} \langle M_\tau^{\varphi_k} \rangle_\tau^2 \right) + C_k(a) (M_\tau^{\varphi_k})^2 \right] \\
& \quad + \sigma^{-1} E^{a,x} \left[M_\tau^a \left((M_\tau^z)^2 - \langle M_\tau^z \rangle_\tau \right) + C(M_\tau^z)^2 \right] \\
& = (2\sigma)^{-1} \sum_k z_k^2 E^{a,x} \left[a_k^2 \left(6^{-1} (M_\tau^{\varphi_k})^4 + 2^{-1} \langle M_\tau^{\varphi_k} \rangle_\tau^2 \right) + C_k(a) (M_\tau^{\varphi_k})^2 \right] \\
& \quad + 2(3\sigma)^{-1} \sum_k a_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^3 \right] + C\sigma^{-1} E^{a,x} \left[(M_\tau^z)^2 \right] \\
& = (2\sigma)^{-1} \sum_k z_k^2 E^{a,x} \left[a_k^2 \left(6^{-1} (M_\tau^{\varphi_k})^4 + 2^{-1} \langle M_\tau^{\varphi_k} \rangle_\tau^2 \right) + C_k(a) (M_\tau^{\varphi_k})^2 \right] \\
& \quad + 2(3\sigma)^{-1} \sum_k a_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^3 \right] + C\sigma^{-1} \sum_k z_k^2 E^{a,x} \left[(M_\tau^{\varphi_k})^2 \right]. \tag{3.102}
\end{aligned}$$

The inequality (2.12) in Lemma 2.1.6, i.e. the inequality

$$E^{a,x} \left[(M_\tau^\varphi)^4 \right] \leq \rho^{(2)} E^{a,x} \left[\langle M_\tau^\varphi \rangle_\tau^2 \right],$$

provides a lower bound on $E^{a,x}[\langle M_\tau^{\varphi_k} \rangle_\tau^2]$, which by substitution into (3.102) yields

$$\begin{aligned}
& \sigma z^\top \nabla^2 \phi^{\sigma,x}(a) z \\
& \geq \sum_k z_k^2 E^{a,x} \left[\left(\frac{1}{12} + \frac{1}{4\rho^{(2)}} \right) (M_\tau^{\varphi_k})^4 a_k^2 + \frac{2}{3} (M_\tau^{\varphi_k})^3 a_k + \left(\frac{1}{2} C_k(a) + C \right) (M_\tau^{\varphi_k})^2 \right].
\end{aligned}$$

We apply the definition (3.94) of the polynomial p_k in order to rewrite the inequality above as

$$\sigma z^\top \nabla^2 \phi^{\sigma,x}(a) z \geq \sum_k z_k^2 E^{a,x} [p_k(M_\tau^{\varphi_k}, a_k; a)]. \tag{3.103}$$

By Lemma 3.3.7, each p_k is nonnegative. Condition (ii) implies that Assumption 3.3.5 holds, which in turn implies that Lemma 3.3.6 holds. In particular, $M_\tau^{\varphi_k}$ assumes nonzero values with strictly positive $\mu^{a,x}$ -probability. Since $M_\tau^{\varphi_k}$ is a continuous random variable, it follows that each expectation in the sum on the right-hand side of (3.103) is strictly positive. Since at least one z_k is nonzero, one of the summands in (3.103) is strictly positive. Thus the Hessian of $\phi^{\sigma,x}$ is positive definite on U . As we have not imposed any conditions on U other than it be open and convex, this implies that the Hessian is positive definite on \mathbb{R}^n .

We now show that the Hessian of $\phi^{\sigma,x}$ is in fact uniformly positive definite. For $k \in \{1, \dots, n\}$, the function

$$\begin{aligned}
h_k &: \mathbb{R}^n \rightarrow \mathbb{R} \\
a &\mapsto E^{a,x} [p_k(M_\tau^{\varphi_k}, a_k; a)]
\end{aligned}$$

is strictly positive (recall that each expectation in the sum on the right-hand side of (3.103) is strictly positive). Furthermore, h_k is continuous, since the $p_k(\cdot, \cdot; a)$ depend smoothly on

a , and since the law $\mu^{a,x}$ depends continuously on a . Recall that the minimum of a finite collection of continuous functions is again a continuous function. Hence, the function

$$h := \min_{1 \leq k \leq n} h_k$$

is both continuous and strictly positive. In particular, for an arbitrary closed, convex subset $\bar{U} \subseteq \mathbb{R}^n$, the minimum value of h on \bar{U} equals $m^x(\bar{U})$, by the definition (3.100). We now show that the value of $m^x(\bar{U})$ must be strictly positive. If the minimum value $m^x(\bar{U})$ were zero, then there would exist some $a' \in \bar{U}$ and $k \in \{1, \dots, n\}$ such that

$$\mu^{a',x}(M_\tau^{\varphi^k} \in \{y \in \mathbb{R} \mid p_k(y, a'_k; a) = 0\}) = 1. \quad (3.104)$$

Since Corollary 3.3.8 guarantees that there is at most one real, nonzero root of $p_k(\cdot, y; a)$ for every $y \in \mathbb{R}$, it follows that (3.104) is consistent with the property that $M_\tau^{\varphi^k}$ has mean zero with respect to $\mu^{a',x}$ if and only if the distribution of $M_\tau^{\varphi^k}$ with respect to $\mu^{a',x}$ is concentrated on the trivial root $x = 0$. This produces a contradiction with the conclusion (3.91) of Lemma 3.3.6. Therefore, the minimum value $m^x(\bar{U})$ is strictly positive. By (3.103) we have, for all $a \in \bar{U}$,

$$\sigma z^\top \nabla^2 \phi^{\sigma,x}(a) z \geq \sum_k z_k^2 E_{\mu(a)}^x [p_k(M_\tau^{\varphi^k}, a_k)] \geq \sum_k z_k^2 m^x(\bar{U}) = m^x(\bar{U}) |z|^2.$$

Thus the Hessian $\nabla^2 \phi^{\sigma,x}(a)$ is uniformly positive definite, by (3.99). By the theory of convex functions, this implies that $\phi^{\sigma,x}$ is strongly convex. \square

Remark 8. In the proof of Theorem 3.3.9, strong convexity of $\phi^{\sigma,x}$ holds under the stated hypotheses, regardless of the dimension n of the approximating subset \mathcal{A} or the dimension d of the state space of the diffusions X^a . The proof does not require that the number of basis functions be bounded. However, since the domain D is assumed to be bounded, and given that increasing n has the effect of reducing the Lebesgue measure of the intersections $S_i \cap D$, then the proof of Theorem 3.3.9 indicates that the value of the parameter m^x will decrease to zero as n increases, because the mean occupation times of the controlled diffusion X^a in each of the supports will decrease, and hence the value of $M_\tau^{\varphi^k}$ for any k will be increasingly concentrated around the initial value of zero. This observation will be relevant when we consider the rate of convergence of a gradient flow in Section §4.1. Note however that if $M_\tau^{\varphi^k}$ is nonconstant (i.e. nonzero), then it is not bounded, so the distribution of $M_\tau^{\varphi^k}$ is supported on all of \mathbb{R} , and hence $m^x(\bar{U})$ will be strictly positive. The essential observation is that neither the dimension n of the approximating subset \mathcal{A} nor the dimension d of the state space of the diffusion affect the strong convexity of $\phi^{\sigma,x}$, provided that conditions (i)-(iv) hold.

In Theorem 3.3.9, we have not imposed the requirement that the basis functions be locally Lipschitz continuous. Therefore, the drift coefficient of the controlled diffusion X^a need not be Lipschitz continuous, and therefore the standard theory on existence and uniqueness of strong solutions to stochastic differential equations cannot guarantee whether a strong solution exists. Zvonkin [67, Theorem 4, Case 2)] proved the existence and uniqueness of solutions in the one-dimensional case, provided that the diffusion coefficient is bounded, α -Hölder continuous for $\alpha \geq 1/2$, and uniformly elliptic, and provided that

the drift is bounded. Veretennikov [62, Theorem 1] generalised Zvonkin's result to the d -dimensional case, with the consequence that when the covariance matrix is the identity matrix, then there exists a unique strong solution to the stochastic differential equation. Since the covariance matrix in (3.1) is proportional to the identity matrix, and since the control c^a is bounded if the basis functions of the approximating set \mathcal{A} are bounded, it holds that existence and uniqueness of solutions is guaranteed, provided that ∇V is also bounded.

The boundedness condition on the basis functions is reasonable even if the optimal control is not bounded, since one often approximates objects in some space by their bounded counterparts (e.g. Lebesgue integrable functions by simple functions). On the other hand, there exist conditions which guarantee that the optimal control is bounded; see Theorem 2.1.23 in Section §2.1.2.

In this chapter, we described a method for constructing strongly convex approximations of the control functional $\bar{\phi}^{\sigma,x}$ defined in (3.4). The approximation is the same as that described in Hartmann and Schütte's paper [35]: restrict the control functional to an approximating subset of feedback controls, spanned by finitely many linearly independent basis functions, and find the coefficients for the linear combination of basis functions that yield the best approximation of the optimal control in the approximating subset. In order to obtain the strong convexity result above, we imposed additional conditions (e.g. boundedness and non-overlapping support conditions), so that the approximating function $\phi^{\sigma,x}$, obtained by restricting the control functional $\bar{\phi}^{\sigma,x}$ to the approximating subset \mathcal{A} , is strongly convex. The non-overlapping supports guarantee that certain continuous local martingales associated to the basis functions are independent. Applying one of the martingale inequalities presented in Lemma 2.1.6 leads to a lower bound on the eigenvalues of the Hessian of the function. The lower bound on g provides some control over the lower bound on the eigenvalues, in the sense that if the work is greater or equal to some value (which depends on the best constants of the martingale inequality), then the eigenvalues are bounded away from zero, provided that the continuous local martingales associated with the basis functions are nonconstant.

The approach we have taken for showing that the approximating function $\phi^{\sigma,x}$ is strongly convex, i.e. the Hessian-based approach, is motivated by the observation that the control functional depends nonlinearly on the feedback control c via the measure $\mu^{c,x}$, where the nonlinear dependence of $\mu^{c,x}$ on c is more complicated than a quadratic nonlinearity. The Hessian-based approach allows us to circumvent the difficulty posed by the nonlinear dependence of $\mu^{c,x}$ on c , since we can bound the Hessian-induced quadratic form from below by expectations of quartic polynomials, each of which factors into the product of two quadratic polynomials. By imposing a lower bound on g (and thus on the path functional W), we can ensure that both quadratic polynomials are nonnegative, and thus ensure that the expectations of the quartic polynomials are strictly positive, provided that the measures $(\mu^{c,x} \circ (M_\tau^{\varphi_i})^{-1})_i$ are supported on all of \mathbb{R} .

Chapter 4

Solution via a gradient flow

In Chapter 3, we analysed the control functional $\bar{\phi}^{\sigma,x}$ of an optimal control problem. The main result of the chapter, Theorem 3.3.9, was that the control functional is strongly convex over an approximating subset \mathcal{A} spanned by finitely many bounded, Borel-measurable functions $\{\varphi_i\}_{1 \leq i \leq n}$ whose supports S_i are subsets of D with strictly positive Lebesgue measure, such that the closures form a partition of the closure of D , and such that the non-overlap condition

$$\overline{S_i} \cap \overline{S_j} \subset \partial S_i \cap \partial S_j, \quad i \neq j \quad (4.1)$$

holds. The strong convexity result followed from analysing the function $\phi^{\sigma,x}$ that corresponded to the restriction of the functional $\bar{\phi}^{\sigma,x}$ to the approximating subset \mathcal{A} .

In this chapter, we apply the results of Chapter 3 in order to characterise the gradient descent algorithm for solving the optimal control problem suggested by Hartmann and Schütte. Since the gradient descent algorithm is an iterative method, it may be viewed as a discrete-time dynamical system. We analyse the algorithm by studying the continuous-time limit.

4.1 The flow associated to a gradient descent algorithm

4.1.1 Exponential convergence to a unique equilibrium

By (2.68) and Theorem 3.3.9, the function $\phi^{\sigma,x} : \mathbb{R}^n \rightarrow [F(\sigma, x), \infty)$ is bounded from below and strongly convex, respectively. Solving the optimal control problem on the approximating subset \mathcal{A} defined in (3.71) corresponds to solving the convex optimisation problem

$$\min_a \phi^{\sigma,x}(a). \quad (4.2)$$

Remark 9. Note that an inequality constraint of the form

$$\sum_i a_i^2 \|\varphi_i\|_\infty^2 \leq C \quad (4.3)$$

for some $C > 0$ may be specified for the convex minimisation problem (4.2), although such an inequality is not strictly necessary, in view of the fact that strong convexity of $\phi^{\sigma,x}$ holds over all of \mathbb{R}^n . Strong duality for the resulting convex minimisation problem may be shown by proving that constraint qualifications, such as Slater's condition, hold (see,

e.g. [10]). However, since neither the primal problem (4.2)–(4.3) nor its dual are easily solved by analytical methods in general, we must resort to a numerical method for solving (4.2).

By Proposition 3.2.1, $\phi^{\sigma,x} : \mathbb{R}^n \rightarrow [F(\sigma, x), \infty)$ is continuously differentiable, and therefore we may use the constant step-size gradient descent algorithm for solving (4.2). Given a step size $t_a > 0$, a number N_{traj} of trajectories to be sampled at each iteration of gradient descent, a number N_{iter} of iterations of gradient descent to perform, and a number n of basis functions $(\varphi_i)_{1 \leq i \leq n}$, the constant step-size gradient descent algorithm is given according to:

```

for  $\ell = 1$  to  $N_{\text{iter}}$  do
  Sample  $N_{\text{traj}}$  trajectories according to the distribution  $\mu^{a^{(\ell-1)},x}$ ;
  Compute the sample means of the estimators  $K^{\sigma,a^{(\ell-1)},1}M_T^{\varphi_i}$  for  $1 \leq i \leq n$ ;
  Estimate  $\phi^{\sigma,x}(a^{(\ell)})$  by the sample mean of  $K^{\sigma,a^{(\ell)},0}$ 
  Obtain the new gradient descent iterate  $a^{(\ell)}$ , by using
      
$$a^{(\ell)} = a^{(\ell-1)} - t_a \nabla \phi^{\sigma,x}(a^{(\ell-1)}), \tag{4.4}$$

  the sample means of  $K^{\sigma,a^{(\ell-1)},1}M_T^{\varphi_i}$ , and Lemma 3.3.1;
end

```

Algorithm 1: Gradient descent algorithm

As we may only estimate the values of the gradient $\nabla_a \phi^{\sigma,x}$ via noise-corrupted measurements, the gradient descent algorithm (4.4) may be implemented as a stochastic approximation algorithm. In the machine learning and stochastic approximation literature, the parameter t_a is also referred to as the ‘learning rate’ or ‘gain’. In many of the early studies on stochastic approximation algorithms, it was desirable that one work with a decreasing sequence $(t_a^{(\ell)})_\ell$ of gains that decreases at each iteration. For example, in Robbins and Monro’s seminal article [56], the convergence of the Robbins-Monro method for stochastic approximation of a root of a function was shown to hold, under the condition that the time discretisation parameter $t_a^{(n)}$ at each iteration was proportional to n^{-1} . The justification for this particular rate of decrease is that the partial sums of gain steps

$$s_N := \sum_{\ell=1}^N t_a^{(\ell)}$$

should form a divergent sequence, in order for the iterates $(a^{(\ell)})_{\ell \in \mathbb{N}}$ to explore the entire parameter space. On the other hand, the gain should also decrease to zero quickly enough, in order to account for the fact that, in all stochastic approximation problems, the function of interest can only be estimated via noise-corrupted measurements; the presence of noise is mitigated by a decreasing gain sequence. We shall not discuss parameter choices further in what follows.

The flow associated to the gradient descent algorithm (4.4) is obtained by taking the continuous-time limit of $t_a \rightarrow 0$:

$$\frac{da}{dt} = -\nabla_a \phi^{\sigma,x}(a). \tag{4.5}$$

Theorem 3.3.9 yields the next

Theorem 4.1.1. *Under the assumptions of Theorem 3.3.9, there exists a unique, asymptotically stable equilibrium a_∞ that satisfies the equation*

$$0 = -\nabla_a \phi^{\sigma,x}(a_\infty), \quad (4.6)$$

and for any initial condition a_0 in a closed, convex subset $\bar{U} \subset \mathbb{R}^n$, it holds that

$$|a_\infty - a_t|^2 \leq |a_0 - a_\infty|^2 \exp(-tm^x(\bar{U})), \quad (4.7)$$

for the strong convexity parameter $m^x(\bar{U})$ defined in (3.100).

Proof. Since the function $\phi^{\sigma,x}$ is strictly convex, there exists a unique global minimum a_∞ at which the gradient $\phi^{\sigma,x}$ vanishes. This proves (4.6). We first show that the equilibrium a_∞ is asymptotically stable using the so-called second method of Lyapunov. Define the Lyapunov function candidate

$$h(a) := \frac{1}{2}|a|^2. \quad (4.8)$$

The candidate function is nonnegative, satisfies $h(a') = 0$ if and only if $a' = 0$, and is radially unbounded. In order to show that the equilibrium is asymptotically stable, we shall show that the time derivative of $h(a_t - a_\infty)$ is negative definite. We shall adapt the proof of convergence of solutions to the overdamped Langevin equation in C. Villani's book [63, Chapter 2]. By (4.5) and (4.6), it follows that

$$\frac{d}{dt}(a_\infty - a_s) = -\nabla \phi^{\sigma,x}(a_\infty) + \nabla \phi^{\sigma,x}(a_s) = \nabla \phi^{\sigma,x}(a_s).$$

Thus, by the chain rule,

$$\frac{d}{dt}h(a_s - a_\infty) = (a_\infty - a_s) \cdot \nabla_a \phi^{\sigma,x}(a_s). \quad (4.9)$$

Taylor's theorem guarantees that there exists some $t \in [0, 1]$ such that $z = ta_s + (1-t)a_\infty$ satisfies

$$\phi^{\sigma,x}(a_\infty) = \phi^{\sigma,x}(a_s) + (a_\infty - a_s) \cdot \nabla \phi^{\sigma,x}(a_s) + \frac{1}{2}(a_\infty - a_s)^\top \nabla^2 \phi^{\sigma,x}(z)(a_\infty - a_s). \quad (4.10)$$

Rearranging (4.10), using the fact that a_∞ is the global minimum of $\phi^{\sigma,x}$, and using the strong convexity (3.99) of $\phi^{\sigma,x}$, we obtain

$$\begin{aligned} (a_\infty - a_s) \cdot \nabla \phi^{\sigma,x}(a_s) &= \phi^{\sigma,x}(a_\infty) - \phi^{\sigma,x}(a_s) - \frac{1}{2}(a_\infty - a_s)^\top \nabla^2 \phi^{\sigma,x}(z)(a_\infty - a_s) \\ &\leq 0 - \frac{1}{2}(a_\infty - a_s)^\top \nabla^2 \phi^{\sigma,x}(z)(a_\infty - a_s) \\ &\leq -\frac{m^x(\bar{U})}{2}|a_\infty - a_s|^2, \end{aligned} \quad (4.11)$$

for any closed, convex set $\bar{U} \subset \mathbb{R}^n$ containing a_s and a_∞ . This proves that, for $a_s \neq a_\infty$,

$$\frac{d}{dt}h(a_s - a_\infty) \leq -m^x(\bar{U})h(a_s - a_\infty) < 0. \quad (4.12)$$

Thus, the time derivative of the Lyapunov function is negative definite, and the unique equilibrium a_∞ is asymptotically stable. Finally, to show (4.7), we integrate (4.12) with respect to the time variable t in order to obtain

$$h(a_t - a_\infty) - h(a_0 - a_\infty) \leq -m^x(\bar{U}) \int_0^t h(a_s - a_\infty) ds.$$

By Gronwall's Lemma, the conclusion follows. \square

Since statements of the form (4.7) define the property of global exponential stability, we have

Corollary 4.1.2. *The unique equilibrium a_∞ is globally exponentially stable.*

Corollary 4.1.3. *For every initial condition a_0 , the solution $a = (a_t)_{t \geq 0}$ exists for all $t > 0$.*

Proof. Let a_0 be arbitrary, and let $v := \phi^{\sigma, x}(a_0)$. Since $\phi^{\sigma, x}$ is convex, the convex hull $\text{conv}(L_v)$ of the level set L_v of v ,

$$L_v := \{a \in \mathbb{R}^n \mid \phi^{\sigma, x}(a) = v\},$$

contains the unique equilibrium a_∞ . In particular, for any $a \in \text{conv}(L_v)$, $\phi^{\sigma, x}(a) \leq v$. Since $\phi^{\sigma, x}$ decreases along every solution of (4.5), it follows that any trajectory initialised in $\text{conv}(L_v)$ does not leave $\text{conv}(L_v)$. This implies that solutions of the flow do not blow up in finite time. \square

Note that (4.7) emphasises the importance of the strong convexity parameter $m^x(\bar{U})$ in determining the rate of convergence to the equilibrium a_∞ . In Remark 8, we observed that, as the dimension n of the approximating subset \mathcal{A} increases, the value of $m^x(\bar{U})$ decreases to zero, although the value of $m^x(\bar{U})$ is strictly positive whenever the supports of the basis functions satisfy Assumption 3.3.5 in Section §3.3.2.

Remark 10. The main reason why we have chosen to study the flow arising from the continuous-time limit of the gradient descent algorithm is that a detailed study of the gradient descent algorithm would require some analysis of stochastic approximation algorithms, which would be beyond the scope of this thesis. Since the pioneering works of Robbins and Monro [56] and Kiefer and Wolfowitz [38], the field has grown steadily, with fundamental contributions regarding convergence and asymptotic distributions by Blum [7], Chung [12], Dvoretzky [26], and Fabian [28], to name just a few. With regards to practical implementations, a shortcoming of the theoretical results concerning stochastic approximation algorithms is the condition that the step sizes, i.e. the t_a in (4.4), be proportional to the inverse of the iteration number, since the step sizes often decrease too rapidly to allow the descent iterates to converge. The innovations due to Spall [59] concerning simultaneous perturbation stochastic approximation, and the method of averaging iterates due to Ruppert, Polyak, and Juditsky [54, 58], provided some solutions around this constraint. Stochastic approximation remains an active field, with many results from stochastic analysis being applied to convergence analysis [60], asymptotic efficiency [52], and stability [3], as well as results from the theory of dynamical systems [5] being applied to the case of constant step sizes. One of the more recent application areas of stochastic

approximation algorithms is machine learning [45], in which one seeks to minimise a ‘loss function’ associated with learning a probability distribution from observed data. Some reviews which may ease one’s foray into what is an increasingly technical field of study include those by Bharath and Borkar [6] and Lai [41], with the section on stochastic approximation in the review [51] being quite accessible.

4.1.2 The unique equilibrium and optimal projection

In this section, we study the unique equilibrium a_∞ defined in Theorem 4.1.1, from the perspective of partial differential equations. In particular, we investigate whether the unique equilibrium yields any useful information about the solution to the Hamilton-Jacobi-Bellman equation (2.85) associated with the optimal control problem. We shall derive results that are meaningful when the unique equilibrium a_∞ is nonzero. In addition to the assumptions introduced so far, we therefore specify the following assumptions. The first assumption, Assumption 4.1.4, is intended to exclude badly chosen approximating subsets \mathcal{A} from further consideration.

Assumption 4.1.4. *Given the approximating subset \mathcal{A} spanned by the basis functions $\{\varphi_i\}_i$, there exists a parametrised control $c^a \in \mathcal{A}$ such that $\bar{\phi}^{\sigma,x}(c^a) < \bar{\phi}^{\sigma,x}(0)$.*

The next assumption is essential to the characterisation of the unique equilibrium in terms of an inner product. Recall that D is bounded and that $\varphi_i \in L^\infty(\bar{D})$ for all i , so that $\varphi_i \in L^2(\bar{D})$ for all i as well.

Assumption 4.1.5. *The function $\nabla_x F(\sigma, \cdot) : D \rightarrow \mathbb{R}^d$ is an element of the Hilbert space $L^2(D; \mathbb{R}^d)$.*

We have the following

Lemma 4.1.6. *Suppose that the random variable W is not $\mu^{0,x}$ -almost surely constant, and suppose that Assumption 4.1.4 holds. Then the unique equilibrium a_∞ is nonzero.*

Proof. The assumption on W guarantees the existence of an admissible feedback control function c such that $\bar{\phi}^{\sigma,x}(c) < \bar{\phi}^{\sigma,x}(0)$. In conjunction with Assumption 4.1.4, it follows that one can always achieve a smaller value of $\bar{\phi}^{\sigma,x}$ by ‘doing something’, i.e. when the parametrised control is nonzero. By continuity of $\phi^{\sigma,x}$, the unique equilibrium must be nonzero. \square

Denote the inner product on $L^2(D; \mathbb{R}^d)$ by $\langle \cdot, \cdot \rangle$ and the associated norm by $\|\cdot\|$. Hence, the non-overlap condition on the basis functions implies that

$$\langle \varphi_i, \varphi_j \rangle = \delta_{ij} \|\varphi_i\|^2. \quad (4.13)$$

For an arbitrary vector $a \in \mathbb{R}^n$, recall that the corresponding parametrised control $c^a \in \mathcal{A}$ is defined by

$$c^a = \sum_i a_i \varphi_i. \quad (4.14)$$

Given the optimal control

$$c_{\text{opt}}^\sigma = -2\varepsilon\sigma \nabla_x F(\sigma, \cdot), \quad (4.15)$$

we define the squared L^2 error function of the parametrised control c^a with respect to c_{opt}^σ by

$$\epsilon(a) := \left\| c_{\text{opt}}^\sigma - c^a \right\|^2 = \int_D |2\varepsilon\sigma \nabla_x F(\sigma, x) + c^a(x)|^2 dx.$$

Given (4.13), we may express the L^2 error ϵ as a quadratic polynomial in a :

$$\epsilon(a) = \|2\varepsilon\sigma \nabla_x F(\sigma, \cdot)\|^2 + 4\varepsilon\sigma \sum_i a_i \langle \nabla_x F(\sigma, \cdot), \varphi_i \rangle + \sum_i a_i^2 \|\varphi_i\|^2.$$

Define $a^* \in \mathbb{R}^n$ by the property that

$$\epsilon(a^*) = \min_{a \in \mathbb{R}^n} \epsilon(a).$$

The vector a^* gives the coordinates of the projection of the optimal control c_{opt}^σ to the approximating subset \mathcal{A} . By taking derivatives of the squared L^2 error, it follows that the i -th component of a^* satisfies

$$a_i^* = -2\varepsilon\sigma \frac{\langle \nabla_x F(\sigma, \cdot), \varphi_i \rangle}{\|\varphi_i\|^2}. \quad (4.16)$$

In order to determine whether the unique equilibrium a_∞ of the flow (4.5) equals the optimal projection coordinate vector a^* , recall that the value function $F(\sigma, x)$ of the optimal control problem

$$\min_c E^{c,x} [K^{\sigma,c,0}] \quad (4.17)$$

subject to

$$dX_t^c = [c - \nabla V](X_t^c) dt + \sqrt{2\varepsilon} dB_t \quad (4.18)$$

where $\tau = \tau(X^c) = \inf \{t > 0 \mid X_t^c \notin D\}$ is the random first exit time from D , uniquely solves the boundary value problem

$$f(x) + LF(\sigma, x) - \varepsilon\sigma |\nabla_x F(\sigma, x)|^2 = 0 \quad x \in D \quad (4.19)$$

$$F(\sigma, x) = g(x) \quad x \in \partial D, \quad (4.20)$$

where the second-order partial differential operator

$$L := \varepsilon\Delta - \nabla V \cdot \nabla$$

is the infinitesimal generator associated to the stochastic differential equation (4.18) for $c \equiv 0$. The optimal control c_{opt}^σ is defined by (4.15) precisely because (4.19) may be rewritten as (see (2.85) in Section §2.2.2)

$$\min_{c(x) \in \mathbb{R}^d} \left\{ f(x) + LF(\sigma, x) + c(x) \cdot \nabla_x F(\sigma, x) + \frac{1}{4\varepsilon\sigma} |c(x)|^2 \right\} = 0 \quad x \in D. \quad (4.21)$$

Given the relation (4.15) between the optimal control and the function $F(\sigma, \cdot)$, we define the projection of the gradient $\nabla_x F(\sigma, \cdot)$ to the approximating subset \mathcal{A} by

$$\nabla_x F^{\sigma, a^*} := - \sum_i \frac{a_i^*}{2\varepsilon\sigma} \varphi_i. \quad (4.22)$$

Observe that this definition ensures that the relation (4.15) holds in the approximating subset, i.e. that

$$-2\varepsilon\sigma\nabla_x F^{\sigma,a^*} = \sum_i a_i^* \varphi_i = c^{a^*},$$

and hence may be considered a reasonable definition. Furthermore, we have

Proposition 4.1.7. *The unique equilibrium minimises the L^2 error, i.e. $a_\infty = a^*$.*

Proof. By substituting $\nabla_x F^{\sigma,a^*}$ for $\nabla_x F(\sigma, \cdot)$ and c^a for c in (4.21), and by using (4.14) and (4.22), we obtain the projection of the Hamilton-Jacobi-Bellman partial differential equation to the approximating subset \mathcal{A} :

$$\min_{c^a(x) \in \mathbb{R}^d} \left\{ f(x) + LF^{\sigma,a^*}(x) + c^a(x) \cdot \nabla_x F^{\sigma,a^*}(x) + \frac{1}{4\varepsilon\sigma} |c^a(x)|^2 \right\} \geq 0 \quad x \in D. \quad (4.23)$$

Note that the equality in (4.21) has been replaced with an inequality in (4.23), because the set of vectors over which the minimisation problem is defined in (4.21) is a subset of the corresponding set in (4.23). For a given $x \in D$, the vector $c^a(x)$ is uniquely determined by the projection coordinate $a \in \mathbb{R}^n$, and hence we may rewrite (4.23) in terms of a minimisation over projection coordinates:

$$\min_{a \in \mathbb{R}^n} \left\{ f(x) + LF^{\sigma,a^*}(x) + c^a(x) \cdot \nabla_x F^{\sigma,a^*}(x) + \frac{1}{4\varepsilon\sigma} |c^a(x)|^2 \right\} \geq 0 \quad x \in D. \quad (4.24)$$

Given the correspondence between the statement (4.21) and the optimal control problem (4.17)–(4.18), it is natural to associate the statement (4.24) to the restricted optimal control problem, i.e. the optimisation problem (4.17) in which the argument c is constrained to the approximating subset \mathcal{A} . This association implies that the unique equilibrium a_∞ of the flow (4.5) solves the minimisation problem (4.24).

For an arbitrary fixed $x \in D$, $f(x) + LF^{\sigma,a^*}(x)$ is constant with respect to a . Hence, the left-hand side of the inequality in (4.24) reduces to the minimisation problem (compare with (2.84))

$$\min_{a \in \mathbb{R}^n} \left\{ c^a(x) \cdot \nabla_x F^{\sigma,a^*}(x) + \frac{1}{4\varepsilon\sigma} |c^a(x)|^2 \right\}.$$

On the other hand, by the non-overlap condition on the basis functions, by the form (4.14) of the parametrised control, and by the definition (4.22) of the projection of the gradient $\nabla_x F(\sigma, \cdot)$ to the approximating subset, we have

$$\begin{aligned} c^a(x) \cdot \left(\nabla_x F^{\sigma,a^*}(x) + \frac{1}{4\varepsilon\sigma} c^a(x) \right) &= \left(\sum_i a_i \varphi_i(x) \right) \cdot \left(\sum_j \frac{a_j - 2a_j^*}{4\varepsilon\sigma} \varphi_j(x) \right) \\ &= \frac{1}{4\varepsilon\sigma} \sum_i a_i (a_i - 2a_i^*) |\varphi_i(x)|^2, \end{aligned} \quad (4.25)$$

where we have used that $\varphi_i(x) \cdot \varphi_j(x) = \delta_{ij} |\varphi_i(x)|^2$ except on sets of Lebesgue measure zero, by the non-overlap condition. The right-hand side of (4.25) is quadratic in the variables $(a_i)_{1 \leq i \leq n}$, and attains its minimum when $a_i = a_i^*$ for all i , thus proving the claim. \square

Proposition 4.1.7 establishes a useful correspondence between the solution of the projected optimal control problem and the solution of the projected Hamilton-Jacobi-Bellman partial differential equation. In particular, if one can iteratively refine the approximating subset so that the error $\|c_{\text{opt}}^\sigma - c^{a^*}\|$ goes to zero in the limit, then solving the optimal control problem corresponds to solving a partial differential equation in the weak sense. Given the non-overlap and the square integrability conditions on the basis functions, the scheme thus described is suggestive of the finite element method. Since neither the dimension d of the state space nor the dimension n of the approximating subset played a role in the existence of the unique equilibrium a_∞ or in the strict convexity of the restricted control functional $\bar{\phi}^{\sigma,x}|_{\mathcal{A}}$, the preceding results indicate that the scheme thus described may be applied to solving boundary value problems defined over high-dimensional state spaces. Note that if the basis functions can be expressed as spatial gradients, i.e. if $\varphi_i = \nabla_x b_i$ for functions $b_i : D \rightarrow \mathbb{R}$ of sufficient regularity, then the function $F(\sigma, \cdot)$ can be approximated without resorting to quadrature, by directly exploiting the relation (4.22) to obtain

$$F^{\sigma, a^*} = - \sum_i \frac{a_i^*}{2\varepsilon\sigma} b_i.$$

For high-dimensional state spaces, a key constraint will be how efficiently expected values can be computed. As is the case for many Monte Carlo methods, the efficiency with which expected values can be computed depends on the variance of the associated estimators. We shall take up the issue of variance reduction in Section §4.2.

4.2 Variance reduction

Recall that, since a_∞ is the global minimum of $\phi^{\sigma,x}$, it holds that $\nabla\phi^{\sigma,x}(a_\infty) = 0$, and hence

$$\forall z \in \mathbb{R}^n, \quad z \cdot \nabla_a \phi^{\sigma,x}(a_\infty) = \text{cov}^{a_\infty,x}(K^{\sigma,a_\infty,1}, M_\tau^z) = 0, \quad (4.26)$$

where $\text{cov}^{a_\infty,x}$ denotes the covariance with respect to the measure $\mu^{a_\infty,x}$, and similarly $\text{var}^{a_\infty,x}$ denotes the variance with respect to $\mu^{a_\infty,x}$. Setting $z = a_\infty$ in the above equation and rearranging yields

$$\text{cov}^{a_\infty,x}(K^{\sigma,a_\infty,0}, M_\tau^{a_\infty}) = -\sigma^{-1} \text{var}^{a_\infty,x}(M_\tau^{a_\infty}). \quad (4.27)$$

Lemma 4.2.1. *At the unique equilibrium a_∞ , the following holds:*

$$\text{var}^{a_\infty,x}(K^{\sigma,a_\infty,0}) = \text{var}^{a_\infty,x}(K^{\sigma,a_\infty,1}) + \text{var}^{a_\infty,x}(\sigma^{-1} M_\tau^{a_\infty}). \quad (4.28)$$

Proof. Substituting (4.26) into

$$\text{var}^{z,x}(K^{\sigma,z,0}) = \text{var}^{z,x}(K^{\sigma,z,1}) - 2\text{cov}^{z,x}(K^{\sigma,z,1}, \sigma^{-1} M_\tau^z) + \text{var}^{z,x}(\sigma^{-1} M_\tau^z) \quad (4.29)$$

yields the desired conclusion. \square

The conclusion (4.28) of Lemma 4.2.1 implies that, at the unique equilibrium a_∞ , the variance of the estimator $K^{\sigma,a_\infty,0}$ that does not include the martingale is strictly larger than the variance of the estimator $K^{\sigma,a_\infty,1}$ that does include the martingale, with the

difference in the variance being proportional to the variance of the martingale at time τ . This motivates the question of whether the variance of $K^{\sigma, a_s, 1}$ is always smaller than the variance of $K^{\sigma, a_s, 0}$, where $a_s \in \mathbb{R}^n$ is any point along a solution $(a_s)_{s \geq 0}$ of the differential equation (4.5).

Variance reduction for solutions initialised at the origin Define the variance gap function $\gamma : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$\gamma(a_s) := \text{var}^{a_s, x}(K^{\sigma, a_s, 0}) - \text{var}^{a_s, x}(K^{\sigma, a_s, 1}).$$

Lemma 4.2.2. *Given the hypotheses of Lemma 4.2.1 and Lemma 4.1.6, the variance gap is strictly positive at the unique equilibrium a_∞ .*

Proof. The assertion follows from the conclusions of Lemma 4.2.1 and Lemma 4.1.6. Since a_∞ is nonzero, and since Assumption 3.3.5 guarantees that $M_\tau^{\varphi_i}$ assumes nonzero values with strictly positive $\mu^{a_\infty, x}$ -probability, the term consisting of the variance of the martingale on the right-hand side of (4.28) is strictly positive. \square

We now consider criteria for which the variance gap is positive along a solution to the flow. For an arbitrary $a_s \in \mathbb{R}^n$, it holds that

$$\frac{d}{dt} \frac{|a_s|^2}{2} = a_s \cdot \frac{da_s}{dt} = a_s \cdot (-\nabla \phi^{\sigma, x}(a_s)) = - \sum_i (a_s)_i E^{a_s, x}[K^{\sigma, a_s, 1} M_\tau^{\varphi_i}].$$

Thus

$$\frac{d}{dt} \frac{|a_s|^2}{2} = -\text{cov}^{a_s, x}(K^{\sigma, a_s, 1}, M_\tau^{a_s}). \quad (4.30)$$

It follows from (4.29) and (4.30) that we can rewrite the variance gap at some point a_s along a solution by

$$\gamma(a_s) = \frac{2}{\sigma} \frac{d}{dt} \frac{|a_s|^2}{2} + \text{var}^{a_s, x}(\sigma^{-1} M_\tau^{a_s}). \quad (4.31)$$

The equation (4.31) proves the next

Proposition 4.2.3. *The norm $|a_s|$ of a solution $a = (a_s)_{s \geq 0}$ is increasing at time $s \geq 0$ if and only if the variance gap γ is larger than the variance of $\sigma^{-1} M_\tau^{a_s}$, i.e.*

$$0 < \frac{d}{dt} |a_s|^2 \iff 0 < \gamma(a_s) - \text{var}^{a_s, x}(\sigma^{-1} M_\tau^{a_s}).$$

Since the unique equilibrium a_∞ is not known a priori, the most straightforward way to apply Proposition 4.2.3 in order to obtain variance reduction along a solution $(a_s)_s$ is to specify an initial condition near or at the origin, since Lemma 4.1.6 guarantees that the solution will move away from the origin. However, this implies that the controlled trajectories will not differ very much from the uncontrolled trajectories, at least for the initial iterations of the gradient descent algorithm. Since sampling uncontrolled trajectories is computationally inefficient, it is of interest to find another approach to variance reduction that applies for all trajectories, not just those initialised near the origin. We consider such an approach next.

Martingale-based control variate Recall that, for the unconstrained control functional $\bar{\phi}^{\sigma,x}$, the optimal control c_{opt}^{σ} has the property that the martingale-based estimator $K^{\sigma,c_{\text{opt}}^{\sigma},1}$ is $\mu^{\sigma,c_{\text{opt}}^{\sigma},x}$ -almost surely equal to the value of $F(\sigma, \cdot)$ at x ,

$$\mu^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},1} = F(\sigma, x)) = 1. \quad (4.32)$$

It follows that (see Corollary 2.2.9)

$$\text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},1}) = 0.$$

Hence, we may calculate the variance of the non-martingale based estimator $K^{\sigma,c_{\text{opt}}^{\sigma},0}$:

$$\begin{aligned} & \text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},0}) \\ &= \text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},1} - \sigma^{-1}M_{\tau}^{\sigma,c_{\text{opt}}^{\sigma}}) \\ &= \text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},1}) - 2\text{cov}^{\sigma,c_{\text{opt}}^{\sigma},x}(K^{\sigma,c_{\text{opt}}^{\sigma},1}, \sigma^{-1}M_{\tau}^{\sigma,c_{\text{opt}}^{\sigma}}) + \text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(\sigma^{-1}M_{\tau}^{\sigma,c_{\text{opt}}^{\sigma}}) \\ &= \text{var}^{\sigma,c_{\text{opt}}^{\sigma},x}(\sigma^{-1}M_{\tau}^{\sigma,c_{\text{opt}}^{\sigma}}), \end{aligned}$$

where the final equality above follows from (4.32). Thus, even at the optimal control, the variance of the non-martingale based estimator $K^{\sigma,c_{\text{opt}}^{\sigma},0}$ with respect to $\mu^{\sigma,c_{\text{opt}}^{\sigma},x}$ is strictly positive. These observations suggest that, at the optimal control c_{opt}^{σ} , the corresponding martingale is a suitable control variate. To fix ideas, we cover the essential ideas behind control variates in the the following paragraph.

Control variates Given a random variable for which one wishes to estimate the mean, a control variate is another random variable whose mean is known and that has nonzero correlation with the desired random variable. The nonzero correlation allows one to reduce the variance in the estimate of the sum of the desired random variable with the control variate. Since one knows the mean of the control variate, this allows one to improve the estimate of the mean of the desired random variables. More precisely, given an E -valued random variable X on a probability space, and given a Borel-measurable function $f : E \rightarrow \mathbb{R}$, a control variate is a random variable $g(X)$ where $g : E \rightarrow \mathbb{R}$ is Borel measurable, such that the mean $E[g(X)]$ is known. If $g(X)$ and $f(X)$ have nonzero covariance, then by defining the function

$$h^{\beta}(X) = f(X) - \beta(g(X) - E[g(X)]), \quad (4.33)$$

it follows that $E[h^{\beta}(X)] = E[f(X)]$ for all values of β , and

$$\begin{aligned} \text{var}(h^{\beta}(X)) &= \text{var}(f(X) + \beta^2\text{var}(g(X)) + 2\text{cov}(f(X), -\beta(g(X) - E[g(X)]))) \\ &= \text{var}(f(X)) + \beta^2\text{var}(g(X)) - 2\text{cov}(f(X), \beta g(X)), \end{aligned}$$

which describes a quadratic polynomial in β . The optimal value of β at which $\text{var}(h^{\beta}(X))$ is minimised is

$$\beta^* = \frac{\text{cov}(f(X), g(X))}{\text{var}(g(X))}, \quad (4.34)$$

and the corresponding optimal value of the variance of $h^{\beta}(X)$ is

$$\text{var}(h^{\beta^*}(X)) = \text{var}(f(X)) \left(1 - (\text{corr}(f(X), g(X)))^2\right), \quad (4.35)$$

where the correlation coefficient is defined by

$$\text{corr}(f(X), g(X)) = \frac{\text{cov}(f(X), g(X))}{(\text{var}(f(X))\text{var}(g(X)))^{1/2}}.$$

In particular, when the correlation coefficient of $f(X)$ and $g(X)$ equals ± 1 , then the variance of $h^{\beta^*}(X)$ equals zero.

Optimal scaling of martingale control variate Now we apply the ideas of control variates to the optimal control problem. We wish to estimate the expected value of $K^{\sigma, c, 0}$ with respect to $\mu^{c, x}$. Since $E^{c, x}[M_\tau^c] = 0$ for any admissible feedback control function c , it holds that M_τ^c is a valid control variate. We now find the optimal value of β (4.34) for the optimal control c_{opt}^σ . Letting $E = \Omega$, $f := K^{\sigma, c_{\text{opt}}^\sigma, 0}$, and $g = M_\tau^{c_{\text{opt}}^\sigma}$, the analogue of (4.33) becomes

$$h^\beta(\omega) = K^{\sigma, c_{\text{opt}}^\sigma, 0}(\omega) - \beta M_\tau^{c_{\text{opt}}^\sigma}(\omega).$$

We observed in Theorem 2.2.8 that the control c is optimal if and only if $K^{\sigma, c, 1}$ is $\mu^{c, x}$ -almost surely constant and equal to $F(\sigma, x)$. The latter statement implies, by the definition of $K^{\sigma, c, \alpha}$, that

$$-\sigma^{-1} M_\tau^{c_{\text{opt}}^\sigma}(\omega) = K^{\sigma, c_{\text{opt}}^\sigma, 0}(\omega) - F(\sigma, x)$$

$\mu^{c_{\text{opt}}^\sigma, x}$ -almost surely, and hence

$$\text{cov}^{c_{\text{opt}}^\sigma, x}(K^{\sigma, c_{\text{opt}}^\sigma, 0}, M_\tau^{c_{\text{opt}}^\sigma}) = -\sigma^{-1} \text{var}^{c_{\text{opt}}^\sigma, x}(M_\tau^{c_{\text{opt}}^\sigma}).$$

Using the above relation in (4.34) yields the optimal value

$$\beta^* = -\sigma^{-1}.$$

The variance of $h^{\beta^*}(X^{c_{\text{opt}}^\sigma})$ equals zero, which is consistent with the conclusions of Theorem 2.2.8.

Now suppose we consider the same questions as above, but for the restricted control functional $\bar{\phi}^{\sigma, x}|_{\mathcal{A}}$. In the preceding discussion, the control was optimal. Since in most practical cases the approximating subset \mathcal{A} will not include the optimal control c_{opt}^σ , we cannot assume that the optimal value of β equals σ^{-1} . Indeed, the preceding discussion made use of a specific property of the optimal control. Hence, for an arbitrary parametrised control $c^z \in \mathcal{A}$, it must be that the optimal value of β depends on $z \in \mathbb{R}^n$. For $0 \neq z \in \mathbb{R}^n$ and $\beta \in \mathbb{R}$, define

$$h^{z, \beta}(\omega) := K^{\sigma, z, 0}(\omega) - \beta M_\tau^z(\omega). \quad (4.36)$$

Note that $h^{z, 0} = K^{\sigma, z, 0}$, $h^{z, -\sigma^{-1}} = K^{\sigma, z, 1}$.

Proposition 4.2.4. *The optimal value $\beta^*(z)$ at which the variance of $h^{z, \beta}$ with respect to $\mu^{z, x}$ attains its minimum is defined by the map*

$$\begin{aligned} \beta^* : \mathbb{R}^n \setminus \{0\} &\rightarrow \mathbb{R} \\ z &\mapsto \frac{\text{cov}^{z, x}(K^{\sigma, z, 0}, M_\tau^z)}{\text{var}^{z, x}(M_\tau^z)}. \end{aligned} \quad (4.37)$$

and the corresponding variance of $h^{z,\beta^*(z)}$ equals

$$\text{var}^{z,x}(h^{z,\beta^*(z)}) = \text{var}^{z,x}(K^{\sigma,z,0}) \left(1 - \frac{(\text{cov}^{z,x}(K^{\sigma,z,0}, M_\tau^z))^2}{\text{var}^{z,x}(K^{\sigma,z,0})\text{var}^{z,x}(M_\tau^z)} \right).$$

Proof. The proposition follows from (4.34) and (4.35). \square

The bound $|\beta^*(z)| \leq (\text{var}^{z,x}(K^{\sigma,z,0})/\text{var}^{z,x}(M_\tau^z))^{1/2}$ holds, by the Cauchy-Schwarz inequality. By Proposition 4.2.4, we have the following algorithm for state-dependent variance reduction in the estimate of the function $\phi^{\sigma,x}(a^{(\ell)})$ at each iterate $a^{(\ell)}$:

for $\ell = 1$ to N_{iter} **do**

Sample N_{traj} trajectories according to the distribution $\mu^{a^{(\ell-1)},x}$;
 Compute the sample means of $(K^{\sigma,a^{(\ell-1)},0} M_\tau^{\varphi_j})_{1 \leq i \leq n}$ and $(M_\tau^{\varphi_i} M_\tau^{\varphi_j})_{1 \leq i,j \leq n}$;
 If $a^{(\ell-1)} \neq 0$, estimate $\phi^{\sigma,x}(a^{(\ell-1)})$ by the sample mean of $h^{a^{(\ell-1)},\beta^*(a^{(\ell-1)})}$, using (4.36) and (4.37). Otherwise, estimate $\phi^{\sigma,x}(0)$ by the sample mean of $K^{\sigma,0,0}$;
 Obtain the new gradient descent iterate $a^{(\ell)}$, by using (4.4), the sample means of $K^{\sigma,a^{(\ell-1)},1} M_\tau^{\varphi_i}$, and Lemma 3.3.1;

end

Algorithm 2: Gradient descent algorithm with martingale-based variance reduction

The next results concern relationships between the growth of $|a_s|$ and covariances of random variables with the martingale $M_\tau^{a_s}$. Note that, by (4.30), we have

$$\frac{d}{dt} \frac{|a_s|^2}{2} = - \left(\beta^*(a_s) + \sigma^{-1} \right) \text{var}^{a_s,x}(M_\tau^{a_s}). \quad (4.38)$$

Lemma 4.2.5. *The norm of the solution at time $s \geq 0$ is increasing if and only if $\beta^*(a_s) + \sigma^{-1}$ is negative and $a_s \neq 0$:*

$$0 \leq \frac{d}{dt} \frac{|a_s|^2}{2} \iff \beta^*(a_s) + \sigma^{-1} \leq 0.$$

Proof. Since $a_s \neq 0$, the variance of $M_\tau^{a_s}$ is strictly positive, and hence the statement follows directly from (4.38). \square

Corollary 4.2.6. *If $a_\infty \neq 0$, then $\beta^*(a_\infty) = -\sigma^{-1}$.*

Corollary shows that there exists a point in $\mathbb{R}^n \setminus \{0\}$ at which the optimal scaling function β^* equals $-\sigma^{-1}$. Lemma 4.2.5 states that the states a_s at which $|a_s|^2$ is stationary are precisely those states for which the optimal scaling function equals $-\sigma^{-1}$. The next result characterises the set of states a_s for which the rate of change of $|a_s|^2$ is stationary:

Proposition 4.2.7. *Let $a = (a_s)_{s \geq 0}$ be an arbitrary solution of the flow (4.5). For an arbitrary time $s \geq 0$, if*

$$0 = \frac{d^2}{dt^2} \frac{|a_s|^2}{2} = - \frac{d}{dt} \text{cov}^{a_s,x}(K^{\sigma,a_s,1}, M_\tau^{a_s}) \quad (4.39)$$

holds, then either

(i) *the state a_s is the unique equilibrium, in which case $\frac{d}{dt}|a_s|^2 = 0$, or*

(ii) the norm of the solution is strictly increasing at time s , i.e. $\frac{d}{dt}|a_s|^2 > 0$.

Proof. Taking the time derivative in (4.30) and applying the chain rule yields

$$\begin{aligned} \frac{d^2}{dt^2} \frac{|a_s|^2}{2} &= \frac{d}{dt} (a_s \cdot \nabla \phi^{\sigma,x}(a_s)) \\ &= \frac{da_s}{dt} \cdot \nabla \phi^{\sigma,x}(a_s) + a_s \cdot \nabla^2 \phi^{\sigma,x}(a_s) (-\nabla \phi^{\sigma,x}(a_s)) \\ &= -\nabla_a \phi^{\sigma,x}(a_s) \cdot \left(\nabla \phi^{\sigma,x}(a_s) + \nabla^2 \phi^{\sigma,x}(a_s) a_s \right). \end{aligned}$$

Therefore, by uniqueness of the gradient, (4.39) holds if and only if either $\nabla \phi^{\sigma,x}(a_s) = 0$ or $\nabla \phi^{\sigma,x}(a_s) + \nabla^2 \phi^{\sigma,x}(a_s) a_s = 0$. Note that it is not possible for both $\nabla \phi^{\sigma,x}(a_s)$ and $\nabla \phi^{\sigma,x}(a_s) + \nabla^2 \phi^{\sigma,x}(a_s) a_s$ to equal zero, since the strict convexity of $\phi^{\sigma,x}$ implies that $\nabla \phi^{\sigma,x}(a_s) = 0$ if and only if $a_s = a_\infty$, and since the invertibility of $\nabla^2 \phi^{\sigma,x}(a_s)$ implies that $\nabla^2 \phi^{\sigma,x}(a_s) a_s = 0$ if and only if $a_s = 0$. Thus we have proven conclusion (i). The condition $\nabla \phi^{\sigma,x}(a_s) + \nabla^2 \phi^{\sigma,x}(a_s) a_s = 0$ holds if and only if

$$-\nabla \phi^{\sigma,x}(a_s) = \nabla^2 \phi^{\sigma,x}(a_s) a_s. \quad (4.40)$$

Taking the inner product of both sides of (4.40) with a_s yields

$$a_s \cdot (-\nabla \phi^{\sigma,x}(a_s)) = a_s \nabla^2 \phi^{\sigma,x}(a_s) a_s \geq m |a_s|^2 > 0.$$

Since

$$\frac{d}{dt} \frac{|a_s|^2}{2} = -a_s \cdot \nabla \phi^{\sigma,x}(a_s),$$

conclusion (ii) follows. \square

In order to understand the significance of Proposition 4.2.7, recall that we observed earlier that we could obtain variance reduction along solutions, provided that the corresponding initial conditions had sufficiently small norm. Proposition 4.2.7 suggests that the restriction to solutions with growing norm was due to the fact that the prefactor σ^{-1} of the random variable $M_\tau^{a_s}$ in the estimator $K^{\sigma,a_s,1}$ was held constant at a value that is optimal for the unique equilibrium a_∞ , but suboptimal for other states a_s along the trajectory. By letting the prefactor vary as a function of the current state of the trajectory, i.e. by setting $\beta(a_s) = \beta^*(a_s)$, where the function β^* is defined in Proposition 4.2.4, we obtain variance reduction along any trajectory, not just those with increasing norm.

In this chapter, we studied the continuous-time limit of the gradient descent algorithm, and used the strong convexity result (Theorem 3.3.9) of Chapter 3 to show that the continuous-time limit of the gradient descent algorithm converges at an exponential rate to the global minimum of the approximating function $\phi^{\sigma,x}$. We exploited the vector space structure of the approximating subset \mathcal{A} and the linearity of the inner product to show that, if both the optimal control and the basis functions of the approximating subset are subsets of the Hilbert space $L^2(D)$, then the non-overlap condition establishes a correspondence between the global minimum and the coordinates of the best approximation of the optimal control in the approximating subset. If each basis function φ_i is the (weak) gradient of a scalar field, then by linearity of the derivative, the coordinates of the best approximation are related to the coordinates of the value function or free energy function $F(\sigma, \cdot)$. We then characterised the variance reduction for estimates of the approximating function that one could attain by state-dependent scaling of a martingale control variate.

Chapter 5

Conclusion

Metastable diffusions constitute an important class of models of statistical-mechanical phenomena involving complex systems, e.g. conformational changes of large molecules in the noisy environment inside a cell. Metastability makes the statistical description of these complex systems challenging: as one must wait on average a long time in order to sample the rare events that are of interest, standard simulations of complex systems will be inefficient. Furthermore, if one wishes to estimate an exponential average that is dominated by values that occur rarely, then the standard Monte Carlo method will converge slowly. These observations motivate the need for importance sampling.

The main purpose of this thesis is the theoretical study of a gradient descent algorithm for importance sampling, proposed by Hartmann and Schütte in 2012. The idea is to first convert the importance sampling problem into a stochastic optimal control problem, in which the solution of the optimal control problem corresponds to the optimal importance sampling measure. Then, by restricting the infinite-dimensional stochastic optimal control problem to an approximating subset spanned by finitely many basis functions, one obtains a finite-dimensional optimisation problem. Given an initial condition, the gradient descent algorithm terminates at the nearest local minimum of the objective function. Two important issues that remained open concerning the gradient descent algorithm were the number of local minima of the objective function, i.e. the number of local minima of the restriction of the control functional to the approximating subset, and the convergence of the gradient descent iterates to the nearest local minimum.

The main result of this thesis is the identification of sufficient conditions for which the objective function is strongly convex. Strong convexity is important because it guarantees exponential convergence of the gradient descent iterates to a unique global minimum of the objective function, thus answering the two open issues mentioned earlier. An important theoretical result that follows from uniqueness is that the global minimum provides the best approximation of the solution to the Hamilton-Jacobi-Bellman equation of the optimal control problem, given the approximating subset. The third result, which is important in the context of importance sampling, is the proof that a certain martingale term, arising from Girsanov's theorem for change of measure, is a suitable control variate. The relevance of this observation to the gradient descent algorithm is that one can identify the optimal scaling for achieving the maximum variance reduction attainable with the martingale control variate.

The strong convexity result was proven using three main ideas. The first idea is to use

Girsanov's theorem in order to obtain expressions for the mixed second partial derivatives of the objective function. The second idea is to construct a collection of pairwise independent continuous local martingales, where independence follows from the non-overlap condition on the supports of the basis functions. The third idea is to bound the spectrum of the Hessian of the objective function away from zero, using Itô's formula, pairwise independence, expressions for the mixed second partial derivatives in terms of fourth-order polynomials of the values of the martingales, and a specific uniform lower bound on the path functional of interest. Results from convex analysis then guarantee that the objective function is strongly convex.

In the proof of strong convexity, we assumed that the basis functions did not change over the iterations of gradient descent. One way to further develop the results above would be to determine a method for iteratively updating the basis functions, in order to improve the approximation quality of the approximating subset. Consider the simple case in which the basis functions are the product of a unit vector with the indicator function of the support, and that after each iteration, we only update the vector. The problem of finding the best basis functions then becomes the problem of finding the best vectors on the unit sphere in \mathbb{R}^d . Intuitively, each best unit vector is proportional to the spatial average over the corresponding support of the optimal control vector field. However, it is not immediately clear how one may efficiently update the basis functions.

Aside from the non-overlap condition, one of the strongest assumptions that we have made in obtaining the result of strong convexity is the almost-sure boundedness of the first passage time of the uncontrolled diffusion. We made this assumption in order to ensure that the value at the first passage time of the Radon-Nikodym derivative of two measures on the space of continuous paths in \mathbb{R}^d is a square integrable random variable. Square integrability of the Radon-Nikodym derivative was key to proving the existence of the first- and second-order variations in the direction of bounded perturbing functions. It would be of interest to remove the almost-sure boundedness assumption, because results from stochastic analysis only guarantee L^1 -boundedness for first passage times for the diffusions we have considered here. Given that well-known conditions from stochastic analysis such as Novikov's and Kazamaki's condition only guarantee that the Radon-Nikodym derivative is uniformly integrable, it appears that establishing the square integrability of the Radon-Nikodym derivative at an integrable stopping time might be nontrivial.

The correspondence between the best approximation of the value function of the optimal control problem and the unique global minimum of the restricted control functional was proven using the assumption that both the optimal control and the basis functions belonged to the Hilbert space $L^2(D)$, and relies on the fact that the approximating subset is in fact a vector space. By considering only those feedback controls which were linear combinations of the basis functions, we could exploit the linearity of the inner product to obtain a quadratic form for which the minimiser could be solved using the calculus.

One appealing feature of the proof is that the dimension of the state space of the controlled diffusion plays no role in establishing the correspondence. Thus, in the context of numerical methods for partial differential equations, one can in principle apply the gradient descent algorithm to solve elliptic boundary value problems defined over high-dimensional spaces. This feature of robustness with respect to dimension is characteristic of Monte Carlo methods. The proof also shows that uniqueness of the global minimum and its relation to best approximation hold, even in the limit of infinitely many basis functions.

An attractive problem suggested by the proof of the second result concerns the uniqueness of the global minimiser when one removes the non-overlap condition on the supports of the basis functions. This problem is motivated by the fact that, since the basis functions have non-overlapping supports and since the domain is bounded, an increase in the number of basis functions leads to a decrease in the convergence rate of the gradient descent iterates to the unique global minimum. Note, however, that the objective function remains strictly convex in the limit of infinitely many basis functions.

The third main result of this thesis concerned a method for variance reduction by state-dependent scaling of a martingale control variate. This result was motivated by the observation that scaling the martingale control variate by the parameter σ^{-1} did not always result in variance reduction in the estimate of the objective function. That the martingale is a suitable control variate at all is due to the fact that the objective function contains the expectation of the quadratic variation of the martingale. This fact follows from the approach we have adopted of using the Cameron-Martin-Girsanov change of measure theorem in order to parametrise the set of importance sampling measures for estimating an exponential average. Given that the partial derivatives of the objective function are expressed as expectations involving martingales, the same or similar control variates could potentially be used to perform variance reduction in the estimates of the partial derivatives.

To close, we consider two directions for future work, more general than the ones suggested so far. In this thesis, we considered the relevance of the optimal control formulation to statistical properties of diffusions that are important in computational biophysics, such as first mean passage times and committor probabilities. We observed that one of the most important statistical properties, namely the free energy difference between two metastable sets on the free energy landscape, cannot be estimated by the stochastic optimal control approach described here. This is because the running cost function that corresponds to the free energy difference does not satisfy the nonnegativity condition required for existence and uniqueness results for linear elliptic boundary value problems, and because the Feynman-Kac representation is used. On the other hand, the Bellman equation allows for real-valued running cost functions. One could improve the applicability of the stochastic optimal control approach in the field of computational biophysics by exploring the consequences of allowing for negative running cost functions.

The final direction for future work that we consider concerns the convex analysis of the unrestricted control functional. If one could show that the unrestricted control functional were convex, then this convexity would be inherited by any discretisation of the control functional. The difficulty in obtaining such a result derives from the nonlinear dependence of the distribution $\mu^{c,x}$ of the paths of the controlled diffusion. This difficulty is what necessitated our use of the Hessian in proving strong convexity: since the first-order partial derivatives contain expectations with respect to $\mu^{c,x}$ of cubic polynomials of the stopped martingales, and since cubic polynomials are not bounded from below, proving the convexity of a C^1 objective function would require fine control on the measure $\mu^{c,x}$. On the other hand, the second-order partial derivatives involve fourth-order polynomials that factorise into two quadratic polynomials; one can bound the work functional so that both quadratic polynomials are nonnegative. This removes the difficulty posed by the measure $\mu^{c,x}$.

Chapter 6

Zusammenfassung

Die Berechnung statistischer Eigenschaften von hochdimensionalen Diffusionsprozessen, zum Beispiel Differenzen der freien Energie zwischen den Konformationen von komplexen Molekülen, kann numerisch sehr aufwändig sein, insbesondere dann, wenn die zu untersuchenden Ereignisse statistisch selten auftreten. Monte-Carlo-Verfahren sind hier prinzipiell geeignet, weil ihre numerische Komplexität nicht (bzw. nur schwach) von der Dimension abhängt, allerdings ist die Konvergenz gewöhnlicher Monte-Carlo-Schätzungen im Falle seltener Ereignisse oft sehr langsam, was die Verfahren ineffizient macht. Importance Sampling ist eine Varianzreduktionsmethode, um Monte-Carlo-Abschätzungen für seltene Ereignisse praktikabel zu machen. Die Idee dabei ist es, Stichproben nach einer veränderten Wahrscheinlichkeitsverteilung zu erzeugen, unter der die seltenen Werte nicht mehr selten sind, und den Schätzer der statistischen Eigenschaften mit dem Likelihood-Quotienten zwischen der ursprünglichen und der veränderten Verteilung entsprechend umzugewichten (gemäß dem Satz von Radon-Nikodym). Hartmann und Schütte haben 2012 gezeigt, dass die Identifikation der optimalen Importance-Sampling-Verteilung auf ein Problem der optimalen Steuerung führt [35]. Eine mögliche Strategie, um das Optimalsteuerungsproblem in hohen Dimensionen zu lösen, ist die Diskretisierung der Steuerung in einem endlichdimensionalen Vektorraum, wodurch sich ein endlichdimensionales Optimierungsproblem ergibt, das z.B. durch das Verfahren des steilsten Abstiegs gelöst werden kann.

Der Schwerpunkt dieser Dissertation ist die Identifizierung von hinreichenden Bedingungen für die gleichmäßige Konvexität des endlichdimensionalen Optimierungsproblems. Aus der gleichmäßigen Konvexität werden dann die Existenz- und Eindeutigkeit von Lösungen des endlichdimensionalen Problems gefolgert und die exponentielle Konvergenz des Abstiegsverfahrens bewiesen. Um den Kreis in Bezug auf das zugrundeliegende Optimalsteuerungsproblems zu schließen, wird gezeigt, dass die eindeutige Lösung des Optimierungsproblems eine Bestapproximation der optimalen Steuerung ist, die sich aus der Lösung der Hamilton-Jacobi-Bellman-Gleichung für die zugehörige Wertefunktion ergibt.

Bibliography

- [1] R. J. ALLEN, D. FRENKEL, AND TEN WOLDE. P. R., *Simulating rare events in equilibrium or nonequilibrium stochastic systems*, J. Chem. Phys., 124 (2006), p. 024102.
- [2] R. J. ALLEN, C. VALERIANI, AND P. REIN TEN WOLDE, *Forward flux sampling for rare event simulations*, Journal of Physics: Condensed Matter, 21 (2009), p. 463102.
- [3] C. ANDRIEU, E. MOULINES, AND P. PRIOURET, *Stability of Stochastic Approximation under Verifiable Conditions*, SIAM J. Control Optim., 44 (2005), pp. 283–312.
- [4] S. ASMUSSEN, D. P. KROESE, AND R. Y. RUBINSTEIN, *Heavy tails, importance sampling and cross-entropy*, Stochastic Models, 21 (2005), pp. 57–76.
- [5] M. BENAÏM AND M. W. HIRSCH, *Stochastic approximation algorithms with constant step size whose average is cooperative*, Ann. Appl. Probab., 9 (1999), pp. 216–241.
- [6] B. BHARATH AND V. S. BORKAR, *Stochastic approximation algorithms: Overview and recent trends*, Sadhana, 24 (1999), pp. 425–452.
- [7] J. R. BLUM, *Multidimensional stochastic approximation methods*, Ann. Math. Statist., 25 (1954), pp. 737–744.
- [8] V. S. BORKAR, *A convex analytic approach to markov decision processes*, 78 (1988), pp. 583–602.
- [9] M. BOUÉ AND P. DUPUIS, *A variational representation for certain functionals of Brownian motion*, Ann. Probab., 26 (1998), pp. 1641–1659.
- [10] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Grundlehren der mathematischen Wissenschaften, Cambridge University Press, Cambridge, 2004.
- [11] G. BUSSI, A. LAIO, AND M. PARRINELLO, *Equilibrium free energies from nonequilibrium metadynamics*, Phys. Rev. Lett., 96 (2006), p. 090601.
- [12] K. L. CHUNG, *On a stochastic approximation method*, Ann. Math. Statist., 25 (1954), pp. 463–483.
- [13] G. CICCOTTI AND M. FERRARIO, *Rare events by constrained molecular dynamics*, Journal of Molecular Liquids, 89 (2000), pp. 1 – 18.
- [14] ———, *Blue moon approach to rare events*, Molecular Simulation, 30 (2004), pp. 787–793.

-
- [15] P. DAI PRA, L. MENEGHINI, AND W. J. RUNGALDIER, *Connections between stochastic control and dynamic games*, Math. Control Signals Syst., 9 (1996), pp. 303–326.
- [16] E. DARVE, D. RODRIGUEZ-GOMEZ, AND A. POHORILLE, *Adaptive biasing force method for scalar and vector free energy calculations*, J. Chem. Phys., 128 (2008), p. 144120.
- [17] B. DAVIS, *On the L^p norms of stochastic integrals and other martingales*, Duke Math. J., 43 (1976), pp. 697–704.
- [18] C. DELLAGO AND P. G. BOLHUIS, *Transition path sampling and other advanced simulation techniques for rare events*, in Advanced Computer Simulation Approaches for Soft Matter Sciences III, Christian Holm and Kurt Kremer, eds., vol. 221 of Advances in Polymer Science, Springer Berlin Heidelberg, 2009, pp. 167–233.
- [19] C. DELLAGO, P. G. BOLHUIS, F. S. CSAJKA, AND D. CHANDLER, *Transition path sampling and the calculation of rate constants*, J. Chem. Phys., 108 (1998).
- [20] N. L. DOLTSINIS, *Free energy and rare events in molecular dynamics*, in Computational Nanoscience: Do it yourself!, D. Marx J. Grotendorst, S. Blügel, ed., John von Neumann Institute for Computing, Jülich, 2006, pp. 375–387.
- [21] R. DU, V. S. PANDE, A. Y. GROSBERG, T. TANAKA, AND E. S. SHAKHNOVICH, *On the transition coordinate for protein folding*, J. Chem. Phys., 108 (1998), pp. 334–350.
- [22] P. DUPUIS, A. D. SEZER, AND H. WANG, *Dynamic importance sampling for queueing networks*, Ann. Appl. Probab., 17 (2007), pp. 1306–1346.
- [23] P. DUPUIS, K. SPILIOPOULOS, AND H. WANG, *Importance sampling for multiscale diffusions*, Multiscale Model. Simul., 10 (2012), pp. 1–27.
- [24] P. DUPUIS AND H. WANG, *Importance sampling, large deviations, and differential games*, Stochastics Stochastics Rep., 76 (2004), pp. 481–508.
- [25] R. DURRETT, *Stochastic calculus : a practical introduction*, Probability and stochastics series, CRC Press, Boca Raton, 1996.
- [26] A. DVORETZKY, *On stochastic approximation*, in Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics, Berkeley, Calif., 1956, University of California Press, pp. 39–55.
- [27] R. ELBER, *Long-timescale simulation methods*, Current Opinion in Structural Biology, 15 (2005), pp. 151 – 156.
- [28] V. FABIAN, *On asymptotic normality in stochastic approximation*, Ann. Math. Statist., 39 (1968), pp. 1327–1332.
- [29] W. H. FLEMING, *Risk sensitive stochastic control and differential games*, Commun. Inf. Syst., 6 (2006), pp. 161–177.

- [30] W. H. FLEMING AND H. M. SONER, *Controlled Markov Processes and Viscosity Solutions*, Springer, New York, 2006.
- [31] E. FOURNIÉ, J. LASRY, J. LEBUCHOUX, P. LIONS, AND N. TOUZI, *Applications of Malliavin calculus to Monte Carlo methods in finance*, Finance and Stochastics, 3 (1999), pp. 391–412.
- [32] P. L. GEISLER, C. DELLAGO, AND D. CHANDLER, *Kinetic pathways of ion pair dissociation in water*, J. Phys. Chem. B, 103 (1999), pp. 3706–3710.
- [33] D. GILBARG AND N. S. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, 2001. (Reprint of the 1998 edition).
- [34] C. HARTMANN, R. BANISCH, M. SARICH, T. BADOWSKI, AND C. SCHÜTTE, *Characterization of rare events in molecular dynamics*, Entropy, 16 (2013), pp. 350–376.
- [35] C. HARTMANN AND C. SCHÜTTE, *Efficient rare event simulation by optimal nonequilibrium forcing*, J. Stat. Mech. Theor. Exp., 2012 (2012), p. P11004.
- [36] C. JARZYNSKI, *Nonequilibrium equality for free energy differences*, Phys. Rev. Lett., 78 (1997), pp. 2690–2693.
- [37] I. KARATZAS AND S. E. SHREVE, *Brownian Motion and Stochastic Calculus*, Springer, Berlin, 1991.
- [38] J. KIEFER AND J. WOLFOWITZ, *Stochastic estimation of the maximum of a regression function*, Ann. Math. Statist., 23 (1952), pp. 462–466.
- [39] S. KULLBACK, *Information Theory and Statistics*, Wiley, New York, 1959.
- [40] S. KULLBACK AND R. A. LEIBLER, *On information and sufficiency*, Ann. Math. Statist., 22 (1951), pp. 79–86.
- [41] T. L. LAI, *Stochastic approximation: invited paper*, Ann. Statist., 31 (2003), pp. 391–406.
- [42] A. LAIO AND M. PARRINELLO, *Escaping free-energy minima*, Proc. Natl. Acad. Sci. U.S.A., 99 (2002), pp. 12562–12566.
- [43] T. LELIÈVRE, M. ROUSSET, AND G. STOLTZ, *Free Energy Computations: A Mathematical Perspective*, Imperial College Press, London, 2010.
- [44] E. J. MCSHANE, *Jensen's inequality*, Bull. Amer. Math. Soc., 43 (1937), pp. 521–527.
- [45] E. MOULINES AND F. R. BACH, *Non-asymptotic analysis of stochastic approximation algorithms for machine learning*, in Advances in Neural Information Processing Systems 24, J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, eds., Curran Associates, Inc., 2011, pp. 451–459.
- [46] A. A. NOVIKOV, *On Stopping Times for a Wiener Process*, Theor. Prob. Appl., 16 (1971), pp. 449–456.

-
- [47] B. K. ØKSENDAL, *Stochastic Differential Equations: An Introduction With Applications*, Springer, Berlin, 2003.
- [48] L. ONSAGER, *Initial recombination of ions*, Phys. Rev., 54 (1938), pp. 554–557.
- [49] A. OWEN AND Y. ZHOU, *Safe and effective importance sampling*, J. Am. Statist. Assoc., 95 (2000), pp. 135–143.
- [50] D. PASSERONE AND M. PARRINELLO, *Action-derived molecular dynamics in the study of rare events*, Phys. Rev. Lett., 87 (2001), p. 108302.
- [51] R. PASUPATHY AND S. KIM, *The stochastic root-finding problem: Overview, solutions, and open questions*, ACM Trans. Model. Comput. Simul., 21 (2011), pp. 19:1–19:23.
- [52] M. PELLETIER, *Asymptotic Almost Sure Efficiency of Averaged Stochastic Algorithms*, SIAM J. Control Optim., 39 (2000), pp. 49–72.
- [53] G. PEŠKIR, *On the exponential Orlicz norms of stopped Brownian motion*, Studia Mathematica, 117 (1996), pp. 253–273.
- [54] B. POLYAK AND A. JUDITSKY, *Acceleration of stochastic approximation by averaging*, SIAM J. Control Optim., 30 (1992), pp. 838–855.
- [55] D. REVUZ AND M. YOR, *Continuous Martingales and Brownian Motion*, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, 2005. Corrected third printing of third edition.
- [56] H. ROBBINS AND S. MONRO, *A stochastic approximation method*, Ann. Math. Statist., 22 (1951), pp. 400–407.
- [57] R. Y. RUBINSTEIN, *Optimization of computer simulation models with rare events*, Eur. J. Oper. Res., 99 (1997), pp. 89 – 112.
- [58] D. RUPPERT, *A newton-raphson version of the multivariate robbins-monro procedure*, The Annals of Statistics, 13 (1985), pp. 236–245.
- [59] J. C. SPALL, *Multivariate stochastic approximation using a simultaneous perturbation gradient approximation*, IEEE T. Automat. Control, 37 (1992), pp. 332–341.
- [60] E. VALKEILA AND A. V. MELNIKOV, *Martingale Models of Stochastic Approximation and Their Convergence*, Theor. Probab. & Appl., 44 (2000), pp. 333–360.
- [61] E. VANDEN-EIJNDEN AND J. WEARE, *Rare event simulation of small noise diffusions*, Comm. Pure Appl. Math., 65 (2012), pp. 1770–1803.
- [62] A. J. VERETENNIKOV, *On strong solutions and explicit formulas for solutions of stochastic integral equations*, Mathematics of the USSR-Sbornik, 39 (1981), p. 387.
- [63] C. VILLANI, *Optimal transport: old and new*, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, 2009.

-
- [64] A. F. VOTER, *Hyperdynamics: Accelerated molecular dynamics of infrequent events*, Phys. Rev. Lett., 78 (1997), pp. 3908–3911.
- [65] ———, *A method for accelerating the molecular dynamics simulation of infrequent events*, J. Chem. Phys., 106 (1997), pp. 4665–4677.
- [66] W. ZHANG, H. WANG, C. HARTMANN, M. WEBER, AND C. SCHÜTTE, *Applications of the cross-entropy method to importance sampling and optimal control of diffusions*, SIAM J. Sci. Comput., 36 (2014), pp. A2654–A2672.
- [67] A. K. ZVONKIN, *A transformation of the phase space of a diffusion process that removes the drift*, Mathematics of the USSR-Sbornik, 22 (1974), p. 129.