

Numerical methods for stochastic optimal control: Applications in rare event estimation and wireless networks

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH Aachen University zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Shyam Mohan Subbiah Pillai, M. Sc.

aus Manama, Bahrain

Berichter: Univ.-Prof. Dr. Raúl Tempone
Asst. Prof. Dr. Nadhir Ben Rached
Univ.-Prof. Dr. Ajay Jasra

Tag der mündlichen Prüfung: 2. Juni 2025

Diese Dissertation ist auf den Internetseiten der Universitätsbibliothek online verfügbar.

Eidesstattliche Erklärung

Shyam Mohan Subbiah Pillai erklärt hiermit, dass diese Dissertation und die darin dargelegten Inhalte die eigenen sind und selbstständig, als Ergebnis der eigenen originären Forschung, generiert wurden.

Hiermit erkläre ich an Eides statt

1. Diese Arbeit wurde vollständig oder größtenteils in der Phase als Doktorand dieser Fakultät und Universität angefertigt;
2. Sofern irgendein Bestandteil dieser Dissertation zuvor für einen akademischen Abschluss oder eine andere Qualifikation an dieser oder einer anderen Institution verwendet wurde, wurde dies klar angezeigt;
3. Wenn immer andere eigene- oder Veröffentlichungen Dritter herangezogen wurden, wurden diese klar benannt;
4. Wenn aus anderen eigenen- oder Veröffentlichungen Dritter zitiert wurde, wurde stets die Quelle hierfür angegeben. Diese Dissertation ist vollständig meine eigene Arbeit, mit der Ausnahme solcher Zitate;
5. Alle wesentlichen Quellen von Unterstützung wurden benannt;
6. Wenn immer ein Teil dieser Dissertation auf der Zusammenarbeit mit anderen basiert, wurde von mir klar gekennzeichnet, was von anderen und was von mir selbst erarbeitet wurde;
7. Teile dieser Arbeit wurden zuvor veröffentlicht und zwar in:
 - [1] Ben Rached, Nadhir and Haji-Ali, Abdul-Lateef and Subbiah Pillai, Shyam Mohan and Tempone, Raúl. "Double-loop importance sampling for McKean–Vlasov stochastic differential equation." *Statistics and Computing* 34, no. 6 (2024): 1-25.
 - [2] Ben Rached, Nadhir and Haji-Ali, Abdul-Lateef and Subbiah Pillai, Shyam Mohan and Tempone, Raúl. "Multilevel importance sampling for rare events associated with the McKean–Vlasov equation." *Statistics and Computing* 35, no. 1 (2025).
 - [3] Ben Rached, Nadhir and Subbiah Pillai, Shyam Mohan and Tempone, Raúl. Optimal Power Procurement for Green Cellular Wireless Networks Under Uncertainty and Chance Constraints. *Entropy* 27, no. 3: 308 (2025).

(Shyam Mohan Subbiah Pillai)

In good and bad times, remember: this too shall pass.

Abstract

This thesis develops numerical methods for non-standard stochastic optimal control (SOC) problems, driven by real-world applications in rare event estimation and energy-efficient wireless networks. The underlying process models are controlled stochastic differential equations. In the first part, we introduce an efficient Monte Carlo (MC) estimator of rare event probabilities associated with the McKean–Vlasov stochastic differential equation (MV-SDE), crucial for analyzing mean-field systems in statistical physics, mathematical finance and many more applications. Using SOC, we derive an optimal importance sampling (IS) measure change that minimises the estimator’s relative statistical error. We then combine IS with hierarchical sampling techniques, like multilevel and multi-index MC, to enhance its computational complexity. The resulting multi-index double loop Monte Carlo (MIDLMC) estimator achieves a significantly improved computational complexity of $\mathcal{O}\left(\text{TOL}_r^{-2} \left(\log \text{TOL}_r^{-1}\right)^2\right)$ for estimating rare event probabilities with a prescribed relative accuracy TOL_r . In the second part, we develop a modelling and numerical framework to solve a chance-constrained SOC problem in cellular wireless networks, where the objective is to compute an optimal short-term power procurement strategy that minimises both operating expenditure and carbon footprint. The model accounts for uncertain renewable energy sources, stochastic wireless channels, and a probabilistic quality-of-service (QoS) constraint, making it a challenging SOC problem. The solution procedure involves a continuous-time Lagrangian relaxation of the QoS constraint, a computationally efficient numerical scheme to solve the Hamilton–Jacobi–Bellman (HJB) partial differential equation (PDE) associated with the relaxed problem, and an optimization framework for the non-smooth dual problem, enabling effective handling of the probabilistic constraint. The proposed numerical procedure provides near-optimal policies for a model cellular base station powered by a hybrid energy system, using German grid and cellular user data. Results demonstrate that our approach delivers solutions in a practical time-frame, emphasizing its computational efficiency and real-world applicability.

Zusammenfassung

In dieser Arbeit werden numerische Methoden für nicht standardisierte **SOC**-Probleme entwickelt, die durch Anwendungen in der Schätzung von Rare Events und energieeffizienter drahtloser Netzwerke motiviert werden. Die grundlegende Prozessmodelle sind kontrollierte stochastische Differentialgleichungen. Im ersten Teil stellen wir einen effizienten **MC**-Schätzer für Rare Event Wahrscheinlichkeiten von McKean–Vlasov stochastische Differentialgleichungen (**MV-SDE**), die für die Analyse von Mean-Field-Systemen in der statistischen Physik, der Finanzmathematik und vielen weiteren Anwendungen wichtig sind. Unter Verwendung von **SOC** leiten wir eine optimale **IS**-Maßänderung ab, die den relativen statistischen Fehler des Schätzers minimiert. Anschließend kombinieren wir **IS** mit hierarchischen **MC**-Methoden wie Multilevel- und Multi-index **MC**, um seine Rechenkomplexität zu reduzieren. Der resultierende **MIDL**MC-Schätzer erreicht eine deutlich verbesserte Rechenkomplexität von $\mathcal{O}\left(\text{TOL}_r^{-2} \left(\log \text{TOL}_r^{-1}\right)^2\right)$ für die Schätzung von Wahrscheinlichkeiten mit einer vorgegebenen relativen Genauigkeit TOL_r . Im zweiten Teil entwickeln wir einen Modellierungs- und numerischen Rahmen zur Lösung eines zufallsbeschränkten **SOC**-Problems in zellularen drahtlosen Netzwerken, bei dem das Ziel darin besteht, eine optimale kurzfristige Strombeschaffungsstrategie zu berechnen, die sowohl die Betriebskosten als auch den CO₂-Fußabdruck minimiert. Das Modell berücksichtigt unsichere erneuerbare Energiequellen, stochastische drahtlose Kanäle und eine probabilistische **QoS**-Bedingung, was es zu einem anspruchsvollen **SOC**-Problem macht. Das Lösungsverfahren beinhaltet eine Lagrange-Relaxation der **QoS**-Bedingung, ein rechnerisch effizientes numerisches Schema zur Lösung der **HJB** partielle Differentialgleichung für den entspannten Problem, und einen Optimierungsrahmen für das nicht-glatte duale Problem, der einen effektiven Umgang mit der probabilistischen Bedingung ermöglicht. Das vorgeschlagene numerische Methode liefert optimale Strategien für eine zellulare Station, die mit einem hybriden Energiesystem betrieben wird, unter Verwendung von deutschen Netz- und zellularen Nutzerdaten. Die Ergebnisse zeigen, dass unser Methode Lösungen in einem praktischen Zeitrahmen liefert, was seine Berechnungseffizienz und seine Anwendbarkeit in der Praxis unterstreicht.

Publications

- Ben Rached, Nadhir and Subbiah Pillai, Shyam Mohan¹ and Tempone, Raúl. "Optimal Power Procurement for Green Cellular Wireless Networks Under Uncertainty and Chance Constraints." *Entropy* 27, no. 3: 308 (2025).
- Ben Rached, Nadhir and Haji-Ali, Abdul-Lateef and Subbiah Pillai, Shyam Mohan¹ and Tempone, Raúl. "Multilevel importance sampling for rare events associated with the McKean–Vlasov equation." *Statistics and Computing* 35, no. 1 (2025).
- Ben Rached, Nadhir and Haji-Ali, Abdul-Lateef and Subbiah Pillai, Shyam Mohan¹ and Tempone, Raúl. "Double-loop importance sampling for McKean–Vlasov stochastic differential equation." *Statistics and Computing* 34, no. 6 (2024): 1-25.

Parts of this work have already been published in a similar form and are reproduced herein: Sections 3.2, 3.3, 4.4, 4.5, and Appendices A.1, A.2, A.3, B.1, B.2 have been published as part of [1]. Section 4.6, and Appendices A.4, B.3, B.4 have been published as part of [2]. Sections 5.2, 5.3, 5.4 and Appendices B.7, B.8, B.9 and most parts of Appendix D have been published as part of [3].

¹Corresponding author

Acknowledgements

First, I would like to express my heartfelt gratitude to my supervisor, Prof. Raúl Tempone, for his unwavering support throughout my research journey and for giving me the opportunity to be a part of the Chair of Mathematics for Uncertainty Quantification, supported by the prestigious Alexander von Humboldt professorship. I am especially thankful to him for introducing me to this exciting and challenging research topic, and for his continuous guidance through scientific advice, insightful remarks, and constructive feedback. I am always awestruck by his remarkable ability to combine meticulous attention to detail with a broad, strategic perspective on our research. I am deeply grateful for the opportunity to participate in world-class conferences and summer schools, and for being introduced to his worldwide network of collaborators. As he often says, I too am looking forward to more.

Next, I would like to thank my scientific advisor, Dr. Nadhir Ben Rached, for his valuable day-to-day guidance and our many fruitful discussions throughout each step of this thesis. It has been a joy to work with someone so passionate about tackling new challenges and driving results. Over the course of my research journey, he has come to feel like an elder brother—offering not only invaluable scientific and career advice, but also bringing a sense of warmth and joy to the work we do. I truly hope this partnership continues for a long time to come.

I extend my special gratitude to my collaborator, Dr. Abdul-Lateef Haji-Ali, for his valuable input and feedback on the first part of this thesis. I am especially thankful for his help in shaping the right research questions, as well as for his insightful scientific and technical comments. I hope to one day acquire his remarkable critical thinking skills. I also wish to thank Prof. Alexander Mitsos from Forschungszentrum Jülich for generously sharing his insights on the second part of this thesis.

I would like to thank all my colleagues at the Chair of Mathematics for Uncertainty Quantification at RWTH Aachen for making me feel at home among them—whether during our lunch breaks, Glühwein sessions, or work trips. Special thanks to my office mate, Sophia Münker, for putting up with me over the past three years (and perhaps even longer!).

I would like to thank my fellow PhD students in the HDS-LEE graduate school. I truly enjoyed the networking opportunities with colleagues who shared various challenges, and broadened my perspectives on mathematical research. The annual retreats have always been a highlight of the year.

I would like to thank Prof. Jesper Ooppelstrup for sharing his years of experience, and helping refine the presentation of the thesis content. Moreover, I want to thank Mathias Horres, Anamika Pandey, Manuel Kleinschmager, Leon Wilkosz, Saifeddine Ben Naamia, Michael Samet, Shreyans Sakhare, and André Carlon for proof-reading.

I would also like to thank my friends in Aachen—my football mates, fellow members of Achoir Joy, the "Prog Ex" crew, and the "Innsbrück" gang—for keeping me sane and grounded, and for helping me blow off some steam during the weekends. Last but not least, I would like to thank my family—Amma, Appa, Raj, Saranya, and Ved—whose unwavering

support and prayers have made it possible for me to reach this point.

This work was performed as part of the Helmholtz School for Data Science in Life, Earth and Energy (HDS-LEE) and received funding from the Helmholtz Association of German Research Centres.

List of Abbreviations

- SDE** stochastic differential equation
- SOC** stochastic optimal control
- MC** Monte Carlo
- HJB** Hamilton–Jacobi–Bellman
- PDE** partial differential equation
- CDF** cumulative distribution function
- PDF** probability density function
- EM** Euler–Maruyama
- KBE** Kolmogorov backward equation
- CLT** Central limit theorem
- iid** independent and identically distributed
- ODE** ordinary differential equation
- IS** importance sampling
- MV-SDE** McKean–Vlasov stochastic differential equation
- QoS** quality-of-service
- LMBM** limited memory bundle method
- SSM** stochastic subgradient method
- 1D** one-dimensional
- DLMC** double loop Monte Carlo
- FDM** finite difference method
- MLMC** multilevel Monte Carlo
- MLDLMC** multilevel double loop Monte Carlo
- MIDLMC** multi-index double loop Monte Carlo
- a.s.** almost-sure
- SNR** signal-to-noise ratio

dB decibel

CFL Courant–Friedrichs–Lewy

List of Figures

1.1	Sample paths of a geometric Brownian motion starting at $X(0) = 100$ simulated using standard MC and IS. Suppose one wishes to estimate $\mathbb{P}[X(T) > K]$, a rare event probability. The horizontal line at level $K = 200$ denotes the rare event threshold. Standard MC samples yield few trajectories crossing K at final time $T = 1$, resulting in MC estimators with high coefficient of variation. In contrast, IS significantly increases the probability of sample paths hitting the rare event region, improving efficiency of MC estimation of rare event probabilities.	2
1.2	Snapshots of a swarm of interacting particles in two-dimensional space at two different times, simulated using the Helbing particle-based pedestrian flow model [4], a MV-SDE in the mean-field limit. Each pedestrian's movement is influenced by both its own dynamics and the collective distribution of the crowd, modelled by mean-field interactions. The evolution of the spatial distribution over time illustrates the collective behaviour emerging from pairwise interactions in the model.	3
4.1	Numerical rate verification for the one-dimensional (1D) Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2$: Figures 4.1a and 4.1b numerically verify Assumptions 4.5.1 and 4.5.2 respectively, implying $\mathcal{O}(P^{-1} + N^{-1})$ error convergence rate for the relative bias of the double loop Monte Carlo (DLMC) estimator.	47
4.2	Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2$: Figures 4.2a and 4.2b numerically verify Assumptions 4.5.3 and 4.5.4 respectively. Both plots imply that the variance of the DLMC estimator converges to $\text{Var}[g(\mathbf{X}(T))]$ asymptotically in the limit $P, N \rightarrow \infty$	48
4.3	Numerical demonstration of variance reduction due to proposed IS scheme for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2$: Figure 4.3a shows reduction in $V_2^{P N}$ by a factor of 6000 with IS, leading to a reduction in the overall variance of the DLMC estimator by around 1000 times as seen in Figure 4.3b. We conclude that the proposed IS scheme reduces $V_2^{P N}$ significantly to reduce the overall variance of the DLMC estimator.	49

4.4 Numerical results of applying Algorithm 2 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2$: Figure 4.4b shows that the proposed DLMC estimator with IS achieves a relative error below the prescribed TOL_r in the sense of (4.14) to estimate a rare event probability of around 2.5×10^{-4} (Figure 4.4a). Figures 4.4c and 4.4d show that $\mathcal{A}_{\text{DLMC}}$ achieves the relative accuracy requirement with a computational work of $\mathcal{O}(\text{TOL}_r^{-4})$. Additionally, Figure 4.4d shows that the proposed IS significantly reduces computational work of the DLMC estimator by a factor of 10^3 50

4.5 Numerical demonstration of variance reduction in the DLMC estimator for level differences due to proposed IS scheme for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}$ for $K = 2.5$: Figure 4.5a shows reduction in $V_{2,\ell}$ by a factor of 100 with IS, leading to a reduction in the overall variance of the DLMC estimator of level differences by around 10 times as seen in Figure 4.5b. We conclude that the proposed IS scheme significantly reduces $\text{Var}[\Delta G_\ell]$, even though the proposed IS scheme minimises $\text{Var}[G_\ell]$. 59

4.6 Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}$ for $K = 2.5$: Figures 4.6a and 4.6b numerically verify Assumptions 4.6.1 and 4.6.2, respectively, using the antithetic sampler (4.56). These plots together determine the rates $\tilde{\alpha} = 1$, $\tilde{w} = 2$, and $\tilde{s} = 1$ 61

4.7 Numerical results of applying Algorithm 3 for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}$ for $K = 2.5$: Figure 4.7a shows that the proposed multilevel double loop Monte Carlo (MLDLMC) estimator with IS achieves a relative error below the prescribed TOL_r in the sense of (4.14) to estimate a rare event quantity of around 3.2×10^{-3} . Figures 4.7b and 4.7c show that $\mathcal{A}_{\text{MLDLMC}}$ achieves the relative accuracy requirement with a computational work of $\mathcal{O}(\text{TOL}_r^{-3})$, one order lesser than that of $\mathcal{A}_{\text{DLMC}}$. Figure 4.7d depicts the independence of the complexity rate on the event rarity. 62

4.8 Schematic illustration of an index set \mathcal{I} with its outer boundary $\partial\mathcal{I}$ as defined in (4.124). 71

4.9 Numerical demonstration of variance reduction for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}(x)$ for $K = 3.5$: Squared coefficient of variation of the DLMC estimator of $\mathbb{E}[\Delta G_{(2,2)}]$ with and without IS with respect to the number of sample paths in the outer loop M_1 for fixed $M_2 = 10^2$. The proposed IS scheme reduces the overall variance $\text{Var}[\Delta G_{(2,2)}]$ by a factor of around 10, even though the proposed IS scheme minimises $\text{Var}[G_\alpha]$ 74

4.10 Numerical demonstration of variance reduction in the DLMC estimator for the first-order mixed differences using the proposed IS scheme on the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}(x)$ for $K = 3.5$: Ratio \mathcal{R} indicates the factor of variance reduction due to the proposed IS scheme. Even though the proposed IS scheme minimises $\text{Var}[G_\alpha]$ as seen in Figure 4.10a, Figure 4.10b depicts significant variance reduction in the mixed differences for coarser multi-indices as well, although the reduction factor reduces as we go finer. 75

4.11	Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$: All plots verify Assumptions 4.7.1 and 4.7.2 using the antithetic sampler (4.83) in the DLMC estimator for the first-order mixed differences. These plots imply the following convergence rates - $\{b_1, b_2\} = \{1, 1\}$, $\{w_1, w_2\} = \{2, 2\}$, and $\{s_1, s_2\} = \{1.5, 2\}$. These plots also demonstrate increased convergence rates of first-order mixed differences, when compared to level differences.	76
4.12	Numerical rate verification on the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$: All plots verify Assumptions 4.7.1 and 4.7.2 for all $\alpha \in \mathbb{N}^2$ using the antithetic sampler (4.83) in the DLMC estimator for the first-order mixed differences. These plots confirm the following convergence rates - $\{b_1, b_2\} = \{1, 1\}$, $\{w_1, w_2\} = \{2, 2\}$, and $\{s_1, s_2\} = \{1.5, 2\}$	77
4.13	Comparing actual and approximate profits of multi-indices in the MIDLMC estimator (4.82) for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$. These plots show that the optimal index set (4.128) is a good proxy for the true optimal index set (4.129).	78
4.14	Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$. Figure 4.14a shows that both MIDLMC and MLDLMC estimators achieve a relative error below TOL_r in the sense of (4.14) to estimate a rare event quantity of around 2.04×10^{-5} . Figure 4.14b reveals that the MIDLMC estimator uses the same maximum P as the MLDLMC estimator, while using finer N discretization to achieve the same relative tolerance.	80
4.15	Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$: Figures 4.15a and 4.15b demonstrate that the MIDLMC and MLDLMC estimators satisfy constraints (4.86) and (4.87), implying optimal choices of discretization parameters and sample numbers for both estimators.	81
4.16	Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{J}}_{\{x>K\}}(x)$ for $K = 3.5$: Figures 4.16a and 4.16b show that $\mathcal{A}_{\text{MIDLMC}}$ achieves the relative accuracy requirement (4.14) with a computational work of $\mathcal{O}\left(\text{TOL}_r^{-2} \left(\log \text{TOL}_r^{-1}\right)^2\right)$, one order lesser than that of $\mathcal{A}_{\text{MLDLMC}}$	82
5.1	Schematic illustration of the power and information flow in a cellular wireless network with multiple base stations. Figure taken from [5].	88
5.2	Schematic illustration of the power flow in a base station in a cellular wireless network.	89
5.3	Schematic illustration of the proposed initialization Algorithm 15.	106
5.4	Schematic illustration of the proposed dual optimization Algorithm 5.	107
5.5	Typical daily cellular user traffic profile, described by (5.54).	108
5.6	Visualizing the spatial distribution of mobile users in urban zones using simple analytical two-dimensional distributions.	109
5.7	Results of numerical simulation of the stochastic differential equation (SDE) in (5.7) governing the Nakagami wireless channel dynamics.	110
5.8	Normalized wind power forecast and real production data in Germany in the region operated by <i>50Hertz</i> in 2024. Mean path and 95% confidence intervals of the SDE in (5.13) calibrated from 2023 data.	111
5.9	Simple characteristic curve of a battery, described by Eqs. (5.14) and (5.15).	112

5.10	Day-ahead spot-price for grid power in Germany on 07.09.2023. Data are from [6].	112
5.11	Numerical verification of the error convergence rate of the PDE discretization with respect to N_x of the proposed upwind finite difference method (FDM) scheme in Section 5.3.1 for the HJB PDE (5.36) with given inputs in Table D.1 and Section 5.4.1.	113
5.12	Slices of the numerical value function v on the axes of time and state variables. For all slices on each component of \mathbf{x} , the rest of the state variables are fixed at 0.	114
5.13	Numerical verification of the dual function's concavity. Dual cost and subgradient values computed with λ^1 constructed from different values of Y_1^1 for $\ell = 1$. Both plots numerically confirm that the finite-dimensional dual maximization Problem 5.2.4 is a convex, non-smooth problem. . . .	116
5.14	Dual cost and subgradient norm computed with λ^1 constructed from Y_1^1 attained at each iteration of Algorithm 15, along with the corresponding quantities using the final output of limited memory bundle method (LMBM) routine [7]. The initialization routine stops within three iterations for the considered problem. The advanced, non-smooth optimization LMBM routine provides an even more optimal point for the upcoming stochastic subgradient method (SSM).	117
5.15	Dual cost and subgradient norm computed with λ^2 constructed from Y^2 at each SSM iteration of Algorithm 5 for $\ell = 2$. The SSM at the given refinement level converges within 20 iterations, suggesting that a good starting point was used.	118
5.16	Dual cost and subgradient norm computed with λ^ℓ constructed from the optimal Y^ℓ obtained from the SSM at each refinement level ℓ of Algorithm 5. The SSM converges within k_{\max} iterations at each refinement level ℓ , indicating the efficiency of the LMBM-boosted initialization procedure and a good choice of step sizes for the SSM.	119
5.17	Evolution of the optimal Lagrange multiplier function $\lambda^\ell(t)$ at each refinement level ℓ in Algorithm 5. The spikes in the Lagrange multiplier function coincide with periods of high cellular traffic.	120
5.18	Optimal power procurement policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5. These curves are the practical outputs received by the base station operator to help make decisions regarding power procurement in the next 24 hours.	121
5.19	MC estimates of the violation in constraint $\mathbb{P}[\phi_{\text{out}}(t) > \phi_{\text{th}}] - \epsilon$ for each $t \in [0, T]$ using optimal controls from solving (a) Problem 5.4.2 and (b) Problem 5.2.4 using Algorithm 5. The unconstrained Problem 5.4.2 provides a solution that does not satisfy the QoS at all. The solution to the dual Problem 5.2.4 satisfies the probabilistic QoS constraint, time-averaged on the Lagrange multiplier discretization.	122
5.20	Optimal dual cost and subgradient norms obtained in 50 independent runs of Algorithm 5 with randomized model and algorithmic parameters as specified in Table D.4. Figure 5.20b demonstrates that Algorithm 5 achieves the prescribed subgradient norm tolerance in all 50 random scenarios, having a wide range of corresponding optimal costs as seen in Figure 5.20a.	125

D.1	Subgradient norm computed with λ^1 constructed from Y_1^1 attained at each iteration of Algorithm 15 for different scenarios, along with the corresponding quantities using the final output of the LMBM routine [7].	169
D.2	Dual cost computed with λ^ℓ constructed from the optimal \mathbf{Y}^ℓ obtained from the SSM at each refinement level ℓ of Algorithm 5 for different scenarios.	170
D.3	Subgradient norm computed with λ^ℓ constructed from the optimal \mathbf{Y}^ℓ obtained from the SSM at each refinement level ℓ of Algorithm 5 for different scenarios.	171
D.4	Optimal base station power consumption ($C_{\text{scaI}}N_u(t)P_{\text{tx}}(t) + C_{\text{offset}}$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.	172
D.5	Optimal battery power usage policy ($P_A(t)$) for transmission, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.	173
D.6	Optimal grid power purchase ($P_F(t)$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.	174
D.7	Optimal power selling ($P_S(t)$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.	175
D.8	Optimal battery charge retention $A(t)$ policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.	176
D.9	Evolution of the optimal Lagrange multiplier function $\lambda^\ell(t)$ at each refinement level ℓ in Algorithm 5 for different scenarios.	177
D.10	MC estimates of the violation in constraint $\mathbb{P}[\phi_{\text{out}}(t) > \phi_{\text{th}}] - \epsilon$ for each $t \in [0, T]$ using optimal controls obtained from Algorithm 5 for different scenarios.	178

List of Tables

4.1	Comparing the average number of samples required for the naive DLMC and proposed DLMC with IS methods to estimate $\mathbb{P}[X(T) > K]$ up to prescribed relative tolerance TOL_r for the 1D Kuramoto model (4.12) for different values of K . The proposed IS scheme drastically reduces the number of samples M_1 and M_2 required to achieve a relative accuracy TOL_r and also ensures the required number of samples does not increase as the event rarity increases.	52
4.2	Summary of verified numerical rates and computational complexities of the various estimators for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}(x)$. In the best case, the MIDLMC estimator has the same computational complexity as that of MC sampling without discretization errors (up to logarithmic terms). However, it requires smoothening of the indicator function and stricter regularity assumptions with respect to Assumptions 4.7.1 and 4.7.2, when compared to the DLMC and MLDLMC estimators.	84
5.1	Comparing the optimal dual costs achieved in various stages of Algorithm 5 with those of Problems 5.4.1 and 5.4.2. The LMBM procedure already gives a decent solution to the dual problem, as there is not much difference between the LMBM solution and the final dual cost obtained. As the Lagrange multiplier becomes finer, we successively obtain more optimal solutions.	120
5.2	Comparing the optimal dual costs (in €) achieved with the dual optimization Algorithm 5 for each scenario with the optimal primal costs of the corresponding unconstrained and almost-sure (a.s.) QoS constrained problems. In all scenarios, the optimal cost of the chance-constrained problem lies in between the unconstrained and a.s. constrained problems.	123
5.3	Comparison of the expected energy balance (in Wh) using the optimal power procurement policy from the optimal value function v obtained from the dual optimization Algorithm 5 for each scenario.	124
D.1	Parameters and coefficients for modeling the cellular base station described in Figure 5.2.	160
D.2	Simulation parameters to run Algorithm 5 and produce numerical results as described in Section 5.4.	167
D.3	Parameters to run the LMBM routine [7] in Algorithm 5 to produce numerical results described in Section 5.4. See [7] for more details.	167
D.4	Distributions of model and algorithmic parameters used to run randomized multi-scenario simulations of Algorithm 5.	168

Contents

Abstract	v
Zusammenfassung	vi
Publications	vii
Acknowledgements	ix
List of Abbreviations	x
List of Figures	xvi
List of Tables	xvii
1 Introduction	1
2 Preliminaries	6
2.1 Common notation	6
2.2 Random variables	6
2.3 Stochastic differential equations	11
2.4 Optimal control of stochastic differential equations	14
2.4.1 Formulation of a stochastic optimal control problem	15
2.4.2 Dynamic programming principle	16
3 Importance sampling via stochastic optimal control	21
3.1 Motivation	21
3.2 Optimal importance sampling for stochastic differential equations	24
3.3 Relating optimal importance sampling control with the Kolmogorov backward equation	29
4 Importance sampling for the McKean–Vlasov stochastic differential equation	31
4.1 McKean–Vlasov stochastic differential equation	32
4.1.1 Strong approximation: The stochastic interacting particle system	33
4.1.2 Examples	34
4.2 Problem statement	35
4.3 Naïve Euler–Maruyama Monte Carlo estimator	35
4.4 Importance sampling using the decoupling approach	36
4.5 Double loop Monte Carlo estimator	39
4.5.1 Error analysis	42
4.5.2 Computational complexity	44
4.5.3 Adaptive algorithm	44

4.5.4	Numerical experiments	45
4.6	Multilevel double loop Monte Carlo estimator	52
4.6.1	Error analysis	54
4.6.2	Computational complexity	55
4.6.3	Adaptive algorithm	57
4.6.4	Numerical experiments	58
4.7	Multi-index double loop Monte Carlo estimator	62
4.7.1	Error analysis	64
4.7.2	Computational complexity	65
4.7.3	Adaptive algorithm	71
4.7.4	Numerical experiments	73
4.8	Conclusions and outlook	83
5	Chance-constrained optimal control of stochastic differential equations	85
5.1	Problem formulation	86
5.2	An example from wireless communications	87
5.2.1	System model	88
5.2.2	Stochastic optimal control formulation	93
5.3	Numerical approach	98
5.3.1	Numerically solving the HJB PDE	98
5.3.2	Estimating subgradients	101
5.3.3	Dual optimization	102
5.4	Numerical experiments and results	107
5.4.1	Description of model cellular base station	107
5.4.2	Numerical results	112
5.5	Conclusions and outlook	124
A	Proofs	127
A.1	Proof of Lemma 3.2.1	127
A.2	Proof of Theorem 3.2.1	130
A.3	Proof of Theorem 4.5.1	135
A.4	Proof of Theorem 4.6.1	136
A.5	Proof of Theorem 4.7.1	138
B	Algorithms	143
B.1	Estimating the bias at level ℓ for adaptive DLMC	143
B.2	Estimating $V_{1,L}$ and $V_{2,L}$ for adaptive DLMC	144
B.3	Estimating level differences for the MLDLMC estimator	145
B.4	Estimating variances for the adaptive MLDLMC algorithm	146
B.5	Estimating first-order mixed differences for the MIDLMC estimator	147
B.6	Estimating variances for the adaptive MIDLMC algorithm	147
B.7	Upwind FDM for HJB PDE (5.36)	148
B.8	MC with EM scheme for subgradient (5.52) estimation	149
B.9	Initialization algorithm for the SSM	150
C	Numerical solvers for control PDEs	151
C.1	FDM solver for SOC-based IS control PDE	151
C.2	Upwind FDM solver for the HJB PDE	154
C.2.1	Convergence analysis	156
C.2.2	Boundary conditions	158

D	Details of numerical problem described in Chapter 5	160
D.1	Base station model parameters	160
D.2	Outage proportion for simple cellular user distributions	160
D.3	Analytical solution of the Hamiltonian (5.37) for simple cellular user distributions	161
D.4	Solution to Problem 5.4.1	163
D.5	Solution to Problem 5.4.2	165
D.6	Simulation parameters in Section 5.4	167
D.7	Sensitivity analysis settings in Section 5.4	168
D.8	Detailed results from multi-scenario simulations in Section 5.4	168
	Bibliography	189

Chapter 1

Introduction

Many real-world systems in engineering and natural sciences evolve continuously in time and are influenced by random factors. A **SDE** provides a powerful framework to describe the evolution of a process in continuous time, accommodating both deterministic trends and random noise. In many applications, one needs to make optimal decisions in such dynamic systems subject to uncertainty. For this purpose, control variables are introduced in an **SDE**, allowing us to influence its behaviour in order to achieve a desired goal. In most applications, the goal is to minimise or maximise an expected cost functional over time. The cost functional is typically the summation of a running cost and a terminal cost. An optimal control is the control variable that minimises/maximises the expected cost functional. **SOC** is an active area of research that combines concepts from optimization, probability theory, and differential equations to tackle this class of problems [8].

The first part of the thesis develops efficient **MC** estimators of rare event probabilities associated with standard **SDEs**, using **SOC**. Accurately estimating rare event probabilities is essential for risk management and decision-making in high-impact fields. Some examples include estimation of financial risks [9], prediction of extreme weather events [10], and establishment of safety domains in aerospace systems [11]. Estimating such probabilities via naive **MC** often results in a prohibitive number of required sample paths due to an exploding coefficient of variation [12] (see Figure 1.1). To address this, we employ **IS** by formulating the optimal change of measure in path space as a **SOC** problem that minimises estimator variance. This leads to a non-standard **SOC** problem with a cost functional, that is the product of a running cost and a terminal cost. The work [13] constructed a theoretically perfect **IS** scheme for **SDEs** using the Funke–Shevlyakov–Hausmann integral representation theorem [14], involving intractable terms. The authors in [15] applied the Donsker–Varadhan variational principle [16] to derive a zero-variance **IS** framework, relying on gradient descent and cross-entropy minimization to obtain tractable, but approximate solutions. In contrast, this thesis derives the optimal **IS** time- and state-dependent control via a direct variance minimization approach, offering a rigorous proof of the corresponding dynamic programming lemma and **HJB PDE** using standard tools in stochastic analysis. Furthermore, we transform the resulting **HJB PDE** into a linear Kolmogorov backward equation (**KBE**), allowing an exact characterisation of the zero-variance **IS** control that is numerically tractable via **PDE** discretization techniques.

Building on this foundation, we extend our **IS** framework to **MV-SDEs** [17], which model interacting particle systems in the mean-field limit and arise in pedestrian dynamics [18] (see Figure 1.2), collective animal behaviour [19], and finance [20]. **MV-SDEs** are a special class of **SDEs** whose drift and diffusion functions depend on the law of the solution. This dependence on the law complicates classical **HJB** formulation for **IS** in **MV-SDEs**. While

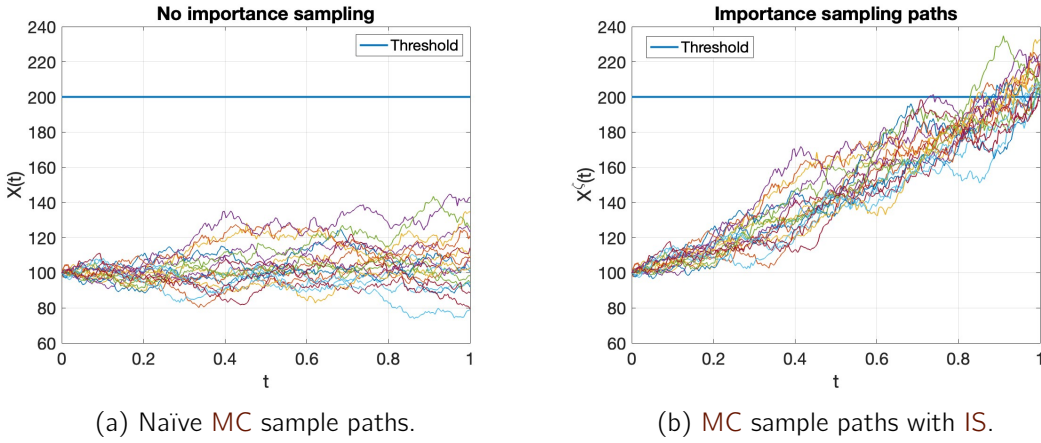


Figure 1.1: Sample paths of a geometric Brownian motion starting at $X(0) = 100$ simulated using standard MC and IS. Suppose one wishes to estimate $\mathbb{P}[X(T) > K]$, a rare event probability. The horizontal line at level $K = 200$ denotes the rare event threshold. Standard MC samples yield few trajectories crossing K at final time $T = 1$, resulting in MC estimators with high coefficient of variation. In contrast, IS significantly increases the probability of sample paths hitting the rare event region, improving efficiency of MC estimation of rare event probabilities.

recent approaches in [21, 22] employ large deviation theory and the Pontryagin principle to construct asymptotically optimal IS controls, this thesis adapts our SOC-based approach to derive a time- and state-dependent zero-variance IS control, using the decoupling framework introduced in [21]. We then propose a novel DLMC estimator with improved computational complexity of $\mathcal{O}(\text{TOL}_r^{-4})$, outperforming the approach in [21] by one complexity order to estimate rare event probabilities with a prescribed relative tolerance of TOL_r . Furthermore, we integrate our SOC-based IS with multilevel Monte Carlo (MLMC) [23] and multi-index MC [24], improving estimator complexity. While previous multilevel IS approaches for standard SDEs [25, 26] use suboptimal level-wise constant controls, we demonstrate that combining our SOC-based IS with multilevel and multi-index MC and analytical smoothing significantly decreases computational work. For MV-SDEs, this results in two new estimators—MLDLMC and MIDLMC—with verifiable error bounds and computational complexity analysis. Numerical experiments on the 1D Kuramoto model [27], a simple MV-SDE example, confirm the reduced complexity from $\mathcal{O}(\text{TOL}_r^{-4})$ to $\mathcal{O}(\text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2)$ in estimating rare events, validating the practical gains of the proposed methodology.

In the second part of the thesis, we shift our focus from using SOC to efficiently solve a forward problem—namely, the estimation of probabilities—to using SOC to solve a control problem with a probability serving as a constraint. This non-standard class of SOC problems, known as chance-constrained SOC, seeks to minimise an expected cost functional subject to a constraint that must hold with high probability [28, 29]. Chance-constrained SOC provides a powerful framework for decision-making under uncertainty while accounting for practical risk tolerances in various domains and striking a balance between robustness and efficiency. Some of its high-impact applications include chance-constrained portfolio selection in insurance [30], accommodation of stochastic renewable energy in electricity markets [31], and development of robust telecommunication networks [32]. Chance (or probabilistic) constraints offer a practical alternative to almost-sure constraints, which




Figure 1.2: Snapshots of a swarm of interacting particles in two-dimensional space at two different times, simulated using the Helbing particle-based pedestrian flow model [4], a **MV-SDE** in the mean-field limit. Each pedestrian’s movement is influenced by both its own dynamics and the collective distribution of the crowd, modelled by mean-field interactions. The evolution of the spatial distribution over time illustrates the collective behaviour emerging from pairwise interactions in the model.

may be either infeasible or overly costly to satisfy [33]. However, these problems do not admit a standard dynamic programming principle with a time- and state-dependent optimal control [34], as the chance constraint imposes a requirement on the state’s joint distribution. To address this, we employ a Lagrangian relaxation of the chance constraint, transforming the problem into a standard **SOC** framework. This enables the derivation of the **HJB PDE** corresponding to the relaxed sub-problem, resulting in a time- and state-dependent optimal control.

Our study is motivated by an application in wireless communication systems, where the objective is to optimise power procurement to minimise both operating expenditure and carbon footprint of cellular wireless networks. These networks operate under uncertainties in wireless channels and renewable energy supply, and must satisfy a **QoS** constraint with high probability at all times. Existing approaches such as [35, 36, 37, 38] are overly conservative, as they optimise for worst-case scenarios. Similarly, chance-constrained methods [39, 5, 40] that employ conservative approximations via Bernstein inequalities [33] lead to overly conservative policies. These works also rely on discrete-time formulations and heuristic solvers, yielding suboptimal solutions. In contrast, this thesis develops a continuous-time chance-constrained **SOC** formulation that decouples the model from its numerical solution, producing a continuous control curve in time. While previous works [41, 42] have addressed continuous-time **SOC** with final-time expectation/probability constraints, our method addresses chance constraints that need to be satisfied at all times. This is more aligned with practical **QoS** requirements in cellular networks. Our numerical solution framework combines a continuous-time Lagrangian relaxation of the chance constraint, an upwind **FDM** for the **HJB** equation corresponding to the relaxed sub-problem, and an effective combination of the **LMBM** for nonsmooth optimization and the **SSM** to navigate the stochasticity of the dual problem. This extends the deterministic optimal control framework in [43] to the stochastic setting. Numerical experiments on a model cellular base station based on the German power system and daily cellular traffic data demonstrate the efficiency and practical relevance of our approach.

Objectives

The central aim of this thesis is to develop numerical methods for solving non-standard **SOC** problems. To this end, the thesis investigates two distinct but thematically connected applications: rare event probability estimation and optimization of green cellular wireless networks. In the first part, the aim is to develop computationally efficient **MC** methods using **SOC** for estimating rare event probabilities of **SDEs**, with a focus on **MV-SDEs**. A key focus is the design of an optimal **IS** strategy through the formulation of a non-standard **SOC** problem that minimises the variance of the **MC** estimator. The thesis rigorously derives the associated **HJB PDE**, thereby establishing a framework for zero-variance **IS** in the path space. This **SOC**-based **IS** approach is applied to **MV-SDEs** using a decoupling technique, introduced in [21], enabling the construction of a time- and state-dependent zero-variance **IS** control. To further improve computational efficiency, the proposed **IS** scheme is combined with hierarchical **MC** techniques, resulting in the novel **DLMC**, **MLDLMC** and **MIDLMC** estimators. These estimators are thoroughly analysed, and their practical effectiveness is demonstrated through numerical experiments on the Kuramoto model [27], highlighting significant reduction in the computational work for rare event probability estimation associated with **MV-SDEs**. In the second part, the thesis develops a continuous-time chance-constrained **SOC** framework for decision-making under uncertainties, with a specific application to power procurement in green cellular networks. This involves formulating a mathematical model that incorporates probabilistic **QoS** constraints over continuous-time, overcoming limitations of existing discrete-time and conservative approaches. The thesis derives tractable solution methods via Lagrangian relaxation, a robust upwind **FDM** solver for the **HJB PDE**, and a combination of the **LMBM** and **SSM** to solve the stochastic non-smooth dual problem. The broader goal is to provide a flexible, computationally efficient framework applicable on various time-discretization schemes, validated through realistic numerical experiments based on energy and cellular traffic data.

Structure of the Thesis

The remainder of this thesis is structured as follows. Chapter 2 introduces some common notation (Section 2.1) and recalls essential concepts in probability, random variables, and **SDEs** (Section 2.2 and 2.3). Section 2.4 provides a brief overview of **SOC** problems and a key technique for solving them.

Chapter 3 motivates the use of **IS** for estimating rare event probabilities in **SDEs** and formulates the problem of optimal **IS** as an **SOC** problem. Starting from a discrete-time framework, Section 3.2 derives the continuous-time formulation in Girsanov's theorem, leading to a time- and state-dependent Markov control from the associated **HJB** equation. We then show, in Section 3.3, how this nonlinear **HJB** can be transformed into a linear **KBE**, enabling numerically tractable variance reduction in **MC** estimators.

In Chapter 4, the **SOC**-based **IS** framework is extended to **MV-SDEs**, a class of **SDEs** whose dynamics depend on the law of the solution. After introducing **MV-SDEs** (Section 4.3) and a naïve **MC** estimator (Section 4.3), we adopt a decoupling approach, introduced in [21], to construct a zero-variance **IS** control in Section 4.4. This leads to the development of a novel **DLMC** estimator in Section 4.5. We then extend this to the multilevel (**MLDLMC**) and multi-index (**MIDLMC**) settings in Sections 4.6 and 4.7 respectively, introducing a correlated antithetic sampler. These methods are validated through numerical experiments on the **1D** Kuramoto model, demonstrating significant computational gains. To our knowledge, this is the first work combining multilevel and

multi-index MC with IS for MV-SDEs. Sections 4.5, 4.6, and 4.7 all include associated relative error analysis, complexity analysis, and an implementable adaptive algorithm.

Chapter 5 turns to SOC problems under chance constraints. We begin with a general formulation in Section 5.1 and then focus on an optimal power procurement problem in green cellular wireless networks in Section 5.2. Section 5.2 includes a continuous-time model of a cellular base station powered by hybrid energy sources, with stochastic dynamics and a probabilistic QoS constraint. Section 5.2 culminates with formulating the relaxed problem and the associated dual optimization problem, enabling us to develop numerical methods to solve it in Section 5.3. Section 5.3 details the numerical scheme: an upwind FDM solver for the HJB PDE, a Euler–Maruyama (EM)-based MC estimator for the subgradient, and a dual optimization algorithm using LMBM-boosted initialization followed by the SSM. Finally, in Section 5.4, we validate the approach using realistic German power grid and mobile traffic data, highlighting the computational efficiency and practical viability of the proposed framework.

Chapter 2

Preliminaries

2.1 Common notation

We denote the set of natural numbers by \mathbb{N} , the set of real numbers by \mathbb{R} , and the set of positive real numbers by \mathbb{R}^+ . We denote the empty set by \emptyset . We indicate a d -dimensional real vector using bold font $\mathbf{x} \in \mathbb{R}^d$, where $d \in \mathbb{N}$. The i^{th} component of \mathbf{x} is denoted by x_i . Unless specifically defined, $\|\mathbf{x}\|$ denotes the Euclidean norm of a vector $\mathbf{x} \in \mathbb{R}^d$ or the Frobenius norm of a matrix $\mathbf{x} \in \mathbb{R}^d \times \mathbb{R}^d$, and $|x|$ denotes the absolute value of a scalar real number $x \in \mathbb{R}$. \mathbf{A}^{tr} denotes the transpose of a matrix \mathbf{A} . ∇x denotes the gradient of a scalar quantity x . $\nabla^2 x$ denotes the Hessian of a scalar quantity x . $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes the Euclidean scalar product of two vectors \mathbf{x}, \mathbf{y} of equal dimension. $\mathbf{A} : \mathbf{B}$ denotes the Frobenius inner product between two matrices \mathbf{A}, \mathbf{B} of equal dimension. The ceiling function $\lceil x \rceil$ denotes the least integer greater than or equal to $x \in \mathbb{R}$. The notation $a \lesssim b$ means that there exists a constant $c \in \mathbb{R}^+$ that is independent of b such that $a \leq cb$. The notation $f = \mathcal{O}(g)$ is called the big-O or Landau notation. It implies that the function f is asymptotically bounded compared to g in the limit $x \rightarrow b$, for some $b \in \mathbb{R} \cup \{-\infty, +\infty\}$. That is,

$$\limsup_{x \rightarrow b} \left| \frac{f(x)}{g(x)} \right| < \infty.$$

Let the set $\mathcal{B} \subset \mathbb{R}^d$. Then, the indicator function is defined by

$$\mathcal{I}_{\{\mathcal{B}\}}(\mathbf{x}) = \begin{cases} 1, & \mathbf{x} \in \mathcal{B}, \\ 0, & \mathbf{x} \notin \mathcal{B}. \end{cases}$$

2.2 Random variables

In this section, we provide a brief introduction to probability theory and random variables based on standard references [44, 45]. This theory contains three main ingredients.

1. The state space: Set of all possible outcomes of an experiment. This is usually denoted by Ω .
2. The events: A subset of Ω , indicating a property which is observed after the experiment is done. We denote by \mathcal{A} the set of all subsets of Ω .
3. The probability: One associates a number denoted by $\mathbb{P}[A]$, called the "probability of A ", for some event $A \in \mathcal{A}$.

Let us now formally introduce some basic concepts in probability theory. We start with the definition of a σ -algebra.

Definition 2.2.1 (σ -algebra). *Let Ω be a non-empty set. A collection \mathcal{F} of subsets of Ω is a σ -algebra on Ω , if*

1. $\emptyset \in \mathcal{F}$ and $\Omega \in \mathcal{F}$;
2. If $A \in \mathcal{F}$ then $A^c \in \mathcal{F}$, where $A^c := \Omega \setminus A$ denotes the complement of A ;
3. If, for some $n \in \mathbb{N}$, $A_1, \dots, A_n \in \mathcal{F}$, then $\cup_{i=1}^n A_i \in \mathcal{F}$ and $\cap_{i=1}^n A_i \in \mathcal{F}$.

The pair (Ω, \mathcal{F}) is called a measurable space.

Consider the following example to aid understanding.

Example 2.2.1. *Let us consider the simple example of a hypothetical, unbiased, 3-sided dice with the numbers 1, 2, 3 written on each face. The experiment is to roll the dice once and observe the number on the visible face. For convenience, we denote the event {Observing the number 1 or 2} as $\{1, 2\}$. Then, we have the following.*

- Set of all possible outcomes $\Omega = \{1, 2, 3\}$.
- Set of all possible subsets of Ω , $\mathcal{A} = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$.
- An example of an event $A \in \mathcal{A}$ is $A = \{1\}$.
- We can assign a probability to the event A , $\mathbb{P}[A] = \frac{1}{3}$.

In this example, one observes that $\mathcal{F}_1 = \{\emptyset, \{1, 2\}, \{3\}, \{1, 2, 3\}\}$ is a σ -algebra on Ω and $\mathcal{F}_2 = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2, 3\}\}$ is not.

Definition 2.2.2 (Generated σ -algebra). *Given a family of sets, $\{A_i\}_{i=1}^n$, there exists a unique σ -algebra, denoted by $\sigma(\{A_i\})$, such that*

1. $\{A_i\}_{i=1}^n \subset \sigma(\{A_i\})$, and
2. if \mathcal{F} is a σ -algebra, $\{A_i\} \subset \mathcal{F} \implies \sigma(\{A_i\}) \subset \mathcal{F}$.

This unique σ -algebra is called the σ -algebra generated by $\{A_i\}_{i=1}^n$.

In Example 2.2.1, the σ -algebra \mathcal{F}_1 is a generated σ -algebra of the event $\{1, 2\}$. That is, $\mathcal{F}_1 = \sigma(\{1, 2\})$. In simple terms, it is the smallest σ -algebra on Ω that contains the event $\{1, 2\}$. Having defined a σ -algebra, we are now equipped to formally define a probability space.

Definition 2.2.3 (Probability space). *A probability space is a triple $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a non-empty set, \mathcal{F} is a σ -algebra on events of Ω , and $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ is a function that satisfies the following.*

1. $\mathbb{P}[\emptyset] = 0$, and $\mathbb{P}[\Omega] = 1$,
2. Let $(A_i)_{i=1}^{\infty}$ be a collection of pairwise disjoint events of \mathcal{F} , that is $A_i \cap A_j = \emptyset$ for $i \neq j$, then $\mathbb{P}[\cup_{i=1}^{\infty} A_i] = \sum_{i=1}^{\infty} \mathbb{P}[A_i]$.

We call \mathbb{P} a probability measure on (Ω, \mathcal{F})

With this definition, a random variable is simply a mapping from the measurable space (Ω, \mathcal{F}) to another measurable space.

Definition 2.2.4 (Random variable). A function $\mathbf{X} : \Omega \rightarrow \mathbb{R}^d$ is defined as a random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ if for all open subsets A of \mathbb{R}^d ,

$$\mathbf{X}^{-1}(A) := \{\omega \in \Omega : \mathbf{X}(\omega) \in A\} \in \mathcal{F}.$$

Equivalently, we say that \mathbf{X} is an \mathcal{F} -measurable function.

Subsequently, we introduce the notation for any open subset $A \in \mathbb{R}^d$

$$\mathbb{P}[\mathbf{X} \in A] := \mathbb{P}[\{\omega \in \Omega : \mathbf{X}(\omega) \in A\}].$$

Definition 2.2.5 (Filtration). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A filtration, denoted by $(\mathcal{F}_t)_{t \geq 0}$, is a collection of sub- σ algebras of \mathcal{F} such that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for all $s \leq t$.

When comparing two random variables, an important property is their independence.

Definition 2.2.6 (Independent random variables). Two \mathbb{R}^d -valued random variables \mathbf{X}, \mathbf{Y} , defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, are independent if and only if for all open sets $A, B \in \mathbb{R}^d$ we have

$$\mathbb{P}[\{\omega \in \Omega : \mathbf{X}(\omega) \in A \text{ and } \mathbf{Y}(\omega) \in B\}] = \mathbb{P}[\mathbf{X} \in A] \mathbb{P}[\mathbf{Y} \in B].$$

An important function related to an \mathbb{R}^d -valued random variable is its cumulative distribution function (CDF).

Definition 2.2.7 (Cumulative distribution function). Given an \mathbb{R}^d -valued random variable \mathbf{X} , defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, the cumulative distribution function is a function $F : \mathbb{R}^d \rightarrow [0, 1]$ such that

$$F(\mathbf{x}) := \mathbb{P}[\mathbf{X} \leq \mathbf{x}], \quad \forall \mathbf{x} \in \mathbb{R}^d.$$

For $\mathbf{x} \in \mathbb{R}^d$, we define $\mathbf{X} \leq \mathbf{x} \iff \{\mathbf{X} \in \mathbb{R}^d : X_i \leq x_i \text{ for all } i = 1, \dots, d\}$.

Definition 2.2.8 (Continuous random variable). A random variable \mathbf{X} is called continuous, if its CDF is absolutely continuous with respect to the Lebesgue measure.

Given the CDF, F , of a continuous \mathbb{R}^d -valued random variable \mathbf{X} , we can define the probability density function (PDF), f .

Definition 2.2.9 (Probability density function). The probability density function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ of a continuous \mathbb{R}^d -valued random variable \mathbf{X} is implicitly defined as

$$F(\mathbf{x}) = \int_{\{\mathbf{y} \in \mathbb{R}^d : \mathbf{y} \leq \mathbf{x}\}} f(\mathbf{y}) d\mathbf{y}, \quad \forall \mathbf{x} \in \mathbb{R}^d.$$

Either the PDF or CDF completely and uniquely defines a continuous random variable. Although these concepts can be applied to non-continuous random variables (e.g. discrete random variables), we only focus on continuous random variables in this thesis. Using the PDF of the random variable, \mathbf{X} , we define the expected value of \mathbf{X} .

Definition 2.2.10 (Expected value of a random variable). The expected (or mean) value of a continuous \mathbb{R}^d -valued random variable \mathbf{X} is defined as

$$\mathbb{E}[\mathbf{X}] := \int_{\mathbb{R}^d} \mathbf{x} f(\mathbf{x}) d\mathbf{x}.$$

We can generalize the expectation definition to any measurable function $g : \mathbb{R}^d \rightarrow \mathbb{R}^{\bar{d}}$ to

$$\mathbb{E}[g(\mathbf{X})] := \int_{\mathbb{R}^d} g(\mathbf{x})f(\mathbf{x})d\mathbf{x}.$$

Definition 2.2.11 (Covariance of a random variable). *The $d \times d$ covariance matrix of a continuous \mathbb{R}^d -valued random variable \mathbf{X} is defined as*

$$\text{Cov}[\mathbf{X}] := \mathbb{E}[\langle \mathbf{X}, \mathbf{X} \rangle] - \langle \mathbb{E}[\mathbf{X}], \mathbb{E}[\mathbf{X}] \rangle.$$

When $d = 1$, the covariance matrix of X is simply the variance of X and is written as

$$\text{Var}[X] := \mathbb{E}[X^2] - (\mathbb{E}[X])^2.$$

Example 2.2.2 (Uniform random variable). *Consider the simplest \mathbb{R} -valued continuous random variable, the uniform random variable, that takes values only on an interval in \mathbb{R} with equal probability. We use the notation $X \sim \mathfrak{U}(a, b)$ to denote that X is a uniform random variable taking values on $[a, b]$. X has the following PDF.*

$$f(x) = \begin{cases} \frac{1}{b-a}, & \text{for } x \in [a, b], \\ 0, & \text{otherwise.} \end{cases}$$

Using Definition 2.2.9, we can compute the CDF, F of X .

$$F(x) = \begin{cases} 0, & \text{for } x < a, \\ \frac{x-a}{b-a}, & \text{for } x \in [a, b], \\ 1, & \text{for } x > b. \end{cases}$$

Using Definition 2.2.10, we conclude that $\mathbb{E}[X] = \frac{a+b}{2}$. From Definition 2.2.11, we can conclude that $\text{Var}[X] = \frac{(b-a)^2}{12}$.

Example 2.2.3 (Normal random variable). *A common \mathbb{R}^d -valued continuous random variable is the d -dimensional normal, or Gaussian, random variable. It is completely defined by its mean, $\boldsymbol{\mu} \in \mathbb{R}^d$, and a positive definite covariance matrix, $\boldsymbol{\Sigma} \in \mathbb{R}^d \times \mathbb{R}^d$. We use the notation $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ to denote that \mathbf{X} is a normal random variable with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. \mathbf{X} has the following PDF.*

$$f(\mathbf{x}) = (2\pi)^{-\frac{d}{2}} \det(\boldsymbol{\Sigma})^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^{\text{tr}} \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})\right),$$

where $\det(\boldsymbol{\Sigma})$ denotes the determinant of the matrix $\boldsymbol{\Sigma}$. Moreover, the random variable $X \sim \mathcal{N}(0, 1)$ is called a one-dimensional standard normal variable.

In this thesis, we are often concerned with a sequence of random variables converging in some sense to another random variable.

Definition 2.2.12 (Convergence of random variables). *We say that a sequence of \mathbb{R}^d -valued random variables $(\mathbf{X}_n)_{n \geq 1}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$:*

- converges in distribution (or converges weakly) to \mathbf{X} defined on $(\Omega, \mathcal{F}, \mathbb{P})$, with CDF F , if

$$\lim_{n \rightarrow \infty} \mathbb{P}[\mathbf{X}_n \leq \mathbf{x}] = F(\mathbf{x})$$

for all $\mathbf{x} \in \mathbb{R}^d$ at which F is continuous. Or equivalently, if $\lim_{n \rightarrow \infty} \mathbb{E} [g(\mathbf{X}_n)] = \mathbb{E} [g(\mathbf{X})]$ for all continuous, bounded functions g on \mathbb{R}^d .

- converges in probability to \mathbf{X} defined on $(\Omega, \mathcal{F}, \mathbb{P})$, if for any $\epsilon > 0$, we have

$$\lim_{n \rightarrow \infty} \mathbb{P} [\{\omega \in \Omega : \|\mathbf{X}_n(\omega) - \mathbf{X}(\omega)\| > \epsilon\}] = 0.$$

- converges in L^p (where $1 \leq p < \infty$) to \mathbf{X} defined on $(\Omega, \mathcal{F}, \mathbb{P})$, if

$$\lim_{n \rightarrow \infty} \mathbb{E} [\|\mathbf{X}_n - \mathbf{X}\|_{\ell^p}^p] = 0.$$

Here, we define $\|\mathbf{x}\|_{\ell^p} = \left(\sum_{i=1}^d |x_i|^p\right)^{\frac{1}{p}}$ for all $\mathbf{x} \in \mathbb{R}^d$.

- converges almost surely to \mathbf{X} defined on $(\Omega, \mathcal{F}, \mathbb{P})$, if

$$\mathbb{P} \left[\left\{ \omega \in \Omega : \lim_{n \rightarrow \infty} \mathbf{X}_n(\omega) = \mathbf{X}(\omega) \right\} \right] = 1.$$

We denote almost sure convergence by writing $\lim_{n \rightarrow \infty} \mathbf{X}_n = \mathbf{X}$ a.s..

Almost sure convergence or L^2 convergence implies convergence in probability which in turn implies convergence in distribution. We conclude this section by defining a particular collection of random variables called stochastic processes.

Definition 2.2.13 (Stochastic process). Given an arbitrary $T \in \mathbb{R}^+$, a stochastic process $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is a function such that $\mathbf{X}(t, \cdot)$ is a random variable in $(\Omega, \mathcal{F}, \mathbb{P})$ for all $t \in (0, T)$.

For convenience, we denote $\mathbf{X}(t) := \mathbf{X}(t, \cdot)$. The t variable is generally associated with the notion of time and T denotes the final (or terminal) time. Next, we define some fundamental notions associated with stochastic processes.

Definition 2.2.14 (Natural filtration of a stochastic process). Given an arbitrary $T \in \mathbb{R}^+$, let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be a stochastic process in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then, the generated σ -algebra

$$\mathcal{F}_t = \sigma(\{\mathbf{X}(s), s \leq t\}), \quad \text{for some } t \in [0, T],$$

is called the natural filtration of the process \mathbf{X} and is denoted by $(\mathcal{F}_t)_{t \geq 0}$.

Definition 2.2.15 (Adapted stochastic process). Given an arbitrary $T \in \mathbb{R}^+$, let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be a stochastic process in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then $(\mathbf{X}(t))_{t \geq 0}$ is said to be adapted to the filtration $(\mathcal{F}_t)_{t \geq 0}$, if for all $t \in [0, T]$, the random variable $\mathbf{X}(t)$ is measurable with respect to \mathcal{F}_t . In other words, \mathbf{X} is an $(\mathcal{F}_t)_{t \geq 0}$ -adapted process.

Definition 2.2.15 implies that a stochastic process is always adapted to its natural filtration. In simple terms, if one observes the paths of the stochastic process \mathbf{X} up to a time $t \geq 0$, one is able to decide if the event $A \in \sigma(\{\mathbf{X}(s), s \leq t\})$ has occurred. One example of a stochastic process is a stochastic differential equation.

2.3 Stochastic differential equations

SDEs are fundamental for mathematical modelling in science, engineering, mathematical finance, and machine learning. They are typically of the form

$$\frac{dS(t)}{dt} = r(t, S(t)) + \text{"noise}(t)"$$

for some evolution function r and a time-dependent "noise(t)" term such that $S(t)$ is well-defined. Itô SDEs [46, 47] are a particular example of the previous formulation, with widespread applications in finance, molecular dynamics, machine learning etc. To formally define an Itô SDE, we first introduce a special type of stochastic process, the Wiener process.

Definition 2.3.1 (Wiener process). *The one-dimensional Wiener process $W : [0, \infty) \times \Omega \rightarrow \mathbb{R}$, also known as the Brownian motion, is a stochastic process that has the following properties.*

1. with probability 1, the mapping $t \rightarrow W(t)$ is continuous and $W(0) = 0$;
2. if $0 = t_0 < t_1 < \dots < t_N = T$, then the increments $W(t_N) - W(t_{N-1}), W(t_{N-1}) - W(t_{N-2}), \dots, W(t_1) - W(t_0)$ are independent;
3. for all $t > s$, the increment $W(t) - W(s)$ has the normal distribution $\mathcal{N}(0, t - s)$.

We now introduce the so-called Itô integral.

Definition 2.3.2 (Itô integral). *Let $f : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ be a Lipschitz continuous function which satisfies*

$$|f(t + \Delta t, W + \Delta W) - f(t, W)| \leq C(\Delta t + |\Delta W|) \quad (2.1)$$

for some constant $C \in \mathbb{R}^+$. Then, for some final time T , the Itô integral is defined by

$$I := \int_0^T f(t, W(t)) dW(t) = \lim_{N \rightarrow \infty} \sum_{n=0}^N f(t_n, W(t_n)) (W(t_{n+1}) - W(t_n)) \quad (\text{limit in } L^2(\mathbb{P})).$$

Here $0 = t_0 < t_1 < \dots < t_N = T$ is any discretization of the interval $[0, T]$ with $\lim_{N \rightarrow \infty} \max_{n=1, \dots, N} (t_{n+1} - t_n) = 0$.

With this, we can define a new stochastic process, the one-dimensional Itô SDE.

Definition 2.3.3 (Itô stochastic differential equation). *Let $W : [0, T] \times \Omega \rightarrow \mathbb{R}$ be a one-dimensional Wiener process on $(\Omega, \mathcal{F}, \mathbb{P})$. A one-dimensional Itô SDE is a stochastic process $X : [0, T] \times \Omega \rightarrow \mathbb{R}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ of the form*

$$X(t) = x_0 + \int_0^t a(s, X(s)) ds + \int_0^t b(s, X(s)) dW(s), \quad (2.2)$$

for a given initial value x_0 , a drift function $a : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, and a diffusion function $b : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$. X is also called an Itô process. For convenience, (2.2) is written in the shorter differential form.

$$\begin{aligned} dX(t) &= a(t, X(t)) dt + b(t, X(t)) dW(t), \\ X(0) &= x_0. \end{aligned}$$

We can extend Definition 2.3.3 to the d -dimensional Itô SDE, \mathbf{X} , given the drift function $\mathbf{a} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, diffusion function $\mathbf{b} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d \times \mathbb{R}^d$, and the d -dimensional Wiener process \mathbf{W} , defined as a d -dimensional vector of d independent one-dimensional Wiener processes.

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}(t, \mathbf{X}(t))dt + \mathbf{b}(t, \mathbf{X}(t))d\mathbf{W}(t), & t \in (0, T] \\ \mathbf{X}(0) = \mathbf{x}_0, \quad \mathbf{x}_0 \in \mathbb{R}^d \end{cases} \quad (2.3)$$

Next, we look at some fundamental results related to SDEs.

Proposition 2.3.1 (Itô formula). *Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be an Itô process satisfying SDE (2.3). Let $g(t, \mathbf{x})$ be a C^2 map from $[0, \infty) \times \mathbb{R}^d$ into \mathbb{R} (i.e., g is twice continuously differentiable on $[0, \infty) \times \mathbb{R}^d$). Define $Y(t) := g(t, \mathbf{X}(t))$. Then $Y : [0, T] \times \Omega \rightarrow \mathbb{R}$ is also an Itô process, satisfying the following Itô SDE.*

$$\begin{cases} dY(t) = \left(\frac{\partial g}{\partial t}(t, \mathbf{X}(t)) + \langle \nabla g(t, \mathbf{X}(t)), \mathbf{a}(t, \mathbf{X}(t)) \rangle \right. \\ \quad \left. + \frac{1}{2} (\mathbf{b}(t, \mathbf{X}(t))\mathbf{b}(t, \mathbf{X}(t))^{tr}) : \nabla^2 g(t, \mathbf{X}(t)) \right) dt \\ \quad + \langle \nabla g(t, \mathbf{X}(t)), \mathbf{b}(t, \mathbf{X}(t))d\mathbf{W}(t) \rangle, \quad t \in (0, T] \\ Y(0) = g(0, \mathbf{x}_0), \quad \mathbf{x}_0 \in \mathbb{R}^d \end{cases} \quad (2.4)$$

The evolution of the joint PDF of an Itô process is given by the Fokker-Planck equation.

Proposition 2.3.2 (Fokker-Planck equation). *Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be an Itô process satisfying SDE (2.3). Then, the conditional joint PDF $\rho(s, \mathbf{y}; t, \mathbf{x})$ of the process $\mathbf{X}(s)$, given that $\mathbf{X}(t) = \mathbf{x}$, for $s > t \geq 0$, is given by the Fokker-Planck equation, or Kolmogorov forward equation.*

$$\begin{cases} -\frac{\partial \rho}{\partial s}(s, \mathbf{y}; t, \mathbf{x}) - \sum_{i=1}^d \frac{\partial}{\partial y_i} (a_i(s, \mathbf{y})\rho(s, \mathbf{y}; t, \mathbf{x})) \\ \quad + \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d \frac{\partial^2}{\partial y_i \partial y_j} \left(\sum_{k=1}^d b_{ik}(s, \mathbf{y})b_{jk}(s, \mathbf{y})\rho(s, \mathbf{y}; t, \mathbf{x}) \right) = 0, \quad (s, \mathbf{y}) \in (t, T] \times \mathbb{R}^d \\ \rho(t, \mathbf{y}; t, \mathbf{x}) = \delta(\mathbf{y} - \mathbf{x}), \quad \mathbf{y} \in \mathbb{R}^d \end{cases}$$

where δ denotes the Dirac distribution.

In large parts of this work, we are interested in estimating $\mathbb{E}[g(\mathbf{X}(T))]$ for some observable g and Itô process \mathbf{X} with starting value \mathbf{x}_0 . The KBE gives a deterministic representation of this expected value.

Proposition 2.3.3 (Kolmogorov backward equation). *Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be an Itô process satisfying SDE (2.3). Let us define the scalar observable $g \in C_0^2(\mathbb{R}^d)$ (i.e., g is twice continuously differentiable with compact support on \mathbb{R}^d). Let us define the cost-to-go function $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ as follows.*

$$u(t, \mathbf{x}) := \mathbb{E}[g(\mathbf{X}(T)) \mid \mathbf{X}(t) = \mathbf{x}], \quad \forall (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$$

Then u solves the following PDE.

$$\begin{cases} \frac{\partial u}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a}(t, \mathbf{x}), \nabla u(t, \mathbf{x}) \rangle \\ \quad + \frac{1}{2} (\mathbf{b}(t, \mathbf{x}) \mathbf{b}(t, \mathbf{x})^{tr}) : \nabla^2 u(t, \mathbf{x}) = 0, & (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d \\ u(T, \mathbf{x}) = g(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^d \end{cases} \quad (2.5)$$

From (2.5), we infer that $\mathbb{E}[g(\mathbf{X}(T))] = u(0, \mathbf{x}_0)$. Standard numerical discretization schemes could be used to solve the linear PDE (2.5). However, they suffer from the "curse of dimensionality". That is, the computational cost required to numerically solve (2.5) up to a prescribed tolerance scales exponentially with the dimension d of the KBE. Moreover, it quickly becomes computationally infeasible to estimate $\mathbb{E}[g(\mathbf{X}(T))]$ using standard numerical schemes up to relative tolerances, which are relevant in rare event estimation (as we will see later). We overcome these issues by using MC, based on simulating approximate paths of the SDE (2.3). Numerically, we produce approximate sample paths of the process \mathbf{X} using the EM time discretization scheme [46].

Definition 2.3.4 (Euler–Maruyama time discretization). *Consider a discretization $0 = t_0 < t_1 < t_2 < \dots < t_N = T$ of the domain $[0, T]$ with N steps. Let $\bar{\mathbf{X}}$ denote the time-discretized version of the process \mathbf{X} corresponding to the SDE (2.3). Then, the EM time discretization of the SDE (2.3) is given as follows.*

- For $n = 0$,

$$\bar{\mathbf{X}}(t_0) = \mathbf{X}(0) = \mathbf{x}_0. \quad (2.6)$$

- For $n = 1, \dots, N - 1$,

$$\bar{\mathbf{X}}(t_{n+1}) = \bar{\mathbf{X}}(t_n) + \mathbf{a}(t_n, \bar{\mathbf{X}}(t_n)) \Delta t_n + \mathbf{b}(t_n, \bar{\mathbf{X}}(t_n)) (\mathbf{W}(t_{n+1}) - \mathbf{W}(t_n)). \quad (2.7)$$

Here $\Delta t_n := t_{n+1} - t_n$, and the i^{th} component of $(\mathbf{W}(t_{n+1}) - \mathbf{W}(t_n))$ is sampled from $\mathcal{N}(0, \Delta t_n)$.

Given M independent samples of the random variable $\bar{\mathbf{X}}(t_N)$ (using (2.7)), the MC approximation of $\mathbb{E}[g(\mathbf{X}(T))]$ is formulated as the following statistical average.

$$\mathbb{E}[g(\mathbf{X}(T))] \approx \mathcal{A}_{\text{MC}} := \frac{1}{M} \sum_{m=1}^M g(\bar{\mathbf{X}}(t_N, \omega^{(m)})). \quad (2.8)$$

There are two sources of error in the above approximation: one arising from the approximation of $\mathbf{X}(T)$ with $\bar{\mathbf{X}}(t_N)$, and the other arising from approximating the expected value $\mathbb{E}[g(\bar{\mathbf{X}}(t_N))]$ with an empirical average. With this, we can bound the global (net) error of the MC estimator as follows.

$$|\mathbb{E}[g(\mathbf{X}(T))] - \mathcal{A}_{\text{MC}}| \leq \underbrace{|\mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\bar{\mathbf{X}}(t_N))]|}_{\text{Bias}} + \underbrace{|\mathbb{E}[g(\bar{\mathbf{X}}(t_N))]| - \mathcal{A}_{\text{MC}}}_{\text{Statistical Error}}. \quad (2.9)$$

For Lipschitz functions \mathbf{a}, \mathbf{b} , the authors in [48, 46] show that the following convergence result holds.

Proposition 2.3.4 (Weak convergence of EM scheme). *Consider the same setting as in Definition 2.3.4. For an observable $g : \mathbf{R}^d \rightarrow \mathbb{R}$ in $C_0^2(\mathbb{R}^d)$,*

$$|\mathbb{E}[g(\mathbf{X}(T)) - g(\bar{\mathbf{X}}(T))]| = \mathcal{O}\left(\max_{1 \leq n \leq N} \Delta t_n\right). \quad (2.10)$$

Proposition 2.3.4 implies that the bias in (2.9) is $\mathcal{O}(N^{-1})$. If we want the bias in (2.9) to be below a prescribed tolerance TOL, then the computational cost of producing one approximate sample path is $\mathcal{O}(\text{TOL}^{-1})$. To analyse the statistical error, we require another important result.

Proposition 2.3.5 (Central limit theorem (CLT) [49]). *Let $\{X_1, \dots, X_M\}$ be a set of $1D$ independent and identically distributed (iid) random variables with $\mathbb{E}[X_i] = 0$ and $\text{Var}[X_i] = 1$ for $i = 1, \dots, M$. Then, in the limit $M \rightarrow \infty$, the random variable $\sum_{i=1}^M \frac{X_i}{\sqrt{M}}$ converges in distribution to the random variable $\xi \sim \mathcal{N}(0, 1)$.*

Proposition 2.3.5 implies that the distribution of the scaled and centered MC estimator, $\sqrt{M}(\mathcal{A}_{\text{MC}} - \mathbb{E}[g(\bar{\mathbf{X}}(t_N))])$, approaches the normal distribution with mean 0 and variance $\text{Var}[g(\bar{\mathbf{X}}(t_N))]$ in the limit $M \rightarrow \infty$. Hence, we have the following probabilistic bound on the statistical error in (2.9).

$$\mathbb{P}\left[\lim_{M \rightarrow \infty} \sqrt{M} \frac{|\mathcal{A}_{\text{MC}} - \mathbb{E}[g(\bar{\mathbf{X}}(t_N))]|}{\sqrt{\text{Var}[g(\bar{\mathbf{X}}(t_N))]}} \leq \Phi^{-1}\left(\frac{1+\epsilon}{2}\right)\right] = \epsilon,$$

for $0 < \epsilon < 1$, where Φ^{-1} is the inverse of the standard normal CDF. For a large number of samples M , the statistical error in (2.9) can be approximated as follows.

$$|\mathbb{E}[g(\bar{\mathbf{X}}(t_N))] - \mathcal{A}_{\text{MC}}| \leq \Phi^{-1}\left(\frac{1+\epsilon}{2}\right) \sqrt{\frac{\text{Var}[g(\bar{\mathbf{X}}(t_N))]}{M}}. \quad (2.11)$$

For example, setting $\epsilon = 0.95$ means $\Phi^{-1}\left(\frac{1+\epsilon}{2}\right) = 1.96$. We need the following number of samples for the MC error to be below a prescribed tolerance TOL with a probability of $1 - \epsilon$.

$$M = \left(\Phi^{-1}\left(\frac{1+\epsilon}{2}\right)\right)^2 \frac{\text{Var}[g(\bar{\mathbf{X}}(t_N))]}{\text{TOL}^2}. \quad (2.12)$$

Hence, we require $\mathcal{O}(\text{TOL}^{-2})$ sample paths, each of cost $\mathcal{O}(\text{TOL}^{-1})$, so that the global error of the MC estimator (2.9) is below a prescribed tolerance TOL with a probability of $1 - \epsilon$. This leads to a computational complexity of $\mathcal{O}(\text{TOL}^{-3})$. Note that the order of computational complexity of the MC estimator is independent of the dimension d of the underlying random variable \mathbf{X} .

2.4 Optimal control of stochastic differential equations

In this section, we provide an introduction to SOC in continuous time based on standard references [8, 50, 51, 52, 34]. First, we will present the basic structure of a SOC problem. Then, we will introduce a key technique used to find its solution.

2.4.1 Formulation of a stochastic optimal control problem

In this section, $(\Omega, \mathcal{F}, \mathbb{P})$ denotes a probability space equipped with the filtration $\{\mathcal{F}_t\}$ and \mathbf{W} denotes a d -dimensional $\{\mathcal{F}_t\}$ -adapted Wiener process. Standard SOC problems consist of the following basic elements.

1. **Time horizon:** In this thesis, we only consider problems with a finite horizon: $[0, T]$ for some $T \in \mathbb{R}^+$.
2. **(Controlled) state process:** The state process is a stochastic process that describes the state of a physical system of interest. The state process takes values in a set $\mathcal{S} \subset \mathbb{R}^d$ called state space. Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathcal{S}$ denote the state process given by the solution of a controlled SDE of the following form.

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}(t, \mathbf{X}(t), \mathbf{U}(t))dt + \mathbf{b}(t, \mathbf{X}(t), \mathbf{U}(t))d\mathbf{W}(t), & t \in (0, T] \\ \mathbf{X}(0) = \mathbf{x}_0, & \mathbf{x}_0 \in \mathbb{R}^d \end{cases} \quad (2.13)$$

Here, the stochastic process \mathbf{U} represents a control. Given a control process \mathbf{U} , the behaviour of the controlled state process \mathbf{X} at each time $t \in [0, T]$ is influenced by modifying the drift and diffusion functions through the value of $\mathbf{U}(t)$.

3. **Control process:** The control process \mathbf{U} in (2.13) is an $\{\mathcal{F}_t\}$ -adapted process that takes values in a set $\mathcal{U} \subset \mathbb{R}^{\bar{d}}$ called the control space. \bar{d} denotes the dimension of the control process \mathbf{U} . Hence, the functions $\mathbf{a} : [0, T] \times \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}^d$ and $\mathbf{b} : [0, T] \times \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}^d \times \mathbb{R}^{\bar{d}}$. In this work, we are particularly interested in the class of deterministic Markov controls. A \mathcal{U} -valued stochastic process \mathbf{U} is called a Markov control if $\mathbf{U}(t) = \mathbf{u}(t, \mathbf{X}(t))$ for some measurable function $\mathbf{u} : [0, T] \times \mathcal{S} \rightarrow \mathcal{U}$. The function \mathbf{u} is called the Markov control function (or policy).
4. **Cost function:** There is some cost associated with the system, which may depend on the system state itself and the control used. The cost is typically expressed as a functional $J_{t, \mathbf{x}}(\mathbf{u})$, denoting the expected total cost starting from the system state \mathbf{x} and time t if the policy \mathbf{u} is implemented. For example, a common cost functional is

$$J_{t, \mathbf{x}}(\mathbf{u}) = \mathbb{E} \left[\int_t^T g(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) ds + h(\mathbf{X}(T)) \mid \mathbf{X}(t) = \mathbf{x} \right], \quad (2.14)$$

where $g : [0, T] \times \mathcal{S} \times \mathcal{U} \rightarrow \mathbb{R}$ is called the running cost function and $h : \mathcal{S} \rightarrow \mathbb{R}$ the terminal cost function. Typically, these functions are assumed to be continuous with at most polynomial growth.

5. **Admissible controls:** Typically, one considers only controls that satisfy certain conditions. An $\{\mathcal{F}_t\}$ -adapted \mathcal{U} -valued stochastic process $\mathbf{U}(t) = \mathbf{u}(t, \mathbf{X}(t))$ is an admissible Markov control if

- There exists a unique solution \mathbf{X} of (2.13). The conditions required for this are:
 - (a) There exists a constant $K_1 \in \mathbb{R}^+$ such that for all $t \in [0, T]$, $\alpha \in \mathcal{U}$, and $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, we have $\|\mathbf{a}(t, \mathbf{x}, \alpha) - \mathbf{a}(t, \mathbf{y}, \alpha)\| + \|\mathbf{b}(t, \mathbf{x}, \alpha) - \mathbf{b}(t, \mathbf{y}, \alpha)\| \leq K_1 \|\mathbf{x} - \mathbf{y}\|$.
 - (b) There exists a constant $K_2 \in \mathbb{R}^+$ such that for all $t \in [0, T]$, $\alpha \in \mathcal{U}$, and $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, we have $\|\mathbf{a}(t, \mathbf{x}, \alpha)\|^2 + \|\mathbf{b}(t, \mathbf{x}, \alpha)\|^2 \leq K_2^2 (1 + \|\mathbf{x}\|^2)$.

(c) There exists some $\delta > 0$ such that for all $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$ the eigenvalues of the matrix $\mathbf{b}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x}))$ are bounded below by δ .

$$\bullet \mathbb{E} \left[\int_0^T |g(t, \mathbf{X}(t), \mathbf{u}(t, \mathbf{X}(t)))| dt + |h(\mathbf{X}(T))| \right] < \infty.$$

Additionally, there could be technical requirements (integrability, smoothness) or physical requirements (constraints on values of the state and control processes). We denote by \mathcal{A} the set of all admissible Markov control functions defined on $[0, T]$.

6. **Value function:** The value function describes the value of the minimum possible cost of the system. We denote it by the function $v : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}$ and is obtained, given state $\mathbf{X}(t) = \mathbf{x}$, by optimizing the cost over all admissible controls.

$$v(t, \mathbf{x}) = \inf_{\mathbf{u} \in \mathcal{A}} J_{t, \mathbf{x}}(\mathbf{u}). \quad (2.15)$$

The goal of a **SOC** problem is to characterize the value function v and find a control \mathbf{U}^* or a policy $\mathbf{u}^* \in \mathcal{A}$ whose cost attains the minimum value. That is, $v(t, \mathbf{x}) = J_{t, \mathbf{x}}(\mathbf{u}^*)$ given the state process $\mathbf{X}(t) = \mathbf{x}$.

In the rest of this thesis, we assume that there exists a unique optimal control that can be attained. That is,

$$v(t, \mathbf{x}) = \inf_{\mathbf{u} \in \mathcal{A}} J_{t, \mathbf{x}}(\mathbf{u}) = \min_{\mathbf{u} \in \mathcal{A}} J_{t, \mathbf{x}}(\mathbf{u}) = J_{t, \mathbf{x}}(\mathbf{u}^*).$$

The conditions required for this are technical and we refer the readers to [51] for more details.

2.4.2 Dynamic programming principle

Next, we introduce one of the main tools that allows us to solve a **SOC** problem, the dynamic programming principle. The dynamic programming lemma expresses the value function recursively as a conditional expectation over future costs.

Proposition 2.4.1 (Dynamic programming lemma). *Let \mathbf{X} be the solution to the controlled SDE (2.13), where $\mathbf{U}(t) = \mathbf{u}(t, \mathbf{X}(t))$ is a Markov control with Markov control function $\mathbf{u} : [t, T] \times \mathcal{S} \rightarrow \mathcal{U}$. Then, the value function $v : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}$, defined in (2.15), satisfies, for all $0 < \delta < T - t$, the following relation:*

$$v(t, \mathbf{x}) = \min_{\mathbf{u} : [t, t+\delta] \times \mathcal{S} \rightarrow \mathcal{U}} \mathbb{E} \left[\int_t^{t+\delta} g(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) ds + v(t + \delta, \mathbf{X}(t + \delta)) \mid \mathbf{X}(t) = \mathbf{x} \right]. \quad (2.16)$$

Proof. Since this is a standard result in continuous-time **SOC**, we refer the readers to [34] for the proof. \square

The idea behind Proposition 2.4.1 is that it is not necessary to optimise the policy \mathbf{u} over the entire time interval $[0, T]$ at once. We can partition the time interval into smaller intervals, and optimise over each individually. This idea becomes particularly useful when we let the partition size go to zero. To see this, let us consider the following limit when $\delta \rightarrow 0$ of (2.16) as follows:

$$0 = \lim_{\delta \rightarrow 0} \min_{\mathbf{u}: [t, t+\delta] \times \mathbf{S} \rightarrow \mathbf{U}} \frac{1}{\delta} \mathbb{E} \left[\int_t^{t+\delta} g(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) ds + v(t + \delta, \mathbf{X}(t + \delta)) - v(t, \mathbf{X}(t)) \middle| \mathbf{X}(t) = \mathbf{x} \right]. \quad (2.17)$$

Applying Itô formula (Proposition 2.3.1) to the value function v , we get

$$\begin{aligned} v(t + \delta, \mathbf{X}(t + \delta)) &= v(t, \mathbf{X}(t)) \\ &+ \int_t^{t+\delta} \left(\frac{\partial v}{\partial s}(s, \mathbf{X}(s)) + \langle \nabla v(s, \mathbf{X}(s)), \mathbf{a}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) \rangle \right. \\ &+ \left. \frac{1}{2} (\mathbf{b}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) \mathbf{b}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s)))^{\text{tr}} : \nabla^2 v(s, \mathbf{X}(s))) \right) ds \\ &+ \int_t^{t+\delta} \langle \nabla v(s, \mathbf{X}(s)), \mathbf{b}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) d\mathbf{W}(s) \rangle. \end{aligned} \quad (2.18)$$

Substituting (2.18) in (2.17) and using the Martingale property of the Itô integral, we get

$$\begin{aligned} 0 &= \lim_{\delta \rightarrow 0} \min_{\mathbf{u}: [t, t+\delta] \times \mathbf{S} \rightarrow \mathbf{U}} \frac{1}{\delta} \mathbb{E} \left[\int_t^{t+\delta} \left(\frac{\partial v}{\partial s}(s, \mathbf{X}(s)) + \langle \nabla v(s, \mathbf{X}(s)), \mathbf{a}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) \rangle \right. \right. \\ &+ \left. \left. \frac{1}{2} (\mathbf{b}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) \mathbf{b}(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s)))^{\text{tr}} : \nabla^2 v(s, \mathbf{X}(s))) \right. \right. \\ &+ \left. \left. g(s, \mathbf{X}(s), \mathbf{u}(s, \mathbf{X}(s))) \right) ds \middle| \mathbf{X}(t) = \mathbf{x} \right]. \end{aligned} \quad (2.19)$$

Given the regularity that we need on v to apply Itô's formula and the regularity of \mathbf{a}, \mathbf{b} to ensure (2.13) has a unique solution, we can apply the limit and use the conditioning on $\mathbf{X}(t) = \mathbf{x}$ to get,

$$\begin{aligned} 0 &= \frac{\partial v}{\partial t}(t, \mathbf{x}) + \min_{\alpha \in \mathbf{U}} \left\{ \langle \nabla v(t, \mathbf{x}), \mathbf{a}(t, \mathbf{x}, \alpha) \rangle \right. \\ &+ \left. \frac{1}{2} (\mathbf{b}(t, \mathbf{x}, \alpha) \mathbf{b}(t, \mathbf{x}, \alpha)^{\text{tr}} : \nabla^2 v(t, \mathbf{x}) + g(t, \mathbf{x}, \alpha)) \right\}. \end{aligned} \quad (2.20)$$

For fixed $(t, \mathbf{x}) \in [0, T] \times \mathbf{S}$, the Markov control function \mathbf{u} appears in the quantity to be minimised through $\alpha \in \mathbf{U}$. That is, suppose we have a solution v of (2.20). Such a procedure will involve finding the minimizing value $\alpha \in \mathbf{U}$ for each $(t, \mathbf{x}) \in [0, T] \times \mathbf{S}$, providing a map between $(t, \mathbf{x}) \in [0, T] \times \mathbf{S}$ and $\alpha \in \mathbf{U}$. This map should coincide with the optimal policy \mathbf{u}^* . Hence, we reduce the computation of the solution to the SOC problem to a pointwise optimization for each $(t, \mathbf{x}) \in [0, T] \times \mathbf{S}$ separately over the set \mathbf{U} , instead of over the set \mathcal{A} of \mathbf{U} -valued Markov control functions \mathbf{u} . This makes finding

the optimal control computationally feasible at least. Next, we need the equation that determines the value function v .

Proposition 2.4.2 (Hamilton–Jacobi–Bellman PDE). *Let \mathbf{X} be the solution to the controlled SDE (2.13), where $\mathbf{U}(t) = \mathbf{u}(t, \mathbf{X}(t))$ is a Markov control with Markov control function $\mathbf{u} : [t, T] \times \mathcal{S} \rightarrow \mathcal{U}$. Then, the value function $v : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}$, defined in (2.15), satisfies the following non-linear HJB PDE:*

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + \mathcal{H}(t, \mathbf{x}, \nabla v(t, \mathbf{x}), \nabla^2 v(t, \mathbf{x})) = 0, & (t, \mathbf{x}) \in [0, T] \times \mathcal{S}, \\ v(T, \mathbf{x}) = h(\mathbf{x}), & \forall \mathbf{x} \in \mathcal{S}, \end{cases} \quad (2.21)$$

with the so-called Hamiltonian function

$$\begin{aligned} \mathcal{H}(t, \mathbf{x}, \nabla v(t, \mathbf{x}), \nabla^2 v(t, \mathbf{x})) = & \min_{\alpha \in \mathcal{U}} \left\{ \langle \nabla v(t, \mathbf{x}), \mathbf{a}(t, \mathbf{x}, \alpha) \rangle \right. \\ & \left. + \frac{1}{2} (\mathbf{b}(t, \mathbf{x}, \alpha) \mathbf{b}(t, \mathbf{x}, \alpha)^{tr}) : \nabla^2 v(t, \mathbf{x}) + g(t, \mathbf{x}, \alpha) \right\}. \end{aligned} \quad (2.22)$$

Proof. We omit the proof of this standard result in this thesis. We refer readers to [51, 34] for the proof. \square

Remark 2.4.1 (Connections to Pontryagin’s Principle). *In the deterministic optimal control setting, when the diffusion \mathbf{b} is zero, the spatial derivative ∇v of the value function in the HJB PDE (2.21) coincides with the adjoint variable in Pontryagin’s principle, a cornerstone of deterministic optimal control theory [53]. For SOC problems, however, Pontryagin’s principle can only be applied in a small noise setting, where the mean path of the SDE is controlled through a time-dependent policy.*

For v to be a unique solution to the PDE (2.21) in the classical sense, v needs to be in the $C^{1,2}([0, T] \times \mathcal{S})$ class of functions. That is, v needs to be once continuously differentiable with respect to t and twice continuously differentiable with respect to \mathbf{x} . Also, v needs to satisfy a polynomial growth condition. That is, $|v(t, \mathbf{x})| \leq K_r(1 + \|\mathbf{x}\|^r)$ for some constants $r, K_r \in \mathbb{R}^+$. The regularity and growth conditions on v imply some conditions on functions $\mathbf{a}, \mathbf{b}, g, h$ and control space \mathcal{U} . We avoid discussion of these technical conditions in this work and refer readers to [51] for more details. Although the above conditions might seem restrictive, it can be relaxed via the notion of viscosity solutions [54]. The corresponding discussion on the existence and uniqueness of viscosity solutions of (2.21) can be found in [52]. The following example of a SOC problem will help readers grasp the above concept better.

Example 2.4.1 (Optimal battery charging). *Consider an energy systems operator who controls a battery that receives energy from a renewable source (solar/wind) and supplies it to a customer. The operator’s goal is to control the battery charging/discharging process to minimise energy production costs under stochastic renewable energy production over the supply time period $[0, T]$.*

1. **Time horizon:** $[0, T]$
2. **(Controlled) state process:** We have a two-dimensional controlled state process $\mathbf{X} = [R, A]$. Let $R : [0, T] \times \Omega \rightarrow \mathbb{R}^+$ denote the incoming renewable power,

modelled by the following Itô SDE:

$$\begin{cases} dR(t) = p(t)dt + \sigma dW(t), & t \in (0, T] \\ R(0) = r_0 \in \mathbb{R}^+ \end{cases} \quad (2.23)$$

Process p represents a deterministic renewable energy inflow and $\sigma > 0$ is the constant diffusion. Let $A : [0, T] \rightarrow [0, \bar{A}]$ denote the battery's state of charge with the battery capacity \bar{A} , modelled by the following ordinary differential equation (ODE):

$$\begin{cases} \frac{dA(t)}{dt} = R(t) - \alpha(t), & t \in (0, T] \\ A(0) = a_0 \in [0, \bar{A}] \end{cases} \quad (2.24)$$

The operator cannot influence the incoming renewable power, but can decide how much power to supply to the customer. Hence, the process α denotes the discharging rate and is the control. That is, by choosing, at time t , the value of $\alpha(t, \mathbf{X}(t))$, the operator controls the amount of power supplied to the customer. The state space is $\mathcal{S} = \mathbb{R}^+ \times [0, \bar{A}]$

3. **Control process:** The control process is defined here by the Markov control function $\alpha : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}^+$. The control space is $\mathcal{U} = \mathbb{R}^+$.
4. **Cost function:** The operator's objective is to minimise the following function

$$J_{t, \mathbf{x}}(\alpha) = \mathbb{E} \left[\int_t^T (c\alpha(s, \mathbf{X}(s))^2 - K(s)\alpha(s, \mathbf{X}(s))) ds + h(A(T)) \mid \mathbf{X}(t) = \mathbf{x} \right]. \quad (2.25)$$

where $c\alpha(s, \mathbf{X}(s))^2$ is a quadratic penalty on discharging to avoid excessive battery wear, $K(s)\alpha(s, \mathbf{X}(s))$ represents the revenue from serving the customer, and $h(A(T))$ represents a penalty for low final battery charge.

5. **Admissible controls:** The admissible set of controls \mathcal{A} are the set of all Markov controls that are \mathbb{R}^+ -valued, and allow a unique solution to (2.24).
6. **Value function:** The value function $v : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}^+$ is obtained, by minimizing the cost over all admissible controls. That is

$$v(t, \mathbf{x}) = \min_{\alpha \in \mathcal{A}} J_{t, \mathbf{x}}(\alpha).$$

The goal of the SOC problem is to determine the Markov control function $\alpha^* \in \mathcal{A}$ whose cost attains the minimum value. That is, $v(t, \mathbf{x}) = J_{t, \mathbf{x}}(\alpha^*)$ given $\mathbf{X}(t) = \mathbf{x} = [r, a]$.

From Proposition 2.4.2, we can show that v satisfies the following HJB equation:

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + p(t) \frac{\partial v}{\partial r}(t, \mathbf{x}) + \frac{\sigma^2}{2} \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) + \min_{v \in \mathbb{R}^+} \left\{ (r - v) \frac{\partial v}{\partial a}(t, \mathbf{x}) + cv^2 - K(t)v \right\} = 0. \\ v(T, \mathbf{x}) = h(a), \quad \forall \mathbf{x} \in \mathcal{S} \end{cases}$$

The optimal control function α^* is characterised by

$$\alpha^*(t, \mathbf{x}) = \arg \min_{v \in \mathbb{R}^+} \left\{ (r - v) \frac{\partial v}{\partial a}(t, \mathbf{x}) + cv^2 - K(t)v \right\},$$

for each $(t, \mathbf{x}) \in [0, T] \times \mathbb{S}$. The optimal control is given by

$$\alpha^*(t, \mathbf{x}) = \frac{1}{2c} \max \left(0, K(t) + \frac{\partial v}{\partial a}(t, \mathbf{x}) \right). \quad (2.26)$$

The optimal value yields the *HJB PDE*:

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + p(t) \frac{\partial v}{\partial r}(t, \mathbf{x}) + r \frac{\partial v}{\partial a}(t, \mathbf{x}) + \frac{\sigma^2}{2} \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) \\ - \frac{1}{4c} \left(\max \left(0, K(t) + \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right)^2 = 0, & (t, \mathbf{x}) \in [0, T] \times \mathbb{S} \\ v(T, \mathbf{x}) = h(a), & \forall \mathbf{x} \in \mathbb{S} \end{cases} \quad (2.27)$$

Hence, by solving *PDE (2.27)* to get the value function v (and its derivatives), and subsequently the optimal control (2.26), one can simulate optimally controlled sample paths of (2.24) that minimise the cost function (2.25).

In Chapter 3, this notion of *SOC* will be used to construct an optimal *IS* estimator for rare event probabilities associated with *SDEs*.

Chapter 3

Importance sampling via stochastic optimal control

In this chapter, we introduce **IS** to efficiently estimate rare event probabilities associated with the solution to standard **SDEs** of the form (2.3) using **MC**. We do this in a systematic manner by posing the problem of deriving an optimal **IS** measure change as a **SOC** problem, and then using the rich solution theory in **SOC** to obtain an optimal Markov control function that minimises the variance of the **MC** estimator. Large parts of this Chapter have already been published in a similar form in [1] and are reproduced herein.

Contributions

- We develop a novel framework (see Section 3.2) to minimise the variance of **MC** estimators of rare-event probabilities through **SOC** of **SDEs**. Although [13, 15] derived an equivalent formulation, we take the novel, but natural, route of variance minimization to obtain an optimal **IS** measure change for **SDEs**.
- For the above novel problem formulation, we state and prove a dynamic programming lemma (see Lemma 3.2.1) for the associated value function with a multiplicative cost structure. In the limit of infinitesimal time intervals, we prove that the dynamic programming lemma converges to a non-linear **PDE** (see Theorem 3.2.1) for the value function, that characterises the optimal control for **IS**. This is done via Itô's formula, Taylor expansion of the exponential function, and bounding the remainder terms from this expansion using standard tools from stochastic analysis (see Appendix A.2). We later link the obtained **PDE** to the **KBE** and the other equivalent formulations in [13, 15] (see Section 3.3).

3.1 Motivation

Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ denote a stochastic process that solves the d -dimensional Itô **SDE** (2.3). Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar-valued function. For a given relative tolerance $\text{TOL}_r > 0$, the objective is to build a computationally efficient **MC** estimator, \mathcal{A}_{MC} , of $\mathbb{P}[g(\mathbf{X}(T)) > K]$, that satisfies the following probabilistic constraint on its relative error:

$$\mathbb{P} \left[\frac{|\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\mathbf{X}(T)) > K]|}{\mathbb{P}[g(\mathbf{X}(T)) > K]} > \text{TOL}_r \right] \leq \epsilon,$$

for some $0 < \epsilon \ll 1$ denoting the confidence level. We assume that $K \in \mathbb{R}$ lies deep

in the right-tail of the distribution of $g(\mathbf{X}(T))$, such that $\mathbb{P}[g(\mathbf{X}(T)) > K]$ is very small. One can rewrite $\mathbb{P}[g(\mathbf{X}(T)) > K]$ as the following expected value: $\mathbb{E}[\mathcal{I}_{\{g(\mathbf{X}(T)) > K\}}]$. One way to approximate $\mathbb{E}[\mathcal{I}_{\{g(\mathbf{X}(T)) > K\}}]$ could be to numerically approximate the solution to the corresponding **KBE** (refer Proposition 2.3.3). It would be computationally infeasible, however, to achieve relative error tolerances with standard numerical schemes (e.g. **FDM**). Moreover, the computational complexity for such schemes scales exponentially with the dimension d of the **KBE**. Hence, **MC** is the preferred method for approximating rare event probabilities. Approximate sample paths $\bar{\mathbf{X}}$ of process \mathbf{X} are produced via **EM** discretization (refer Definition 2.3.4), and the expected value is approximated by the following empirical average over M iid realizations of $\bar{\mathbf{X}}(T)$:

$$\mathbb{P}[g(\mathbf{X}(T)) > K] \approx \mathcal{A}_{\text{MC}} = \frac{1}{M} \sum_{m=1}^M \mathcal{I}_{\{g(\bar{\mathbf{X}}(T, \omega^{(m)})) > K\}}.$$

The relative statistical error of \mathcal{A}_{MC} is defined as $\frac{|\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]|}{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]}$. We wish to control the statistical error in the following sense.

$$\mathbb{P}\left[\frac{|\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]|}{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]} > \text{TOL}_r\right] \leq \epsilon. \quad (3.1)$$

We rewrite (3.1) as follows.

$$\mathbb{P}\left[\frac{|\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]|}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}} > \frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}}\right] \leq \epsilon. \quad (3.2)$$

The Berry–Esseen theorem [45] provides the following bound for the probability in the right-hand side of (3.2).

$$\begin{aligned} & \left| \mathbb{P}\left[\frac{\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}} > \frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}}\right] \right. \\ & \quad \left. - \Phi\left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}}\right) \right| \\ & \leq \frac{1}{\sqrt{M}} C_{\text{BE}} \left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}} , \frac{\mathbb{E}[\mathcal{A}_{\text{MC}}^3]}{\text{Var}[\mathcal{A}_{\text{MC}}]^{\frac{3}{2}}} \right), \end{aligned} \quad (3.3)$$

where Φ is the standard normal **CDF** and C_{BE} is a bounded, positive constant that depends on the required relative error tolerance and the third moment of \mathcal{A}_{MC} . Using (3.3), we can obtain the following bound for the probability in the left-hand side of (3.2).

$$\begin{aligned} & \mathbb{P}\left[\frac{|\mathcal{A}_{\text{MC}} - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]|}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}} > \frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}}\right] \\ & \leq 2 \left(1 - \Phi\left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}}\right) \right) \\ & \quad + \frac{2}{\sqrt{M}} C_{\text{BE}} \left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{\text{MC}}]}} , \frac{\mathbb{E}[\mathcal{A}_{\text{MC}}^3]}{\text{Var}[\mathcal{A}_{\text{MC}}]^{\frac{3}{2}}} \right). \end{aligned} \quad (3.4)$$

By imposing the upper bound in (3.4) to be less than ϵ instead of (3.1), we obtain the following condition that the relative statistical error of \mathcal{A}_{MC} needs to satisfy.

$$\begin{aligned} \Phi \left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{MC}]}} \right) &= 1 - \frac{\epsilon}{2} \\ &+ \frac{1}{\sqrt{M}} C_{BE} \left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{MC}]}} , \frac{\mathbb{E}[\mathcal{A}_{MC}^3]}{\text{Var}[\mathcal{A}_{MC}]^{\frac{3}{2}}} \right). \end{aligned} \quad (3.5)$$

One needs to solve (3.5) to find the optimal M that satisfies the relative statistical error tolerance TOL_r in the sense of (3.1). In the limit $\text{TOL}_r \rightarrow 0$, the CLT (refer Proposition 2.3.5) can be used to instead give us the following simpler condition that the relative statistical error of \mathcal{A}_{MC} needs to satisfy.

$$\Phi \left(\frac{\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \text{TOL}_r}{\sqrt{\text{Var}[\mathcal{A}_{MC}]}} \right) = 1 - \frac{\epsilon}{2}. \quad (3.6)$$

This leads to a direct formulation of the number of samples for the relative statistical error to be below a prescribed relative tolerance TOL_r in the sense of (3.1).

$$M = \left(\Phi^{-1} \left(\frac{2 - \epsilon}{2} \right) \right)^2 \frac{\text{Var}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}]}{\text{TOL}_r^2 (\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K])^2}, \quad (3.7)$$

where Φ^{-1} is the inverse of the standard normal CDF. We can further simplify the term $\text{Var}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}]$ by using Definition 2.2.11 of the variance of a random variable.

$$\begin{aligned} \text{Var}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}] &= \mathbb{E} \left[\left(\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}} \right)^2 \right] - \mathbb{E}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}]^2 \\ &= \mathbb{E}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}] - \mathbb{E}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}]^2 \\ &= \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] (1 - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]). \end{aligned} \quad (3.8)$$

Inserting (3.8) in (3.7), we get the following optimal number of samples to achieve relative tolerance TOL_r :

$$M = \left(\Phi^{-1} \left(\frac{2 - \epsilon}{2} \right) \right)^2 \frac{(1 - \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K])}{\text{TOL}_r^2 \mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]}. \quad (3.9)$$

In this thesis, we use the CLT to obtain probabilistic bounds on the relative statistical error of all our MC estimators. We refer readers to [55] for a detailed comparison of Berry–Essen and CLT bounds on the statistical error of MC estimators. Equation (3.9) implies that the number of required samples scales inversely with the rare event probability $\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]$ to be estimated. As the probability reduces, the number of samples required by a crude MC estimator to achieve a relative statistical error tolerance TOL_r quickly becomes computationally infeasible. For example, if $\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] \approx 10^{-5}$, then one would require around 4×10^9 sample paths to even achieve 1% relative tolerance with 95% confidence. Hence, crude MC provides poor estimates of rare event probabilities. From (3.7), it is evident that one needs to have low values of $\frac{\text{Var}[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}]}{(\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K])^2}$ (often referred to as the coefficient of variation), for a fixed $\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]$, to have a computationally

efficient MC estimator. IS is a popular technique widely used to reduce the variance of MC estimators of rare event probabilities [13, 12]. When used appropriately, IS can reduce the optimal number of sample paths (3.7) to achieve TOL_r by reducing the estimator's coefficient of variation. This technique involves transforming the underlying probability measure of the random variable and drawing samples according to a new measure. We first introduce the IS technique for estimating expectations associated with a continuous random variable.

Proposition 3.1.1 (Importance sampling for a continuous random variable). *Let \mathbf{Y} be an \mathbb{R}^d -valued continuous random variable with PDF $\rho_{\mathbf{Y}}$. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar-valued function. Let \mathbf{Z} be an auxiliary \mathbb{R}^d -valued continuous random variable with PDF $\rho_{\mathbf{Z}}$, such that the support of $g\rho_{\mathbf{Y}}$ is a subset of the support of $\rho_{\mathbf{Z}}$. That is,*

$$\rho_{\mathbf{Z}}(\mathbf{x}) = 0 \implies g(\mathbf{x})\rho_{\mathbf{Y}}(\mathbf{x}) = 0, \quad \forall \mathbf{x} \in \mathbb{R}^d.$$

Suppose we want to estimate $\mathbb{E}[g(\mathbf{Y})]$. Then, we can rewrite this expectation as follows:

$$\begin{aligned} \mathbb{E}[g(\mathbf{Y})] &= \int_{\mathbb{R}^d} g(\mathbf{x})\rho_{\mathbf{Y}}(\mathbf{x})d\mathbf{x} = \int_{\mathbb{R}^d \setminus \mathcal{D}_{\mathbf{Y}}} g(\mathbf{x})\rho_{\mathbf{Y}}(\mathbf{x})d\mathbf{x} \\ &= \int_{\mathbb{R}^d \setminus \mathcal{D}_{\mathbf{Z}}} g(\mathbf{x}) \underbrace{\frac{\rho_{\mathbf{Y}}(\mathbf{x})}{\rho_{\mathbf{Z}}(\mathbf{x})}}_{:=L(\mathbf{x})} \rho_{\mathbf{Z}}(\mathbf{x})d\mathbf{x} = \mathbb{E}[g(\mathbf{Z})L(\mathbf{Z})], \end{aligned}$$

where $\mathcal{D}_{\mathbf{Y}} := \{\mathbf{x} \in \mathbb{R}^d : g(\mathbf{x})\rho_{\mathbf{Y}}(\mathbf{x}) = 0\}$ and $\mathcal{D}_{\mathbf{Z}} := \{\mathbf{x} \in \mathbb{R}^d : \rho_{\mathbf{Z}}(\mathbf{x}) = 0\}$. L denotes the likelihood factor defined by the following mapping:

$$L : \mathbb{R}^d \rightarrow \mathbb{R} : \mathbf{x} \mapsto \frac{\rho_{\mathbf{Y}}(\mathbf{x})}{\rho_{\mathbf{Z}}(\mathbf{x})}.$$

By a clever choice of the auxiliary random variable \mathbf{Z} (or equivalently PDF $\rho_{\mathbf{Z}}$), one can make $\text{Var}[g(\mathbf{Z})L(\mathbf{Z})] \ll \text{Var}[g(\mathbf{Y})]$, even though $\mathbb{E}[g(\mathbf{Y})] = \mathbb{E}[g(\mathbf{Z})L(\mathbf{Z})]$. Such a choice will be highly beneficial for any MC estimator of $\mathbb{E}[g(\mathbf{Y})]$, ensuring lower statistical error for a given number of MC samples. The notion of IS for continuous random variables can be extended to Itô SDEs and is described in Section 3.2.

3.2 Optimal importance sampling for stochastic differential equations

Let us now go back to the problem of estimating $\mathbb{P}[g(\mathbf{X}(T)) > K]$ for an Itô process \mathbf{X} using MC. In this Section, the objective is to methodically derive an optimal IS measure change for SDE (2.3) that minimises the variance (or coefficient of variation) of the MC estimator \mathcal{A}_{MC} of the rare event probability $\mathbb{P}[g(\mathbf{X}(T)) > K]$, based on the SOC theory introduced in Section 2.4. We pose the problem of finding the optimal IS measure as a SOC problem, resulting in a time- and state-dependent optimal Markov control function.

For easy understanding of measure change in SDEs, we start with a discrete-time measure change on the EM-discretization of SDE (2.3). For this, we rewrite the EM-discretization in Definition 2.3.4 as follows:

$$\begin{cases} \bar{\mathbf{X}}(t_{n+1}) = \bar{\mathbf{X}}(t_n) + \mathbf{a}(t_n, \bar{\mathbf{X}}(t_n)) \Delta t_n \\ \quad + \mathbf{b}(t_n, \bar{\mathbf{X}}(t_n)) \sqrt{\Delta t_n} \boldsymbol{\varepsilon}_n, \quad \forall n = 0, \dots, N-1. \\ \bar{\mathbf{X}}(t_0) = \mathbf{X}(0) = \mathbf{x}_0 \end{cases} \quad (3.10)$$

Here $\boldsymbol{\varepsilon}_n$ is a d -dimensional standard normal random variable such that $\sqrt{\Delta t_n} \boldsymbol{\varepsilon}_n$ is a d -dimensional Wiener increment. For each time step $n = 0, \dots, N-1$, we perform the following mean-shift measure change of the random variable $\boldsymbol{\varepsilon}_n$:

$$\hat{\boldsymbol{\varepsilon}}_n = \sqrt{\Delta t_n} \boldsymbol{\zeta}_n + \boldsymbol{\varepsilon}_n, \quad (3.11)$$

where $\boldsymbol{\zeta}_n := \boldsymbol{\zeta}(t_n, \bar{\mathbf{X}}(t_n))$ are discrete-time evaluations of $\boldsymbol{\zeta}$, an \mathbb{R}^d -valued Markov control function. Since we know the PDFs of $\boldsymbol{\varepsilon}_n$ and $\hat{\boldsymbol{\varepsilon}}_n$, the likelihood factor L_n arising from this measure change at time step n is

$$\begin{aligned} L_n(\hat{\boldsymbol{\varepsilon}}_n) &= \exp\left\{-\frac{1}{2}\|\hat{\boldsymbol{\varepsilon}}_n\|^2\right\} \exp\left\{\frac{1}{2}\|\hat{\boldsymbol{\varepsilon}}_n - \sqrt{\Delta t_n} \boldsymbol{\zeta}_n\|^2\right\} \\ &= \exp\left\{\frac{1}{2}\Delta t_n \|\boldsymbol{\zeta}_n\|^2 - \sqrt{\Delta t_n} \langle \hat{\boldsymbol{\varepsilon}}_n, \boldsymbol{\zeta}_n \rangle\right\}. \end{aligned} \quad (3.12)$$

Inserting (3.11) in (3.12) and writing L_n as a function of $\boldsymbol{\varepsilon}_n$, we get

$$L_n(\boldsymbol{\varepsilon}_n) = \exp\left\{-\frac{1}{2}\Delta t_n \|\boldsymbol{\zeta}_n\|^2 - \sqrt{\Delta t_n} \langle \boldsymbol{\varepsilon}_n, \boldsymbol{\zeta}_n \rangle\right\}. \quad (3.13)$$

Aggregating the measure change over all time steps, the overall likelihood factor is $\prod_{n=0}^{N-1} L_n(\boldsymbol{\varepsilon}_n)$. We can rewrite our quantity of interest $\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]$ as follows:

$$\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K] = \mathbb{E}\left[\mathcal{I}_{\{g(\bar{\mathbf{X}}(T)) > K\}}\right] = \mathbb{E}\left[\mathcal{I}_{\{g(\bar{\mathbf{X}}^\zeta(T)) > K\}} \prod_{n=0}^{N-1} L_n(\boldsymbol{\varepsilon}_n)\right]. \quad (3.14)$$

The dynamics of the new discrete-time process $\bar{\mathbf{X}}^\zeta$ is derived by inserting the measure change (3.11) in (3.10), resulting in

$$\begin{cases} \bar{\mathbf{X}}^\zeta(t_{n+1}) = \bar{\mathbf{X}}^\zeta(t_n) + \left(\mathbf{a}(t_n, \bar{\mathbf{X}}^\zeta(t_n)) \right. \\ \quad \left. + \mathbf{b}(t_n, \bar{\mathbf{X}}^\zeta(t_n)) \boldsymbol{\zeta}(t_n, \bar{\mathbf{X}}^\zeta(t_n)) \right) \Delta t_n \\ \quad + \mathbf{b}(t_n, \bar{\mathbf{X}}^\zeta(t_n)) \sqrt{\Delta t_n} \boldsymbol{\varepsilon}_n, \quad \forall n = 0, \dots, N-1 \\ \bar{\mathbf{X}}^\zeta(t_0) = \bar{\mathbf{X}}(t_0) = \mathbf{x}_0 \end{cases} \quad (3.15)$$

The objective of IS here is to minimise the variance of the MC estimator of $\mathbb{P}[g(\bar{\mathbf{X}}(T)) > K]$, or equivalently its second moment. That is, we seek to solve the following minimization problem.

$$\min_{\{\boldsymbol{\zeta}_n\}_{n=0}^{N-1}} \mathbb{E}\left[\mathcal{I}_{\{g(\bar{\mathbf{X}}^\zeta(T)) > K\}} \prod_{n=0}^{N-1} (L_n(\boldsymbol{\varepsilon}_n))^2 \mid \bar{\mathbf{X}}^\zeta(0) = \mathbf{x}_0\right]$$

$$\begin{aligned}
&= \min_{\{\zeta_n\}_{n=0}^{N-1}} \mathbb{E} \left[\mathcal{I}_{\{g(\bar{\mathbf{X}}^\zeta(T)) > K\}} \prod_{n=0}^{N-1} \exp \left\{ -\Delta t_n \|\zeta_n\|^2 \right. \right. \\
&\quad \left. \left. - 2\sqrt{\Delta t_n} \langle \varepsilon_n, \zeta_n \rangle \right\} \mid \bar{\mathbf{X}}^\zeta(0) = \mathbf{x}_0 \right] \\
&= \min_{\{\zeta_n\}_{n=0}^{N-1}} \mathbb{E} \left[\mathcal{I}_{\{g(\bar{\mathbf{X}}^\zeta(T)) > K\}} \exp \left\{ -\sum_{n=0}^{N-1} \Delta t_n \|\zeta_n\|^2 \right. \right. \\
&\quad \left. \left. - 2 \sum_{n=0}^{N-1} \sqrt{\Delta t_n} \langle \varepsilon_n, \zeta_n \rangle \right\} \mid \bar{\mathbf{X}}^\zeta(0) = \mathbf{x}_0 \right]. \tag{3.16}
\end{aligned}$$

In the limit $\max_{n=0, \dots, N-1} \Delta t_n \rightarrow 0$, (3.16) can be posed as the following continuous-time SOC problem.

1. **Time horizon:** $[0, T]$ for some $T \in \mathbb{R}^+$.
2. **(Controlled) state process:** The controlled state $\mathbf{X}^\zeta : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ solves the following Itô SDE.

$$\begin{cases} d\mathbf{X}^\zeta(t) = \left(\mathbf{a}(t, \mathbf{X}^\zeta(t)) + \mathbf{b}(t, \mathbf{X}^\zeta(t)) \zeta(t, \mathbf{X}^\zeta(t)) \right) dt \\ \quad + \mathbf{b}(t, \mathbf{X}^\zeta(t)) d\mathbf{W}(t), \quad 0 < t \leq T, \\ \mathbf{X}^\zeta(0) = \mathbf{X}(0) = \mathbf{x}_0 \end{cases}. \tag{3.17}$$

Observe that (3.15) is the EM-discretization of (3.17). The state space is $\mathbf{S} = \mathbb{R}^d$.

3. **Control process:** The control process is defined by the Markov control function $\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$. The control space is $\mathbf{U} = \mathbb{R}^d$.
4. **Cost function:** The objective is to minimise the second moment of the MC estimator of $\mathbb{P}[g(\mathbf{X}(T)) > K]$ with IS.

$$\begin{aligned}
J_{t,\mathbf{x}}(\zeta) &= \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{X}^\zeta(T)) > K\}} \exp \left\{ -\int_t^T \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right. \right. \\
&\quad \left. \left. - 2 \int_t^T \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right\} \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]. \tag{3.18}
\end{aligned}$$

5. **Admissible controls:** The admissible set of controls \mathcal{A} are the set of all Markov control functions that are \mathbb{R}^d -valued, and allow a unique solution to (3.17).
6. **Value function:** The value function $v : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is obtained by minimizing the cost over all admissible controls. That is

$$v(t, \mathbf{x}) = \min_{\zeta \in \mathcal{A}} J_{t,\mathbf{x}}(\zeta). \tag{3.19}$$

The goal of the SOC problem is to determine the Markov control function $\zeta^* \in \mathcal{A}$ whose cost attains the minimum value. That is $v(t, \mathbf{x}) = J_{t,\mathbf{x}}(\zeta^*)$ given $\mathbf{X}^\zeta(t) = \mathbf{x}$. Such SOC-based IS schemes have been developed in various other contexts, including stochastic

reaction networks [56], and sums of independent random variables [57]. We extend this scheme to standard Itô SDEs in this Section.

Before proceeding with solving the SOC problem, we comment on the proposed measure change for SDEs described above. Although the measure change defined by (3.11) may at first seem ad hoc or heuristic, this measure change in continuous-time corresponds precisely to the Girsanov theorem for measure change in SDEs. The Girsanov theorem provides a framework for changing the probability measure in SDEs, by altering its drift while preserving the Brownian motion under the new measure.

Proposition 3.2.1 (Girsanov theorem for measure change in SDEs [47]). *Let \mathbb{P} be a probability measure defined on the measurable space (Ω, \mathcal{F}) . Let $\mathbf{W}^{\mathbb{P}}$ be a d -dimensional Wiener process in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be a stochastic process driven by the Wiener process $\mathbf{W}^{\mathbb{P}}$ in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$.*

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}(t, \mathbf{X}(t))dt + \mathbf{b}(t, \mathbf{X}(t))d\mathbf{W}^{\mathbb{P}}(t), & t \in (0, T], \\ \mathbf{X}(0) = \mathbf{x}_0, & \mathbf{x}_0 \in \mathbb{R}^d \end{cases} \quad (3.20)$$

Let \mathbb{Q} be another probability measure defined on the same measurable space (Ω, \mathcal{F}) defined by the so-called Radon–Nikodym derivative:

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_0^t = \exp \left\{ - \int_0^t \langle \boldsymbol{\zeta}(s, \mathbf{X}(s)), d\mathbf{W}^{\mathbb{P}}(s) \rangle - \frac{1}{2} \int_0^t \|\boldsymbol{\zeta}(s, \mathbf{X}(s))\|^2 ds \right\}, \quad (3.21)$$

where $\boldsymbol{\zeta}$ is an \mathcal{F} -adapted Markov process. Then, $\mathbf{W}^{\mathbb{Q}}$ is a d -dimensional Wiener process in the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$ and is given by

$$d\mathbf{W}^{\mathbb{Q}}(t) = d\mathbf{W}^{\mathbb{P}}(t) - \boldsymbol{\zeta}(t, \mathbf{X}(t))dt. \quad (3.22)$$

Let $\mathbf{X}^{\boldsymbol{\zeta}}$ denote the process \mathbf{X} in the new probability space $(\Omega, \mathcal{F}, \mathbb{Q})$. Then, from (3.22), $\mathbf{X}^{\boldsymbol{\zeta}} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ solves the following Itô SDE.

$$\begin{cases} d\mathbf{X}^{\boldsymbol{\zeta}}(t) = \left(\mathbf{a}(t, \mathbf{X}^{\boldsymbol{\zeta}}(t)) + \mathbf{b}(t, \mathbf{X}^{\boldsymbol{\zeta}}(t)) \boldsymbol{\zeta}(t, \mathbf{X}^{\boldsymbol{\zeta}}(t)) \right) dt \\ \quad + \mathbf{b}(t, \mathbf{X}^{\boldsymbol{\zeta}}(t)) d\mathbf{W}^{\mathbb{Q}}(t), & 0 < t \leq T, \\ \mathbf{X}^{\boldsymbol{\zeta}}(0) = \mathbf{X}(0) = \mathbf{x}_0 \end{cases} \quad (3.23)$$

Note that the overall likelihood factor $\prod_{n=0}^{N-1} L_n(\boldsymbol{\epsilon}_n)$ defined in (3.13) is a time-discretized version of the Radon–Nikodym derivative (3.21) and the discrete dynamics (3.15) is the EM-discretization of (3.23). We refer readers to [58] for a formal proof of Proposition 3.2.1 as a limit of the discrete time formulation described above.

Let us now get back to solving the SOC problem of optimal IS for SDEs. Note that the cost function (3.18) possesses a multiplicative structure of a running cost times a terminal cost, instead of an additive structure of a running cost plus a terminal cost as seen in (2.14) in Section 2.4. It is unclear if a dynamic programming equation holds for such a non-standard SOC problem formulation. Therefore, we first state and prove a dynamic programming lemma for the value function v (3.19).

Lemma 3.2.1 (Dynamic programming lemma for optimal IS in standard SDEs). *Let process \mathbf{X} solve the SDE (2.3), and controlled process $\mathbf{X}^{\boldsymbol{\zeta}}$ solve the SDE (3.17), where*

$\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the Markov control function satisfying Novikov's condition $\mathbb{E} \left[\exp \left\{ \frac{1}{2} \int_0^T \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right\} \right] < \infty$. Assume the value function $v : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, defined in (3.19), belongs to the $C^{1,2}([0, T] \times \mathbb{R}^d)$ class of functions. Then, v satisfies the following dynamic programming relation for all $0 < \delta < T - t$,

$$v(t, \mathbf{x}) = \min_{\zeta: [t, t+\delta] \times \mathbb{R}^d \rightarrow \mathbb{R}^d} \mathcal{J}_{t, \mathbf{x}}(\zeta), \quad (3.24)$$

where

$$\mathcal{J}_{t, \mathbf{x}}(\zeta) = \mathbb{E} \left[\exp \left\{ - \int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds - 2 \int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right\} v(t + \delta, \mathbf{X}^\zeta(t + \delta)) \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]. \quad (3.25)$$

Proof. See Appendix A.1. □

Similar to the procedure in Section 2.4.2, one can consider Lemma 3.2.1 in the limit $\delta \rightarrow 0$ and derive the PDE that determines the value function v and characterizes the optimal Markov control function ζ .

Theorem 3.2.1 (Optimal control for IS in standard SDEs). *Let process \mathbf{X} solve the SDE (2.3), and controlled process \mathbf{X}^ζ solve the SDE (3.17), where $\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the Markov control function satisfying the condition $\mathbb{E} \left[\exp \left\{ \frac{1}{2} \int_0^T \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right\} \right] < \infty$. In addition, assume the value function $v : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, defined in (3.19), belongs to the $C^{1,2}([0, T] \times \mathbb{R}^d)$ class of functions, and $v(t, \mathbf{x}) \neq 0$ for all $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$. Then, v satisfies the following non-linear PDE:*

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a}(t, \mathbf{x}), \nabla v(t, \mathbf{x}) \rangle + \frac{1}{2} \nabla^2 v(t, \mathbf{x}) : (\mathbf{b}(t, \mathbf{x}) \mathbf{b}(t, \mathbf{x})^{tr}) \\ - \frac{1}{4v(t, \mathbf{x})} \|\mathbf{b}(t, \mathbf{x})^{tr} \nabla v(t, \mathbf{x})\|^2 = 0, & (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d, \\ v(T, \mathbf{x}) = \mathcal{I}_{\{g(\mathbf{x}) > K\}}, & \forall \mathbf{x} \in \mathbb{R}^d, \end{cases} \quad (3.26)$$

with an optimal Markov control function that minimises the cost function in (3.18) defined as follows.

$$\zeta^*(t, \mathbf{x}) = \frac{1}{2} \mathbf{b}(t, \mathbf{x})^{tr} \nabla \log v(t, \mathbf{x}), \quad \forall (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d. \quad (3.27)$$

Proof. See Appendix A.2. □

In Theorem 3.2.1, the regularity assumptions on the value function v might be restrictive in many cases. However, these assumptions can be relaxed by introducing the notion of viscosity solutions [54]. We refer readers to [52] for a formal proof.

Remark 3.2.1 (Equivalent HJB formulation). *Previous approaches [15, 59] formulated the problem of optimal IS for SDEs via SOC through a variational characterization of thermodynamic free energy, minimizing the following function:*

$$\gamma(t, \mathbf{x}) = -\frac{1}{2} \log v(t, \mathbf{x})$$

$$= \inf_{\zeta \in \mathcal{A}} \mathbb{E} \left[\int_t^T \frac{1}{2} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds - \log \left(\mathcal{I}_{\{g(\mathbf{X}^\zeta(T)) > K\}} \right) \middle| \mathbf{X}^\zeta(t) = \mathbf{x} \right],$$

with the function v defined in (3.19) and the Itô process \mathbf{X}^ζ following the SDE (3.17), leading to the same HJB PDE (3.26).

Various approaches have been proposed to numerically solve (3.26) and obtain an approximate control ζ^* . [60] solved the d -dimensional HJB PDE (3.26) using least-squares regression, whereas [61] solved it using model-reduction techniques for higher dimensions. Neural networks have also been employed to solve the HJB PDE in higher dimensions with stochastic gradient [15] and cross-entropy [62] learning methods for the above SOC formulation. In this thesis, however, we use the equivalence of the nonlinear PDE (3.26) and the linear KBE to solve the SOC problem.

3.3 Relating optimal importance sampling control with the Kolmogorov backward equation

Recall that the value function v , defined in (3.19), minimises the second moment of the MC estimator with IS of the probability $\mathbb{P}[g(\mathbf{X}(T)) > K]$. In this section, we demonstrate that the proposed measure change with the optimal control defined in Theorem 3.2.1 results in a MC estimator with zero variance. To see this, we define the new function $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$v(t, \mathbf{x}) = u^2(t, \mathbf{x}), \quad \forall (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d. \quad (3.28)$$

Replacing v and its derivatives with u and its derivatives in (3.26) using (3.28), we get the following PDE that solves for the function u :

$$\begin{cases} \frac{\partial u}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a}(t, \mathbf{x}), \nabla u(t, \mathbf{x}) \rangle \\ + \frac{1}{2} \nabla^2 u(t, \mathbf{x}) : (\mathbf{b}(t, \mathbf{x}) \mathbf{b}(t, \mathbf{x})^{\text{tr}}) = 0, & (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d \\ u(T, \mathbf{x}) = \mathcal{I}_{\{g(\mathbf{x}) > K\}}, & \mathbf{x} \in \mathbb{R}^d \end{cases} \quad (3.29)$$

with optimal control

$$\zeta^*(t, \mathbf{x}) = \mathbf{b}(t, \mathbf{x})^{\text{tr}} \nabla \log u(t, \mathbf{x}). \quad (3.30)$$

Equation (3.29) is exactly the KBE (refer Proposition 2.3.3) that solves for the following conditional expectation.

$$u(t, \mathbf{x}) = \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{X}(T)) > K\}} \middle| \mathbf{X}(t) = \mathbf{x} \right],$$

where the Itô process \mathbf{X} follows the SDE (2.3). Hence, (3.28) implies that the optimal second moment of the MC estimator with IS is equal to the square of the first moment of the MC estimator without (or with) IS. Then, from the Definition 2.2.11 of the variance of a random variable, it follows that the optimal IS measure change described above leads to a zero-variance estimator. The idea is to solve the linear KBE (3.29), instead of the non-linear HJB PDE (3.26), to get the optimal control function ζ^* that minimises the variance of the MC estimator with IS. For the purpose of IS, it is sufficient to roughly solve (3.29) to obtain substantial variance reduction. Hence, we develop a MC estimator combined with IS using the optimal control ζ^* from solving (3.29) to estimate the rare

event probability $\mathbb{P}[g(\mathbf{X}(T)) > K]$. This is computationally much cheaper than directly solving *KBE* (3.29) to estimate the probability up to relative tolerances, especially in higher dimensions ($d \gg 1$).

Remark 3.3.1 (Using *KBE* for *IS*). *Derivation of a MC estimator with optimal IS that involves the solution to the KBE is a classical result [13]. As IS produces an unbiased MC estimator, it is sufficient to roughly approximate u and its derivatives to get an approximate but useful control. The authors in [63, 64] have applied this in the context of IS for Markov processes and SDEs.*

In Chapter 4, we perform optimal *IS*, as explained in this section, for a challenging class of *SDEs* known as the *MV-SDE*.

Chapter 4

Importance sampling for the McKean–Vlasov stochastic differential equation

In this chapter, we use the SOC-based IS measure change introduced in Chapter 3 to construct a computationally efficient estimator of rare event probabilities associated with a challenging class of SDEs, known as the MV-SDE. We further enhance the computational complexity of this estimator using multilevel and multi-index MC techniques. Large parts of this Chapter have already been published in a similar form in [1, 2] and are reproduced herein.

Contributions

- Based on the decoupling approach introduced in [21], we develop a novel MC estimator for approximating rare event probabilities associated with MV-SDEs. We develop an SOC-based IS scheme to derive a time- and state-dependent IS control function for the resulting decoupled MV-SDE, resulting in a zero-variance estimator (see Section 4.4).
- We develop a novel DLMC estimator for approximating rare event probabilities associated with MV-SDEs up to a prescribed relative error tolerance TOL_r , along with detailed relative bias and statistical error analysis (see Section 4.5). The DLMC estimator with the proposed SOC-based IS scheme is proven to have an optimal computational complexity of $\mathcal{O}(\text{TOL}_r^{-4})$ with a bounded coefficient of variation. We formulate a robust DLMC algorithm that iteratively chooses the optimal parameters.
- We combine IS with MLMC to formulate a novel MLDLMC estimator for approximating rare event probabilities associated with MV-SDEs up to a prescribed relative error tolerance TOL_r (see Section 4.6). We develop a novel antithetic sampler to enhance variance convergence rates of the level difference DLMC estimators. The MLDLMC estimator with the proposed SOC-based IS scheme and the antithetic sampler is proven to have an optimal computational complexity of $\mathcal{O}(\text{TOL}_r^{-3})$ with a bounded coefficient of variation for the considered example. We formulate a robust MLDLMC algorithm that iteratively chooses the optimal parameters.
- We combine IS with multi-index MC to formulate a novel MIDLMC estimator for approximating rare event probabilities associated with MV-SDEs up to a prescribed

relative error tolerance TOL_r (see Section 4.7). We develop a novel antithetic sampler to enhance variance convergence rates of the first-order mixed difference **DLMC** estimators. The **MIDLMC** estimator with the proposed **SOC**-based **IS** scheme and the antithetic sampler is proven to have an optimal computational complexity of $\mathcal{O}\left(\text{TOL}_r^{-2}(\log \text{TOL}_r^{-1})^2\right)$ with a bounded coefficient of variation for the considered example. We formulate a robust **MIDLMC** algorithm that iteratively chooses the optimal parameters and index set.

4.1 McKean–Vlasov stochastic differential equation

MV-SDEs are a special class of **SDEs** whose drift and diffusion functions depend on the process' law itself [17]. In this thesis, we consider a broad class of **MV-SDEs** that arise from the mean-field limit of stochastic interacting particle systems with pairwise interaction kernels [65]. We consider the probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, where \mathcal{F}_t is the filtration of a standard Wiener process. For functions $\mathbf{a} : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$, $\mathbf{b} : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^{d \times d}$, $\kappa_1 : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, and $\kappa_2 : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, we consider the following Itô **SDE** for the McKean–Vlasov stochastic process $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$:

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}\left(\mathbf{X}(t), \int_{\mathbb{R}^d} \kappa_1(\mathbf{X}(t), \mathbf{x}) \mu_t(d\mathbf{x})\right) dt \\ \quad + \mathbf{b}\left(\mathbf{X}(t), \int_{\mathbb{R}^d} \kappa_2(\mathbf{X}(t), \mathbf{x}) \mu_t(d\mathbf{x})\right) d\mathbf{W}(t), \quad t > 0 \\ \mathbf{X}(0) = \mathbf{x}_0 \sim \mu_0 \in \mathcal{P}(\mathbb{R}^d), \end{cases} \quad (4.1)$$

where $\mathbf{W} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ is a standard d -dimensional Wiener process with mutually independent components. μ_t is the law of the stochastic process $\mathbf{X}(t)$. $\mathbf{x}_0 \in \mathbb{R}^d$ is a random initial state with distribution $\mu_0 \in \mathcal{P}(\mathbb{R}^d)$, where $\mathcal{P}(\mathbb{R}^d)$ is the space of probability measures on \mathbb{R}^d . **MV-SDEs** arise from the mean-field behaviour of stochastic interacting particle systems, commonly used in diverse applications in pedestrian dynamics [18], collective animal behavior [19], oscillator systems [27, 66], biological interactions [67], and financial mathematics [20]. We assume the following standard conditions for the existence and uniqueness of solutions to (4.1) [65, 68, 69, 70].

- κ_1, κ_2 are Lipschitz continuous, i.e. there exist constants $C_{\kappa_1}, C_{\kappa_2} \in \mathbb{R}^+$ such that for all $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in \mathbb{R}^d$,

$$|\kappa_1(\mathbf{x}, \mathbf{y}) - \kappa_1(\mathbf{x}', \mathbf{y}')| \leq C_{\kappa_1} (\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{y} - \mathbf{y}'\|), \quad (4.2)$$

$$|\kappa_2(\mathbf{x}, \mathbf{y}) - \kappa_2(\mathbf{x}', \mathbf{y}')| \leq C_{\kappa_2} (\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{y} - \mathbf{y}'\|). \quad (4.3)$$

- Functions \mathbf{a}, \mathbf{b} are Lipschitz continuous, i.e. there exist constants $C_{\mathbf{a}}, C_{\mathbf{b}} \in \mathbb{R}^+$ such that for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ and all $\mathbf{y}, \mathbf{y}' \in \mathbb{R}$,

$$\|\mathbf{a}(\mathbf{x}, \mathbf{y}) - \mathbf{a}(\mathbf{x}', \mathbf{y}')\| \leq C_{\mathbf{a}} (\|\mathbf{x} - \mathbf{x}'\| + |\mathbf{y} - \mathbf{y}'|), \quad (4.4)$$

$$\|\mathbf{b}(\mathbf{x}, \mathbf{y}) - \mathbf{b}(\mathbf{x}', \mathbf{y}')\| \leq C_{\mathbf{b}} (\|\mathbf{x} - \mathbf{x}'\| + |\mathbf{y} - \mathbf{y}'|). \quad (4.5)$$

- Functions \mathbf{a}, \mathbf{b} are bounded, i.e there exists a constant $C_{\text{bound}} \in \mathbb{R}$ such that for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ and $\mathbf{y}, \mathbf{y}' \in \mathbb{R}$,

$$\|\mathbf{a}(\mathbf{x}, \mathbf{y})\| + \|\mathbf{b}(\mathbf{x}', \mathbf{y}')\| \leq C_{\text{bound}} < \infty. \quad (4.6)$$

- There exists a positive constant $C_{\text{DG}} \geq 1$, such that for all $y \in \mathbb{R}$ and all $\mathbf{x}, \zeta \in \mathbb{R}^d$

$$\frac{1}{C_{\text{DG}}} \|\zeta\|^2 \leq \langle \mathbf{b} \mathbf{b}^{\text{tr}}(\mathbf{x}, y) \zeta, \zeta \rangle \leq C_{\text{DG}} \|\zeta\|^2. \quad (4.7)$$

More recent existence and uniqueness results in [71, 72, 73] relax some of the above assumptions, but we do not consider these extensions in this thesis. Under the above conditions, the time evolution of the deterministic mean-field law μ_t of the McKean–Vlasov process is given by the following nonlinear Fokker–Planck equation [74].

$$\begin{cases} \frac{\partial \mu(t, \mathbf{x})}{\partial t} + \sum_{i=1}^d \frac{\partial}{\partial x_i} \left(a_i \left(\mathbf{x}, \int_{\mathbb{R}^d} \kappa_1(\mathbf{x}, \mathbf{z}) \mu(t, \mathbf{z}) d\mathbf{z} \right) \mu(t, \mathbf{x}) \right) \\ - \sum_{i=1}^d \sum_{j=1}^d \frac{1}{2} \frac{\partial^2}{\partial x_i \partial x_j} \left(\sum_{k=1}^d b_{ik} b_{jk} \left(\mathbf{x}, \int_{\mathbb{R}^d} \kappa_2(\mathbf{x}, \mathbf{z}) \mu(t, \mathbf{z}) d\mathbf{z} \right) \right) \cdot \\ \mu(t, \mathbf{x}) \end{cases} = 0, \quad (t, \mathbf{x}) \in (0, T) \times \mathbb{R}^d \quad (4.8)$$

The well-posedness of the PDE (4.8) has been shown in [75, 76] subject to conditions (4.2)–(4.7). Equation (4.8) represents a nonlinear PDE due to the dependency of \mathbf{a} and \mathbf{b} on $\mu(t, \mathbf{x})$. Equation (4.8) is also an integro-differential equation with nonlocal terms because the drift and diffusion depend on an integral over \mathbb{R}^d with respect to $\mu(t, \mathbf{x})$. It is computationally prohibitive to solve such a PDE numerically to access solutions to the MV-SDE (4.1) up to relative error tolerances, particularly in high dimensions ($d \gg 1$). This motivates the use of strong approximations that converge strongly (in the L^p sense) to the MV-SDE (4.1).

4.1.1 Strong approximation: The stochastic interacting particle system

A strong approximation to the solution of the MV-SDE (4.1) is obtained by solving a system of P exchangeable Itô SDEs, also known as a stochastic interacting particle system with pairwise interaction kernels [65]. For $p = 1, \dots, P$, the process $\mathbf{X}_p^P : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ solves the following SDE:

$$\begin{cases} d\mathbf{X}_p^P(t) = \mathbf{a} \left(\mathbf{X}_p^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_1(\mathbf{X}_p^P(t), \mathbf{X}_j^P(t)) \right) dt \\ + \mathbf{b} \left(\mathbf{X}_p^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{X}_p^P(t), \mathbf{X}_j^P(t)) \right) d\mathbf{W}_p(t), \quad t > 0 \\ \mathbf{X}_p^P(0) = (\mathbf{x}_0)_p \sim \mu_0 \in \mathcal{P}(\mathbb{R}^d), \end{cases} \quad (4.9)$$

where $\{(\mathbf{x}_0)_p\}_{p=1}^P$ are iid random variables sampled from the initial distribution, μ_0 , and $\{\mathbf{W}_p : [0, T] \times \Omega \rightarrow \mathbb{R}^d\}_{p=1}^P$ are mutually independent d -dimensional Wiener processes, which are also independent of $\{(\mathbf{x}_0)_p\}_{p=1}^P$. Equation (4.9) approximates the mean-field law, μ_t , in (4.1) by an empirical law based on particles $\{\mathbf{X}_p^P\}_{p=1}^P$.

$$\mu_t(d\mathbf{x}) \approx \mu_t^P(d\mathbf{x}) = \frac{1}{P} \sum_{j=1}^P \delta_{\mathbf{X}_j^P(t)}(d\mathbf{x}), \quad (4.10)$$

where the particles $\{\mathbf{X}_p^P\}_{p=1}^P$ are identically distributed, but not mutually independent due to the interaction kernels in the drift and diffusion functions. Strong convergence (convergence in L^p -sense) of the particle system to the **MV-SDE** (4.1) follows from the results in [77, 78, 79], under the assumptions on \mathbf{a}, \mathbf{b} stated in this thesis. The high dimensionality of the Fokker–Planck **PDE**, satisfied by the joint probability density of the particle system, motivates the use of **MC** methods, which do not suffer from the curse of dimensionality.

4.1.2 Examples

Gaussian McKean–Vlasov process

A simple **1D** example of a **MV-SDE** is the following:

$$\begin{cases} dX(t) = (-a_1 X(t) + a_2 \mathbb{E}[X(t)]) dt + b dW(t), & t > 0 \\ X(0) = x_0 \sim \mathcal{N}(\mu, \sigma^2), \end{cases} \quad (4.11)$$

$a_1, a_2 \in \mathbb{R}^+$ are constants, the diffusion $b \in \mathbb{R}$ is constant, and x_0 represents an initial value sampled from the normal distribution $\mathcal{N}(\mu, \sigma^2)$. In addition, W represents a **1D** Wiener process, and x_0, W are mutually independent. The drift depends on the law of $X(t)$ via its expectation. This simple example has the analytical solution

$$X(t) \sim \mathcal{N}\left(\mu e^{(a_2 - a_1)t}, \sigma^2 e^{-2a_1 t} + \frac{b^2}{2a_1}(1 - e^{-2a_1 t})\right).$$

Hence, a strong approximation is not required to solve this **MV-SDE**, even in higher dimensions. Moreover, **MV-SDE** (4.11) can be solved (semi-)analytically even if the drift depends linearly on higher moments (e.g. $\mathbb{E}[X(t)^2]$) or quantiles (e.g. $\mathbb{P}[X(t) > \alpha(t)]$).

Kuramoto synchronisation model

The Kuramoto model is a **1D** example of a **MV-SDE** describing synchronization in statistical physics to model the behaviour of large sets of coupled oscillators. This **SDE** has been used to model cooperative synchronisation phenomena like synchronous flashing of fireflies [27], neuronal firing synchrony in the brain [80], and synchronisation of candle flame arrays [66]. It is a system of P fully connected, synchronized oscillators. We consider the following Itô **SDE** for the phase $X_p^P : [0, T] \times \Omega \rightarrow \mathbb{R}$ of each oscillator:

$$\begin{cases} dX_p^P(t) = \left(\xi_p + \frac{1}{P} \sum_{q=1}^P \sin(X_p^P(t) - X_q^P(t)) \right) dt + \sigma dW_p(t), & t > 0 \\ X_p^P(0) = (x_0)_p \sim \mu_0 \in \mathcal{P}(\mathbb{R}), \end{cases} \quad (4.12)$$

where $\{\xi_p\}_{p=1}^P$ denotes **iid** random variables sampled from a prescribed distribution and models the natural frequency of each oscillator. The diffusion $\sigma \in \mathbb{R}$ is constant, and $\{(x_0)_p\}_{p=1}^P$ represents **iid** random variables sampled from a prescribed distribution μ_0 . In addition, $\{W_p\}_{p=1}^P$ represents mutually independent **1D** Wiener processes, and $\{\xi_p\}_{p=1}^P$, $\{(x_0)_p\}_{p=1}^P$, and $\{W_p\}_{p=1}^P$ are mutually independent. This coupled particle system converges strongly (in the L^p sense) in the mean-field limit as the number of oscillators tends

to infinity. In this limit, each particle satisfies the following **MV-SDE**:

$$\begin{cases} dX(t) = \left(\xi + \int_{\mathbb{R}} \sin(X(t) - x) \mu_t(dx) \right) dt + \sigma dW(t), & t > 0 \\ X(0) = x_0 \sim \mu_0 \in \mathcal{P}(\mathbb{R}), \end{cases} \quad (4.13)$$

where $X(t)$ is the McKean–Vlasov process at time t , ξ represents a random variable sampled from some prescribed distribution, and μ_t is the law of $X(t)$. A common quantity of interest for this model is the order parameter, $\mathbb{E} [\sin(X(t))]^2 + \mathbb{E} [\cos(X(t))]^2$, which quantifies the synchronization level of the oscillators. An order parameter close to one indicates synchronization, implying strong coupling between the oscillators, while an order parameter near zero suggests that the oscillators are not synchronized.

4.2 Problem statement

We let $T > 0$ be some finite terminal time and $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ denote the McKean–Vlasov process (4.1). We let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a given scalar observable. The objective is to build a computationally efficient **MC** estimator \mathcal{A}_{MC} for $\mathbb{E} [g(\mathbf{X}(T))]$ with a given relative tolerance $\text{TOL}_r > 0$ that satisfies

$$\mathbb{P} \left[\frac{|\mathcal{A}_{\text{MC}} - \mathbb{E} [g(\mathbf{X}(T))]|}{|\mathbb{E} [g(\mathbf{X}(T))]|} \geq \text{TOL}_r \right] \leq \nu, \quad (4.14)$$

for a given confidence level determined by $0 < \nu \ll 1$. The high dimensionality of the **KBE** corresponding to the stochastic particle system (4.9) makes **PDE**-based numerical solutions of $\mathbb{E} [g(\mathbf{X}(T))]$ up to a relative accuracy TOL_r computationally infeasible. This motivates employing **MC** methods to overcome the curse of dimensionality. In this thesis, we are specifically interested in computing rare event probabilities associated with the final state $\mathbf{X}(T)$, implying that the function g is an indicator function or a smoothed version of it. Let \mathcal{D} denote a d -dimensional failure region deep in the tails of $\mathbf{X}(T)$. Then $g(\mathbf{x}) = \mathcal{I}_{\{\mathcal{D}\}}(\mathbf{x})$.

4.3 Naïve Euler–Maruyama Monte Carlo estimator

A naïve **MC** method to approximate $\mathbb{E} [g(\mathbf{X}(T))]$ would be to generate approximate sample paths of the interacting particle system (4.9) by using the **EM** discretization scheme (see Definition 2.3.4). Previous works have investigated **MC** methods using this numerical scheme for bounded, Lipschitz drift/diffusion functions and smooth observables [81, 82, 83], and they were able to achieve $\mathcal{O}(\text{TOL}^{-4})$ computational complexity for a prescribed absolute error tolerance TOL .

Proposition 4.3.1 (Euler–Maruyama time discretization for stochastic particle systems). *Consider a discretization $0 = t_0 < t_1 < t_2 < \dots < t_N = T$ of the domain $[0, T]$ with N time steps. Let $\mathbf{X}_p^{P|N}$ denote the time discretized version of the process \mathbf{X}_p^P corresponding to the **SDE** (4.9). Then, the **EM** time discretization of the **SDE** (4.9) is given for all $p = 1, \dots, P$ as follows.*

- For $n = 0$,

$$\mathbf{X}_p^{P|N}(t_0) = \mathbf{X}_p^P(0) = (\mathbf{x}_0)_p. \quad (4.15)$$

- For $n = 1, \dots, N - 1$,

$$\begin{aligned} \mathbf{X}_p^{P|N}(t_{n+1}) &= \mathbf{X}_p^{P|N}(t_n) + \mathbf{a} \left(\mathbf{X}_p^{P|N}(t_n), \frac{1}{P} \sum_{j=1}^P \kappa_1(\mathbf{X}_p^{P|N}(t_n), \mathbf{X}_j^{P|N}(t_n)) \right) \Delta t_n \\ &\quad + \mathbf{b} \left(\mathbf{X}_p^{P|N}(t_n), \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{X}_p^{P|N}(t_n), \mathbf{X}_j^{P|N}(t_n)) \right) \Delta \mathbf{W}_p(t_n). \end{aligned} \quad (4.16)$$

Here $\Delta t_n = t_{n+1} - t_n$, and $\Delta \mathbf{W}_p(t_n) = \sqrt{\Delta t_n} \boldsymbol{\varepsilon}_n$ are independent Wiener increments with $\boldsymbol{\varepsilon}_n$ being a d -dimensional standard normal variable.

Given M independent samples of the P -particle system $\{\mathbf{X}_p^{P|N}(t_N)\}_{p=1}^P$, a naïve MC approximation of $\mathbb{E}[g(\mathbf{X}(T))]$ is formulated as the following empirical average [82]:

$$\mathbb{E}[g(\mathbf{X}(T))] \approx \mathcal{A}_{\text{MC}} := \frac{1}{M} \sum_{m=1}^M \frac{1}{P} \sum_{p=1}^P g(\mathbf{X}_p^{P|N}(t_N, \omega_{1:P}^{(m)})), \quad (4.17)$$

where $\omega_{z:P}^{(m)} := \{\omega_q^{(m)}\}_{q=z}^P$ and, for each q , $\omega_q^{(m)}$ denotes the m^{th} sample of all random variables used to compute $\mathbf{X}_q^{P|N}$. For rare event probability estimation, however, the computational feasibility of naïve MC rapidly diminishes due to the problem of exploding coefficient of variation (see Section 3.1 and [12]). Hence, we combine the IS variance reduction technique with MC methods to produce computationally feasible estimates of rare event probabilities associated with MV-SDEs.

4.4 Importance sampling using the decoupling approach

Chapter 3 presented an SOC-based approach to obtain the optimal IS measure change for standard SDEs. However, there are two main difficulties while applying this for MV-SDEs. The drift and diffusion functions in the MV-SDE are inherently tied to $\{\mu_t : t \in [0, T]\}$, the law of the process itself. Changing the measure alters these coefficients as well. If, instead, we use the P -particle approximation (4.9) and apply the SOC-based IS approach in Chapter 3, we get a HJB PDE in Pd dimensions, that is computationally infeasible to solve numerically. We overcome this by utilizing a decoupling approach introduced in [21]. The key idea here is to transform the MV-SDE (4.1) into a standard SDE by conditioning on an empirical approximation of the mean-field law $\{\mu_t : t \in [0, T]\}$. The authors in [21] employed large deviations and the Pontryagin principle to solve approximately for a suboptimal IS measure change. Instead, we harness the optimal IS measure change for standard SDEs introduced in Chapter 3 to build an efficient MC estimator.

The decoupling approach first replaces the deterministic mean-field law $\{\mu_t : t \in [0, T]\}$ with an empirical approximation (4.10) using the interacting particle system (4.9). Conditioned on this empirical law computed beforehand, a decoupled MV-SDE is introduced and the IS change of measure is applied to it. In this way, the empirical estimation of the law $\{\mu_t : t \in [0, T]\}$ and the IS measure change is decoupled. The decoupling approach comprises the following steps.

1. We approximate the MV-SDE law $\{\mu_t : t \in [0, T]\}$ using the empirical measure $\{\mu_t^P : t \in [0, T]\}$ in (4.10) using particles $\{\mathbf{X}_p^P(t) : t \in [0, T]\}_{p=1}^P$ satisfying (4.9).

2. Given $\{\mathbf{X}_p^P(t)\}_{p=1}^P \sim \mu_t^P : t \in [0, T]$, we define the decoupled McKean–Vlasov process $\bar{\mathbf{X}}^P : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ as the solution to the following Itô SDE.

$$\begin{cases} d\bar{\mathbf{X}}^P(t) = \mathbf{a} \left(\bar{\mathbf{X}}^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_1(\bar{\mathbf{X}}^P(t), \mathbf{X}_j^P(t)) \right) dt \\ \quad + \mathbf{b} \left(\bar{\mathbf{X}}^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_2(\bar{\mathbf{X}}^P(t), \mathbf{X}_j^P(t)) \right) d\bar{\mathbf{W}}(t), \quad t \in [0, T] \\ \bar{\mathbf{X}}^P(0) = \bar{\mathbf{x}}_0 \sim \mu_0, \quad \bar{\mathbf{x}}_0 \in \mathbb{R}^d, \end{cases} \quad (4.18)$$

where superscript P indicates that the drift and diffusion functions in (4.18) are computed using $\{\mu_t^P : t \in [0, T]\}$ derived from the stochastic P -particle system. Functions \mathbf{a} and \mathbf{b} are the same as defined in Section 4.1. In addition, $\bar{\mathbf{W}} : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ is a standard d -dimensional Wiener process independent of the Wiener processes $\{\mathbf{W}_p\}_{p=1}^P$ used in (4.9), and $\bar{\mathbf{x}}_0 \in \mathbb{R}^d$ is a random initial state sampled from μ_0 as defined in (4.1) and is independent from $\{(\mathbf{x}_0)_p\}_{p=1}^P$ in (4.9). Thus, (4.18) is a standard SDE with random coefficients.

3. We introduce a copy space [21] to distinguish the decoupled MV-SDE (4.18) from the stochastic P -particle system. We suppose (4.9) is defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We define a copy space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$; hence, we define (4.18) on the product space $(\Omega, \mathcal{F}, \mathbb{P}) \times (\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$. Thus, \mathbb{P} is a probability measure generated by the randomness of $\{\mu_t^P : t \in [0, T]\}$, and $\bar{\mathbb{P}}$ denotes the measure generated by the randomness of the Wiener process driving (4.18) conditioned on $\{\mu_t^P : t \in [0, T]\}$.
4. We approximate the quantity of interest as the following nested expectation.

$$\begin{aligned} \mathbb{E}[g(\mathbf{X}(T))] &\approx \mathbb{E}_{\mathbb{P} \otimes \bar{\mathbb{P}}}[g(\bar{\mathbf{X}}^P(T))] \\ &= \mathbb{E}_{\mathbb{P}} \left[\mathbb{E}_{\bar{\mathbb{P}}} [g(\bar{\mathbf{X}}^P(T)) \mid \{\mu_t^P : t \in [0, T]\}] \right]. \end{aligned} \quad (4.19)$$

Henceforth, we omit the probability measures while denoting the expectation to simply mean that $\mathbb{E}[g(\bar{\mathbf{X}}^P(T))]$ denotes an expectation over all randomness in the decoupled MV-SDE (4.18).

Since (4.18) is a standard SDE for a given empirical law $\{\mathbf{X}_p^P(t)\}_{p=1}^P \sim \mu_t^P : t \in [0, T]$, we follow the procedure in Section 3.2 to derive the optimal IS measure change. First, we formulate the HJB PDE that provides the optimal IS control for the decoupled MV-SDE conditioned on the empirical law $\{\mathbf{X}_p^P(t)\}_{p=1}^P \sim \mu_t^P : t \in [0, T]$.

Corollary 4.4.1 (HJB PDE for decoupled MV-SDE). *Let the decoupled process $\bar{\mathbf{X}}^P$ follow the dynamics (4.18) and let the controlled process $\bar{\mathbf{X}}_\xi^P : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ follow*

the given controlled dynamics with the Markov control function $\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$:

$$\left\{ \begin{array}{l} d\bar{\mathbf{X}}_\zeta^P(t) = \left(\mathbf{a} \left(\bar{\mathbf{X}}_\zeta^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_1(\bar{\mathbf{X}}_\zeta^P(t), \mathbf{X}_j^P(t)) \right) \right. \\ \quad \left. + \mathbf{b} \left(\bar{\mathbf{X}}_\zeta^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_2(\bar{\mathbf{X}}_\zeta^P(t), \mathbf{X}_j^P(t)) \right) \zeta(t, \bar{\mathbf{X}}_\zeta^P(t)) \right) dt \\ \quad + \mathbf{b} \left(\bar{\mathbf{X}}_\zeta^P(t), \frac{1}{P} \sum_{j=1}^P \kappa_2(\bar{\mathbf{X}}_\zeta^P(t), \mathbf{X}_j^P(t)) \right) d\mathbf{W}(t), \quad 0 < t < T \\ \bar{\mathbf{X}}_\zeta^P(0) = \bar{\mathbf{X}}^P(0) = \bar{\mathbf{x}}_0 \sim \mu_0. \end{array} \right. \quad (4.20)$$

We use (4.9) to obtain $\{\mathbf{X}_p^P(t)\}_{p=1}^P \sim \mu_t^P : t \in [0, T]$ in (4.18) and (4.20). We assume that the policy ζ satisfies the Novikov's condition (see Lemma 3.2.1).

$$\mathbb{E} \left[\exp \left\{ \frac{1}{2} \int_0^T \|\zeta(s, \bar{\mathbf{X}}_\zeta^P(s))\|^2 ds \right\} \mid \{\mu_t^P : t \in [0, T]\} \right] < \infty.$$

The value function $v : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ that minimises the second moment of the MC estimator of $\mathbb{E} [g(\bar{\mathbf{X}}^P(T)) \mid \{\mu_t^P : t \in [0, T]\}]$ can be defined as

$$v(t, \mathbf{x}) = \min_{\zeta \in \mathcal{Z}} \mathbb{E} \left[g^2(\bar{\mathbf{X}}_\zeta^P(T)) \exp \left\{ - \int_t^T \|\zeta(s, \bar{\mathbf{X}}_\zeta^P(s))\|^2 ds \right. \right. \\ \left. \left. - 2 \int_t^T \langle \zeta(s, \bar{\mathbf{X}}_\zeta^P(s)), d\mathbf{W}(s) \rangle \right\} \mid \bar{\mathbf{X}}_\zeta^P(t) = \mathbf{x}, \{\mu_t^P : t \in [0, T]\} \right]. \quad (4.21)$$

Here, $\mathcal{Z} = \{f \in C^1([0, T] \times \mathbb{R}^d)\}$ is a set of admissible deterministic d -dimensional Markov controls. Define a new function $\gamma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, such that $v(t, \mathbf{x}) = \exp\{-2\gamma(t, \mathbf{x})\}$. Then, γ satisfies the following nonlinear HJB equation:

$$\left\{ \begin{array}{l} \frac{\partial \gamma}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_1(\mathbf{x}, \mathbf{X}_j^P(t)) \right), \nabla \gamma(t, \mathbf{x}) \rangle \\ \quad + \frac{1}{2} \nabla^2 \gamma(t, \mathbf{x}) : \left(\mathbf{b} \mathbf{b}^{\text{tr}} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{x}, \mathbf{X}_j^P(t)) \right) \right) \\ \quad - \frac{1}{4} \left\| \mathbf{b}^{\text{tr}} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{x}, \mathbf{X}_j^P(t)) \right) \nabla \gamma(t, \mathbf{x}) \right\|^2 \\ \quad = 0, \quad (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d \\ \gamma(T, \mathbf{x}) = -\log |g(\mathbf{x})|, \quad \mathbf{x} \in \mathbb{R}^d, \end{array} \right. \quad (4.22)$$

with optimal control

$$\zeta^*(t, \mathbf{x}) = -\mathbf{b}^{\text{tr}} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{x}, \mathbf{X}_j^P(t)) \right) \nabla \gamma(t, \mathbf{x}), \quad (4.23)$$

which minimises the second moment (4.21), conditioned on $\{\mu_t^P : t \in [0, T]\}$.

Remark 4.4.1 (Zero-variance control for MV-SDEs). *If we somehow have access to the true mean-field law $\{\mu_t : t \in [0, T]\}$, the controls given by (4.23) result in a zero-variance estimator of $\mathbb{E}[g(\mathbf{X}(T))]$ for MV-SDEs. However, in practice, we can only approximate μ_t by an empirical estimate.*

In Corollary 4.4.1, we control the decoupled McKean–Vlasov process $\bar{\mathbf{X}}^P$ (4.18) instead of the particles $\{\mathbf{X}_p^P\}_{p=1}^P$ from the interacting particle system (4.9) because the optimal control PDE is d -dimensional instead of Pd -dimensional. As shown in Section 3.3, we can recover the linear KBE with the change in variable $v(t, \mathbf{x}) = u(t, \mathbf{x})^2$, and obtain the following zero-variance control PDE for a given $\{\mu_t^P : t \in [0, T]\}$.

$$\begin{cases} \frac{\partial u}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_1(\mathbf{x}, \mathbf{X}_j^P(t)) \right), \nabla u(t, \mathbf{x}) \rangle \\ + \frac{1}{2} \nabla^2 u(t, \mathbf{x}) : (\mathbf{b} \mathbf{b}^{\text{tr}}) \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{x}, \mathbf{X}_j^P(t)) \right) \\ = 0, \quad (t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d \\ u(T, \mathbf{x}) = |g(\mathbf{x})|, \quad \mathbf{x} \in \mathbb{R}^d, \end{cases} \quad (4.24)$$

with optimal control

$$\zeta^*(t, \mathbf{x}) = \mathbf{b}^{\text{tr}} \left(\mathbf{x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{x}, \mathbf{X}_j^P(t)) \right) \nabla \log u(t, \mathbf{x}). \quad (4.25)$$

The likelihood brought about by the IS measure change described in Corollary 4.4.1 can be written as the following Radon–Nikodym derivative.

$$\begin{aligned} \mathbb{L}_{[0, T]}^P &= \frac{d\bar{\mathbb{Q}}}{d\bar{\mathbb{P}}} \Big|_0^T \\ &= \exp \left\{ - \int_0^T \langle \zeta(s, \bar{\mathbf{X}}_\zeta^P(s)), d\mathbf{W}^{\bar{\mathbb{Q}}}(s) \rangle - \frac{1}{2} \int_0^T \|\zeta(s, \bar{\mathbf{X}}_\zeta^P(s))\|^2 ds \right\}. \end{aligned} \quad (4.26)$$

Here $\bar{\mathbb{Q}}$ denotes the shifted probability measure of the copy space, such that (4.20) is defined on the product space $(\Omega, \mathcal{F}, \mathbb{P}) \times (\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{Q}})$. With this measure change, the nested expectation (4.19) can be rewritten as follows.

$$\begin{aligned} \mathbb{E}[g(\mathbf{X}(T))] &\approx \mathbb{E}_{\mathbb{P}} \left[\mathbb{E}_{\mathbb{P}} \left[g(\bar{\mathbf{X}}^P(T)) \mid \{\mu_t^P : t \in [0, T]\} \right] \right] \\ &= \mathbb{E}_{\mathbb{P}} \left[\mathbb{E}_{\bar{\mathbb{Q}}} \left[g(\bar{\mathbf{X}}_\zeta^P(T)) \mathbb{L}_{[0, T]}^P \mid \{\mu_t^P : t \in [0, T]\} \right] \right]. \end{aligned} \quad (4.27)$$

We approximate the nested expectation (4.27) using a DLMC estimator.

4.5 Double loop Monte Carlo estimator

The DLMC estimator for the nested expectation (4.27) is constructed through the following steps:

1. Approximate the mean-field law μ_t in (4.1) by the empirical law $\{\mathbf{X}_p^P(t)\}_{p=1}^P \sim \mu_t^P : t \in [0, T]$ using (4.9). In practice, we obtain a time-discretized empirical law with

N time steps from the EM discretization $\{\mathbf{X}_p^{P|N}\}_{p=1}^P$ of the particle system (4.16). Let $\bar{\mu}^{P|N}$ be the discrete law obtained.

$$\bar{\mu}^{P|N}(t_n) = \frac{1}{P} \sum_{p=1}^P \delta_{\mathbf{x}_p^{P|N}(t_n)}, \quad \forall n = 0, \dots, N. \quad (4.28)$$

Let $\omega_{1:P}$ denote the P underlying sets of random variables used to generate a realization of $\mathbf{X}_p^{P|N}$ using (4.9).

2. The empirical law must be defined for all $t \in [0, T]$ for the decoupled MV-SDE to be well-defined. For this, we use the EM continuous-time extension of the time-discretized stochastic particle system $\{\mathbf{X}_p^{P|N}\}_{p=1}^P$. Given $\{\mathbf{X}_p^{P|N}\}_{p=1}^P$ for all time steps $n = 0, \dots, N$, we have

$$\begin{aligned} \mathbf{X}_p^{P|N}(t) &= \mathbf{X}_p^{P|N}(t_n) + \mathbf{a} \left(\mathbf{X}_p^{P|N}(t_n), \frac{1}{P} \sum_{j=1}^P \kappa_1(\mathbf{X}_p^{P|N}(t_n), \mathbf{X}_j^{P|N}(t_n)) \right) (t - t_n) \\ &+ \mathbf{b} \left(\mathbf{X}_p^{P|N}(t_n), \frac{1}{P} \sum_{j=1}^P \kappa_2(\mathbf{X}_p^{P|N}(t_n), \mathbf{X}_j^{P|N}(t_n)) \right) (\mathbf{W}(t) - \mathbf{W}(t_n)), \quad t_n < t < t_{n+1}. \end{aligned} \quad (4.29)$$

The continuous-time extension $\mu^{P|N}$ of the discrete empirical law $\bar{\mu}^{P|N}$ is then defined as follows

$$\mu^{P|N}(t) = \frac{1}{P} \sum_{j=1}^P \delta_{\mathbf{x}_j^{P|N}(t)}, \quad \forall t \in [0, T]. \quad (4.30)$$

3. Given $\mu^{P|N}$ from (4.30), we solve (4.24) and use (4.25) to obtain the optimal policy ζ for IS.
4. Given the empirical law $\mu^{P|N}$ from (4.30) and control function $\zeta : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, we generate sample paths of the controlled decoupled MV-SDE (4.20). For this, we use the EM discretization of SDE (4.20) with N time steps to obtain approximate sample paths denoted by $\{\bar{\mathbf{X}}_\zeta^{P|N}(t_n)\}_{n=1}^N$. Let $\bar{\omega}$ denote the set of random variables used to generate one of these sample paths.
5. The nested expectation (4.27) is approximated by the DLMC estimator $\mathcal{A}_{\text{DLMC}}$ using the following nested sample average.

$$\mathcal{A}_{\text{DLMC}} = \frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} g \left(\bar{\mathbf{X}}_\zeta^{P|N}(t_N) \right) \mathbb{L}^{P|N} \left(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)} \right). \quad (4.31)$$

where $\mathbb{L}^{P|N}$ denotes the time-discretized approximation of the likelihood (4.26) defined as

$$\begin{aligned} \mathbb{L}^{P|N} \left(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)} \right) &= \prod_{n=0}^{N-1} \exp \left\{ -\frac{1}{2} \Delta t_n \left\| \zeta(t_n, \bar{\mathbf{X}}_\zeta^{P|N}(t_n)) \right\|^2 \right. \\ &\quad \left. - \langle \Delta \mathbf{W}(t_n), \zeta(t_n, \bar{\mathbf{X}}_\zeta^{P|N}(t_n)) \rangle \right\} \left(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)} \right). \end{aligned} \quad (4.32)$$

M_1 denotes the number of independent realizations of $\mu^{P|N}$ in the DLMC estimator, and $\omega_{1:P}^{(m_1)}$ denotes the m_1^{th} realization of $\omega_{1:P}$. For each realization of $\mu^{P|N}$, M_2 denotes the number of sample paths of the decoupled MV-SDE, and $\bar{\omega}^{(m_2)}$ denotes the m_2^{th} realization of $\bar{\omega}$. Further, $\Delta\mathbf{W}(t_n)$ are the Wiener increments driving the dynamics of the time-discretized decoupled MV-SDE (4.20).

The above DLMC estimator requires solving the optimal control PDE (4.24) for each $\mu^{P|N}$ realization, making it computationally cumbersome. To avoid this, we independently obtain a sufficiently accurate empirical law realization off-line using a sufficiently large number of particles and time steps and use this law to solve the control PDE (4.24). This approach is motivated by the convergence of the empirical law to the mean-field law $\{\mu_t : t \in [0, T]\}$ as the number of particles tends to infinity. This approach has two main advantages: we do not need to solve the original KBE (4.24) for the value function u to satisfy a relative tolerance TOL_r , and we don't require high accuracy of the optimal IS control ζ for the purpose of IS. Algorithm 1 depicts the above procedure in a pseudo-algorithmic manner.

Algorithm 1: DLMC algorithm with IS for MV-SDEs

Off-line:

Generate one realisation of law $\{\mathbf{X}_p^{P|N}\}_{p=1}^{\bar{P}} \sim \mu^{\bar{P}|N}$ with \bar{P} -particle system and \bar{N} time steps using (4.16) with some large \bar{P}, \bar{N} ;

Given $\mu^{\bar{P}|N}$, solve PDE (4.24) to obtain IS control ζ ;

Inputs: P, N, M_1, M_2, ζ ;

for $m_1 = 1, \dots, M_1$ **do**

Generate realization of random variables $\omega_{1:P}^{(m_1)}$;

Generate realization of law $\{\mathbf{X}_p^{P|N}\}_{p=1}^P \sim \mu^{P|N}(\omega_{1:P}^{(m_1)})$ with P -particle system and N time steps using (4.16);

for $m_2 = 1, \dots, M_2$ **do**

Generate realization of random variables $\bar{\omega}^{(m_2)}$;

Given $\mu^{P|N}(\omega_{1:P}^{(m_1)})$ and ζ , generate a sample path of (4.20) with N time steps;

Compute $g(\bar{\mathbf{X}}_{\zeta}^{P|N}(t_N))(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)})$;

Compute $\mathbb{L}^{P|N}(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)})$ using (4.32);

end

Approximate $\mathbb{E}\left[g(\bar{\mathbf{X}}_{\zeta}^{P|N}(t_N))\mathbb{L}^{P|N} \mid \mu^{P|N}(\omega_{1:P}^{(m_1)})\right]$ by

$$\frac{1}{M_2} \sum_{m_2=1}^{M_2} g(\bar{\mathbf{X}}_{\zeta}^{P|N}(t_N))\mathbb{L}^{P|N}(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)});$$

end

Approximate $\mathbb{E}\left[g(\bar{\mathbf{X}}^{P|N}(T))\right]$ by

$$\frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} g(\bar{\mathbf{X}}_{\zeta}^{P|N}(t_N))\mathbb{L}^{P|N}(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)});$$

4.5.1 Error analysis

We bound the relative error of the DLMC estimator, $\mathcal{A}_{\text{DLMC}}$, as

$$\frac{|\mathcal{A}_{\text{DLMC}} - \mathbb{E}[g(\mathbf{X}(T))]|}{|\mathbb{E}[g(\mathbf{X}(T))]|} \leq \underbrace{\frac{|\mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^{P|N}(T))]|}{|\mathbb{E}[g(\mathbf{X}(T))]|}}_{=\epsilon_b, \text{ Relative bias}} + \underbrace{\frac{|\mathbb{E}[g(\tilde{\mathbf{X}}^{P|N}(T))] - \mathcal{A}_{\text{DLMC}}|}{|\mathbb{E}[g(\mathbf{X}(T))]|}}_{=\epsilon_s, \text{ Relative statistical error}}. \quad (4.33)$$

Although $\mathcal{A}_{\text{DLMC}}$ should satisfy a prescribed relative error tolerance TOL_r in the sense of (4.14), we impose more restrictive conditions, which can be expressed as follows:

$$\text{Bias constraint: } \epsilon_b \leq \theta \text{TOL}_r, \quad (4.34)$$

$$\text{Statistical constraint: } \mathbb{P}[\epsilon_s \leq (1 - \theta)\text{TOL}_r] > 1 - \nu, \quad (4.35)$$

for a given tolerance splitting parameter $\theta \in (0, 1)$ and confidence level determined by $0 < \nu \ll 1$. Let us first analyse the estimator bias, which can be split into two terms:

$$\begin{aligned} \epsilon_b &= \frac{|\mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^{P|N}(T))]|}{|\mathbb{E}[g(\mathbf{X}(T))]|} \quad (4.36) \\ &\leq \underbrace{\frac{|\mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^P(T))]|}{|\mathbb{E}[g(\mathbf{X}(T))]|}}_{\text{Relative decoupling error}} + \underbrace{\frac{|\mathbb{E}[g(\tilde{\mathbf{X}}^P(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^{P|N}(T))]|}{|\mathbb{E}[g(\mathbf{X}(T))]|}}_{\text{Relative time discretization error}}, \end{aligned}$$

Assumption 4.5.1 (Decoupling error). *There exists a constant $C_1 > 0$ independent of P , such that*

$$\left| \mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^P(T))] \right| \leq C_1 P^{-1}.$$

Assumption 4.5.1 is motivated by the weak convergence of the stochastic particle system (4.9) to the MV-SDE (4.1) under the standard conditions (4.2)-(4.7) on the drift and diffusion functions. Since this is a well-known result for drift and diffusion functions satisfying (4.2)-(4.7), we omit its proof in this thesis. We refer readers to the proofs in [84, 85].

Assumption 4.5.2 (Time discretization error). *There exists a constant $C_2 > 0$ independent of N , such that*

$$\left| \mathbb{E}[g(\tilde{\mathbf{X}}^P(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^{P|N}(T))] \right| \leq C_2 N^{-1}.$$

Assumption 4.5.2 is motivated by the weak convergence of the EM time-discretization scheme for standard SDEs (see Proposition 2.3.4) under the standard conditions (4.2)-(4.7) on the drift and diffusion functions. The constants C_1, C_2 depend on T and the regularity and boundedness of the functions \mathbf{a}, \mathbf{b} and g and their derivatives. This is a standard result for the EM-scheme, and we refer readers to [48] for the proof. By substituting Assumptions 4.5.1 and 4.5.2 into (4.36), the bias bound can be expressed as

$$\epsilon_b |\mathbb{E}[g(\mathbf{X}(T))]| \leq \frac{C_1}{P} + \frac{C_2}{N}. \quad (4.37)$$

Analogous to the procedure in Section 3.1, we can use the CLT to satisfy the probabilistic constraint (4.35) on the relative statistical error of $\mathcal{A}_{\text{DLMC}}$, yielding the following constraint on the estimator's variance.

$$\frac{\text{Var}[\mathcal{A}_{\text{DLMC}}]}{|\mathbb{E}[g(\mathbf{X}(T))]|^2} \leq \left(\frac{(1-\theta)\text{TOL}_r}{C_\nu} \right)^2, \quad (4.38)$$

where C_ν is the $(1 - \frac{\nu}{2})$ -quantile of the standard normal distribution. Alternately, one could employ the approach in [55], which uses the Berry–Esseen theorem to derive non-asymptotic bounds on the relative statistical error. While this method can, in principle, be applied to the estimator $\mathcal{A}_{\text{DLMC}}$, it introduces additional complexity in the adaptive determination of the sample sizes M_1 and M_2 needed to satisfy the constraint in (4.35). Therefore, we proceed with the simpler asymptotic bound (4.38), based on the CLT. For notational convenience, we denote by $Y^{P|N}$ the random variable $g(\bar{\mathbf{X}}_\zeta^{P|N}(T)) \mathbb{L}^{P|N}$, and we denote by $Y_{m_1, m_2}^{P|N} := g(\bar{\mathbf{X}}_\zeta^{P|N}(T)) \mathbb{L}^{P|N}(\omega_{1:P}^{(m_1)} \times \bar{\omega}^{(m_2)})$ a random sample of $Y^{P|N}$. Thus, the DLMC estimator variance can be expressed as

$$\text{Var}[\mathcal{A}_{\text{DLMC}}] = \text{Var} \left[\frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} Y_{m_1, m_2}^{P|N} \right] \quad (4.39)$$

$$= \frac{1}{M_1} \text{Var} \left[\frac{1}{M_2} \sum_{m_2=1}^{M_2} Y_{1, m_2}^{P|N} \right]. \quad (4.40)$$

Then, by using the law of total variance, we obtain

$$\begin{aligned} \text{Var}[\mathcal{A}_{\text{DLMC}}] &= \frac{1}{M_1} \text{Var} \left[\mathbb{E} \left[\frac{1}{M_2} \sum_{m_2=1}^{M_2} Y_{1, m_2}^{P|N} \middle| \mu^{P|N} \right] \right] \\ &\quad + \frac{1}{M_1} \mathbb{E} \left[\text{Var} \left[\frac{1}{M_2} \sum_{m_2=1}^{M_2} Y_{1, m_2}^{P|N} \middle| \mu^{P|N} \right] \right] \\ &= \frac{1}{M_1} \text{Var} \left[\mathbb{E} \left[Y^{P|N} \middle| \mu^{P|N} \right] \right] \\ &\quad + \frac{1}{M_1 M_2} \mathbb{E} \left[\text{Var} \left[Y^{P|N} \middle| \mu^{P|N} \right] \right], \end{aligned} \quad (4.41)$$

where $\{Y_{1, m_2}^{P|N}\}_{m_2=1}^{M_2}$ denotes independent samples of $Y^{P|N}$ conditioned on $\mu^{P|N}$.

Assumption 4.5.3. *There exists a constant, $C_3 > 0$, which is independent of N and P , such that*

$$V_1^{P|N} = \text{Var} \left[\mathbb{E} \left[Y^{P|N} \middle| \mu^{P|N} \right] \right] \leq C_3 P^{-1}.$$

Assumption 4.5.3 is motivated by the strong convergence of the stochastic particle system (4.9) to the MV-SDE (4.1) [77]. As the outer variance is related to the randomness in the empirical law $\mu^{P|N}$, $V_1^{P|N}$ vanishes with the increase in the number of particles.

Assumption 4.5.4. *There exists a constant, $C_4 > 0$, which is independent of N and P , such that*

$$V_2^{P|N} = \mathbb{E} \left[\text{Var} \left[Y^{P|N} \middle| \mu^{P|N} \right] \right] \leq C_4 < \infty.$$

Assumption 4.5.4 is motivated by the bounded conditional variance after the optimal IS measure change in the dynamics (4.18). By substituting Assumptions 4.5.3 and 4.5.4

in (4.41), (4.38) can be expressed as

$$\frac{C_3}{PM_1} + \frac{C_4}{M_1M_2} \leq \left(\frac{(1-\theta)\text{TOL}_r |\mathbb{E}[G(\mathbf{X}(T))]|}{C_v} \right)^2. \quad (4.42)$$

4.5.2 Computational complexity

To analyse the computational cost of the DLMC estimator (4.31), we assume first a naïve method (with computational cost $\mathcal{O}(P)$) to compute the empirical means in the drift and diffusion functions in (4.9) and (4.18). Then, the computational cost required to generate one realization of the empirical law using the EM scheme (4.16) is $\mathcal{O}(NP^2)$. The computational cost required to generate one realization of the decoupled McKean–Vlasov process (4.20) using the EM scheme is $\mathcal{O}(PN)$. The computational cost to solve the optimal control PDE (4.24) and get the optimal IS control function is considered an off-line cost in this analysis. Then, the net computational cost of the DLMC estimator (4.31) is

$$\mathcal{W} := \mathcal{O}(M_1NP^2 + M_1M_2NP). \quad (4.43)$$

Under the constraints in (4.34) and (4.42), we determine optimal parameters P, N, M_1 , and M_2 that minimise the computational cost (4.43). This can be presented as the following optimization problem:

$$\left\{ \begin{array}{l} \min_{\{P, N, M_1, M_2\}} \quad \mathcal{W} = M_1NP^2 + M_1M_2NP \\ \text{s.t.} \quad \frac{C_1}{P} + \frac{C_2}{N} \approx \theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|, \\ \quad \quad C_v^2 \left(\frac{C_3}{PM_1} + \frac{C_4}{M_1M_2} \right) \approx (1-\theta)^2 \text{TOL}_r^2 |\mathbb{E}[g(\mathbf{X}(T))]|^2, \end{array} \right. \quad (4.44)$$

with the solution formulated in Theorem 4.5.1.

Theorem 4.5.1 (Optimal DLMC complexity). *Consider the DLMC estimator in (4.31), obtained from Algorithm 1. For any $\text{TOL}_r > 0$, there exist optimal parameters P, N, M_1, M_2 such that (4.34) and (4.42) hold. The optimal computational work of this estimator is*

$$\mathcal{W} = M_1NP^2 + M_1M_2NP = \mathcal{O}(\text{TOL}_r^{-4}). \quad (4.45)$$

Proof. Refer to Appendix A.3. □

The double loop approach for this problem crucially guarantees $\mathcal{O}(\text{TOL}_r^{-4})$ complexity for the proposed estimator. From (4.44) and Theorem 4.5.1, one can deduce that the use of only a single realization of the empirical law ($M_1 = 1$) results in an increase in complexity to $\mathcal{O}(\text{TOL}_r^{-5})$. Moreover, in comparison to the naïve MC estimator in [82], the proposed IS scheme considerably reduces the estimator's coefficient of variation for rare event probabilities.

4.5.3 Adaptive algorithm

We formulate a DLMC algorithm that adaptively selects optimal parameters P, N, M_1 , and M_2 . For computational convenience, we restrict our choices for P, N by implementing the following hierarchies:

- $P_\ell = P_0 \times 2^\ell$, $\ell = 0, \dots, L$;
- $N_\ell = N_0 \times 2^\ell$, $\ell = 0, \dots, L$.

To build an adaptive algorithm, we must estimate the bias and variances $V_{1,\ell} := V_1^{P_\ell|N_\ell}$ and $V_{2,\ell} := V_2^{P_\ell|N_\ell}$ for some level ℓ in a cheap and robust manner. First, we select a level $\ell = L$ that satisfies the bias constraint (4.34). For convenience of the following analysis, we denote $G = g(\mathbf{X}(T))$ and its discretization at level ℓ as $G_\ell = g\left(\tilde{\mathbf{X}}_\zeta^{P_\ell|N_\ell}(T)\right) \mathbb{L}^{P_\ell|N_\ell}$. The likelihood factor, $\mathbb{L}^{P_\ell|N_\ell}$, at level ℓ is computed using (4.32). Then, the relative bias for level ℓ can be expressed as

$$\epsilon_b = \frac{|\mathbb{E}[G - G_\ell]|}{|\mathbb{E}[G]|} \leq \frac{1}{|\mathbb{E}[G]|} \left(\frac{C_1}{P_\ell} + \frac{C_2}{N_\ell} \right). \quad (4.46)$$

We use the Richardson extrapolation technique [86] to estimate the bias at level ℓ as

$$|\mathbb{E}[G - G_\ell]| \approx 2|\mathbb{E}[G_{\ell+1} - G_\ell]|. \quad (4.47)$$

We use a robust DLMC estimate of $\mathbb{E}[\Delta G_{\ell+1}] = \mathbb{E}[G_{\ell+1} - G_\ell]$ using Algorithm 6 with \hat{M}_1, \hat{M}_2 samples. Algorithm 6 in Appendix B.1 incorporates an antithetic sampler [82] to estimate the bias using sufficiently correlated samples of $G_{\ell+1}$ and G_ℓ . Given that a level $\ell = L$ satisfies (4.34), we find optimal parameters M_1 and M_2 that satisfy (4.38). This corresponds to solving the following optimization problem.

$$\begin{cases} \min_{\{M_1, M_2\}} & \mathcal{W} = M_1 N_L P_L^2 + M_1 M_2 N_L P_L \\ & C_v^2 \left(\frac{V_{1,L}}{M_1} + \frac{V_{2,L}}{M_1 M_2} \right) \approx (1 - \theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2, \end{cases} \quad (4.48)$$

where $V_{1,L} := V_1^{P_L|N_L}$ and $V_{2,L} := V_2^{P_L|N_L}$. By solving (4.48), we get

$$M_1 = \left(V_{1,L} + \sqrt{\frac{V_{1,L} V_{2,L}}{P_L}} \right) \frac{C_v^2}{(1 - \theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2}, \quad M_2 = \sqrt{\frac{V_{2,L} P_L}{V_{1,L}}}. \quad (4.49)$$

In principle, M_1 and M_2 obtained in (4.49) are real-valued. In practice, we use the next highest integer values for M_1 and M_2 . From (4.49), one observes that we must estimate $V_{1,L}$ and $V_{2,L}$ to obtain a DLMC estimator that satisfies (4.38). The proposed adaptive algorithm uses a heuristic DLMC estimator for these variances, as stated in Algorithm 7 in Appendix B.2, with \bar{M}_1 and \bar{M}_2 samples.

The proposed adaptive algorithm updates a heuristic estimate for $\mathbb{E}[g(\mathbf{X}(T))]$ (the quantity of interest) to check the estimator's relative error at each step. To initialize the adaptive algorithm, we produce an initial rough DLMC estimate by using IS with \tilde{M}_1 and \tilde{M}_2 samples. With this, we have all the components required to formulate the adaptive DLMC Algorithm 2.

4.5.4 Numerical experiments

We implement the DLMC Algorithm 2 and demonstrate its effectiveness over naïve MC while estimating rare event probabilities. The following numerical experiments are based on the 1D Kuramoto model (4.12) with $\sigma = 0.4$, $T = 1$, $(x_0)_p \sim \mathcal{N}(0, 0.2)$, and $\xi_p \sim \mathcal{U}(-0.2, 0.2)$ for all $p = 1, \dots, P$. We use DLMC with IS to estimate $\mathbb{P}[X(T) > K] =$

Algorithm 2: Adaptive DLMC algorithm with IS for MV-SDEs**Off-line:**

Generate one realisation of law $\{\mathbf{X}_p^{\bar{P}|\bar{N}}\}_{p=1}^{\bar{P}} \sim \mu^{\bar{P}|\bar{N}}$ with \bar{P} -particle system and \bar{N} time steps using (4.16) with sufficiently large \bar{P}, \bar{N} ;

Given $\mu^{\bar{P}|\bar{N}}$, solve PDE (4.24) to obtain IS control function ζ ;

Input: $P_0, N_0, \text{TOL}_r, \zeta$;

$\ell = 0$;

Estimate $\hat{\alpha} = \mathbb{E}[G_0]$ using $P_0, N_0, \tilde{M}_1, \tilde{M}_2, \zeta$ in **Algorithm 1**;

while Bias $> \theta \text{TOL}_r \hat{\alpha}$ **do**

$P_\ell = P_0 \times 2^\ell, \quad N_\ell = N_0 \times 2^\ell$;

Estimate $V_{1,\ell}$ and $V_{2,\ell}$ using $P_\ell, N_\ell, \tilde{M}_1, \tilde{M}_2, \zeta$ in **Algorithm 7**;

Compute optimal M_1, M_2 with estimated $V_{1,\ell}, V_{2,\ell}$ using (4.49);

Estimate Bias = $\frac{1}{\hat{\alpha}} |2\mathbb{E}[G_{\ell+1} - G_\ell]|$ with $P_\ell, N_\ell, \tilde{M}_1, \tilde{M}_2, \zeta$ in **Algorithm 6**;

Update $\hat{\alpha} = \mathbb{E}[G_\ell]$ using $P_\ell, N_\ell, M_1, M_2, \zeta$ in **Algorithm 1**;

$\ell \leftarrow \ell + 1$;

end

$\mathcal{A}_{\text{DLMC}} = \hat{\alpha}$.

$\mathbb{E} \left[\mathcal{I}_{\{X(T) > K\}} \right]$, where the process X satisfies the 1D Kuramoto MV-SDE (4.13) and $K \in \mathbb{R}$ lies deep in the right tail of the distribution of $X(T)$. The SOC-based IS scheme for the decoupled MV-SDE was obtained by numerically solving the optimal control PDE (4.24) with a FDM solver on a grid of 100 time discretization and 200 space discretization points, and using (4.25) to get the optimal IS control function ζ . Appendix C.1 presents details of the numerical solver.

We begin by verifying the assumptions on the bias (Assumptions 4.5.1, 4.5.2) and the variance (Assumptions 4.5.3, 4.5.4) of the DLMC estimator. However, we do not have the exact value of $\mathbb{E}[g(X(T))]$ to verify Assumption 4.5.1. Instead, we use Richardson extrapolation [86] with the Assumption 4.5.1 with P and $2P$ particles. This results in the following statement.


$$\left| \mathbb{E} \left[g(\bar{X}^{2P}(T)) \right] - \mathbb{E} \left[g(\bar{X}^P(T)) \right] \right| \leq \frac{C_1}{2P}. \quad (4.50)$$

Figure 4.1a plots $\left| \mathbb{E} \left[g(\bar{X}^{2P|N}(T)) \right] - \mathbb{E} \left[g(\bar{X}^{P|N}(T)) \right] \right|$ using a DLMC estimator with a large N for various values of P and verifies (4.50), implying that Assumption 4.5.1 holds. Figure 4.1a verifies the $\mathcal{O}(P^{-1})$ convergence order of the decoupling error in Assumption 4.5.1. Similarly, we use Richardson extrapolation with N and $2N$ time steps, resulting in the following.


$$\left| \mathbb{E} \left[g(\bar{X}^{P|2N}(T)) \right] - \mathbb{E} \left[g(\bar{X}^{P|N}(T)) \right] \right| \leq \frac{C_2}{2N}. \quad (4.51)$$

Figure 4.1b plots $\left| \mathbb{E} \left[g(\bar{X}^{P|2N}(T)) \right] - \mathbb{E} \left[g(\bar{X}^{P|N}(T)) \right] \right|$ using a DLMC estimator with a large P for various value of N and verifies (4.51), implying that Assumption 4.5.2 holds. Figure 4.1b verifies the $\mathcal{O}(N^{-1})$ convergence order of the EM discretization scheme in Assumption 4.5.1.

Figure 4.2 numerically verifies Assumptions 4.5.3 and 4.5.4. We estimate the variances and expected values in Assumptions 4.5.3 and 4.5.4 with sample means and variances computed in Algorithm 7. Figure 4.2a verifies the $\mathcal{O}(P^{-1})$ convergence of $V_1^{P|N}$ and



(a) Verifying Assumption 4.5.1. DLMC estimate of $\left| \mathbb{E} \left[g(\bar{X}^{2P|N}(T)) \right] - \mathbb{E} \left[g(\bar{X}^{P|N}(T)) \right] \right|$ by using Algorithm 1 with inputs $N = 2^7, M_1 = 10^2$, and $M_2 = 10^3$.




(b) Verifying Assumption 4.5.2. DLMC estimate of $\left| \mathbb{E} \left[g(\bar{X}^{P|2N}(T)) \right] - \mathbb{E} \left[g(\bar{X}^{P|N}(T)) \right] \right|$ by using Algorithm 1 with inputs $P = 80, M_1 = 10^2$, and $M_2 = 10^3$.


Figure 4.1: Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}$ for $K = 2$: Figures 4.1a and 4.1b numerically verify Assumptions 4.5.1 and 4.5.2 respectively, implying $\mathcal{O}(P^{-1} + N^{-1})$ error convergence rate for the relative bias of the DLMC estimator.

Figure 4.2b verifies the boundedness of $V_2^{P|N}$ with IS.

Figure 4.3 depicts the results of two numerical experiments conducted to verify the variance reduction of the proposed IS scheme for MV-SDEs. In the first experiment shown in Figure 4.3a, we verify variance reduction on the MC estimator of the inner expectation




(a) Verifying Assumption 4.5.3. DLMLC estimate of $V_1^{P|N}$ using Algorithm 7 with inputs $N = 2^7$, $M_1 = 10^2$, and $M_2 = 10^4$.




(b) Verifying Assumption 4.5.4. DLMLC estimate of $V_2^{P|N}$ by using Algorithm 7 with inputs $N = 2^8$, $M_1 = 10^2$, and $M_2 = 10^4$.

Figure 4.2: Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}$ for $K = 2$: Figures 4.2a and 4.2b numerically verify Assumptions 4.5.3 and 4.5.4 respectively. Both plots imply that the variance of the DLMLC estimator converges to $\text{Var}[g(\mathbf{X}(T))]$ asymptotically in the limit $P, N \rightarrow \infty$.

in (4.27) conditioned on an empirical law, $\mu^{P|N}$. To obtain Figure 4.3a, we acquire one realization of $\mu^{P|N}$ empirically by using the stochastic P -particle system (4.9) with $P = 200$ and $N = 32$. We use this law not only to numerically solve the optimal control PDE (4.24) and obtain the optimal IS control ζ , but also as an input to all realizations of the decoupled MV-SDE (4.18). We simulate sample paths of the decoupled MV-SDE using the EM scheme with $N = 32$ time steps. Figure 4.3a compares the squared coefficients of variation of the MC estimator of the inner conditional expectation in (4.27) with and without IS



(a) MC estimator squared coefficient of variation, conditioned on a fixed empirical $\mu^{P|N}$, with respect to number of sample paths M of the decoupled MV-SDE (4.18) conditioned on $\mu^{P|N}$.




(b) DLMC estimator squared coefficient of variation with respect to the number of sample paths in the inner loop M_2 .


Figure 4.3: Numerical demonstration of variance reduction due to proposed IS scheme for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2$: Figure 4.3a shows reduction in $V_2^{P|N}$ by a factor of 6000 with IS, leading to a reduction in the overall variance of the DLMC estimator by around 1000 times as seen in Figure 4.3b. We conclude that the proposed IS scheme reduces $V_2^{P|N}$ significantly to reduce the overall variance of the DLMC estimator.

with respect to the number of sample paths M of the decoupled MV-SDE. The plots verify that the squared coefficient of variation for the estimator reduces approximately 6000-fold with IS. In the second experiment shown in Figure 4.3b, we use $\bar{P} = 200$ particles and $\bar{N} = 100$ time steps to estimate an empirical $\mu^{P|N}$, and subsequently obtain the optimal IS control. Then, we set $P = 100, N = 32, M_1 = 10^3$, and varied M_2 as


inputs to Algorithm 1. Figure 4.3b shows the squared coefficient of variation of the DLMC estimator with respect to M_2 . We observe that the DLMC estimator with IS displays a remarkably reduced coefficient of variation (approximately 1000-fold).



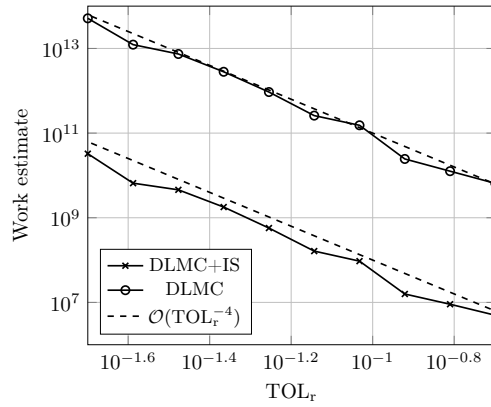
(a) DLMC estimator $\mathcal{A}_{\text{DLMC}}$ values with respect to prescribed relative tolerance TOL_r .



(b) Relative error of estimator $\mathcal{A}_{\text{DLMC}}$ with respect to prescribed relative tolerance TOL_r .



(c) Computational runtime of Algorithm 2 with respect to prescribed relative tolerance TOL_r .



(d) Computational work of estimator $\mathcal{A}_{\text{DLMC}}$ with respect to prescribed relative tolerance TOL_r .

Figure 4.4: Numerical results of applying Algorithm 2 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}$ for $K = 2$: Figure 4.4b shows that the proposed DLMC estimator with IS achieves a relative error below the prescribed TOL_r in the sense of (4.14) to estimate a rare event probability of around 2.5×10^{-4} (Figure 4.4a). Figures 4.4c and 4.4d show that $\mathcal{A}_{\text{DLMC}}$ achieves the relative accuracy requirement with a computational work of $\mathcal{O}(\text{TOL}_r^{-4})$. Additionally, Figure 4.4d shows that the proposed IS significantly reduces computational work of the DLMC estimator by a factor of 10^3 .

Finally, we apply Algorithm 2 with $\bar{P} = 1000$ particles and $\bar{N}_1 = 100$ time steps to estimate one realization of the empirical law $\mu^{\bar{P}|\bar{N}}$ using the stochastic particle system (4.16) and obtaining the optimal IS control ζ . We use $P_0 = 5, N_0 = 4$, and various relative error tolerances TOL_r as inputs to Algorithm 2. The following heuristics were employed to ensure the robustness of the proposed algorithm.

- We use $\tilde{M}_1 = 10^3$ and $\tilde{M}_2 = 10^2$ to obtain an initial rough estimate for the required quantity to aid in quantifying the required relative tolerance.

- Algorithm 7 is employed to estimate $V_{1,\ell}$ and $V_{2,\ell}$ with $\bar{M}_1 = 50$ and $\bar{M}_2 = 10^3$ for the first three levels, i.e., $\ell = 1, 2, 3$. In addition, Assumptions 4.5.3 and 4.5.4 are employed for the subsequent levels to extrapolate $V_{1,\ell}$ and $V_{2,\ell}$.
- Algorithm 6 is used to estimate bias, with $\hat{M}_1 = \max(M_1, 100)$ and $\hat{M}_2 = \max(M_2, 50)$, where M_1, M_2 are the computed optimal sample sizes. The estimated bias is then compared with the extrapolated bias from the last two levels, and the maximum of the three values is selected, ensuring bias estimate robustness for $\ell > 3$.

Figure 4.4a illustrates that $K = 2$ corresponds to a probability of approximately 2.53×10^{-4} , a sufficiently rare event probability. Figure 4.4b shows the relative error of the DLMC estimator for different runs of Algorithm 2 over various prescribed relative error tolerances. We use a reference DLMC approximation computed with $\text{TOL}_r = 1.5\%$. Figure 4.4c shows that the computational runtime closely follows the predicted theoretical rate of $\mathcal{O}(\text{TOL}_r^{-4})$ for small relative tolerances. Additionally, we compare the estimated computational work, given by (4.43), for the DLMC estimator with and without IS. Because running a crude DLMC is computationally infeasible for rare events, we use a heuristic estimate of the computational cost of crude DLMC without actually running the algorithm. For this, we approximate the sample variance of the DLMC estimator without IS as follows.

$$\begin{aligned} \text{Var} \left[\mathcal{I}_{\{X(T) > K\}} \right] &= \mathbb{E} \left[\left(\mathcal{I}_{\{X(T) > K\}} \right)^2 \right] - \mathbb{E} \left[\mathcal{I}_{\{X(T) > K\}} \right]^2 \\ &= \mathbb{E} \left[\mathcal{I}_{\{X(T) > K\}} \right] - \mathbb{E} \left[\mathcal{I}_{\{X(T) > K\}} \right]^2 \\ &\approx \mathbb{E} \left[\mathcal{I}_{\{X(T) > K\}} \right]. \end{aligned}$$

This variance is then estimated by the DLMC estimator $\mathcal{A}_{\text{DLMC}}$ with IS. This sample variance is used to estimate the number of samples M_1, M_2 required to estimate the rare event probability without IS. Figure 4.4d provides numerical evidence that the IS estimator reduced the computational cost to achieve a prescribed relative error tolerance, by multiple orders (three orders of magnitude in this case). This implies that the DLMC estimator with the proposed IS scheme dramatically reduces the computational cost associated with estimating rare event probabilities.

Table 4.1 shows the average number of samples required to reach a given relative tolerance with and without IS for different values of K . For the DLMC estimator without IS, the number of samples required to satisfy a given relative tolerance increases as the probability reduces (see Section 3.1). The proposed SOC-based IS scheme reduces both M_1 and M_2 . In fact, the required number of samples for IS is considerably lower than that without IS, and it remains of the same order regardless of the event rarity. Thus, the proposed IS measure change results in a bounded estimator coefficient of variation, regardless of the value of K . In conclusion, the proposed DLMC estimator with the optimal IS measure change drastically reduces the computational work required to estimate rare event probabilities, with a computational complexity rate of $\mathcal{O}(\text{TOL}_r^{-4})$. The presence of the discretization parameters P and N hints towards the use of MLMC coupled with IS to further reduce the computational complexity of the DLMC estimator.

K	$\mathcal{A}_{\text{DLMC}}$			DLMC+IS		DLMC	
		TOL_r	ℓ	M_1	M_2	M_1	M_2
1	5.6×10^{-2}	20%	3	59	12	157	65
		10%	4	111	27	284	131
		5%	5	167	58	456	297
1.5	6.8×10^{-3}	20%	4	41	24	230	325
		10%	5	64	48	386	700
		5%	6	105	102	662	1491
2	2.53×10^{-4}	20%	5	17	48	477	3629
		10%	6	33	90	913	6929
		5%	7	57	186	1711	14864

Table 4.1: Comparing the average number of samples required for the naïve DLMC and proposed DLMC with IS methods to estimate $\mathbb{P}[X(T) > K]$ up to prescribed relative tolerance TOL_r for the 1D Kuramoto model (4.12) for different values of K . The proposed IS scheme drastically reduces the number of samples M_1 and M_2 required to achieve a relative accuracy TOL_r and also ensures the required number of samples does not increase as the event rarity increases.

4.6 Multilevel double loop Monte Carlo estimator

MLMC was introduced as an improvement to MC for SDEs in [23]. MLMC is based on generalized control variates and improves efficiency when an approximation of the solution is computed based on a discretization parameter [87]. Most MLMC simulations are performed at cheaper, coarse levels, with relatively few simulations applied at the costlier, fine levels. The authors in [88] introduce a MLMC estimator with EM-discretization of complexity $\mathcal{O}\left(\text{TOL}^{-2} \left(\log \text{TOL}^{-1}\right)^2\right)$ to estimate functionals associated with a particle system arising from the stochastic representation of the Landau–Fokker–Planck PDE, with the assumption that the mean-field law is constant in the considered time scale. Subsequently, the work [82] developed an MLMC estimator for MV-SDEs with a partitioning sampler with a computational complexity of $\mathcal{O}\left(\text{TOL}^{-3}\right)$ for smooth observables and bounded, Lipschitz drift and diffusion functions. The authors in [89] obtain the same complexity for MLMC in MV-SDEs with any general dependence on the mean-field law with an antithetic estimator. However, even MLMC methods become extremely expensive in the context of rare events due to the same exploding coefficient of variation problem seen in Section 3.1. Hence, we combine IS with MLMC to reduce the estimator’s coefficient of variation in the rare event regime by extending the DLMC estimator in Section 4.5 to the multilevel setting. Combining IS with MLMC has been previously explored in other contexts. The works [25, 26] employ suboptimal level-wise constant controls for standard SDEs, whereas the work [90] uses MLMC with IS to estimate functionals of ergodic SDEs that do not satisfy the contractivity condition. The authors in [91] also construct a level-dependent IS scheme to avoid catastrophic coupling of level differences while implementing MLMC for stochastic reaction networks. In this Section, we combine IS with MLMC to estimate rare events associated with MV-SDEs using the novel MLDLMC estimator.

Two discretization parameters (P and N) are required to generate sample paths of the decoupled MV-SDE (4.18). We introduce the level ℓ that couples both discretization

parameters. The geometric sequence of levels is defined for $\ell = 0, 1, \dots, L$ as follows:

$$\begin{aligned} P_\ell &= P_0 \times 2^\ell, \\ N_\ell &= N_0 \times 2^\ell. \end{aligned} \quad (4.52)$$

We propose the following method to couple IS with a MLDLMC estimator. We solve the same optimal control PDE described in Section 4.4 off-line using one realization of the stochastic particle system (4.9) with a large number of particles \bar{P} and time steps \bar{N} , and obtain the optimal IS control function ζ . We apply the same policy ζ to generate sample paths of the controlled decoupled MV-SDE (4.20) at all discretization levels ℓ . We recall the notation $G = g(\mathbf{X}(T))$, and $G_\ell = g(\bar{\mathbf{X}}_\zeta^{P_\ell|N_\ell}(T))\mathbb{L}^{P_\ell|N_\ell}$. \mathbf{X} denotes the McKean–Vlasov process following (4.1) and $\bar{\mathbf{X}}_\zeta^{P_\ell|N_\ell}$ denotes the approximate, EM-discretized decoupled MV-SDE process (4.18) conditioned on the empirical law $\mu^{P_\ell|N_\ell}$ (4.10). The likelihood factor $\mathbb{L}^{P_\ell|N_\ell}$ at level ℓ is computed using (4.32). MLMC uses a telescoping sum [87] for a given level $L \in \mathbb{N}$:

$$\mathbb{E}[G_L] = \sum_{\ell=0}^L \mathbb{E}[G_\ell - G_{\ell-1}], \quad G_{-1} = 0. \quad (4.53)$$

We approximate each expectation in (4.53) independently using a DLMC method, resulting in the MLDLMC estimator.

$$\mathcal{A}_{\text{MLDLMC}} = \sum_{\ell=0}^L \frac{1}{M_{1,\ell}} \sum_{i=1}^{M_{1,\ell}} \frac{1}{M_{2,\ell}} \sum_{j=1}^{M_{2,\ell}} (G_\ell - \mathcal{G}_{\ell-1})(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}), \quad (4.54)$$

where $\mathcal{G}_{\ell-1}$ is a random variable correlated to G_ℓ such that $\mathcal{G}_{-1} = 0$ and $\mathbb{E}[\mathcal{G}_{\ell-1}] = \mathbb{E}[G_{\ell-1}]$, ensuring $\mathbb{E}[\mathcal{A}_{\text{MLDLMC}}] = \mathbb{E}[G_L]$. In addition, $\omega_{1:P_\ell}^{(\ell,i)}$ refers to the i^{th} realization of the P_ℓ underlying sets of random variables used to estimate $\mu^{P_\ell|N_\ell}$ at level ℓ , and $\bar{\omega}^{(\ell,j)}$ denotes the j^{th} realization of the random variables in (4.18) for the given $\mu^{P_\ell|N_\ell}$ realization at level ℓ . Samples of $G_\ell(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)})$ and $\mathcal{G}_{\ell-1}(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)})$ must be sufficiently correlated to ensure that the MLDLMC estimator has better computational complexity than the DLMC estimator. We explore two correlated samplers motivated by the MLMC estimators in [82].

1. *Naïve Sampler.* We use the first $P_{\ell-1}$ random variables out of each set of P_ℓ random variables to obtain empirical law $\mu_{\ell-1} := \mu^{P_{\ell-1}|N_{\ell-1}}$ to generate a sample of $\mathcal{G}_{\ell-1}$ using (4.9). Given $\mu_{\ell-1}$, we solve (4.18) at level $\ell - 1$ using the same $\bar{\omega}$ as for level ℓ and compute the quantity of interest:

$$\begin{aligned} \mathcal{G}_{\ell-1}(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) &= \bar{\mathcal{G}}_{\ell-1}(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) \\ &:= G_{\ell-1}(\omega_{1:P_{\ell-1}}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}). \end{aligned} \quad (4.55)$$

2. *Antithetic Sampler.* We split the P_ℓ sets of random variables into 2 iid groups of size $P_{\ell-1}$ each. Then, to generate a sample of $\mathcal{G}_{\ell-1}$, we use each group to independently simulate (4.9) and obtain empirical laws $\mu_{\ell-1}^{(1)}$ and $\mu_{\ell-1}^{(2)}$. Given the two empirical laws, we solve (4.18) at level $\ell - 1$ independently for each law using the same $\bar{\omega}$ as for level ℓ . The quantity of interest is computed for each group and averaged over

the 2 groups:

$$\begin{aligned} \mathcal{G}_{\ell-1}(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) &= \hat{\mathcal{G}}_{\ell-1}(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) \\ &:= \frac{1}{2} \left(G_{\ell-1}(\omega_{1:P_{\ell-1}}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) + G_{\ell-1}(\omega_{(P_{\ell-1}+1):P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)}) \right). \end{aligned} \quad (4.56)$$

4.6.1 Error analysis

We bound the relative error of $\mathcal{A}_{\text{MLDLMC}}$ as follows:

$$\frac{|\mathbb{E}[G] - \mathcal{A}_{\text{MLDLMC}}|}{|\mathbb{E}[G]|} \leq \underbrace{\frac{|\mathbb{E}[G] - \mathbb{E}[G_L]|}{|\mathbb{E}[G]|}}_{=\epsilon_b, \text{ Relative bias}} + \underbrace{\frac{|\mathbb{E}[G_L] - \mathcal{A}_{\text{MLDLMC}}|}{|\mathbb{E}[G]|}}_{=\epsilon_s, \text{ Relative statistical error}}, \quad (4.57)$$

Although $\mathcal{A}_{\text{MLDLMC}}$ should satisfy a prescribed relative error tolerance TOL_r in the sense of (4.14), we impose more restrictive conditions.

$$\bar{\epsilon}_b \leq \theta \text{TOL}_r, \quad (4.58)$$

$$\mathbb{P}[\epsilon_s \geq (1 - \theta)\text{TOL}_r] \leq \nu, \quad (4.59)$$

for a given tolerance splitting parameter $\theta \in (0, 1)$ and confidence level determined by $0 < \nu \ll 1$.

Assumption 4.6.1 (MLDLMC bias). *There exists a constant $\tilde{\alpha} > 0$ such that*

$$|\mathbb{E}[G] - \mathbb{E}[G_\ell]| \lesssim 2^{-\tilde{\alpha}\ell},$$

where constant $\tilde{\alpha}$ denotes the bias convergence rate with respect to level ℓ .

Assumption 4.6.1 is a direct implication of (4.37) arising from Assumptions 4.5.1 and 4.5.2. The finest level $L \in \mathbb{N}$ is chosen to satisfy the bias constraint (4.58). The statistical error constraint (4.59) can be rewritten as follows:

$$\mathbb{P}[\epsilon_s \geq (1 - \theta)\text{TOL}_r |\mathbb{E}[G|]] = \mathbb{P}\left[\frac{\epsilon_s}{\sqrt{\text{Var}[\mathcal{A}_{\text{MLDLMC}}]}} \geq \frac{(1 - \theta)\text{TOL}_r |\mathbb{E}[G|]}{\sqrt{\text{Var}[\mathcal{A}_{\text{MLDLMC}}]}}\right] \leq \nu. \quad (4.60)$$

By assuming the normality (at least asymptotically) of the estimator $\mathcal{A}_{\text{MLDLMC}}$, we obtain the following condition for the estimator variance:

$$\frac{\text{Var}[\mathcal{A}_{\text{MLDLMC}}]}{|\mathbb{E}[G]|^2} \leq \left(\frac{(1 - \theta)\text{TOL}_r}{C_\nu}\right)^2, \quad (4.61)$$

where C_ν is the $(1 - \frac{\nu}{2})$ -quantile of the standard normal distribution. The asymptotic normality of MLMC estimators can be demonstrated using the Lindeberg–Feller CLT [92]. The proposed estimator variance can be expressed as

$$\text{Var}[\mathcal{A}_{\text{MLDLMC}}] = \sum_{\ell=0}^L \frac{1}{M_{1,\ell}} \text{Var}\left[\frac{1}{M_{2,\ell}} \sum_{j=1}^{M_{2,\ell}} (\Delta G_\ell)^{(1,j)}\right], \quad (4.62)$$

where $(\Delta G_\ell)^{(i,j)} := (G_\ell - \mathcal{G}_{\ell-1})(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)})$. Using the law of total variance,

$$\begin{aligned} \text{Var}[\mathcal{A}_{\text{MLDLMC}}] &= \sum_{\ell=0}^L \frac{1}{M_{1,\ell}} \left(\underbrace{\text{Var}[\mathbb{E}[\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}]]}_{=V_{1,\ell}} + \frac{1}{M_{2,\ell}} \underbrace{\mathbb{E}[\text{Var}[\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}]]}_{=V_{2,\ell}} \right) \\ &= \sum_{\ell=0}^L \left(\frac{V_{1,\ell}}{M_{1,\ell}} + \frac{V_{2,\ell}}{M_{1,\ell}M_{2,\ell}} \right), \end{aligned} \quad (4.63)$$

where laws $\{\mu_\ell, \mu_{\ell-1}\}$ are coupled by the same sets of random variables $\omega_{1:P_\ell}^{(\ell,\cdot)}$ as described in the naïve (4.55) and antithetic (4.56) samplers.

Assumption 4.6.2 (MLDLMC variance). *There exist constants $\tilde{w} > 0$ and $\tilde{s} > 0$ such that,*

$$V_{1,\ell} = \text{Var}[\mathbb{E}[\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}]] \lesssim 2^{-\tilde{w}\ell}, \quad (4.64)$$

$$V_{2,\ell} = \mathbb{E}[\text{Var}[\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}]] \lesssim 2^{-\tilde{s}\ell}. \quad (4.65)$$

where constants \tilde{w} and \tilde{s} are the convergence rates for $V_{1,\ell}$ and $V_{2,\ell}$, respectively, with respect to ℓ .

4.6.2 Computational complexity

For each level ℓ , we must estimate $\mathbb{E}[\Delta G_\ell] = \mathbb{E}[G_\ell - G_{\ell-1}]$. In this thesis, we compute the empirical means in the drift and diffusion functions in (4.9) and (4.18) using a naïve method with computational cost $\mathcal{O}(P)$. The computational cost to generate one realization of the random variable G_ℓ comprises two parts (see Section 4.5): i) generation of a realization of the empirical law μ_ℓ using the EM scheme with N_ℓ time steps ($\mathcal{O}(N_\ell P_\ell^2)$), and ii) the generation of one sample path of the decoupled McKean–Vlasov process (4.20) using the EM scheme ($\mathcal{O}(P_\ell N_\ell)$). Hence, the net computational cost of the proposed MLDLMC estimator (4.54) is

$$\mathcal{W}[\mathcal{A}_{\text{MLDLMC}}] \lesssim \sum_{\ell=0}^L (M_{1,\ell} P_\ell^2 N_\ell + M_{1,\ell} M_{2,\ell} P_\ell N_\ell). \quad (4.66)$$

Let some level L satisfy the bias constraint (4.58). We aim to compute the optimal parameters $\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L$ which minimise the work (4.66) while satisfying (4.61), which can be posed as the following optimization problem:

$$\begin{cases} \arg \min_{\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L} \sum_{\ell=0}^L (M_{1,\ell} P_\ell^2 N_\ell + M_{1,\ell} M_{2,\ell} P_\ell N_\ell) \\ \text{s.t. } C_v^2 \left(\sum_{\ell=0}^L \left(\frac{V_{1,\ell}}{M_{1,\ell}} + \frac{V_{2,\ell}}{M_{1,\ell} M_{2,\ell}} \right) \right) \approx (1 - \theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2 \end{cases}. \quad (4.67)$$

The solution of (4.67) in \mathbb{R}^+ is found using the Lagrangian multiplier method $\forall \ell = 0, \dots, L$. These are,

$$\begin{aligned}
\mathcal{M}_{1,\ell} &= \frac{C_v^2}{(1-\theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2} \frac{\sqrt{V_{1,\ell}}}{\sqrt{P_\ell^2 N_\ell}} \left(\sum_{j=0}^L \sqrt{P_j N_j} (\sqrt{V_{1,j} P_j} + \sqrt{V_{2,j}}) \right), \\
\tilde{\mathcal{M}}_\ell &= \mathcal{M}_{1,\ell} \mathcal{M}_{2,\ell} \\
&= \frac{C_v^2}{(1-\theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2} \frac{\sqrt{V_{2,\ell}}}{\sqrt{P_\ell N_\ell}} \left(\sum_{j=0}^L \sqrt{P_j N_j} (\sqrt{V_{1,j} P_j} + \sqrt{V_{2,j}}) \right).
\end{aligned} \tag{4.68}$$

In practice, we only use natural numbers for $\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L$. Hence, we use the following quasi-optimal solution to (4.67):

$$M_{1,\ell} = \lceil \mathcal{M}_{1,\ell} \rceil, \quad M_{2,\ell} = \left\lceil \frac{\tilde{\mathcal{M}}_\ell}{\lceil \mathcal{M}_{1,\ell} \rceil} \right\rceil. \tag{4.69}$$

We obtain the optimal computational complexity of the proposed **MLDLMC** estimator using (4.69).

Theorem 4.6.1 (Optimal **MLDLMC** complexity). *Let $G_\ell^{(i,j)} := G_\ell(\omega_{1:P_\ell}^{(\ell,i)} \times \bar{\omega}^{(\ell,j)})$ be a sample of G_ℓ for every $\ell \in \mathbb{N}$. Consider the **MLDLMC** estimator (4.54) with $G_{-1}^{(i,j)} = 0$. Let Assumptions 4.6.1 and 4.6.2 hold. Let the constants $\tilde{\alpha}$, $\tilde{\omega}$, $\tilde{\mathfrak{s}}$ from Assumptions 4.6.1 and 4.6.2 be such that $\tilde{\alpha} \geq \frac{1}{2} \min(\tilde{\omega}, 1 + \tilde{\mathfrak{s}}, 3)$. Then, there exists an optimal L and sequences $\{M_{1,\ell}\}_{\ell=0}^L$ and $\{M_{2,\ell}\}_{\ell=0}^L$ such that*

$$\mathbb{P} \left[\frac{|\mathcal{A}_{\text{MLDLMC}} - \mathbb{E}[G]|}{|\mathbb{E}[G]|} \geq \text{TOL}_r \right] \leq \nu, \tag{4.70}$$

and

$$\mathcal{W}[\mathcal{A}_{\text{MLDLMC}}] \lesssim \text{TOL}_r^{-2 - \max(0, \frac{3-\tilde{\omega}}{\tilde{\alpha}}, \frac{2-\tilde{\mathfrak{s}}}{\tilde{\alpha}})} (\log \text{TOL}_r^{-1})^{2\mathcal{J}}, \tag{4.71}$$

where

$$\mathcal{J} := \begin{cases} 1, & \text{if } \min(\tilde{\omega}, 1 + \tilde{\mathfrak{s}}) = 3, \\ 0, & \text{otherwise.} \end{cases}$$

Proof. See Appendix A.4. □

The complexity rate in (4.71) is independent of θ and the dimension, d , of the **MV-SDE**. This is because the optimal control **PDE** (4.24) is solved off-line only once to get the **IS** control function ζ . These parameters, however, affect the associated constant in (4.71). Further optimization of the **MLMC** algorithm with respect to these parameters has been investigated in [93, 94].

Remark 4.6.1. *In many applications, the variance and the second moments of the level differences in the **MLDLMC** estimator are of the same order. In this case, we can demonstrate that $V_{1,\ell} \leq \mathbb{E}[\mathbb{E}[\Delta G_\ell^2 | \{\mu_\ell, \mu_{\ell-1}\}]] \approx V_{2,\ell}$, implying $\tilde{\omega} \geq \tilde{\mathfrak{s}}$. This would simplify the analysis presented in Appendix A.4.*

Algorithm 8 in Appendix B.3 implements the proposed **IS** scheme in the level-difference estimator for the proposed **MLDLMC** method and can be easily modified for any other correlated sampler, such as the naïve sampler. Next, we build an adaptive **MLDLMC** algorithm that sequentially chooses parameters $L, \{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L$ satisfying constraints (4.58) and (4.59). The bias and variances $V_{1,\ell}$ and $V_{2,\ell}$ corresponding to level ℓ must be estimated cheaply and robustly to develop such an algorithm.

4.6.3 Adaptive algorithm

Estimating bias at level ℓ

The bias for level ℓ can be approximated using Richardson extrapolation [86]:

$$|\mathbb{E}[G] - \mathbb{E}[G_\ell]| \approx (1 - 2^{-\tilde{\alpha}})^{-1} |\mathbb{E}[G_{\ell+1} - G_\ell]|. \quad (4.72)$$

Then, we use Algorithm 8 with at least \underline{M}_1 and \underline{M}_2 samples to obtain a DLMC estimate of the bias. To ensure robust bias estimates at all levels, we actually use \hat{M}_1 and \hat{M}_2 samples in Algorithm 8 defined as follows:

$$\begin{aligned} \hat{M}_1 &= \max(M_{1,\ell}, \underline{M}_1), \\ \hat{M}_2 &= \max(M_{2,\ell}, \underline{M}_2). \end{aligned} \quad (4.73)$$

To ensure reliable bias estimates at higher levels ($\ell > 3$), we compare the DLMC bias estimator with the extrapolated bias from two previous levels using Assumption 4.6.1:

$$|\mathbb{E}[G] - \mathbb{E}[G_\ell]| \approx \max \left(\frac{|\mathbb{E}[\Delta G_{\ell+1}]|}{1 - 2^{-\tilde{\alpha}}}, \frac{|\mathbb{E}[G] - \mathbb{E}[G_{\ell-1}]|}{2^{\tilde{\alpha}}}, \frac{|\mathbb{E}[G] - \mathbb{E}[G_{\ell-2}]|}{2^{2\tilde{\alpha}}} \right), \quad \ell > 3. \quad (4.74)$$

Estimating $V_{1,\ell}, V_{2,\ell}$ at level ℓ

To compute the optimal number of samples required to satisfy the statistical error constraint (4.61) using (4.69), we require cheap and robust empirical estimates of the variance terms $V_{1,\ell}$ and $V_{2,\ell}$ for each level ℓ . For this, we apply the DLMC algorithm (Algorithm 9 in Appendix B.4) with appropriately chosen values of \tilde{M}_1 and \tilde{M}_2 . Algorithm 9 could become computationally expensive at higher levels. We exploit Assumption 4.6.2 to avoid this overload and extrapolate variances for higher levels. For levels $\ell > 3$,

$$\begin{aligned} V_{1,\ell} &\approx \max \left(\frac{V_{1,\ell-1}}{2^{\tilde{w}}}, \frac{V_{1,\ell-2}}{2^{2\tilde{w}}} \right), \\ V_{2,\ell} &\approx \max \left(\frac{V_{2,\ell-1}}{2^{\tilde{s}}}, \frac{V_{2,\ell-2}}{2^{2\tilde{s}}} \right). \end{aligned} \quad (4.75)$$

Relative error control

To control the relative bias and statistical errors, we require a heuristic estimate of the quantity of interest $\mathbb{E}[G]$ itself. This estimate is continuously updated at each level ℓ . At level $\ell = 0$, we use the DLMC algorithm (Algorithm 8) with \tilde{M}_1 and \tilde{M}_2 samples to obtain an initial robust but cheap estimate of $\mathbb{E}[G]$. For subsequent levels, we apply the MLDLMC estimator (4.54) with optimal values for $\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L$ to update the estimate. Putting all this together, we propose the adaptive MLDLMC algorithm (Algorithm 3) for rare event observables associated with MV-SDEs. The IS control function ζ in Algorithm 3 is obtained off-line by generating one realization of the empirical law $\mu^{\tilde{P}|\tilde{N}}$ with large \tilde{P} and \tilde{N} and then numerically solving the control PDE (4.24) given $\mu^{\tilde{P}|\tilde{N}}$.

Algorithm 3: Adaptive MLDLMC algorithm with IS**Off-line:**

Generate one realisation of law $\{\{\mathbf{X}_p^{\bar{P}|\bar{N}}\}_{p=1}^{\bar{P}} \sim \mu^{\bar{P}|\bar{N}}\}$ with \bar{P} -particle system and \bar{N} time steps using (4.16) with sufficiently large \bar{P}, \bar{N} ;

Given $\mu^{\bar{P}|\bar{N}}$, solve PDE (4.24) to obtain IS control function ζ ;

Input: $P_0, N_0, \text{TOL}_r, \zeta, \{\bar{M}_1, \bar{M}_2\}, \{\tilde{M}_1, \tilde{M}_2\}, \{\underline{M}_1, \underline{M}_2\}$;

Estimate $\bar{G} = \mathbb{E}[G_0]$ with $P_0, N_0, \bar{M}_1, \bar{M}_2, \zeta$ using **Algorithm 8**;

Estimate and store $V_{1,0}, V_{2,0}$ with $P_0, N_0, \bar{M}_1, \bar{M}_2, \zeta$ using **Algorithm 9**;

Set $L = 1$;

while $\epsilon_b > \theta \text{TOL}_r \bar{G}$ **do**

$P_L = P_0 \times 2^L, \quad N_L = N_0 \times 2^L$;

Estimate and store $V_{1,L}, V_{2,L}$ with $P_L, N_L, \tilde{M}_1, \tilde{M}_2, \zeta$ using **Algorithm 9**;

Compute optimal $\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L$ using (4.69);

Estimate ϵ_b using (4.72) with $P_{L+1}, N_{L+1}, \tilde{M}_1, \tilde{M}_2, \zeta$ using (4.73) and

Algorithm 8;

Reevaluate $\bar{G} = \mathbb{E}[G_0] + \sum_{\ell=1}^L \mathbb{E}[\Delta G_\ell]$ with $\{M_{1,\ell}, M_{2,\ell}\}_{\ell=0}^L, \zeta$ using

Algorithm 8 for each ℓ ;

$L \leftarrow L + 1$;

end

$\mathcal{A}_{\text{MLDLMC}} = \bar{G}$.

4.6.4 Numerical experiments

We provide numerical evidence for the assumptions and computational complexity rates derived in Section 4.6. The results outlined below focus on the Kuramoto model (4.12) with the following settings: $\sigma = 0.4, T = 1, (x_0)_p \sim \mathcal{N}(0, 0.2)$, and $\zeta_p \sim \mathfrak{U}(-0.2, 0.2)$ for all $p = 1, \dots, P$. We set the parameters as follows: $\theta = 0.5$, and $\nu = 0.05$. Moreover, we set the hierarchies for the MLDLMC estimator as


$$P_\ell = 5 \times 2^\ell, \quad N_\ell = 4 \times 2^\ell. \quad (4.76)$$

We use MLDLMC with IS to estimate $\mathbb{P}[X(T) > K] \approx \mathbb{E}[\tilde{\mathcal{I}}_{\{X(T) > K\}}]$, where the process X satisfies the 1D Kuramoto MV-SDE (4.13) and $K \in \mathbb{R}$ lies deep in the right tail of the distribution of $X(T)$. The function $\tilde{\mathcal{I}}$ is a Lipschitz piece-wise linear approximation of the indicator function \mathcal{I} defined as follows.


$$\tilde{\mathcal{I}}_{\{x > K\}}(x) = \begin{cases} 0 & , \quad x < K - \beta, \\ \frac{1}{2} + \frac{x-K}{2\beta} & , \quad K - \beta < x < K + \beta, \\ 1 & , \quad x > K + \beta. \end{cases} \quad (4.77)$$

In the limit $\beta \rightarrow 0$, the function $\tilde{\mathcal{I}}_{\{x > K\}}(x)$ converges to the indicator function $\mathcal{I}_{\{x > K\}}(x)$. This approximation helps avoid poor variance convergence rates of level differences with non-Lipschitz observables in MLMC [95]. We fix $\beta = 0.5$ heuristically for our numerical experiments. The SOC-based IS scheme for the decoupled MV-SDE was obtained by numerically solving the optimal control PDE (4.24) with a FDM solver (see Appendix C.1) and using (4.25) to get the optimal IS control function ζ .

Figure 4.5 depicts the results of two numerical experiments conducted to verify the variance reduction of the proposed IS scheme in the DLMC estimator of the level differences. In Figure 4.5a, we verify variance reduction on the MC estimator of $\mathbb{E}[\Delta G_\ell | \mu^{\bar{P}|\bar{N}}]$. For



(a) Experiment 1: Squared coefficient of variation of the MC estimator of $\mathbb{E}[\Delta G_3 | \mu^{\bar{P}|\bar{N}}]$ with and without IS with respect to the number of sample paths M .



(b) Experiment 2: Squared coefficient of variation of the DLMC estimator of $\mathbb{E}[\Delta G_3]$ with and without IS with respect to the number of sample paths in the inner loop M_2 .

Figure 4.5: Numerical demonstration of variance reduction in the DLMC estimator for level differences due to proposed IS scheme for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}$ for $K = 2.5$: Figure 4.5a shows reduction in $V_{2,\ell}$ by a factor of 100 with IS, leading to a reduction in the overall variance of the DLMC estimator of level differences by around 10 times as seen in Figure 4.5b. We conclude that the proposed IS scheme significantly reduces $\text{Var}[\Delta G_\ell]$, even though the proposed IS scheme minimises $\text{Var}[G_\ell]$.

this, we run the DLMC Algorithm 8 with $M_1 = 1$ and $\ell = 3$ for different values of M_2 . To generate Figure 4.5a, we acquire one realization of the empirical law $\mu^{\bar{P}|\bar{N}}$ by employing the stochastic P -particle system (4.9) with $\bar{P} = 200$ particles and $\bar{N} = 100$ time steps. We use this law not only to numerically solve the optimal control PDE (4.24)


and obtain the IS control ζ , but also to generate all realizations of the random variable ΔG_ℓ . Figure 4.5a compares the squared coefficient of variation of the MC estimator with and without IS versus the number of samples M of ΔG_ℓ . The results verify that the proposed IS scheme reduces the squared coefficient of variation approximately 100-fold. In Figure 4.5b, we verify the variance reduction in the DLMC estimator of $\mathbb{E}[\Delta G_\ell]$ with the proposed IS scheme using the DLMC Algorithm 8 with $\ell = 3$ and $M_1 = 10^3$ for different values of M_2 . To generate Figure 4.5b, we obtain the IS control ζ by solving the optimal control PDE (4.24) off-line using an independent realization of the empirical law $\mu^{\bar{P}|\bar{N}}$ with $\bar{P} = 1000$ particles and $\bar{N} = 100$ time steps. Figure 4.5b depicts a significantly reduced squared coefficient of variation (approximately one order of magnitude) with IS. The variance of the DLMC estimator of $\mathbb{E}[\Delta G_\ell]$ is given by $\frac{V_{1,\ell}}{M_{1,\ell}} + \frac{V_{2,\ell}}{M_{1,\ell}M_{2,\ell}}$ (see Equation (4.62)). We notice convergence of the squared coefficient of variation as the second term vanishes with a large value of M_2 .

Next, we numerically verify Assumptions 4.6.1 and 4.6.2 to determine constants $\tilde{\alpha}$, \tilde{w} , and \tilde{s} . Figure 4.6 verifies Assumptions 4.6.1 and 4.6.2 with $\tilde{\alpha} = 1$ and $\{\tilde{w}, \tilde{s}\} = \{2, 1\}$ for the antithetic sampler (4.56). These rates imply a computational complexity of $\mathcal{O}(\text{TOL}_r^{-3})$ for the proposed MLDLMC estimator with IS according to Theorem 4.6.1. The computational complexity rate of the MLDLMC estimator with the antithetic sampler is one order better than that with the naïve sampler because the variance convergence rates $\{\tilde{w}, \tilde{s}\} = \{1, 1\}$ for the sampler are worse than those of the antithetic sampler [2], implying a MLDLMC computational complexity rate of $\mathcal{O}(\text{TOL}_r^{-4})$, the same as the DLMC estimator with IS.


Lastly, we implement the proposed adaptive MLDLMC Algorithm 3 with inputs $\{\tilde{M}_1, \tilde{M}_2\} = \{1000, 100\}$, $\{\tilde{M}_1, \tilde{M}_2\} = \{25, 100\}$, and $\{\underline{M}_1, \underline{M}_2\} = \{100, 50\}$. We use $\bar{P} = 1000$ particles and $\bar{N} = 100$ time steps to independently obtain one empirical realization of $\mu^{\bar{P}|\bar{N}}$ to compute the IS control function ζ . Figure 4.7 presents the results of Algorithm 3 and numerically verifies the complexity rates obtained in Theorem 4.6.1. We observe that $K = 2.5$ corresponds to an expected value $\approx 3.2 \times 10^{-3}$, a sufficiently rare event. Figures 4.7a, 4.7b, and 4.7c correspond to $K = 2.5$. Figure 4.7a presents the exact relative error of the proposed MLDLMC estimator, estimated using a reference MLDLMC approximation with $\text{TOL}_r = 1\%$, for multiple runs of Algorithm 3 with various prescribed relative error tolerances. Figure 4.7a verifies that the MLDLMC estimator with IS satisfies the error constraint (4.14) with 95% confidence (corresponding to $\nu = 0.05$). The runtimes in Figure 4.7b include the estimation time of the bias, $V_{1,\ell}$ and $V_{2,\ell}$, in adaptive Algorithm 3. Figure 4.7b confirms that the average computational runtime closely follows the predicted theoretical rate $\mathcal{O}(\text{TOL}_r^{-3})$ for the entire range of relative tolerances. Figure 4.7c depicts the average computational cost estimate for the MLDLMC estimator over the prescribed TOL_r values. The cost estimate is computed using the following proxy.

$$\text{Computational cost}[\mathcal{A}_{\text{MLDLMC}}] \approx \sum_{\ell=0}^L (M_{1,\ell} P_\ell^2 N_\ell + M_{1,\ell} M_{2,\ell} P_\ell N_\ell). \quad (4.78)$$

Figures 4.7b and 4.7c both display a complexity of $\mathcal{O}(\text{TOL}_r^{-3})$ for the MLDLMC estimator with IS and the antithetic sampler, achieving one order complexity reduction compared with the DLMC estimator with the same IS scheme. Figure 4.7d plots the average computational cost estimate (4.78) for various values of K over a range of relative error tolerances. This verifies that the complexity rates of the MLDLMC estimator are independent of the K value, or the probability of the rare event itself. In conclusion, the



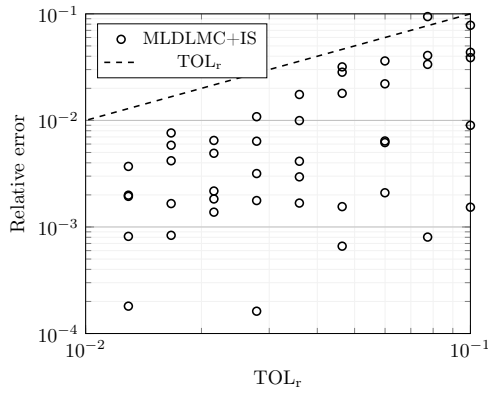
(a) Verifying Assumption 4.6.1. DLMC estimate of bias (4.72) using Algorithm 8 with inputs $M_1 = 10^2$ and $M_2 = 10^4$ versus level ℓ .




(b) Verifying Assumption 4.6.2. DLMC estimates of $V_{1,\ell}$ and $V_{2,\ell}$ using Algorithm 9 with inputs $M_1 = 10^2$ and $M_2 = 10^4$ versus level ℓ .

Figure 4.6: Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}$ for $K = 2.5$: Figures 4.6a and 4.6b numerically verify Assumptions 4.6.1 and 4.6.2, respectively, using the antithetic sampler (4.56). These plots together determine the rates $\tilde{\alpha} = 1$, $\tilde{w} = 2$, and $\tilde{s} = 1$.


proposed MLDLMC estimator with the proposed IS scheme reduced the computational complexity rate from $\mathcal{O}(\text{TOL}_r^{-4})$ for the DLMC estimator to $\mathcal{O}(\text{TOL}_r^{-3})$ for the treated example, under certain assumptions that are numerically verified. The presence of the two discretization parameters P and N for simulating MV-SDEs motivates extending this estimator to a multi-index MC setting to further reduce the computational complexity.




(a) Relative error of the estimator $\mathcal{A}_{\text{MLMLMC}}$ with respect to the prescribed relative tolerance TOL_r .



(b) Average computational runtime in seconds (s) of Algorithm 3 with respect to the prescribed relative tolerance TOL_r .



(c) Average computational cost estimate of the estimator $\mathcal{A}_{\text{MLMLMC}}$ with respect to the prescribed relative tolerance TOL_r .



(d) Average computational cost estimate with respect to the prescribed relative tolerance TOL_r for various K values.

Figure 4.7: Numerical results of applying Algorithm 3 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}$ for $K = 2.5$: Figure 4.7a shows that the proposed MLMLMC estimator with IS achieves a relative error below the prescribed TOL_r in the sense of (4.14) to estimate a rare event quantity of around 3.2×10^{-3} . Figures 4.7b and 4.7c show that $\mathcal{A}_{\text{MLMLMC}}$ achieves the relative accuracy requirement with a computational work of $\mathcal{O}(\text{TOL}_r^{-3})$, one order lesser than that of $\mathcal{A}_{\text{DLMLC}}$. Figure 4.7d depicts the independence of the complexity rate on the event rarity.

4.7 Multi-index double loop Monte Carlo estimator

The multi-index MC method was introduced in [24] to address problems involving multiple (two or more) discretization parameters. This approach builds on the efficiency of MLMC but requires mixed regularity with respect to the discretization parameters. While the multi-index method has been applied to non-rare event observables associated with particle systems in [82], the goal in this section is to combine the multi-index technique with the SOC-based IS scheme for rare events introduced in Section 4.4. The objective is to develop an efficient MC estimator for rare event probabilities associated with MV-SDEs that surpasses the performance of the MLMLMC estimator in Section 4.6. This gives rise to the novel MIDLMC estimator, representing the first attempt in literature to integrate IS with multi-index MC to the best of our knowledge.

Two discretization parameters (P, N) are required to generate sample paths of the decoupled **MV-SDE** (4.20). We introduce the multi-index $\alpha := (\alpha_1, \alpha_2) \in \mathbb{N}^2$, which defines the discretization parameters:

$$\begin{aligned} P_{\alpha_1} &= P_0 \times 2^{\alpha_1}, \\ N_{\alpha_2} &= N_0 \times 2^{\alpha_2}. \end{aligned} \quad (4.79)$$

P_0 and N_0 are the minimum number of particles and time steps, respectively, used to generate approximate sample paths of the decoupled **MV-SDE** (4.20). We propose the following method to couple **IS** with a **MIDLMC** estimator: we solve the same optimal control **PDE** (4.24) described in Section 4.4 using one realization of the stochastic particle system (4.9) with a large number of particles \bar{P} and time steps \bar{N} , and obtain the optimal **IS** control function ζ using (4.25). We apply the same policy ζ to generate sample paths of the decoupled **MV-SDE** (4.20) at all discretization multi-indices α . We recall the notation $G = g(\mathbf{X}(T))$, where \mathbf{X} denotes the McKean–Vlasov process following (4.1). Let $G_\alpha := g\left(\bar{\mathbf{X}}_\zeta^{P_{\alpha_1}|N_{\alpha_2}}(T)\right) \mathbb{L}^{P_{\alpha_1}|N_{\alpha_2}}$, where $\bar{\mathbf{X}}_\zeta^{P_{\alpha_1}|N_{\alpha_2}}$ denotes the **EM**-discretized controlled decoupled **MV-SDE** process (4.20) conditioned on the empirical law $\mu^{P_{\alpha_1}|N_{\alpha_2}}$ (4.10). The likelihood factor $\mathbb{L}^{P_{\alpha_1}|N_{\alpha_2}}$ at multi-index α is computed using (4.32). Following [24], we define the first-order mixed difference for this setting:

$$\Delta G_\alpha := \left(G_{(\alpha_1, \alpha_2)} - G_{(\alpha_1-1, \alpha_2)} \right) - \left(G_{(\alpha_1, \alpha_2-1)} - G_{(\alpha_1-1, \alpha_2-1)} \right). \quad (4.80)$$

The multi-index **MC** method is based on the following telescoping sum:

$$\mathbb{E}[G] = \sum_{\alpha \in \mathbb{N}^2} \mathbb{E}[\Delta G_\alpha], \quad (4.81)$$

where we set $G_{(-1,0)} = 0$, $G_{(0,-1)} = 0$, and $G_{(-1,-1)} = 0$. We let $\Delta \mathcal{G}_\alpha$ be a random variable such that $\mathbb{E}[\Delta \mathcal{G}_\alpha] = \mathbb{E}[\Delta G_\alpha]$. In the trivial case, $\Delta \mathcal{G}_\alpha = \Delta G_\alpha$. We could also choose $\Delta \mathcal{G}_\alpha$ cleverly, such that $\text{Var}[\Delta \mathcal{G}_\alpha] \ll \text{Var}[\Delta G_\alpha]$. Each expectation in (4.81) is approximated using a **DLMC** estimator, giving rise to the **MIDLMC** estimator:

$$\mathcal{A}_{\text{MIDLMC}}(\mathcal{I}) = \sum_{\alpha \in \mathcal{I}} \frac{1}{M_{1,\alpha}} \sum_{m_1=1}^{M_{1,\alpha}} \frac{1}{M_{2,\alpha}} \sum_{m_2=1}^{M_{2,\alpha}} \Delta \mathcal{G}_\alpha \left(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right), \quad (4.82)$$

where $\mathcal{I} \subset \mathbb{N}^2$ is an appropriately chosen index set and $\{M_{1,\alpha}, M_{2,\alpha}\}$ are the integer numbers of samples in the inner and outer loops, respectively, of the **DLMC** estimator for each $\alpha \in \mathcal{I}$. Similar to Section 4.6, we construct an antithetic estimator $\Delta \mathcal{G}_\alpha$ defined as

$$\begin{aligned} \Delta \mathcal{G}_\alpha \left(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right) &:= \left(\left(G_{(\alpha_1, \alpha_2)} - \mathcal{G}_{(\alpha_1-1, \alpha_2)} \right) \right. \\ &\quad \left. - \left(G_{(\alpha_1, \alpha_2-1)} - \mathcal{G}_{(\alpha_1-1, \alpha_2-1)} \right) \right) \left(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right), \end{aligned} \quad (4.83)$$

where $\mathcal{G}_{(\alpha_1-1, \alpha_2)}$ is highly correlated to $G_{(\alpha_1, \alpha_2)}$ and is defined as

$$\mathcal{G}_{(\alpha_1-1, \alpha_2)} \left(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right) = \frac{1}{2} \left(G_{(\alpha_1-1, \alpha_2)} \left(\omega_{1:P_{\alpha_1-1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right) \right.$$

$$+ G_{(\alpha_1-1, \alpha_2)} \left(\omega_{(P_{\alpha_1-1}+1):P_\alpha}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)} \right). \quad (4.84)$$

In (4.84), we split the P_{α_1} sets of random variables, used to generate $G_{(\alpha_1, \alpha_2)}$, into two iid groups of size P_{α_1-1} , and use each group to generate a realization of the empirical law independently. Conditioned on each realization of the empirical law, we generate approximate sample paths of $\bar{\mathbf{X}}_\zeta^{P_{\alpha_1-1}|N_{\alpha_2}}$ using the same $\bar{\omega}$ as for $G_{(\alpha_1, \alpha_2)}$ before averaging the quantity of interest over the two groups.

4.7.1 Error analysis

We bound the relative error of $\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})$ as

$$\begin{aligned} \frac{|\mathbb{E}[G] - \mathcal{A}_{\text{MIDLMC}}(\mathcal{I})|}{|\mathbb{E}[G]|} &\leq \underbrace{\frac{|\mathbb{E}[G] - \mathbb{E}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})]|}{|\mathbb{E}[G]|}}_{:=\epsilon_b, \text{ Relative bias}} \\ &+ \underbrace{\frac{|\mathbb{E}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})] - \mathcal{A}_{\text{MIDLMC}}(\mathcal{I})|}{|\mathbb{E}[G]|}}_{:=\epsilon_s, \text{ Relative statistical error}}. \end{aligned} \quad (4.85)$$

Instead of satisfying a prescribed relative error tolerance TOL_r in the sense of (4.14), we split the required accuracy between the relative bias and statistical errors using parameter $\theta \in (0, 1)$ and impose the following stricter constraints:

$$\epsilon_b \leq (1 - \theta)\text{TOL}_r, \quad (4.86)$$

$$\mathbb{P}[\epsilon_s \geq \theta\text{TOL}_r] \leq \nu. \quad (4.87)$$

for a confidence level determined by $0 < \nu \ll 1$.

Assumption 4.7.1. For all $\alpha \in \mathcal{I}$, the absolute value of the expected value of $\Delta\mathcal{G}_\alpha$ satisfies

$$|\mathbb{E}[\Delta\mathcal{G}_\alpha]| \leq Q_B \times 2^{-\alpha_1 b_1 - \alpha_2 b_2}, \quad (4.88)$$

for constants $Q_B > 0$ and $b_1, b_2 > 0$.

Assumption 4.7.1 is a direct implication of (4.37) arising from Assumptions 4.5.1 and 4.5.2. It relies on the mixed first-order difference term $\Delta\mathcal{G}_\alpha$ giving a product of the convergence rates in Assumptions 4.5.1 and 4.5.2 instead of a sum. b_1 and b_2 denote the weak convergence rates of the mixed first-order differences with respect to the number of particles and number of time steps respectively. The index set $\mathcal{I} \subset \mathbb{N}^2$ is chosen to satisfy the bias constraint (4.86). Using the asymptotic normality of the multi-index estimator [24], the statistical error constraint (4.87) can be expressed as the following constraint on the variance of the estimator:

$$\frac{\text{Var}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})]}{|\mathbb{E}[G]|^2} \leq \left(\frac{\theta\text{TOL}_r}{C_\nu} \right)^2, \quad (4.89)$$

where C_ν is the $(1 - \frac{\nu}{2})$ -quantile of the standard normal distribution. Due to the independence of the DLMC estimators corresponding to each α in (4.82), the MIDLMC estimator

variance can be written as

$$\text{Var}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})] = \sum_{\alpha \in \mathcal{I}} \frac{1}{M_{1,\alpha}} \text{Var} \left[\frac{1}{M_{2,\alpha}} \sum_{m_2=1}^{M_{2,\alpha}} \Delta \mathcal{G}_\alpha \left(\omega_{1:P_{\alpha_1}}^{(\alpha,1)} \times \bar{\omega}^{(\alpha,m_2)} \right) \right]. \quad (4.90)$$

Using the law of total variance,

$$\begin{aligned} \text{Var}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})] &= \sum_{\alpha \in \mathcal{I}} \frac{1}{M_{1,\alpha}} \left(\underbrace{\text{Var} \left[\mathbb{E} \left[\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha,1)} \right] \right]}_{=V_{1,\alpha}} \right. \\ &\quad \left. + \frac{1}{M_{2,\alpha}} \underbrace{\mathbb{E} \left[\text{Var} \left[\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha,1)} \right] \right]}_{=V_{2,\alpha}} \right) \\ &= \sum_{\alpha \in \mathcal{I}} \left(\frac{V_{1,\alpha}}{M_{1,\alpha}} + \frac{V_{2,\alpha}}{M_{1,\alpha} M_{2,\alpha}} \right). \end{aligned} \quad (4.91)$$

Assumption 4.7.2. For all $\alpha \in \mathcal{I}$, the variance terms in (4.91) satisfy

$$V_{1,\alpha} := \text{Var} \left[\mathbb{E} \left[\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha,1)} \right] \right] \lesssim 2^{-\alpha_1 w_1 - \alpha_2 w_2}, \quad (4.92)$$

$$V_{2,\alpha} := \mathbb{E} \left[\text{Var} \left[\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha,1)} \right] \right] \lesssim 2^{-\alpha_1 s_1 - \alpha_2 s_2}, \quad (4.93)$$

for constants $w_i > 0$ and $s_i > 0$ for $i = 1, 2$.

Assumption 4.7.2 is motivated from Assumption 4.6.2. To the best of our knowledge, there are currently no proofs of these assumptions for **MV-SDEs**, but we verify them numerically in Section 4.7.4.

4.7.2 Computational complexity

For each multi-index $\alpha \in \mathbb{N}^2$, we estimate $\mathbb{E}[\Delta G_\alpha]$ that requires four evaluations of the random variable G_α for different combinations of discretization parameters (4.80). We assume a naïve method (with computational cost $\mathcal{O}(P)$) to compute the empirical means in the drift and diffusion functions in (4.9) and (4.20). The computational cost to generate one realization of the random variable G_α comprises two parts (see Section 4.5): generation of a realization of the empirical law $\mu^{P_{\alpha_1} | N_{\alpha_2}}$ using the **EM** scheme with N_{α_2} time steps ($\mathcal{O}(N_{\alpha_2} P_{\alpha_2}^2)$), and the generation of one sample path of the decoupled **MV-SDE** (4.20) using the **EM** scheme with N_{α_2} time steps ($\mathcal{O}(P_{\alpha_1} N_{\alpha_2})$). Thus, we can bound the total computational cost of the **MIDLMC** estimator:

$$\mathcal{W}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})] \lesssim \sum_{\alpha \in \mathcal{I}} M_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2 + M_{1,\alpha} M_{2,\alpha} N_{\alpha_2} P_{\alpha_1}. \quad (4.94)$$

For a given index set \mathcal{I} , we find the optimal $M_{1,\alpha}$ and $M_{2,\alpha}$ for all $\alpha \in \mathcal{I}$ that minimises the total computational cost (4.94) subject to the statistical error constraint (4.89):

$$\begin{cases} \arg \min_{\{M_{1,\alpha}, M_{2,\alpha}\}_{\alpha \in \mathcal{I}}} \sum_{\alpha \in \mathcal{I}} M_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2 + M_{1,\alpha} M_{2,\alpha} N_{\alpha_2} P_{\alpha_1} \\ \text{s.t.} \quad \left(\sum_{\alpha \in \mathcal{I}} \frac{V_{1,\alpha}}{M_{1,\alpha}} + \frac{V_{2,\alpha}}{M_{1,\alpha} M_{2,\alpha}} \right) \approx \left(\frac{\theta \text{TOL}_r |\mathbb{E}[G]|}{C_v} \right)^2. \end{cases} \quad (4.95)$$

The solution to (4.95) in \mathbb{R}^+ using the Lagrange multiplier method for constrained optimization yields

$$\begin{aligned}\mathcal{M}_{1,\alpha} &:= \left(\frac{C_v}{\theta \text{TOL}_r \mathbb{E}[G]} \right)^2 \sqrt{\frac{V_{1,\alpha}}{N_{\alpha_2} P_{\alpha_1}^2}} \sum_{\beta \in \mathcal{I}} \left(\sqrt{V_{1,\beta} N_{\beta_2} P_{\beta_1}^2} + \sqrt{V_{2,\beta} N_{\beta_2} P_{\beta_1}} \right), \\ \tilde{\mathcal{M}}_\alpha &:= \mathcal{M}_{1,\alpha} \mathcal{M}_{2,\alpha} \\ &= \left(\frac{C_v}{\theta \text{TOL}_r \mathbb{E}[G]} \right)^2 \sqrt{\frac{V_{2,\alpha}}{N_{\alpha_2} P_{\alpha_1}}} \sum_{\beta \in \mathcal{I}} \left(\sqrt{V_{1,\beta} N_{\beta_2} P_{\beta_1}^2} + \sqrt{V_{2,\beta} N_{\beta_2} P_{\beta_1}} \right).\end{aligned}\quad (4.96)$$

In practice, we use natural numbers for $\{\mathcal{M}_{1,\alpha}, \mathcal{M}_{2,\alpha}\}_{\alpha \in \mathcal{I}}$. Hence, we use the following quasi-optimal solution to (4.95):

$$M_{1,\alpha} = \lceil \mathcal{M}_{1,\alpha} \rceil, \quad M_{2,\alpha} = \left\lceil \frac{\tilde{\mathcal{M}}_\alpha}{\lceil \mathcal{M}_{1,\alpha} \rceil} \right\rceil. \quad (4.97)$$

Using (4.97) we bound the estimator cost as

$$\begin{aligned}\mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})] &\lesssim \sum_{\alpha \in \mathcal{I}} \left((\mathcal{M}_{1,\alpha} + 1) N_{\alpha_2} P_{\alpha_1}^2 + (\mathcal{M}_{1,\alpha} + 1) \left(\frac{\tilde{\mathcal{M}}_\alpha}{\lceil \mathcal{M}_{1,\alpha} \rceil} + 1 \right) N_{\alpha_2} P_{\alpha_1} \right) \\ &\leq \underbrace{\sum_{\alpha \in \mathcal{I}} (\mathcal{M}_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2 + \tilde{\mathcal{M}}_\alpha N_{\alpha_2} P_{\alpha_1})}_{:= W_1(\mathcal{I}) \text{ (objective function in (4.95))}} \\ &\quad + \underbrace{\sum_{\alpha \in \mathcal{I}} (P_{\alpha_1}^2 N_{\alpha_2} + P_{\alpha_1} N_{\alpha_2})}_{:= W_2(\mathcal{I}) \text{ (cost of one sample per multi-index)}} \\ &\quad + \underbrace{\sum_{\alpha \in \mathcal{I}} \mathcal{M}_{1,\alpha} N_{\alpha_2} P_{\alpha_1}}_{:= W_3(\mathcal{I})} + \underbrace{\sum_{\alpha \in \mathcal{I}} \frac{\tilde{\mathcal{M}}_\alpha}{\lceil \mathcal{M}_{1,\alpha} \rceil} N_{\alpha_2} P_{\alpha_1}}_{:= W_4(\mathcal{I})}.\end{aligned}\quad (4.98)$$

Using $P_{\alpha_1} > 0$ we see that,

$$\begin{aligned}W_3(\mathcal{I}) &= \sum_{\alpha \in \mathcal{I}} \mathcal{M}_{1,\alpha} N_{\alpha_2} P_{\alpha_1} \leq \sum_{\alpha \in \mathcal{I}} \mathcal{M}_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2 \leq W_1(\mathcal{I}), \\ W_4(\mathcal{I}) &= \sum_{\alpha \in \mathcal{I}} \frac{\tilde{\mathcal{M}}_\alpha}{\lceil \mathcal{M}_{1,\alpha} \rceil} N_{\alpha_2} P_{\alpha_1} \leq \sum_{\alpha \in \mathcal{I}} \frac{\tilde{\mathcal{M}}_\alpha}{\max(1, \mathcal{M}_{1,\alpha})} N_{\alpha_2} P_{\alpha_1} \\ &\leq \sum_{\alpha \in \mathcal{I}} \tilde{\mathcal{M}}_\alpha N_{\alpha_2} P_{\alpha_1} \leq W_1(\mathcal{I}).\end{aligned}$$

Hence, we can rewrite (4.98) as follows,

$$\mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})] \lesssim W_1(\mathcal{I}) + W_2(\mathcal{I}). \quad (4.99)$$

Substituting (4.96) in (4.99), we obtain

$$\begin{aligned}\mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})] &\lesssim \left(\frac{C_v}{\theta \text{TOL}_r \mathbb{E}[G]} \right)^2 \left(\sum_{\alpha \in \mathcal{I}} \sqrt{V_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2} + \sqrt{V_{2,\alpha} N_{\alpha_2} P_{\alpha_1}} \right)^2 \\ &\quad + \sum_{\alpha \in \mathcal{I}} (P_{\alpha_1}^2 N_{\alpha_2} + P_{\alpha_1} N_{\alpha_2}).\end{aligned}\quad (4.100)$$

Under Assumptions 4.7.1 and 4.7.2, the total computational cost (4.100) to reach a prescribed relative error tolerance TOL_r can be estimated as

$$\begin{aligned} \mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})] &\lesssim \left(\frac{C_v}{\theta \text{TOL}_r \mathbb{E}[G]} \right)^2 \left(\underbrace{\sum_{\alpha \in \mathcal{I}} (\exp\langle \bar{\mathbf{g}}, \alpha \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle)}_{=:\tilde{W}_1(\mathcal{I})} \right)^2 \\ &\quad + \underbrace{\sum_{\alpha \in \mathcal{I}} \exp\langle \boldsymbol{\lambda}, \alpha \rangle}_{=:\tilde{W}_2(\mathcal{I})}. \end{aligned} \quad (4.101)$$

We define the vectors $\bar{\mathbf{g}}, \bar{\bar{\mathbf{g}}}, \boldsymbol{\lambda} \in \mathbb{R}^2$ with the following entries:

$$\bar{\mathbf{g}} := \log(2) \left[\frac{2-w_1}{2}, \frac{1-w_2}{2} \right]^{\text{tr}}, \quad (4.102)$$

$$\bar{\bar{\mathbf{g}}} := \log(2) \left[\frac{1-s_1}{2}, \frac{1-s_2}{2} \right]^{\text{tr}} \quad (4.103)$$

$$\boldsymbol{\lambda} := \log(2) [2, 1]^{\text{tr}}. \quad (4.104)$$

$\tilde{W}_2(\mathcal{I})$ represents the total computational cost required to generate one sample for each multi-index $\alpha \in \mathcal{I}$, which is the minimum cost for a MIDLMC estimator and it is essential to ensure that it does not dominate the first term in the bound (4.101). For now, we assume that the first term dominates the second term, allowing us to focus solely on $\tilde{W}_1(\mathcal{I})$. We will explicitly outline the conditions that guarantee the dominance of the first term in Theorem 4.7.1.

Optimal index set

One of the main objectives of this work is to motivate a choice for the set of multi-indices $\mathcal{I} := \mathcal{I}(\text{TOL}_r)$ that minimises $\mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})]$. We obtain the following bound for the relative bias, subject to the following relative bias constraint:

$$\epsilon_b(\mathcal{I}) := \frac{1}{|\mathbb{E}[G]|} \left| \sum_{\alpha \notin \mathcal{I}} \mathbb{E}[\Delta \mathcal{G}_\alpha] \right| \leq \frac{1}{|\mathbb{E}[G]|} \sum_{\alpha \notin \mathcal{I}} |\mathbb{E}[\Delta \mathcal{G}_\alpha]|. \quad (4.105)$$

We choose the multi-index set \mathcal{I} to satisfy the following relative bias constraint.

$$\frac{1}{|\mathbb{E}[G]|} \sum_{\alpha \notin \mathcal{I}} |\mathbb{E}[\Delta \mathcal{G}_\alpha]| \leq (1-\theta) \text{TOL}_r. \quad (4.106)$$

Using (4.88), we rewrite (4.106) as

$$\tilde{B}(\mathcal{I}) := \sum_{\alpha \notin \mathcal{I}} \exp\langle -\boldsymbol{\rho}, \alpha \rangle \leq \frac{(1-\theta) \text{TOL}_r |\mathbb{E}[G]|}{Q_B}, \quad (4.107)$$

where we define the vector $\boldsymbol{\rho} := \log(2) [b_1, b_2]^{\text{tr}} \in \mathbb{R}^2$. $\tilde{B}(\mathcal{I})$ can be considered an upper bound on the absolute bias of $\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})$. To determine the optimal \mathcal{I} , we solve

the following optimization problem:

$$\min_{\mathcal{I} \in \mathbb{N}^2} \mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I})] \quad \text{such that} \quad \epsilon_b(\mathcal{I}) \leq (1 - \theta)\text{TOL}_r. \quad (4.108)$$

To simplify the problem, we work with the bounds on the bias (4.107) and the work (4.101). We assume that the first term in (4.101) dominates the second. Subsequently, we solve the following simplified problem to obtain a quasi-optimal set \mathcal{I} :

$$\min_{\mathcal{I} \in \mathbb{N}^2} \tilde{W}_1(\mathcal{I}) \quad \text{such that} \quad \tilde{B}(\mathcal{I}) \leq \frac{(1 - \theta)\text{TOL}_r |\mathbb{E}[G]|}{Q_B}. \quad (4.109)$$

$\tilde{W}_1(\mathcal{I})$ is defined in (4.101) and $\tilde{B}(\mathcal{I})$ is defined in (4.107). For convenience, we refer to the objective \tilde{W}_1 as the ‘‘work’’ and the constraint function \tilde{B} as the ‘‘error’’ in the rest of this section. To solve the optimization problem (4.109), we closely follow the procedure in [24]. First, we write the total error associated with index set \mathcal{I} from (4.107) as follows:

$$\tilde{B}(\mathcal{I}) = \sum_{\alpha \notin \mathcal{I}} \exp\langle -\rho, \alpha \rangle = \sum_{\alpha \notin \mathcal{I}} \mathcal{E}_\alpha. \quad (4.110)$$

where $\mathcal{E}_\alpha := \exp\langle -\rho, \alpha \rangle$ denotes the error contribution of multi-index α . Next, we write the total work using (4.101), as follows:

$$\tilde{W}_1(\mathcal{I}) = \sum_{\alpha \in \mathcal{I}} (\exp\langle \bar{\mathbf{g}}, \alpha \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle) = \sum_{\alpha \in \mathcal{I}} \omega_\alpha. \quad (4.111)$$

where $\omega_\alpha := \exp\langle \bar{\mathbf{g}}, \alpha \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle$ denotes the work contribution of multi-index α . Next, we assign a ‘‘profit’’ indicator to each multi-index and add only the most profitable multi-indices to the index set \mathcal{I} . We define the profit of a multi-index α as follows:

$$\mathcal{P}_\alpha := \frac{\mathcal{E}_\alpha}{\omega_\alpha}. \quad (4.112)$$

Using (4.110),(4.111), we obtain the following expression for the profit:

$$\mathcal{P}_\alpha = \frac{\exp\langle -\rho, \alpha \rangle}{\exp\langle \bar{\mathbf{g}}, \alpha \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle}. \quad (4.113)$$

Following [24], we define the optimal index set for a given level $v \in \mathbb{R}^+$ as the following level set:

$$\mathcal{I}(v) := \{\alpha \in \mathbb{N}^2 : \mathcal{P}_\alpha \geq v\}. \quad (4.114)$$

Inserting (4.113) into (4.114), we obtain the optimal index set for the set threshold $L \in \mathbb{R}^+$ for the MIDLMC estimator:

$$\mathcal{I}(L) := \left\{ \alpha \in \mathbb{N}^2 : \exp\langle \underbrace{(\bar{\mathbf{g}} + \rho)}_{=: \bar{\delta}}, \alpha \rangle + \exp\langle \underbrace{(\bar{\bar{\mathbf{g}}} + \rho)}_{=: \bar{\bar{\delta}}}, \alpha \rangle \leq L \right\}. \quad (4.115)$$

where $\bar{\delta} := [\bar{\delta}_1, \bar{\delta}_2]^{\text{tr}}$, $\bar{\bar{\delta}} := [\bar{\bar{\delta}}_1, \bar{\bar{\delta}}_2]^{\text{tr}} \in (\mathbb{R}^+)^2$ are normalized weighting vectors. These are written explicitly as follows:

$$\bar{\delta} = \left[\frac{\left(\frac{2-w_1}{2} + b_1\right)}{C_{\bar{\delta}}}, \frac{\left(\frac{1-w_2}{2} + b_2\right)}{C_{\bar{\delta}}} \right]^{\text{tr}},$$

$$\begin{aligned}
C_{\bar{\delta}} &= \left(\frac{2-w_1}{2} + b_1 \right) + \left(\frac{1-w_2}{2} + b_2 \right), \\
\bar{\delta} &= \left[\frac{\left(\frac{1-s_1}{2} + b_1 \right)}{C_{\bar{\delta}}}, \frac{\left(\frac{1-s_2}{2} + b_2 \right)}{C_{\bar{\delta}}} \right]^{\text{tr}}, \\
C_{\bar{\delta}} &= \left(\frac{1-s_1}{2} + b_1 \right) + \left(\frac{1-s_2}{2} + b_2 \right).
\end{aligned} \tag{4.116}$$

To ensure admissible index sets according to the classical multi-index method [24], we need that $\bar{\delta}_1, \bar{\delta}_2, \bar{\delta}_1, \bar{\delta}_2 > 0$, leading to the following condition on the convergence rates:

$$2b_1 \geq \max(w_1 - 1, s_1) - 1, \quad 2b_2 \geq \max(w_2, s_2) - 1. \tag{4.117}$$

The goal of the following theorem is: i) to choose an optimal multi-index $\mathcal{I}(L)$ defined in (4.115) that minimises the computational work of the MIDLMC estimator while satisfying the relative bias constraint (4.106), and ii) to choose the optimal number of samples per multi-index in the obtained set $\mathcal{I}(L)$ to get an optimal MIDLMC estimator that also satisfies the variance constraint (4.89). For this theorem, we introduce the following notation:

$$\begin{aligned}
\chi_{11} &:= C_{\bar{\delta}} \log(2) \max \left(\frac{2-w_1}{2-w_1+2b_1}, \frac{1-w_2}{1-w_2+2b_2} \right), \\
\chi_{12} &:= C_{\bar{\delta}} \log(2) \max \left(\frac{1-s_1}{2-w_1+2b_1}, \frac{1-s_2}{1-w_2+2b_2} \right), \\
\chi_{21} &:= C_{\bar{\delta}} \log(2) \max \left(\frac{2-w_1}{1-s_1+2b_1}, \frac{1-w_2}{1-s_2+2b_2} \right), \\
\chi_{22} &:= C_{\bar{\delta}} \log(2) \max \left(\frac{1-s_1}{1-s_1+2b_1}, \frac{1-s_2}{1-s_2+2b_2} \right), \\
\eta_1 &:= C_{\bar{\delta}} \log(2) \min \left(\frac{2b_1}{2-w_1+2b_1}, \frac{2b_2}{1-w_2+2b_2} \right), \\
\eta_2 &:= C_{\bar{\delta}} \log(2) \min \left(\frac{2b_1}{1-s_1+2b_1}, \frac{2b_2}{1-s_2+2b_2} \right), \\
e_1 &:= \begin{cases} 2, & \frac{2-w_1}{2b_1} = \frac{1-w_2}{2b_2} \\ 1, & \text{otherwise} \end{cases}, \quad e_2 := \begin{cases} 2, & \frac{1-s_1}{2b_1} = \frac{1-s_2}{2b_2} \\ 1, & \text{otherwise} \end{cases}, \\
\aleph_1 &:= \begin{cases} 2, & \frac{1-s_1}{2-w_1+2b_1} = \frac{1-s_2}{1-w_2+2b_2} \\ 1, & \text{otherwise} \end{cases}, \quad \aleph_2 := \begin{cases} 2, & \frac{2-w_1}{1-s_1+2b_1} = \frac{1-w_2}{1-s_2+2b_2} \\ 1, & \text{otherwise} \end{cases}, \\
d_1 &:= \begin{cases} 2, & w_1 = 2 \text{ and } w_2 = 1 \\ 1, & w_1 = 2 \text{ or } w_2 = 1 \\ 0, & \text{otherwise} \end{cases}, \quad d_2 := \begin{cases} 2, & s_1 = 1 \text{ and } s_2 = 1 \\ 1, & s_1 = 1 \text{ or } s_2 = 1 \\ 0, & \text{otherwise} \end{cases}, \\
\varsigma &:= \min \left(\max \left(0, \frac{\chi_{11}}{\eta_1}, \frac{\chi_{12}}{\eta_1} \right), \max \left(0, \frac{\chi_{21}}{\eta_2}, \frac{\chi_{22}}{\eta_2} \right) \right),
\end{aligned}$$

$$\varrho := \begin{cases} \max(d_1, d_2), & \zeta = 0 \\ (e_1 - 1) \left(1 + \frac{\chi_{11}}{\eta_1}\right), & \zeta = \frac{\chi_{11}}{\eta_1} \\ (\aleph_1 - 1) + (e_1 - 1) \frac{\chi_{12}}{\eta_1}, & \zeta = \frac{\chi_{12}}{\eta_1} \\ (\aleph_2 - 1) + (e_2 - 1) \frac{\chi_{21}}{\eta_2}, & \zeta = \frac{\chi_{21}}{\eta_2} \\ (e_2 - 1) \left(1 + \frac{\chi_{22}}{\eta_2}\right), & \zeta = \frac{\chi_{22}}{\eta_2} \end{cases},$$

$$\Psi := \min \left(\frac{2}{2 - w_1 + 2b_1'}, \frac{1}{1 - w_2 + 2b_2'}, \frac{2}{1 - s_1 + 2b_1'}, \frac{1}{1 - s_2 + 2b_2'} \right).$$

Theorem 4.7.1 (Optimal MIDLMC complexity). *We consider the optimal multi-index set given by (4.115) with optimal vectors $\bar{\delta}$ and $\bar{\delta}$, as defined in (4.116), satisfying the condition (4.117). Under Assumptions 4.7.1 and 4.7.2, the optimal $\bar{L} \in \mathbb{R}_+$ satisfies the bias constraint (4.107) in the following sense:*

$$\lim_{\text{TOL}_r \downarrow 0} \frac{\tilde{B}(\mathcal{I}(\bar{L}))}{\frac{(1-\theta)\text{TOL}_r \mathbb{E}[G]}{Q_B}} \leq 1. \quad (4.118)$$

Then, the total computational cost of the MIDLMC estimator (4.82) $\mathcal{W}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I}(\bar{L}))]$, subject to constraint (4.89) satisfies the following:

$$\limsup_{\text{TOL}_r \downarrow 0} \frac{\mathcal{W}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I}(\bar{L}))]}{\text{TOL}_r^{-2-2\zeta} \left(\log \text{TOL}_r^{-1}\right)^{2\varrho}} \leq C_{\text{work}} < \infty, \quad (4.119)$$

given that the following condition holds:

$$\Psi \leq 1 + \zeta. \quad (4.120)$$

Proof. See Appendix A.5. □

To demonstrate the superiority of the MIDLMC estimator over the MLDLMC estimator, we examine the case where $w_1 = w_2 = s_1 = s_2 = s > 0$, and $b_1 = b_2 = b > 0$. This implies that the bias and variances converge at the same rate in the P and N directions. From Theorem 4.7.1, as $\text{TOL}_r \rightarrow 0$, we asymptotically have

$$\mathcal{W}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I}(\bar{L}))] = \begin{cases} \mathcal{O}(\text{TOL}_r^{-2}), & s > 2 \\ \mathcal{O}(\text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2), & s = 2. \\ \mathcal{O}(\text{TOL}_r^{-2 - \frac{2-s}{b}}), & s < 2 \end{cases}. \quad (4.121)$$

In contrast, the MLDLMC estimator has the following complexity rates (from Theorem 4.6.1):

$$\mathcal{W}[\mathcal{A}_{\text{MLDLMC}}] = \begin{cases} \mathcal{O}(\text{TOL}_r^{-2}), & s > 3 \\ \mathcal{O}(\text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2), & s = 3. \\ \mathcal{O}(\text{TOL}_r^{-2 - \frac{3-s}{b}}), & s < 3 \end{cases}. \quad (4.122)$$

Comparing (4.121) and (4.122), the MIDLMC estimator has better complexity rates than

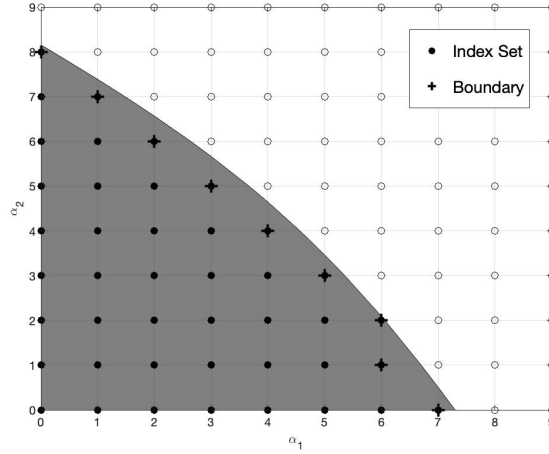


Figure 4.8: Schematic illustration of an index set \mathcal{I} with its outer boundary $\partial\mathcal{I}$ as defined in (4.124).

the **MLDLMC** estimator in all regimes of variance convergence rate s . In addition, the condition (4.120) for the **MIDLMC** estimator simplifies to $2b \geq s$. For the multilevel case, we have the more restrictive condition $2b \geq \min(s, 3)$ (see Theorem 4.6.1). Moreover, we also deduce that the complexity of the **MIDLMC** estimator is the same as that of a **MLDLMC** estimator with only P as the discretization parameter. However, the **MIDLMC** estimator requires mixed regularity (in the sense of Assumptions 4.7.1 and 4.7.2) concerning P and N , whereas the **MLDLMC** estimator only requires ordinary regularity.

Algorithm 10 in Appendix B.5 implements the proposed **IS** scheme to estimate $\mathbb{E}[\Delta G_{\alpha}]$ for each $\alpha \in \mathcal{I}$, required for the **MIDLMC** estimator (4.82). Next, we develop an adaptive **MIDLMC** algorithm that sequentially selects the index set \mathcal{I} and determines the optimal number of samples $\{M_{1,\alpha}, M_{2,\alpha}\}_{\alpha \in \mathcal{I}}$ to satisfy the error constraints specified in (4.86) and (4.87). The adaptive algorithm relies on cost-effective, reliable estimates of the relative bias and variances $\{V_{1,\alpha}, V_{2,\alpha}\}_{\alpha \in \mathcal{I}}$.

4.7.3 Adaptive algorithm

Estimating the bias of index set \mathcal{I}

The optimal index set (4.115) is defined using the set threshold L . Given that $\mathcal{I}(L) \subset \mathcal{I}(L+1)$, we apply the following heuristic estimate for the absolute bias corresponding to index set $\mathcal{I}(L)$.

$$|\mathbb{E}[G - \mathcal{A}_{\text{MIDLMC}}(\mathcal{I}(L))]| \approx \sum_{\alpha \in \partial\mathcal{I}(L)} |\mathbb{E}[\Delta G_{\alpha}]|, \quad (4.123)$$

where $\partial\mathcal{I}(L)$ is the outer boundary of the index set $\mathcal{I}(L)$, defined as follows:

$$\partial\mathcal{I}(L) = \{\alpha \in \mathcal{I}(L) : \alpha + (1, 0) \notin \mathcal{I}(L) \text{ or } \alpha + (0, 1) \notin \mathcal{I}(L)\}. \quad (4.124)$$

Figure 4.8 visualises $\partial\mathcal{I}(L)$ according to (4.124). Because $\partial\mathcal{I}(L) \subset \mathcal{I}(L)$, we use already-computed nested averages using the optimal $\{M_{1,\alpha}, M_{2,\alpha}\}_{\alpha \in \partial\mathcal{I}(L)}$ in (4.82) to estimate each expectation in (4.123).

Estimating variances $\{V_{1,\alpha}, V_{2,\alpha}\}$

Computationally cheap, robust empirical estimates of $\{V_{1,\alpha}, V_{2,\alpha}\}$ for all $\alpha \in \mathcal{I}(L)$ are necessary to compute the optimal number of samples to satisfy the variance constraint (4.89) of the estimator using (4.97). Therefore, DLMC Algorithm 11 in Appendix B.6 is used with appropriately chosen \bar{M}_1, \bar{M}_2 . However, estimating $\{V_{1,\alpha}, V_{2,\alpha}\}$ for all $\alpha \in \mathcal{I}(L)$ using Algorithm 11 can become computationally burdensome when $\mathcal{I}(L)$ is even moderately large. To alleviate this computational overload, we leverage Assumption 4.7.2 and employ an extrapolation approach to estimate $\{V_{1,\alpha}, V_{2,\alpha}\}$ for deeper multi-indices. Specifically, we employ Algorithm 11 solely to estimate $\{V_{1,\alpha}, V_{2,\alpha}\}$ for the small full tensor index set $\{0, 1, 2\} \times \{0, 1, 2\}$. Then, we employ the extrapolation Algorithm 12 in Appendix B.6, applying Assumption 4.7.2 to estimate $\{V_{1,\alpha}, V_{2,\alpha}\}$ for the remaining multi-indices. To further alleviate the computational burden, we only estimate sample variances $\{V_{1,\alpha}, V_{2,\alpha}\}$ for the newly added multi-indices in each iteration (i.e., for $\alpha \in \mathcal{I}(L+1) - \mathcal{I}(L)$). As $\mathcal{I}(L) \subset \mathcal{I}(L+1)$ by construction, we can reuse the estimates of $\{V_{1,\alpha}, V_{2,\alpha}\}$ for multi-indices from previous iterations.

Relative error control

To meet the relative error constraints (4.86) and (4.89) for a given relative error tolerance TOL_r , the adaptive algorithm requires a heuristic estimate of $\mathbb{E}[G]$. Algorithm 10 with appropriately selected \bar{M}_1 and \bar{M}_2 gives an initial estimate \bar{G} for $\mathbb{E}[\Delta G_{(0,0)}]$. In subsequent iterations, the multi-index estimator (4.82) with optimal values of $\{M_{1,\alpha}, M_{2,\alpha}\}_{\alpha \in \mathcal{I}(L)}$ updates \bar{G} and the absolute error tolerances. Combining these components, we present the adaptive MIDLMC Algorithm 4 for rare events associated with MV-SDEs. The IS control function ζ in Algorithm 4 is computed off-line by generating one realization of the empirical law $\mu^{\bar{P}|\bar{N}}$ with some large \bar{P} and \bar{N} and then numerically solving the control PDE (4.24) given $\mu^{\bar{P}|\bar{N}}$, and obtaining ζ using (4.25).

Algorithm 4: Adaptive MIDLMC with IS

User Provided Input: $P_0, N_0, \text{TOL}_r, \zeta, \nu, \theta, \{\bar{M}_1, \bar{M}_2\}, \{\tilde{M}_1, \tilde{M}_2\};$

Input from Pilots: $\{b_1, b_2\}, \{w_1, w_2\}, \{s_1, s_2\};$

Estimate $\bar{G} = \mathbb{E}[\Delta G_{(0,0)}]$ using **Algorithm 10** with $\bar{M}_1, \bar{M}_2, \zeta;$

Estimate and store $\{V_{1,\alpha}, V_{2,\alpha}\}$ for $\alpha \in [0, 1, 2] \times [0, 1, 2]$ using **Algorithm 11** with $\bar{M}_1, \bar{M}_2, \zeta;$

Set $L = 1;$

while $\epsilon_b(\mathcal{I}(L)) > (1 - \theta)\text{TOL}_r$ **do**

Generate index set $\mathcal{I}(L)$ from (4.115);

Estimate and store $\{V_{1,\alpha}, V_{2,\alpha}\}_{\alpha \in \mathcal{I}(L)}$ using **Algorithm 12**;

Compute optimal sample sizes $\{M_{1,\alpha}, M_{2,\alpha}\}_{\alpha \in \mathcal{I}(L)}$ using (4.97);

Reevaluate $\bar{G} = \sum_{\alpha \in \mathcal{I}(L)} \mathbb{E}[\Delta G_\alpha]$ as per (4.82) with $\{M_{1,\alpha}, M_{2,\alpha}\}, \zeta$ using **Algorithm 10** for each $\alpha;$

$\epsilon_b(\mathcal{I}(L)) \approx \frac{1}{|\bar{G}|} \sum_{\alpha \in \partial \mathcal{I}(L)} |\mathbb{E}[\Delta G_\alpha]|;$

$L \leftarrow L + 1;$

end

$\mathcal{A}_{\text{MIDLMC}}(\mathcal{I}(L)) = \bar{G}.$

In Algorithm 4, the rates $\{b_1, b_2\}$, $\{w_1, w_2\}$, and $\{s_1, s_2\}$ associated with Assumptions 4.7.1 and 4.7.2 must be estimated a priori, to determine the optimal index set $\mathcal{I}(L)$

using (4.115). Reliable pilot runs are required to obtain these estimates. Such pilot runs are not required for the adaptive DLMC Algorithm 2 or MLDLMC Algorithm 3.

4.7.4 Numerical experiments

We provide numerical evidence for the assumptions and computational complexity rates derived in Section 4.7. The results focus on the 1D Kuramoto model (4.12) with the following settings: $\sigma = 0.4$, $T = 1$, $x_0 \sim \mathcal{N}(0, 0.2)$, and $\zeta \sim \mathfrak{U}(-0.2, 0.2)$. We set the parameters $\theta = 0.5$ and $\nu = 0.05$. In addition, we set the following hierarchies for the MIDLMC estimator.

$$P_{\alpha_1} = 5 \times 2^{\alpha_1}, \quad N_{\alpha_2} = 4 \times 2^{\alpha_2}. \quad (4.125)$$

We use MIDLMC with IS to estimate $\mathbb{P}[X(T) > K] \approx \mathbb{E}[\tilde{\mathcal{I}}_{\{X(T) > K\}}]$, where the process X satisfies the Kuramoto MV-SDE (4.13) and $K \in \mathbb{R}$ lies deep in the right tail of the distribution of $X(T)$. The function $\tilde{\mathcal{I}}$ is a smooth, differentiable approximation of the indicator function \mathcal{I} defined as follows.

$$\tilde{\mathcal{I}}_{\{x > K\}}(x) = \frac{1}{2} \left(1 + \tanh \left(\frac{x - K}{\beta} \right) \right). \quad (4.126)$$

In the limit $\beta \rightarrow 0$, the function $\tilde{\mathcal{I}}_{\{x > K\}}(x)$ converges to the indicator function $\mathcal{I}_{\{x > K\}}(x)$. This approximation helps avoid poor variance convergence rates of the first-order mixed differences in multi-index MC [24]. We fix $\beta = \frac{1}{3}$ and $K = 3.5$ for our numerical experiments, corresponding to $\mathbb{E}[\tilde{\mathcal{I}}_{\{X(T) > K\}}] \approx 2.04 \times 10^{-5}$, a sufficiently rare event. We assess two algorithms: the adaptive MLDLMC Algorithm 3 and the adaptive MIDLMC algorithm 4 with the index set defined in (4.115). For both methods, we generate one realization of the particle system (4.9) off-line with $\bar{P} = 1000$ particles and $\bar{N} = 100$ time steps to compute the IS control function ζ . Given the empirical law, the SOC-based IS scheme for the decoupled MV-SDE is obtained by numerically solving the optimal control PDE (4.24) with a FDM solver (see Appendix C.1) and using (4.25) to get the optimal IS control function ζ .

We start by verifying the variance reduction for the DLMC estimator of the mixed difference $\mathbb{E}[\Delta G_{\alpha}]$. Figure 4.9 illustrates the squared coefficient of variation for the estimator of $\mathbb{E}[\Delta G_{\alpha}]$ with and without IS, plotted as a function of M_1 with a fixed M_2 . The results in Figure 4.9 demonstrate a significantly reduced squared coefficient of variation (around one order of magnitude) for the fixed $\alpha = (2, 2)$. Both variances decrease with $\mathcal{O}(M_1^{-1})$, and the proposed IS scheme decreases the associated constant.

Subsequently, we examine the effect of the IS scheme on all multi-indices $\alpha \in \mathbb{N}^2$. To evaluate this, we analyse the following ratio:

$$\mathcal{R} := \frac{\text{Var}[\mathcal{A}_{\text{DLMC}}^{\text{IS}}]}{\text{Var}[\mathcal{A}_{\text{DLMC}}]}. \quad (4.127)$$

Here, $\mathcal{A}_{\text{DLMC}}^{\text{IS}}$ and $\mathcal{A}_{\text{DLMC}}$ denote the DLMC estimators of some quantity of interest (either $\mathbb{E}[G_{\alpha}]$ or $\mathbb{E}[\Delta G_{\alpha}]$) with and without IS respectively. Figure 4.10b presents a contour plot illustrating the values of the ratio \mathcal{R} for the DLMC estimator of $\mathbb{E}[\Delta G_{\alpha}]$ across various $\alpha \in \mathbb{N}^2$. Figure 4.10a depicts the same plot for the DLMC estimator of $\mathbb{E}[G_{\alpha}]$, representing the quantity of interest using discretization parameters defined by α . We observe that the variance reduction achieved by the IS control function ζ , obtained via SOC as explained in Section 4.4, diminishes with deeper multi-indices for the first-order




Figure 4.9: Numerical demonstration of variance reduction for the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x>K\}}(x)$ for $K = 3.5$: Squared coefficient of variation of the DLMC estimator of $\mathbb{E}[\Delta G_{(2,2)}]$ with and without IS with respect to the number of sample paths in the outer loop M_1 for fixed $M_2 = 10^2$. The proposed IS scheme reduces the overall variance $\text{Var}[\Delta G_{(2,2)}]$ by a factor of around 10, even though the proposed IS scheme minimises $\text{Var}[G_\alpha]$.

mixed differences estimator. However, as expected, the variance reduction for the DLMC estimator of $\mathbb{E}[G_\alpha]$ is consistently strong at all discretizations.

Figures 4.11 and 4.12 confirm Assumptions 4.7.1 and 4.7.2 while numerically determining the convergence rates $\{b_1, b_2\}$, $\{w_1, w_2\}$, and $\{s_1, s_2\}$. Figure 4.11a plots the convergence of $\mathbb{E}[\Delta G_\alpha]$, obtaining the rates $b_1 = 1$ and $b_2 = 1$. Additionally, Figures 4.11b and 4.11c display the convergence of $V_{1,\alpha}$ and $V_{2,\alpha}$, respectively, for the first-order mixed differences, yielding the rates $w_1 = 2$, $w_2 = 2$, $s_1 = 1.5$, and $s_2 = 2$. These convergence rates are compared to those of the level differences, $\mathbb{E}[\Delta G_\ell] = \mathbb{E}[G_\ell - G_{\ell-1}]$, used in the MLDLMC estimator (see Section 4.6). The first-order mixed differences of the MIDLMC estimator exhibit higher convergence rates than the level differences of the multi-level estimator.

Figure 4.12 additionally reveals that Assumptions 4.7.1 and 4.7.2 are satisfied for sufficiently fine discretizations. Using the rates in Figures 4.11 and 4.12, we can explicitly determine the optimal index set:

$$\mathcal{I}(L) = \left\{ \alpha \in \mathbb{N}^2 : \exp\left(\frac{2}{3}\alpha_1 + \frac{1}{3}\alpha_2\right) + \exp\left(\frac{2}{5}\alpha_1 + \frac{3}{5}\alpha_2\right) \leq L \right\}. \quad (4.128)$$

We insert the above rates in Theorem 4.7.1 and expect the work complexity of MIDLMC with the index set (4.128) to be $\mathcal{O}\left(\text{TOL}_r^{-2} \left(\log \text{TOL}_r^{-1}\right)^2\right)$. In addition, the condition (4.120) for MIDLMC is satisfied. For the corresponding multilevel convergence rates in Figure 4.11, we expect a work complexity of $\mathcal{O}\left(\text{TOL}_r^{-3}\right)$ for the MLDLMC estimator (see Theorem 4.6.1).

Figure 4.13b plots the index set (4.128) for different values of L . We compare this to





(a) Contour plot of \mathcal{R} for the DLMC estimator of $\mathbb{E}[G_\alpha]$.(b) Contour plot of \mathcal{R} for the DLMC estimator of $\mathbb{E}[\Delta G_\alpha]$.


Figure 4.10: Numerical demonstration of variance reduction in the DLMC estimator for the first-order mixed differences using the proposed IS scheme on the 1D Kuramoto model (4.12) with $g(x) = \tilde{\mathcal{I}}_{\{x > K\}}(x)$ for $K = 3.5$: Ratio \mathcal{R} indicates the factor of variance reduction due to the proposed IS scheme. Even though the proposed IS scheme minimises $\text{Var}[G_\alpha]$ as seen in Figure 4.10a, Figure 4.10b depicts significant variance reduction in the mixed differences for coarser multi-indices as well, although the reduction factor reduces as we go finer.



(a) Verifying Assumption 4.7.1 using DLMC Algorithm 10 with inputs $M_1 = 10^3$ and $M_2 = 10^3$.




(b) Verifying (4.92) from Assumption 4.7.2 using Algorithm 11 with inputs $M_1 = 10^2$ and $M_2 = 10^4$.




(c) Verifying (4.93) from Assumption 4.7.2 using Algorithm 11 with inputs $M_1 = 10^2$ and $M_2 = 10^4$.


Figure 4.11: Numerical rate verification for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}(x)$ for $K = 3.5$: All plots verify Assumptions 4.7.1 and 4.7.2 using the antithetic sampler (4.83) in the DLMC estimator for the first-order mixed differences. These plots imply the following convergence rates - $\{b_1, b_2\} = \{1, 1\}$, $\{w_1, w_2\} = \{2, 2\}$, and $\{s_1, s_2\} = \{1.5, 2\}$. These plots also demonstrate increased convergence rates of first-order mixed differences, when compared to level differences.



(a) Contour plots of sample mean $|\mathbb{E}[\Delta G_\alpha]|$ of the first-order mixed differences estimator using Algorithm 10 with inputs $M_1 = 10^3$ and $M_2 = 10^2$.




(b) Contour plots of sample $V_{1,\alpha}$ of the first-order mixed differences estimator using Algorithm 10 with inputs $M_1 = 10^2$ and $M_2 = 10^3$.




(c) Contour plots of sample $V_{2,\alpha}$ of the first-order mixed differences estimator using Algorithm 10 with inputs $M_1 = 10^2$ and $M_2 = 10^3$.

Figure 4.12: Numerical rate verification on the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}(x)$ for $K = 3.5$: All plots verify Assumptions 4.7.1 and 4.7.2 for all $\alpha \in \mathbb{N}^2$ using the antithetic sampler (4.83) in the DLMC estimator for the first-order mixed differences. These plots confirm the following convergence rates - $\{b_1, b_2\} = \{1, 1\}$, $\{w_1, w_2\} = \{2, 2\}$, and $\{s_1, s_2\} = \{1.5, 2\}$.



(a) Contour plot of numerical profits of index set (4.129).



(b) Contour plot of optimal index sets defined in (4.115).

Figure 4.13: Comparing actual and approximate profits of multi-indices in the MIDLMC estimator (4.82) for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}(x)$ for $K = 3.5$. These plots show that the optimal index set (4.128) is a good proxy for the true optimal index set (4.129).

the index set obtained with actual numerical profit computed as follows:

$$\mathcal{I}(v) = \left\{ \alpha \in \mathbb{N}^2 : \mathcal{P}_\alpha = \frac{\mathcal{E}_\alpha}{\omega_\alpha} \approx \frac{|\mathbb{E}[\Delta G_\alpha]|}{\sqrt{V_{1,\alpha} P_{\alpha_1}^2 N_{\alpha_2}} + \sqrt{V_{2,\alpha} P_{\alpha_1} N_{\alpha_2}}} \geq v \right\}. \quad (4.129)$$

We use DLMC estimates of $|\mathbb{E}[\Delta G_\alpha]|$, $V_{1,\alpha}$, $V_{2,\alpha}$ computed in Figure 4.12 to approximate the numerical profits in (4.129) to produce Figure 4.13a, revealing that (4.128) is a reasonable approximation to (4.129) for sufficiently fine discretizations.


Figures 4.14 to 4.16 depict the results of running the adaptive MIDLMC Algorithm 4 with IS and comparing them with the results from the adaptive MLDLMC Algorithm 3. We employ the parameter values $\{\bar{M}_1, \bar{M}_2\} = \{10^3, 10^2\}$ and $\{\tilde{M}_1, \tilde{M}_2\} = \{25, 100\}$. Algorithm 4 and Algorithm 3 are independently executed five times each for different prescribed relative tolerances TOL_r , and the combined results are plotted. Figure 4.14a displays the relative error achieved by the proposed MIDLMC estimator with IS for various prescribed relative tolerances TOL_r , estimated using a reference MIDLMC approximation with $\text{TOL}_r = 1\%$. Each marker represents a separate run of the corresponding adaptive algorithm. Figure 4.14a confirms that the multi-index and multilevel estimators satisfy the prescribed TOL_r , in the sense of (4.14). Figure 4.14b illustrates the maximum discretization level required for the number of particles and time steps in order for the multilevel and multi-index estimators to satisfy the relative bias constraint (4.86). Figure 4.14b reveals that the maximum number of particles is approximately the same for the multilevel and multi-index estimators.

Figure 4.15 presents a plot of the error estimates computed during the execution of Algorithm 4. Figure 4.15a displays the estimated relative bias at the final iteration of Algorithm 4, calculated using (4.123). Furthermore, Figure 4.15b provides the estimated relative statistical error of the estimator, $\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})$. As per the algorithm design, both relative errors should remain below $\frac{\text{TOL}_r}{2}$ for Algorithm 4 to terminate. Moreover, Figure 4.15b verifies that the number of samples chosen in (4.97) for the estimator, $\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})$, is optimal.


Figure 4.16a presents the average computational runtime of Algorithms 4 and 3 for various relative error tolerances. Additionally, Figure 4.16b displays the average computational cost of both estimators for different values of TOL_r , computed using (4.94). That is,

$$\text{Computational cost}[\mathcal{A}_{\text{MIDLMC}}(\mathcal{I})] \approx \sum_{\alpha \in \mathcal{I}(L)} (M_{1,\alpha} N_{\alpha_2} P_{\alpha_1}^2 + M_{1,\alpha} M_{2,\alpha} N_{\alpha_2} P_{\alpha_1}). \quad (4.130)$$

Figure 4.16 also highlights that the adaptive MLDLMC algorithm failed to produce estimates of the quantity of interest up to $\text{TOL}_r = 1\%$, whereas the MIDLMC algorithm achieved it within a fixed computational budget. Furthermore, Figures 4.16a and 4.16b provide further evidence supporting the superiority of the proposed MIDLMC estimator over the MLDLMC estimator. The multi-index estimator demonstrates a complexity reduction of approximately one order (up to logarithmic terms) compared to the multilevel estimator. Moreover, Figure 4.16 also indicates that the complexity rates derived in Theorem 4.7.1 are reasonably accurate.




(a) Relative error of estimators $\mathcal{A}_{\text{MIDLMC}}$ and $\mathcal{A}_{\text{MLDLMC}}$ with respect to prescribed relative tolerance TOL_r .




(b) Maximum discretization level of P and N for the estimators $\mathcal{A}_{\text{MIDLMC}}$ and $\mathcal{A}_{\text{MLDLMC}}$ with respect to prescribed relative tolerance TOL_r .

Figure 4.14: Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}(x)$ for $K = 3.5$. Figure 4.14a shows that both MIDLMC and MLDLMC estimators achieve a relative error below TOL_r in the sense of (4.14) to estimate a rare event quantity of around 2.04×10^{-5} . Figure 4.14b reveals that the MIDLMC estimator uses the same maximum P as the MLDLMC estimator, while using finer N discretization to achieve the same relative tolerance.




(a) Estimated relative bias of estimators $\mathcal{A}_{\text{MIDLMC}}$ and $\mathcal{A}_{\text{MLDLMC}}$ with respect to prescribed relative tolerance TOL_r .




(b) Estimated relative statistical error of estimators $\mathcal{A}_{\text{MIDLMC}}$ and $\mathcal{A}_{\text{MLDLMC}}$ with respect to prescribed relative tolerance TOL_r .

Figure 4.15: Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x > K\}}(x)$ for $K = 3.5$: Figures 4.15a and 4.15b demonstrate that the MIDLMC and MLDLMC estimators satisfy constraints (4.86) and (4.87), implying optimal choices of discretization parameters and sample numbers for both estimators.



(a) Average computational runtime of Algorithms 4 and 3 with respect to prescribed relative tolerance TOL_r .



(b) Average computational cost estimate of estimators \mathcal{A}_{MIDLDC} and \mathcal{A}_{MLDLDC} with respect to prescribed relative tolerance TOL_r .

Figure 4.16: Numerical results of applying Algorithms 3 and 4 for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{I}_{\{x>K\}}(x)$ for $K = 3.5$: Figures 4.16a and 4.16b show that \mathcal{A}_{MIDLDC} achieves the relative accuracy requirement (4.14) with a computational work of $\mathcal{O}\left(TOL_r^{-2} \left(\log TOL_r^{-1}\right)^2\right)$, one order lesser than that of \mathcal{A}_{MLDLDC} .

4.8 Conclusions and outlook

In Chapter 4, we used **SOC** to derive an optimal **IS** measure change for estimating rare event probabilities associated with **MV-SDEs**. We successfully combined this **IS** scheme with **MC**, **MLMC**, and multi-index **MC** methods for **MV-SDEs** to get computationally efficient **DLMC**, **MLDLMC**, and **MIDLMC** estimators of rare events. We showed both theoretically and numerically, under certain assumptions which could be verified numerically, the computational advantage of these **MC** estimators. The use of hierarchical sampling techniques improved the computational complexity rates from $\mathcal{O}(\text{TOL}_r^{-4})$ for the **DLMC** estimator to $\mathcal{O}(\text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2)$ for the **MIDLMC** estimator, while ensuring a bounded coefficient of variation for rare event estimation associated with the **1D** Kuramoto model (4.12). Table 4.2 summarises the achieved bias and variance convergence rates, and computational complexities of the various estimators introduced in this Chapter for estimating rare event probabilities associated with the **1D** Kuramoto model (4.12).

Extensions of this work could include the following.

1. It would be valuable to investigate the performance of the **IS** scheme and the various **MC** estimators for high-dimensional problems. Obtaining a good **IS** optimal control in high dimensions might become computationally infeasible using **FDM** to solve the optimal control **PDE** (4.24). One might then revert to model reduction techniques like Markovian projection [96] or machine learning based methods [64] to approximate the **IS** control function.
2. Due to poor variance convergence rates of the level differences [95] and first-order mixed differences [24] estimators in the **MLDLMC** and **MIDLMC** estimators of probabilities, we approximated the indicator function \mathcal{I} with smoothed versions $\tilde{\mathcal{I}}$ and $\tilde{\mathcal{J}}$ respectively. This creates an additional bias while estimating rare event probabilities, that we do not control in this work. One can revert to numerical smoothing techniques like adaptive sampling [97] or Brownian bridge conditioning [98] to avoid this additional bias, enabling more accurate estimates of rare event probabilities.
3. The combination of multilevel and multi-index **MC** with the proposed **IS** scheme explained in Sections 4.6 and 4.7, although still efficient, can be further optimised. By balancing the costs of the numerical **PDE** solver for **IS** and the sampling costs for the **DLMC** estimators, one can derive the optimal discretization parameters for the **PDE** solver and the empirical law that it is conditioned on. It would be interesting to investigate how much computational gain this would provide. Furthermore, creating unbiased estimators based on a simple randomization idea introduced in [99] could reduce the sources of error and help simplify the complexity analysis of our estimators.

While the primary objective in Chapter 4 was the accurate estimation of rare event probabilities, such probabilities also play a crucial role in constrained optimization problems in finance and engineering, where they appear as probabilistic constraints that must be satisfied. In Chapter 5, we extend our focus from the forward problem (estimation) to the control problem, formulating and solving an **SOC** problem in which a probability serves as a constraint.

	Bias	Variance	Complexity
DLMC	$ \mathbb{E}[g(\tilde{\mathbf{X}}(T))] - \mathbb{E}[g(\tilde{\mathbf{X}}^{P N}(T))] $ $= \mathcal{O}(P^{-1} + N^{-1})$	$V_1^{P N} = \mathcal{O}(P^{-1})$ $V_2^{P N} = \mathcal{O}(1)$	$\mathcal{O}(\text{TOL}_r^{-4})$
MLDLMC	$ \mathbb{E}[G - G_\ell] = \mathcal{O}(2^{-\ell})$	$V_{1,\ell} = \mathcal{O}(2^{-2\ell})$ $V_{2,\ell} = \mathcal{O}(2^{-\ell})$	$\mathcal{O}(\text{TOL}_r^{-3})$
MIDLMC	$ \mathbb{E}[\Delta\mathcal{G}_\alpha] = \mathcal{O}(2^{-\alpha_1 - \alpha_2})$	$V_{1,\alpha} = \mathcal{O}(2^{-2\alpha_1 - 2\alpha_2})$ $V_{2,\alpha} = \mathcal{O}(2^{-2\alpha_1 - 1.5\alpha_2})$	$\mathcal{O}(\text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2)$

Table 4.2: Summary of verified numerical rates and computational complexities of the various estimators for the 1D Kuramoto model (4.12) with $g(x) = \mathcal{J}_{\{x > K\}}(x)$. In the best case, the MIDLMC estimator has the same computational complexity as that of MC sampling without discretization errors (up to logarithmic terms). However, it requires smoothening of the indicator function and stricter regularity assumptions with respect to Assumptions 4.7.1 and 4.7.2, when compared to the DLMC and MLDLMC estimators.

Chapter 5

Chance-constrained optimal control of stochastic differential equations

In this chapter, we study **SOC** problems with a probability constraint. More precisely, we aim to solve a **SOC** problem where the probability that a function of the state and control stays below a threshold is high. We refer to this class of **SOC** problems as chance-constrained **SOC** problems. This study is motivated by an example from the telecommunications industry, where we seek an optimal power procurement schedule for green cellular wireless networks under uncertainty and chance constraints. Large parts of this Chapter have already been published in a similar form in [3] and are reproduced herein.

Contributions

- We propose a novel time-continuous optimization framework for power procurement in green cellular wireless networks, subject to **SDE** dynamics and **QoS** chance constraints. Compared to the discrete-time formulation in previous studies [39, 5, 40], the proposed approach decouples the model development from the numerical approximation, enhancing the model fidelity (see Remark 5.2.1). This formulation also yields a continuous control curve over time, allowing its application for any time discretization scheme and eliminating the need for ad hoc interpolations [43].
- We apply Lagrangian relaxation to the probabilistic **QoS** constraint, transforming the problem into a standard continuous-time **SOC** problem. Previous studies have explored numerical methods for continuous-time **SOC** with final-time chance constraints using Lagrangian relaxation [41, 100, 42] or reformulation as a stochastic target problem [101, 102]. However, the proposed approach is novel in addressing a chance constraint that must be satisfied at every time point, and is implemented within the context of cellular wireless networks.
- We develop an iterative algorithm to optimise the dual function within a finite-dimensional function class numerically. Each iteration involves solving the **HJB PDE** to compute the dual function value and its noisy subgradient. The proposed approach extends the work in [43] on the deterministic continuous-time optimization of coupled hydrothermal power systems to the stochastic setting.
- With access only to noisy realizations of the subgradient, we combine the **LMBM** method [103] with the **SSM** [104] to solve the dual optimization problem. We iteratively refine the Lagrange multiplier to ensure compliance with the chance constraint.

5.1 Problem formulation

Chance-constrained **SOC** is a subclass of **SOC** problems, first introduced in [28, 29]. This approach allows constraints to be violated with a certain prescribed probability or risk level. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space equipped with the filtration $\{\mathcal{F}_t\}$ and a d -dimensional $\{\mathcal{F}_t\}$ -adapted Wiener process \mathbf{W} . A chance-constrained **SOC** problem has the following elements.

1. **Time horizon:** We consider a fixed time horizon $[0, T]$ for some $T \in \mathbb{R}^+$.
2. **(Controlled) state process:** The controlled state process takes values in the state space $\mathbb{S} \subset \mathbb{R}^d$. Let $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{S}$ denote the state process given by the solution of the following controlled Itô **SDE**.

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}(t, \mathbf{X}(t), \Phi(t))dt + \mathbf{b}(t, \mathbf{X}(t))d\mathbf{W}(t), & t \in (0, T] \\ \mathbf{X}(0) = \mathbf{x}_0, \quad \mathbf{x}_0 \in \mathbb{R}^d \end{cases} \quad (5.1)$$

Here, the stochastic process Φ represents a control. In this thesis, we consider problems where the control process appears only in the drift function \mathbf{a} of the **SDE** (5.1).

3. **Control process:** The control process Φ in (5.1) is an $\{\mathcal{F}_t\}$ -adapted process that takes values in the control space $\mathbb{U} \subset \mathbb{R}^{\bar{d}}$, where \bar{d} is the dimension of the control process. In this work, we are particularly interested in the class of deterministic Markov controls, that is $\Phi(t) := \phi(t, \mathbf{X}(t))$, where $\phi : [0, T] \times \mathbb{S} \rightarrow \mathbb{U}$ is the Markov control function or policy.
4. **Cost function:** The objective is to minimise some cost functional of the following form.

$$J_{t, \mathbf{x}}(\phi) = \mathbb{E} \left[\int_t^T g(s, \mathbf{X}(s), \phi(s, \mathbf{X}(s))) ds + h(\mathbf{X}(T)) \mid \mathbf{X}(t) = \mathbf{x} \right]. \quad (5.2)$$

where $g : [0, T] \times \mathbb{S} \times \mathbb{U} \rightarrow \mathbb{R}$ is the running cost function and $h : \mathbb{S} \rightarrow \mathbb{R}$ the terminal cost function.

5. **Admissible controls:** We consider only controls that satisfy certain physical requirements that are expressed via *a.s.* constraints for all $t \in [0, T]$ of the following form.

$$\begin{aligned} f_{\text{eq}}(t, \mathbf{X}(t), \Phi(t)) &= \mathbf{0}, \\ f_{\text{in}}(t, \mathbf{X}(t), \Phi(t)) &\leq \mathbf{0}. \end{aligned} \quad (5.3)$$

Here $f_{\text{eq}} : [0, T] \times \mathbb{S} \times \mathbb{U} \rightarrow \mathbb{R}^{n_{\text{eq}}}$ and $f_{\text{in}} : [0, T] \times \mathbb{S} \times \mathbb{U} \rightarrow \mathbb{R}^{n_{\text{in}}}$ are vector valued convex functions. n_{eq} and n_{in} represent the number of equality and inequality *a.s.* constraints in the system. Let $\mathcal{A}(t, \mathbf{X}(t))$ denote the set of controls that satisfy the constraints (5.3) at time t for given state $\mathbf{X}(t)$. Subsequently, let \mathcal{A} denote the set of all admissible policies ϕ satisfying the constraints for all $t \in [0, T]$ and for all realizations of \mathbf{X} .

6. **Chance constraint:** The state and the control process must also satisfy a probability

constraint at all points in time of the following form

$$\mathbb{P} [\mathbf{f}_{\text{prob}}(t, \mathbf{X}(t), \Phi(t)) \geq \mathbf{f}_{\text{th}}] \leq \epsilon, \quad \forall t \in [0, T]. \quad (5.4)$$

$\mathbf{f}_{\text{prob}} : [0, T] \times \mathbf{S} \times \mathbf{U} \rightarrow \mathbb{R}$ denotes a real valued convex function, $\mathbf{f}_{\text{th}} \in \mathbb{R}$ denotes a threshold value that the function \mathbf{f}_{prob} must avoid going above with a probability of $1 - \epsilon$ for some $0 < \epsilon \ll 1$.

The chance constraint (5.4) can be viewed as a relaxation of the requirement to enforce the *a.s.* constraint $\mathbf{f}_{\text{prob}}(t, \mathbf{X}(t), \Phi(t)) \leq \mathbf{f}_{\text{th}}$ at all times $t \in [0, T]$ and for all realizations of $\mathbf{X}(t)$, which could be extremely costly or even impossible. Section 5.2 explains in detail why this approach is well-suited for optimization in wireless communication systems, where the associated QoS constraint is commonly defined as a chance constraint [105, 39, 5].

5.2 An example from wireless communications

There is a growing need for building robust and novel techniques for optimizing modern-day communication network systems. Since 2021, worldwide mobile broadband traffic has increased by an average of 19.6% annually, reaching 10^9 terabytes in 2023 [106]. The 4G network coverage increased from 41% of the world population in 2015 to 92% of the world population in 2024 [106]. The estimated transmission network electricity consumption in 2022 was 240 to 340 TWh, an increase of 64% from 2015 [107]. Next-generation transmission technology (e.g., 5G) requires even more power to ensure sufficient mobile network coverage [108]. A consequence of the high energy demand is increased greenhouse gas emissions. Data-transmission networks accounted for around 330 Mt CO₂-equivalent in 2020, around 0.9% of all energy-related emissions [107]. About 83% of the energy consumed by cellular networks came from fossil fuels in 2021, whereas only 9% came from renewables [109]. This problem motivates the coupling of cellular wireless networks with renewable energy to reduce both the costs and carbon footprint of the telecommunications industry. Hybrid energy systems provide a consistent power supply to wireless networks by combining renewable energy sources (e.g., solar and wind) [110]. However, modeling these systems and deriving an optimal power-procurement policy can be challenging, due to the inherent stochasticity of solar/wind power and wireless channels [111].

Previous studies have considered the optimal power-procurement problem in a discrete-time setting for general energy systems under an uncertain energy demand and incoming renewable energy [112, 113, 114]. The corresponding optimization problem for wireless networks has additional difficulties due to the uncertainty in wireless channels and probabilistic constraints due to the network QoS [115, 116, 117, 118]. The solution method for this problem in the current literature can be generally categorized into robust optimization and stochastic optimization. The robust optimization approach constructs a bounded uncertainty set encompassing all realizations of the stochastic variables and optimises the worst-case objective function value [35, 36, 37, 38]. Robust optimization, while ensuring feasibility under the worst-case scenario, often leads to overly conservative policies that could underutilize renewable energy resources. The stochastic optimization approach assumes a known probability distribution for the random variables and optimises the expected value of the objective function over the probability distribution of the uncertain parameters [119, 120]. The state-of-the-art stochastic optimization approach for wireless networks is chance-constrained optimization. The standard approach to solving discrete chance-constrained problems is to replace the chance constraint with a conservative but tractable approximation using a Bernstein inequality [33, 105], which provides a bound

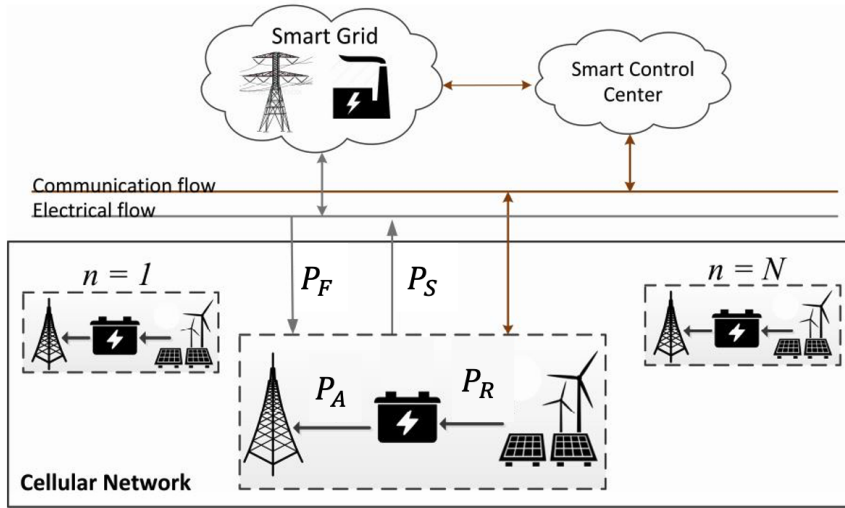


Figure 5.1: Schematic illustration of the power and information flow in a cellular wireless network with multiple base stations. Figure taken from [5].

on the tail probabilities. However, the computational complexity of the method in the discrete-time setting increases exponentially for large datasets and complex problems. For example, the authors in [39, 5, 40] solved the resulting discrete optimization problem using heuristic techniques, yielding suboptimal solutions. In this Chapter, we formulate and solve a finite horizon, continuous-time **SOC** problem to derive a near-optimal policy minimizing both the expected operating expenditure and environmental impact of the wireless network. The **SOC** problem is subject to demand constraints, a time-pointwise **QoS** chance constraint, and the stochastic dynamics and capacity limits of energy systems and wireless channels.

5.2.1 System model

This work considers a cellular network of multiple non-interacting base stations, each powered by renewable energy sources (such as solar panels or wind turbines) and battery storage. Figure 5.1 depicts the schematic network.

Figure 5.2 illustrates the electrical power flow in each base station, equipped with a battery that stores the power generated from renewables (P_R). The power transmitted by the base station to serve cellular users consists of power procured from the battery (P_A) and bought from the grid (P_F). Any extra energy stored in the battery could be sold back to the grid (P_S) for revenue. We aim to optimise the energy procurement at each base station for a one-day operation cycle.

Base station model

The instantaneous power procured by the mobile operator from the traditional grid to power the base station is denoted by $\{P_F(t) \in \mathbb{R}^+ : t \in [0, T]\}$, where 0 and T denote the time at the beginning and end of the current day, respectively. The corresponding unit price is denoted by $\{K_b(t) \in \mathbb{R}^+ : t \in [0, T]\}$. This price can vary during the day depending on strategies followed by stakeholders in the energy market. Moreover, the base station has its own internal source of power (i.e., renewable power generator) whose generated power is denoted by $\{P_R(t) \in \mathbb{R}^+ : t \in [0, T]\}$. This power is assumed to be free of charge.




Figure 5.2: Schematic illustration of the power flow in a base station in a cellular wireless network.

The internal power source is assumed to have a maximum power output capacity denoted by $\bar{P}_R \in \mathbb{R}^+$. The base station has a battery, allowing the mobile operator to store the incoming renewable power. As the base station is interconnected with the traditional grid, the mobile operator chooses to use the stored renewable power to run the base station or sell the extra energy back to the grid for revenue. The instantaneous power drawn from the battery to run the base station is denoted by $\{P_A(t) \in \mathbb{R}^+ : t \in [0, T]\}$. The instantaneous power from the battery sold back to the grid is denoted by $\{P_S(t) \in \mathbb{R}^+ : t \in [0, T]\}$, and the corresponding unit price is denoted by $\{K_s(t) \in \mathbb{R}^+ : t \in [0, T]\}$. We reasonably assume $K_s(t) \leq K_b(t)$ for all $t \in [0, T]$. The instantaneous power transmitted by the base station to serve its users is denoted by $\{P_{\text{tx}}^{\text{tot}}(t) \in \mathbb{R}^+ : t \in [0, T]\}$. We assume the total transmitted power is equally divided among each average mobile user connected to the network. The instantaneous power transmitted by the base station per average user is denoted by $\{P_{\text{tx}}(t) \in \mathbb{R}^+ : t \in [0, T]\}$ and $P_{\text{tx}}^{\text{tot}}(t) = N_u(t)P_{\text{tx}}(t)$ for all $t \in [0, T]$, where $\{N_u(t) \in \mathbb{R}^+ : t \in [0, T]\}$ is the number of connected mobile users at time t . The instantaneous power balance equation for the base station is

$$P_A(t) + P_F(t) = C_{\text{scal}}N_u(t)P_{\text{tx}}(t) + C_{\text{offset}}, \quad 0 \leq t \leq T, \quad (5.5)$$

where $C_{\text{scal}} \in \mathbb{R}^+$ represents a factor that scales with the transmitted power due to the amplifier and feeder losses in the network, and $C_{\text{offset}} \in \mathbb{R}^+$ models the offset power required to operate the base station irrespective of the transmitted power [5]. We assume that the base station can transmit a maximum power of $\bar{P}_{\text{tx}} \in \mathbb{R}^+$.

$$P_{\text{tx}}^{\text{tot}}(t) \leq \bar{P}_{\text{tx}}, \quad 0 \leq t \leq T. \quad (5.6)$$

The term $N_u(t)$ in (5.5) is determined by the daily mobile user traffic profile, heavily influencing the power demand in the base station.

Cellular network model

According to the power law, the power received by the user decays with increased distance from the base station. Moreover, the transmitted power is uncertain due to atmospheric conditions, natural and human-made environmental obstacles, and interference from other signals [116]. Hence, the instantaneous power received by a user u at position $\mathbf{x}_u(t) := [x_u(t), y_u(t)]^{\text{tr}} \in \mathbb{R}^2$ from a base station at $\mathbf{x}_{\text{BS}} := [x_{\text{BS}}, y_{\text{BS}}]^{\text{tr}} \in \mathbb{R}^2$ is given as $P_{\text{rx}}^u(t) := P_{\text{tx}}(t)\zeta(t)\kappa\|\mathbf{x}_u(t) - \mathbf{x}_{\text{BS}}\|^{-\eta}$, where $\kappa, \eta \in \mathbb{R}^+$ denotes the path loss constant and exponent, respectively. The stochastic process $\zeta : [0, T] \times \Omega \rightarrow [\underline{\zeta}, \infty)$ models a wireless fading channel, that is greater than $\underline{\zeta} \in \mathbb{R}^+$ a.s..

This work assumes that the channel follows Nakagami fading [121]. The Nakagami fading model is a good data-fit model for cellular systems in urban and suburban areas [122]. In this model, the instantaneous power is gamma distributed, a Pearson type III distribution [123]. The Pearson diffusion [124] naturally models the continuous-time Nakagami fading channel, where process ζ lies in $[\underline{\zeta}, \infty)$ for all $t \in [0, T]$. This model also has the advantage of analytical tractability, providing closed-form solutions for the invariant distribution of the process ζ and offering better insight into its behaviour. This approach makes the model choice computationally more efficient than other Markov-based models (e.g., discrete-state Markov chains), whose complexity increases quadratically as the number of possible states increases. Hence, the Pearson diffusion is employed to model the Nakagami fading channel. Thus, this work constructs an SDE for the process ζ (5.7) with a mean-reverting linear drift and a squared diffusion, which is linear in the state, and whose invariant distribution is the shifted-gamma distribution [124] with shape parameter $\mu > 1$ and scale parameter $\theta \in \mathbb{R}^+$. We consider the probability space $\{\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, P\}$, where \mathcal{F}_t is the filtration of the standard 1D Wiener process $W : [0, T] \times \Omega \rightarrow \mathbb{R}$:

$$\begin{cases} d\zeta(t) = -\theta(\zeta(t) - \underline{\zeta} - \mu) dt + \sqrt{2\theta(\zeta(t) - \underline{\zeta})} dW(t), & t > 0 \\ \zeta(0) \sim \mu_0^{\underline{\zeta}}, \end{cases} \quad (5.7)$$

where $\mu_0^{\underline{\zeta}}$ is the invariant distribution of the process ζ . The network performance is often characterized by the offered QoS, determined by the signal-to-noise ratio (SNR) threshold SNR_{th} . The SNR, in decibel (dB), at user u at $\mathbf{x}_u(t)$ served by the base station can be expressed as follows:

$$\text{SNR}^u(t) := 10 \log_{10} \left(\frac{P_{\text{rx}}^u(t)}{\sigma_0} \right) = 10 \log_{10} \left(\frac{P_{\text{tx}}(t)\zeta(t)\kappa\|\mathbf{x}_u(t) - \mathbf{x}_{\text{BS}}\|^{-\eta}}{\sigma_0} \right). \quad (5.8)$$

The ambient noise level is denoted by $\sigma_0 \in \mathbb{R}^+$. The QoS of the base station is determined by the proportion of users whose SNR is above SNR_{th} . A user u is defined to be in outage at time t if $\text{SNR}^u(t) < \text{SNR}_{\text{th}}$. The proportion of users in outage is denoted by $\{\phi_{\text{out}}(t) \in [0, 1] : t \in [0, T]\}$ and can be expressed as follows:

$$\phi_{\text{out}}(t) := \frac{1}{N_u(t)} \sum_{u=1}^{N_u(t)} \mathcal{I}_{\{\text{SNR}^u(t) < \text{SNR}_{\text{th}}\}}, \quad (5.9)$$

where $\phi_{\text{out}}(t)$ can be interpreted as an empirical average over a user distribution. The following approximation assumes that each user is identical and that sufficient number of

mobile users exist at all times:

$$\begin{aligned}
\phi_{\text{out}}(t) &= \frac{1}{N_u(t)} \sum_{u=1}^{N_u(t)} \mathcal{I}_{\{\text{SNR}^u(t) < \text{SNR}_{\text{th}}\}} \\
&= \frac{1}{N_u(t)} \sum_{u=1}^{N_u(t)} \mathcal{I}_{\left\{10 \log_{10} \left(\frac{P_{\text{tx}}(t) \zeta(t) \kappa \|\mathbf{x}_u(t) - \mathbf{x}_{\text{BS}}\|^{-\eta}}{\sigma_0} \right) < \text{SNR}_{\text{th}} \right\}} \\
&\approx \int_{\mathbb{R}^2} \mathcal{I}_{\left\{ \|\mathbf{z} - \mathbf{x}_{\text{BS}}\| > \left(\frac{P_{\text{tx}}(t) \zeta(t) \kappa}{\sigma_0 10^{-\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{1}{\eta}} \right\}} \rho_z(t) \mathrm{d}z, \tag{5.10}
\end{aligned}$$

where $\rho_z(t)$ denotes the distribution of mobile users at time t . This approximation is justified by the law of large numbers [49]: as the number of users grow, the empirical sum (5.9) over individual users converges to the expected value under the user distribution. In (5.10), $\phi_{\text{out}}(t)$ is the complementary CDF of the user distribution. This quantity can be computed (even analytically) for many well-known distributions (see Appendix D.2).

The objective of the network operator is to ensure $\phi_{\text{out}}(t) < \phi_{\text{th}}$ for all $t \in [0, T]$, where $0 < \phi_{\text{th}} \ll 1$ represents a QoS threshold ratio. Achieving this objective in an a.s. sense is infeasible due to possibly huge optimal costs. Instead, we introduce the following chance constraint for the network QoS.

$$\mathbb{P}[\phi_{\text{out}}(t) \geq \phi_{\text{th}}] \leq \epsilon, \quad 0 \leq t \leq T, \tag{5.11}$$

where $1 - \epsilon$, for $0 < \epsilon \ll 1$, is the confidence level of satisfying the constraint $\phi_{\text{out}}(t) < \phi_{\text{th}}$. The network operator also generates revenue from servicing the users throughout the day:

$$\mathcal{R}_1 := \mathbb{E} \left[\int_0^T K_{\text{net}}(t) N_u(t) (1 - \phi_{\text{out}}(t)) \mathrm{d}t \right], \tag{5.12}$$

where $\{K_{\text{net}}(t) \in \mathbb{R}^+ : t \in [0, T]\}$ denotes the unit price paid by users to connect to the network, and $N_u(t)(1 - \phi_{\text{out}}(t))$ represents the number of users that are connected to the network (not in outage).

Renewable power model

We apply a data-driven parametric SDE to model the stochastic renewable energy production, whose solution defines a stochastic process that models the error in a given forecast for renewable power [125]. Let the process $R : [0, T] \times \Omega \rightarrow [0, 1]$ denote the normalized generated renewable power.

$$\begin{cases} \mathrm{d}R(t) = (\dot{p}(t) - \theta(t)(R(t) - p(t))) \mathrm{d}t + \sqrt{2\alpha\theta_0 R(t)(1 - R(t))} \mathrm{d}W(t), & t > 0 \\ R(0) \sim \mu_0^R, \end{cases} \tag{5.13}$$

where $\{p(t) \in [0, 1] : t \in [0, T]\}$ is a deterministic forecast for the normalized power provided by an official source and $\{\dot{p}(t) : t \in [0, T]\}$ is its time derivative. In addition, $\{\theta(t) : t \in [0, T]\}$ denotes the mean-reversion time-varying parameter ensuring that the process R is the unique strong solution to (5.13) for all $t \in [0, T]$ with range $[0, 1]$,

a.s.. Moreover, parameters $\alpha, \theta_0 \in \mathbb{R}^+$ in the diffusion function, are inferred from discrepancies between historical real-world observations and their forecasts. In addition, $W : [0, T] \times \Omega \rightarrow \mathbb{R}$ is a standard 1D Wiener process. The drift function in (5.13) is designed such that the process reverts to its mean $p(t)$, with a time-varying speed $\theta(t)$ proportional to the deviation of the process $R(t)$ from its mean, and it tracks the time derivative $\dot{p}(t)$. The diffusion function in (5.13) is designed so that the process R avoids exiting from the range $[0, 1]$. The net renewable power generated is then given by $P_R(t) = \bar{P}_R R(t)$ for all $t \in [0, T]$.

Battery model

This work uses a simple model for the battery based on a circuit model of the state of the charge of a lithium-ion battery [126]. We assume that it has no energy loss when used and that no associated operational costs or aging effects exist. Process $A : [0, T] \rightarrow [0, 1]$ denotes the normalized charge in the battery:

$$\begin{cases} dA(t) = \frac{(P_R(t) - P_S(t) - P_A(t))}{\bar{A}} dt \\ A(0) = A_0 \\ -\underline{P}_A(A(t)) \leq P_A(t) + P_S(t) - P_R(t) \leq \bar{P}_A(A(t)), \end{cases} \quad (5.14)$$

where $\bar{A} \in \mathbb{R}^+$ represents the maximum battery charge capacity, \underline{P}_A denotes the maximum power the battery can absorb, and \bar{P}_A indicates the maximum power that the battery can supply. These quantities are characteristics of the battery. Generally, the battery can supply more power than it can absorb; hence, $\underline{P}_A < \bar{P}_A$ for all $t \in [0, T]$. In addition, $A_0 \in [0, 1]$ represents the normalized charge in the battery at the start of the day. The requirement that $0 \leq A(t) \leq 1$ for all $t \in [0, T]$ is enforced via constraints on charging and discharging at the extremes:

$$\begin{cases} P_R(t) - P_S(t) - P_A(t) \geq 0, & \text{if } A(t) = 0, \\ P_R(t) - P_S(t) - P_A(t) \leq 0, & \text{if } A(t) = 1. \end{cases} \quad (5.15)$$

The network operator generates revenue by selling some energy stored in the battery throughout the day:

$$\mathcal{R}_2 := \mathbb{E} \left[\int_0^T K_s(t) P_S(t) dt \right]. \quad (5.16)$$

While the battery model (5.14) is idealized, realistic battery efficiency losses can be included in the model by incorporating simple energy loss terms in the objective function of the optimal control problem. For example, this adjustment could include Ohmic efficiency losses [127] as a polynomial function of the battery discharge $(P_A(t) + P_S(t) - P_R(t))$, and state-of-charge losses [128] as an exponential function of the normalized battery charge A , which is outside the current scope of the thesis.

Grid power model

The network operator incurs the following cost of buying power from the traditional grid throughout the day:

$$\mathcal{C}_1 := \mathbb{E} \left[\int_0^T K_b(t) P_F(t) dt \right]. \quad (5.17)$$

The power from the traditional grid is primarily generated by fossil fuel stations, and

has a highly negative environmental effect. This work does not impose an upper limit on how much power could be bought from the grid. To restrict the green cellular network from buying substantial amounts of grid power, we also impose an additional "environmental" cost of buying fossil fuel-based power [129].

$$\mathcal{C}_2 := \mathbb{E} \left[\int_0^T (C_1 P_F(t) + C_2 P_F(t)^2) dt \right], \quad (5.18)$$

where $C_1 \in \mathbb{R}^+$ and $C_2 \in \mathbb{R}^+$ denote the pollutant emission coefficients of grid power.

Running horizon framework

Although this work aims to formulate an optimal power-procurement problem for a one-day operation cycle, such a short-term approach could lead to spurious results at the end of the day. For example, any optimal policy that minimises costs would have operators selling all stored energy at the end of the day to generate revenue. This scenario is unrealistic because the operator would prefer to keep some charge in the battery for the next day. This work devises a running horizon framework to avoid producing such impractical optimal policies and still work in a short-term finite horizon. This framework continuously solves a finite-horizon optimal control problem in a rolling manner, and consists of two facets. First, we add fictitious revenue from storing battery charge at the end of the day. More specifically, we introduce a third revenue source, as follows:

$$\mathcal{R}_3 := \mathbb{E} [P_K \bar{A}A(T)], \quad (5.19)$$

where $P_K \in \mathbb{R}^+$ represents the fictitious cost per unit stored battery charge. For example, publicly available deterministic forecasts of normalized renewable power and day-ahead energy price forecasts can be applied to determine a suitable value of P_K .

Next, we formulate the optimal control problem for a two-day operation cycle (instead of one) but apply the optimal power-procurement policy only for one day. This approach requires two-day-ahead forecasts of renewable power p , energy prices K_b, K_s , and cellular user demand N_{it} . The optimal power-procurement policy is updated daily, as and when the operators receive the two-day-ahead forecasts of these quantities. Regarding the problem formulation, the only change is that the time horizon is now denoted by $[0, T]$, where 0 denotes beginning of the day, and T denotes the end of the second day (instead of the end of the current day). Henceforth, this notation is applied while formulating the optimization problem.

5.2.2 Stochastic optimal control formulation

We gather all the elements of the SOC problem for a cellular base station.

1. **Time horizon:** $[0, T]$ where 0 denotes the start of the current day and T denotes the end of next day.
2. **(Controlled) state process:** The stochastic process $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{R}^3$ denotes the state process with the components $\mathbf{X} = [A, R, \zeta]$, A being the normalized state of battery charge, R being the normalized incoming renewable power, and ζ being the instantaneous wireless fading channel. Each component follows the SDE dynamics (5.14), (5.13), and (5.7) respectively. We combine these dynamics in the

form of the following 3-dimensional Itô SDE.

$$\begin{cases} d\mathbf{X}(t) = \mathbf{a}(t, \mathbf{X}(t), \Phi(t))dt + \mathbf{b}(t, \mathbf{X}(t))d\mathbf{W}(t), & t \in (0, T] \\ \mathbf{X}(0) \sim \mu_0^{\mathbf{X}}, \quad \mu_0^{\mathbf{X}} \in \mathcal{P}(\mathbb{R}^3) \end{cases} \quad (5.20)$$

Here the drift and diffusion functions are defined as follows.

$$\mathbf{a}(t, \mathbf{X}(t), \Phi(t)) = \begin{bmatrix} \frac{(\bar{P}_R R(t) - P_S(t) - P_A(t))}{\dot{p}(t) - \theta(t)(R(t) - p(t))} \\ -\theta(\xi(t) - \underline{\xi} - \mu) \end{bmatrix},$$

$$\mathbf{b}(t, \mathbf{X}(t)) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \sqrt{2\alpha\theta_0 R(t)(1-R(t))} & 0 \\ 0 & 0 & \sqrt{2\theta(\xi(t) - \underline{\xi})} \end{bmatrix}.$$

Moreover, \mathbf{W} is a standard 3-dimensional Wiener process and $\mu_0^{\mathbf{X}}$ is the initial distribution of the process \mathbf{X} belonging to $\mathcal{P}(\mathbb{R}^3)$, the class of probability distributions in \mathbb{R}^3 . The state space is $\mathcal{S} = [0, 1] \times [0, 1] \times [\underline{\xi}, \infty)$.

3. **Control process:** The control process is denoted by Φ and has 4 components $\Phi := [P_A, P_F, P_S, P_{\text{tx}}]$. Φ is an \mathcal{F}_t -adapted Markov control process, implying that $\Phi(t)$ is a deterministic function ϕ of the time t and the current state $\mathbf{X}(t)$. That is, $\Phi(t) := \phi(t, \mathbf{X}(t))$, where ϕ is the Markov control function or policy.
4. **Admissible controls:** The policy ϕ must satisfy the following equality a.s. constraint arising from the power balance equation (5.5) of the base station.

$$P_A(t) + P_F(t) - C_{\text{scal}} N_u(t) P_{\text{tx}}(t) - C_{\text{offset}} = 0, \quad 0 \leq t \leq T. \quad (5.21)$$

The policy ϕ must also satisfy the following inequality a.s. constraints arising from the capacity constraints of the transmission and energy systems introduced before.

$$-P_A(t) \leq 0, \quad (5.22)$$

$$-P_F(t) \leq 0, \quad (5.23)$$

$$-P_{\text{tx}}(t) \leq 0 \quad (5.24)$$

$$P_{\text{tx}}(t) - \frac{\bar{P}_{\text{tx}}}{N_u(t)} \leq 0, \quad (5.25)$$

$$-P_S(t) \leq 0, \quad (5.26)$$

$$\bar{P}_R R(t) - P_A(t) - P_S(t) - \underline{P}_A(A(t)) \leq 0, \quad (5.27)$$

$$P_A(t) + P_S(t) - \bar{P}_R R(t) - \bar{P}_A(A(t)) \leq 0. \quad (5.28)$$

Let $\mathcal{A}(t, \mathbf{X}(t))$ denote the set of controls that satisfy the constraints (5.21)-(5.28) at time t for given state $\mathbf{X}(t)$. Subsequently, let \mathcal{A} denote the set of all admissible policies ϕ satisfying the constraints for all $t \in [0, T]$ and for all realizations of \mathbf{X} .

$$\mathcal{A} := \left\{ \phi : \phi \text{ is } \mathcal{F}_t\text{-adapted and Markovian, } \phi(t, \mathbf{X}(t, \omega)) \in \mathcal{A}(t, \mathbf{X}(t, \omega)), \right. \\ \left. \forall \omega \in \Omega, \quad \forall t \in [0, T] \right\}. \quad (5.29)$$

5. **Chance constraint:** The state and the control process must satisfy the following network QoS constraint (5.11) at all points in time.

$$\mathbb{P}[\phi_{\text{out}}(t, \mathbf{X}(t), \boldsymbol{\phi}(t, \mathbf{X}(t))) \geq \phi_{\text{th}}] \leq \epsilon, \quad \forall t \in [0, T]. \quad (5.30)$$

where $0 < \epsilon \ll 1$ determines the confidence level of satisfying the QoS constraint (5.11). In modern-day cellular networks, base station operators require ϵ to be around 10^{-4} to 10^{-5} to ensure high levels of network performance. In general, ϵ is chosen to ensure that the probability of violating the QoS constraint is very low.

6. **Cost function:** The network operator seeks to minimise both operating costs and environmental penalty (5.18) by introducing a trade-off parameter $w \in [0, 1]$ to balance these objectives. The net expected cost to be minimised is then written as

$$\begin{aligned} \mathcal{U} &:= w(\mathcal{C}_1 - \mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + (1-w)\mathcal{C}_2 \\ &= \mathbb{E} \left[\int_0^T \left(w \left(K_b(t)P_F(t) - K_s(t)P_S(t) \right. \right. \right. \\ &\quad \left. \left. \left. - K_{\text{net}}(t)N_u(t)(1 - \phi_{\text{out}}(t, \boldsymbol{\zeta}(t), P_{\text{tx}}(t))) \right) \right. \right. \\ &\quad \left. \left. + (1-w)(C_1P_F(t) + C_2P_F(t)^2) \right) dt - wP_K \bar{A}A(T) \right] \\ &= \mathbb{E} \left[\int_0^T g(t, \mathbf{X}(t), \boldsymbol{\phi}(t, \mathbf{X}(t))) dt + h(\mathbf{X}(T)) \right], \end{aligned}$$

where we group all the running costs into the function g and the terminal costs into the function h . The objective is to minimise the following cost functional.

$$J_{t,\mathbf{x}}(\boldsymbol{\phi}) = \mathbb{E} \left[\int_t^T g(s, \mathbf{X}(s), \boldsymbol{\phi}(s, \mathbf{X}(s))) ds + h(\mathbf{X}(T)) \mid \mathbf{X}(t) = \mathbf{x} \right]. \quad (5.31)$$

Problem 5.2.1 (Primal problem). *Given current time t and current state $\mathbf{X}(t) = \mathbf{x}$, and deterministic forecasts for daily cellular user traffic profile $N_u(s)$, energy spot prices $K_b(s)$ and $K_s(s)$, renewable energy forecast $p(s)$, and mobile network usage price $K_{\text{net}}(s)$ for $t \leq s \leq T$, we solve the following:*

$$\boldsymbol{\phi}^* = \arg \min_{\boldsymbol{\phi} \in \mathcal{A}} J_{t,\mathbf{x}}(\boldsymbol{\phi}), \quad (5.32)$$

where the minimization is subject to the dynamics (5.7), (5.13), and (5.14) and the chance constraint (5.30) for the network QoS.

The chance constraint (5.30) does not allow us to define a value function that is a function of time t and state \mathbf{x} as done in (2.15) in Section 2.4. This is because constraint (5.30) controls the joint distribution of the Markovian dynamics and controls, and not each realization of the state/controls. The standard dynamic programming principle (see Section 2.4.2) does not apply if the value function additionally depends on the state's distribution. While it is theoretically possible to extend the principle to incorporate constraints on the state distribution, such an approach introduces significant challenges. From a computational perspective, tracking the full state distribution—or

even just its moments—leads to an exponential increase in the effective state space. Theoretically, the associated Hamiltonian may no longer be convex, and the existence of a unique minimum cannot be guaranteed. We overcome this problem using Lagrangian relaxation of the probabilistic constraint (5.30) for all $t \in [0, T]$, yielding a continuous-time Lagrangian relaxation approach. For this approach, we define the time-continuous deterministic Lagrange multiplier function $\lambda : [0, T] \rightarrow \mathbb{R}^+$, associated with the relaxed constraint (5.30). Using the classical Lagrangian relaxation technique for constrained optimization problems [130], we define the Lagrangian functional as follows:

$$\begin{aligned} \mathcal{L}_{t,\mathbf{x}}(\boldsymbol{\phi}, \lambda) &= J_{t,\mathbf{x}}(\boldsymbol{\phi}) + \int_t^T \lambda(s) (\mathbb{P}[\phi_{\text{out}}(s, \mathbf{X}(s), \boldsymbol{\phi}(s, \mathbf{X}(s))) \geq \phi_{\text{th}} \mid \mathbf{X}(t) = \mathbf{x}] - \epsilon) ds \\ &= J_{t,\mathbf{x}}(\boldsymbol{\phi}) + \mathbb{E} \left[\int_t^T \lambda(s) \left(\mathcal{I}_{\{\phi_{\text{out}}(s, \mathbf{X}(s), \boldsymbol{\phi}(s, \mathbf{X}(s))) \geq \phi_{\text{th}}\}} - \epsilon \right) ds \middle| \mathbf{X}(t) = \mathbf{x} \right]. \end{aligned} \quad (5.33)$$

Problem 5.2.2 (Relaxed problem). *Given current time t and current state $\mathbf{X}(t) = \mathbf{x}$, the Lagrange multiplier function λ , and deterministic forecasts for daily cellular user traffic profile $N_u(s)$, energy spot prices $K_b(s)$ and $K_s(s)$, renewable energy forecast $p(s)$, and mobile network usage price $K_{\text{net}}(s)$ for $t \leq s \leq T$, we solve the following:*

$$\boldsymbol{\phi}^*(\lambda) = \arg \min_{\{\boldsymbol{\phi} \in \mathcal{A}\}} \mathcal{L}_{t,\mathbf{x}}(\boldsymbol{\phi}, \lambda), \quad (5.34)$$

where the minimization is subject to the dynamics (5.7), (5.13), and (5.14).

For a given Lagrange multiplier function λ , this relaxation of the probabilistic constraint allows us to use standard SOC theory (Section 2.4) and define the value function $v : [0, T] \times \mathcal{S} \rightarrow \mathbb{R}$ that describes the value of the minimum possible Lagrangian for the system.

$$v(t, \mathbf{x}) = \min_{\boldsymbol{\phi} \in \mathcal{A}} \mathcal{L}_{t,\mathbf{x}}(\boldsymbol{\phi}, \lambda). \quad (5.35)$$

The goal of the relaxed SOC problem is to characterize the value function v and find a policy $\boldsymbol{\phi}^* \in \mathcal{A}$ whose Lagrangian (5.33) attains the minimum value for a given λ . That is $v(t, \mathbf{x}) = \mathcal{L}_{t,\mathbf{x}}(\boldsymbol{\phi}^*(\lambda), \lambda)$ for given λ and $\mathbf{X}(t) = \mathbf{x}$. We denote the components of $\mathbf{x} := [a, r, \chi]$, where a , r , and χ denote the variables corresponding to components of the state vector $A(t)$, $R(t)$, and $\xi(t)$, respectively. We apply the dynamic programming principle (see Section 2.4.2) to obtain the following second-order nonlinear HJB PDE which solves for the value function v in (5.35).

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + \mathcal{H}(t, \mathbf{x}, \nabla v, \nabla^2 v; \lambda) = 0, \\ v(T, \mathbf{x}) = -wP_K \bar{A}a, \quad \forall \mathbf{x} \in \mathcal{S}, \end{cases} \quad (5.36)$$

The Hamiltonian \mathcal{H} associated with Problem 5.2.2 in (5.36) is defined as follows:

$$\begin{aligned} \mathcal{H}(t, \mathbf{x}, \nabla v, \nabla^2 v; \lambda) &= \min_{\boldsymbol{\phi} \in \mathcal{A}(t, \mathbf{x})} \left[\frac{(\bar{P}_R r - P_S - P_A)}{\bar{A}} \frac{\partial v}{\partial a} + (\dot{p}(t) - \theta(t)(r - p(t))) \frac{\partial v}{\partial r} \right. \\ &\quad \left. + \alpha \theta_0 r (1 - r) \frac{\partial^2 v}{\partial r^2} - \theta (\chi - \underline{\xi} - \mu) \frac{\partial v}{\partial \chi} + \theta (\chi - \underline{\xi}) \frac{\partial^2 v}{\partial \chi^2} \right] \end{aligned}$$

$$+ g(t, \mathbf{x}, \boldsymbol{\phi}) + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}) > \phi_{\text{th}}\}} - \epsilon \right) \Big]. \quad (5.37)$$

We numerically solve PDE (5.36) for a given Lagrange multiplier function λ using an upwind FDM scheme. Section 5.3 details the numerical scheme. To get the optimal λ , we solve the dual problem associated with the primal Problem 5.2.1.

Problem 5.2.3 (Dual problem). *We determine*

$$\lambda^* = \arg \max_{\lambda \in \Lambda} \Theta(\lambda), \quad (5.38)$$

where the dual function $\Theta(\lambda)$ is given by

$$\Theta(\lambda) := \mathcal{L}_{0, \mathbf{x}(0)}(\boldsymbol{\phi}^*(\lambda), \lambda), \quad (5.39)$$

and $\boldsymbol{\phi}^*(\lambda)$ solves Problem 5.2.2 for a given λ . Λ denotes the set of all positive real-valued functions of time t .

Problem 5.2.3 is an infinite-dimensional optimization problem because we optimise over a set of functions Λ . To make this problem tractable, we discretize $\lambda(t)$ into a piecewise constant function to approximate the infinite-dimensional dual problem, allowing for practical numerical optimization. This approach implies satisfying the constraint (5.30) on average over a finite number of subintervals in $[0, T]$. That is, the integral of the pointwise violation of the constraint (5.30) over each subinterval must be zero. We consider the time discretization $0 = t_0 < t_1 < \dots < t_{\ell-1} < t_\ell = T$ of the time domain $[0, T]$. The Lagrange multiplier is approximated as follows:

$$\lambda(t) \approx \lambda^\ell(t) := \sum_{i=1}^{\ell} Y_i^\ell \mathcal{I}_{\{t \in [t_{i-1}, t_i]\}}(t), \quad \forall t \in [0, T], \quad (5.40)$$

where $\{Y_i^\ell\}_{i=1}^{\ell}$ denotes components of the vector of Lagrange multipliers \mathbf{Y}^ℓ . For the given time discretization, we define

$$\Lambda^\ell := \{\lambda^\ell : \text{with } \lambda^\ell(t) \text{ defined in (5.40) and } Y_i^\ell \in \mathbb{R}^+ \forall i \in [1, \dots, \ell]\}. \quad (5.41)$$

With this approximation, the finite-dimensional approximation of Problem 5.2.3 is formulated.

Problem 5.2.4 (Finite-dimensional dual problem). *For a given refinement level ℓ , we determine the following:*

$$\bar{\lambda}^\ell = \arg \max_{\lambda^\ell \in \Lambda^\ell} \Theta(\lambda^\ell), \quad (5.42)$$

where the dual function Θ is defined in (5.39).

The following Lagrangian is written for a given $\lambda^\ell \in \Lambda^\ell$:

$$\mathcal{L}_{0, \mathbf{x}(0)}(\boldsymbol{\phi}, \lambda^\ell) = J_{0, \mathbf{x}(0)}(\boldsymbol{\phi}) + \sum_{i=1}^{\ell} Y_i^\ell \left(\mathfrak{D}\Theta(\lambda^\ell) \right)_i \quad (5.43)$$

with

$$\left(\mathfrak{D}\Theta(\lambda^\ell) \right)_i := \mathbb{E} \left[\int_{t_{i-1}}^{t_i} \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}(t), \boldsymbol{\phi}(t, \mathbf{x}(t))) \geq \phi_{\text{th}}\}} - \epsilon \right) dt \right], \quad \forall i \in [1, \dots, \ell], \quad (5.44)$$

where $(\mathfrak{D}\Theta(\lambda^\ell))_i$ denotes the i^{th} component of an ℓ -dimensional subgradient vector $\mathfrak{D}\Theta$ of the dual function Θ evaluated at $\lambda^\ell \in \Lambda^\ell$. To solve Problem 5.2.4, the optimal ℓ -dimensional vector $\mathbf{Y}^\ell \in (\mathbb{R}^+)^{\ell}$ that maximises the dual function Θ must be found. To ensure a unique solution to Problem 5.2.4, we analyse its regularity and convexity. We begin by examining conditions for the relaxed problem 5.2.2 to have a unique solution. In the controlled SDE (5.20), the drift function \mathbf{a} is linear in both the state \mathbf{X} and control Φ , while the squared diffusion \mathbf{b}^2 is uncontrolled, independent across components, and quadratic in \mathbf{X} . The running cost g and terminal cost h remain bounded due to the physical constraints on Φ (Constraints (5.21) to (5.28)) and the finite state space \mathcal{S} . To ensure a minimum for the Hamiltonian (5.37), we verify that the admissible control set \mathbf{A} is defined by linear constraints in Φ , while constraints (5.27) and (5.28) are typically linear in \mathbf{X} and Φ . If ϕ_{out} is convex in ξ and P_{tx} , then the running cost g is convex. In the relaxed problem (5.33), the indicator function \mathcal{I} introduces discontinuities, making the existence of a minimum dependent on the function ϕ_{out} . For common mobile user traffic distributions such as those described in Appendix D.2, the corresponding ϕ_{out} ensures a convex Hamiltonian with a unique minimum. These conditions guarantee a unique solution to the relaxed Problem 5.2.2. Since λ is positive and the relaxed problem 5.2.2 has a unique solution, the dual function Θ is a point-wise minimum of an affine function in λ . Hence, the dual function is concave with a unique maximum [130]. A piecewise constant approximation of λ (5.40) leads to a non-smooth but still concave Θ . Hence, Problem 5.2.4 is a convex, ℓ -dimensional optimization problem that is not generally smooth. Therefore, we apply subgradient methods to solve Problem 5.2.4. Moreover, the choice of ℓ and the corresponding time discretization grid is set to control the pointwise violation in the chance constraint (5.30). Section 5.3 details this approach.

Remark 5.2.1 (Discrete vs. continuous-time formulation). *The optimal solution to a discrete-time formulation leads to a discrete set of control decisions at each time step, dependent on the selected discretization. Applying the dynamic programming principle to a discrete-time formulation leads to a stochastic optimization problem at each time step. Finer grids substantially increase the computational work, making the problem computationally infeasible without guaranteeing convergence of the optimal policies. In a continuous-time formulation, the limit of the dynamic programming principle exists as the time step becomes infinitesimally small, leading to the derivation of the associated HJB PDE. Numerical discretization methods are applied just once to solve this PDE and obtain a solution at all time points. This approach facilitates discretization error control, ensuring convergence to the true optimal solution, and enables the flexible application of adaptive and higher-order numerical approximation schemes.*

5.3 Numerical approach

5.3.1 Numerically solving the HJB PDE

We employ the explicit Euler upwind finite-difference scheme to numerically approximate the solution to the HJB equation in (5.36). This method provides a convergent numerical scheme under certain conditions that can be proven for the considered problem [131]. We consider the following time and space discretization of $(t, \mathbf{x}) \in [0, T] \times \mathcal{S}$ with the following finite-difference grid $\tau := \tau_t \times \tau_a \times \tau_r \times \tau_\chi$:

$$\tau_t : 0 = t_0 < t_1 < \dots < t_{N_t-1} < t_{N_t} = T, \quad \Delta t = \frac{T}{N_t}, \quad (5.45)$$

$$\begin{aligned}\tau_a : 0 = a_0 < a_1 < \dots < a_{N_1-1} < a_{N_1} = 1, \quad \Delta a = \frac{1}{N_1}, \\ \tau_r : 0 = r_0 < r_1 < \dots < r_{N_2-1} < r_{N_2} = 1, \quad \Delta r = \frac{1}{N_2}, \\ \tau_\chi : 0 = \chi_0 < \chi_1 < \dots < \chi_{N_3-1} < \chi_{N_3} = 1, \quad \Delta \chi = \frac{1}{N_3}.\end{aligned}$$

For the dimension χ , we approximate the domain $[\underline{\xi}, \infty)$ by $[\underline{\xi}, \bar{\chi}]$. The upper limit $\bar{\chi}$ is chosen such that $\mathbb{P}[\xi(t) \geq \bar{\chi}]$ is negligible for all $t \in [0, T]$. We then scale $[\underline{\xi}, \bar{\chi}]$ to $[0, 1]$.

$$\begin{aligned}t_n &= n\Delta t, \quad n = 0, \dots, N_t, \\ a_i &= i\Delta a, \quad i = 0, \dots, N_1, \\ r_j &= j\Delta r, \quad j = 0, \dots, N_2, \\ \chi_k &= k\Delta \chi, \quad k = 0, \dots, N_3.\end{aligned}$$

The upwind FDM involves using the following finite differences to approximate the derivatives in (5.36).

$$\begin{aligned}\Delta_t^- v(t, \mathbf{x}) &= \frac{v(t, \mathbf{x}) - v(t - \Delta t, \mathbf{x})}{\Delta t} = \frac{\partial v}{\partial t}(t, \mathbf{x}) + \mathcal{O}(\Delta t), \\ \Delta_a^+ v(t, \mathbf{x}) &= \frac{v(t, [a + \Delta a, r, \chi]) - v(t, [a, r, \chi])}{\Delta a} = \frac{\partial v}{\partial a}(t, \mathbf{x}) + \mathcal{O}(\Delta a), \\ \Delta_a^- v(t, \mathbf{x}) &= \frac{v(t, [a, r, \chi]) - v(t, [a - \Delta a, r, \chi])}{\Delta a} = \frac{\partial v}{\partial a}(t, \mathbf{x}) + \mathcal{O}(\Delta a), \\ \Delta_r^+ v(t, \mathbf{x}) &= \frac{v(t, [a, r + \Delta r, \chi]) - v(t, [a, r, \chi])}{\Delta r} = \frac{\partial v}{\partial r}(t, \mathbf{x}) + \mathcal{O}(\Delta r), \\ \Delta_r^- v(t, \mathbf{x}) &= \frac{v(t, [a, r, \chi]) - v(t, [a, r - \Delta r, \chi])}{\Delta r} = \frac{\partial v}{\partial r}(t, \mathbf{x}) + \mathcal{O}(\Delta r), \\ \Delta_\chi^+ v(t, \mathbf{x}) &= \frac{v(t, [a, r, \chi + \Delta \chi]) - v(t, [a, r, \chi])}{\Delta \chi} = \frac{\partial v}{\partial \chi}(t, \mathbf{x}) + \mathcal{O}(\Delta \chi), \\ \Delta_\chi^- v(t, \mathbf{x}) &= \frac{v(t, [a, r, \chi]) - v(t, [a, r, \chi - \Delta \chi])}{\Delta \chi} = \frac{\partial v}{\partial \chi}(t, \mathbf{x}) + \mathcal{O}(\Delta \chi), \\ \Delta_r^2 v(t, \mathbf{x}) &= \frac{v(t, [a, r + \Delta r, \chi]) - 2v(t, [a, r, \chi]) + v(t, [a, r - \Delta r, \chi])}{\Delta r^2} \\ &= \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) + \mathcal{O}(\Delta r^2), \\ \Delta_\chi^2 v(t, \mathbf{x}) &= \frac{v(t, [a, r, \chi + \Delta \chi]) - 2v(t, [a, r, \chi]) + v(t, [a, r, \chi - \Delta \chi])}{\Delta \chi^2} \\ &= \frac{\partial^2 v}{\partial \chi^2}(t, \mathbf{x}) + \mathcal{O}(\Delta \chi^2).\end{aligned}\tag{5.46}$$

We introduce the following notation for some real-valued function f .

$$f^+(t, \mathbf{x}, \boldsymbol{\phi}) = \max(f(t, \mathbf{x}, \boldsymbol{\phi}), 0), \quad f^-(t, \mathbf{x}, \boldsymbol{\phi}) = \max(-f(t, \mathbf{x}, \boldsymbol{\phi}), 0)\tag{5.47}$$

We introduce the following notation to indicate the evaluation of the following functions on the numerical grid τ (5.45).

$$v_{(i,j,k)}^n := v(t_n, [a_i, r_j, \chi_k]) = v(n\Delta t, [i\Delta a, j\Delta r, k\Delta \chi]),\tag{5.48}$$

$$\begin{aligned}
\Delta_a^+ v |_{(i,j,k)}^n &:= \frac{v_{(i+1,j,k)}^n - v_{(i,j,k)}^n}{\Delta a}, \\
\Delta_a^- v |_{(i,j,k)}^n &:= \frac{v_{(i,j,k)}^n - v_{(i-1,j,k)}^n}{\Delta a}, \\
\mathbf{a} |_{(i,j,k)}^n &:= \mathbf{a}(t_n, [a_i, r_j, \chi_k], (\boldsymbol{\phi}^*)_{(i,j,k)}^n), \\
\mathbf{b} |_{(i,j,k)}^n &:= \mathbf{b}(t_n, [a_i, r_j, \chi_k]), \\
g^\lambda |_{(i,j,k)}^n &:= g\left(t_n, [a_i, r_j, \chi_k], (\boldsymbol{\phi}^*)_{(i,j,k)}^n\right) + \lambda(t_n) \left(\mathcal{I}_{\{\phi_{\text{out}}(t_n, [a_i, r_j, \chi_k], (\boldsymbol{\phi}^*)_{(i,j,k)}^n) > \phi_{\text{th}}\}} - \epsilon\right).
\end{aligned}$$

Optimal controls at grid point $(t_n, [a_i, r_j, \chi_k])$, denoted by $(\boldsymbol{\phi}^*)_{(i,j,k)}^n$, are computed using the known derivatives of the numerical value function as follows:

$$\begin{aligned}
(\boldsymbol{\phi}^*)_{(i,j,k)}^n &= \arg \min_{\boldsymbol{\phi} \in \mathbf{A}(t_n, [a_i, r_j, \chi_k])} \left[a_1^+(t_n, [a_i, r_j, \chi_k], \boldsymbol{\phi}) \Delta_a^+ v |_{(i,j,k)}^n \right. \\
&\quad + a_1^-(t_n, [a_i, r_j, \chi_k], \boldsymbol{\phi}) \Delta_a^- v |_{(i,j,k)}^n + g(t_n, [a_i, r_j, \chi_k], \boldsymbol{\phi}) \\
&\quad \left. + \lambda(t_n) \left(\mathcal{I}_{\{\phi_{\text{out}}(t_n, [a_i, r_j, \chi_k], \boldsymbol{\phi}) > \phi_{\text{th}}\}} - \epsilon\right) \right]. \tag{5.49}
\end{aligned}$$

Then, the explicit update rule for the upwind finite-difference scheme is given as follows:

$$\begin{aligned}
v_{(i,j,k)}^{n-1} &= v_{(i,j,k)}^n \left(1 - \left(b_{2,2} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{\Delta r^2} - \left(b_{3,3} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{\Delta \chi^2} \right. \\
&\quad \left. - \frac{\Delta t}{\Delta a} |a_1 |_{(i,j,k)}^n| - \frac{\Delta t}{\Delta r} |a_2 |_{(i,j,k)}^n| - \frac{\Delta t}{\Delta \chi} |a_3 |_{(i,j,k)}^n| \right) \\
&\quad + v_{(i+1,j,k)}^n \frac{\Delta t}{\Delta a} a_1^+ |_{(i,j,k)}^n + v_{(i-1,j,k)}^n \frac{\Delta t}{\Delta a} a_1^- |_{(i,j,k)}^n \\
&\quad + v_{(i,j+1,k)}^n \left(\left(b_{2,2} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{2\Delta r^2} + a_2^+ |_{(i,j,k)}^n \frac{\Delta t}{\Delta r} \right) \\
&\quad + v_{(i,j-1,k)}^n \left(\left(b_{2,2} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{2\Delta r^2} + a_2^- |_{(i,j,k)}^n \frac{\Delta t}{\Delta r} \right) \\
&\quad + v_{(i,j,k+1)}^n \left(\left(b_{3,3} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{2\Delta \chi^2} + a_3^+ |_{(i,j,k)}^n \frac{\Delta t}{\Delta \chi} \right) \\
&\quad + v_{(i,j,k-1)}^n \left(\left(b_{3,3} |_{(i,j,k)}^n\right)^2 \frac{\Delta t}{2\Delta \chi^2} + a_3^- |_{(i,j,k)}^n \frac{\Delta t}{\Delta \chi} \right) + \Delta t g^\lambda |_{(i,j,k)}^n. \\
v_{(i,j,k)}^{N_t} &= -w P_K \bar{A} a_i, \quad \forall i, j, k.
\end{aligned} \tag{5.50}$$

The convergence of the upwind FDM numerical scheme (5.50) is analysed in Section C.2.1 in Appendix C.2, establishing conditions on functions \mathbf{a} , \mathbf{b} , g that guarantees convergence. In general, boundary conditions must be imposed on the HJB PDE (5.36) to obtain a unique numerical solution. However, boundary conditions at the boundaries $a = 0, 1$, $r = 0, 1$, and $\chi = \underline{\chi}$ are unnecessary because the components of the drift function \mathbf{a} naturally have a sign at the boundaries ensuring the use of an interior point to approximate first-order derivatives. Similarly, components of the diffusion function \mathbf{b} are naturally 0 at these boundaries, ensuring no exterior points are necessary to approximate second-order derivatives. At the boundary $\chi = \bar{\chi}$, we impose a nonreflective boundary condition to indicate that the actual domain is not bounded at $\chi = \bar{\chi}$. Refer to Section C.2.2 in

Appendix C.2 for more details about boundary conditions. We impose the Courant–Friedrichs–Lewy (CFL) condition at each grid point in τ to ensure the stability of the numerical scheme.

$$\Delta t \left(\frac{|a_1 |_{(i,j,k)}^n|}{\Delta a} + \frac{|a_2 |_{(i,j,k)}^n|}{\Delta r} + \frac{|a_3 |_{(i,j,k)}^n|}{\Delta \chi} + \frac{(b_{2,2} |_{(i,j,k)}^n)^2}{\Delta r^2} + \frac{(b_{3,3} |_{(i,j,k)}^n)^2}{\Delta \chi^2} \right) \leq 1. \quad (5.51)$$

Section C.2.1 in Appendix C.2 derives the CFL condition in (5.51). Although the HJB PDE (5.36) is nonlinear, the numerical scheme given by Eqs. (5.49) to (5.50) is explicit because the value function v at time t_{n-1} is completely determined by its value and derivatives at time t_n . However, this relationship is nonlinear and cannot be vectorized. Moreover, the constrained optimization problem in (5.49) must be solved at every grid point in τ . Depending on the structure of $\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi})$, (5.49) can be solved analytically or numerically. Linear interpolation is applied to extend the numerical solution from the grid τ to the entire domain $[0, T] \times \mathbb{S}$. Next, the dual function $\Theta(\lambda)$ is approximated by $\bar{\Theta}(\lambda) := v(0, \mathbf{X}(0))$ for a given λ . Algorithm 13 in Appendix B.7 presents a pseudo-algorithm of the numerical solver.

Remark 5.3.1 (Discretization error of the PDE solver). *The upwind FDM scheme is a first-order explicit scheme with a discretization error of $\mathcal{O}(\Delta t + \Delta a + \Delta r + \Delta \chi)$ (see Section C.2.1 in Appendix C.2). When solving the overall optimization problem up to a relative tolerance of TOL with respect to the dual function value $\Theta(\lambda)$, the PDE discretization error must be $\mathcal{O}(|\Theta(\lambda)|\text{TOL})$. Thus, $\Delta a, \Delta r, \Delta \chi$ must be $\mathcal{O}(|\Theta(\lambda)|\text{TOL})$, so that the work per time step is $\mathcal{O}(|\Theta(\lambda)|^{-3}\text{TOL}^{-3})$. Although the upwind FDM scheme uses a first-order time stepping scheme, the CFL condition (5.51) for stability requires $\Delta t = \mathcal{O}(\Delta r^2 + \Delta \chi^2)$, implying the computational work of $\mathcal{O}(|\Theta(\lambda)|^{-5}\text{TOL}^{-5})$ to solve the HJB PDE once.*

5.3.2 Estimating subgradients

Upon solving the HJB PDE (5.36) numerically and obtaining an approximate value function v in the domain $[0, T] \times \mathbb{S}$, we can numerically simulate optimally controlled paths of the state variable \mathbf{X} forward in time, from the initial condition $\mathbf{X}(0)$ using dynamics (5.20). We use the EM time discretization of the SDEs in (5.20) (see Definition 2.3.4). We consider the discretization $0 = \bar{t}_0 < \bar{t}_1 < \dots < \bar{t}_{\bar{N}_t} = T$ of the time domain $[0, T]$ with \bar{N}_t time steps. We denote the EM time-discretized version of the controlled state process \mathbf{X} by $\bar{\mathbf{X}}^{\bar{N}_t}$. The optimal controls $\boldsymbol{\phi}^*(\bar{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n))$ at each EM timestep are obtained by minimizing (5.49) at the current time \bar{t}_n and state $\bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n)$, which may be a point outside the grid τ . The values of the numerical derivatives $\Delta_a^\pm v$ (5.46) of the value function v at those points are linearly interpolated from the computed derivative values at the nearest points in grid τ . Although the dynamics (5.20) are designed to remain in the state space \mathbb{S} a.s., the EM discretization may occasionally produce trajectories that exit \mathbb{S} with low probability. To address such cases, the approximate SDE paths are reflected at the boundaries of the state space. The subgradient of Θ at $\lambda^\ell \in \Lambda^\ell$ is computed according to (5.44), approximating the expected value using MC sampling with M_{SG} iid sample paths of the optimally controlled process $\bar{\mathbf{X}}^{\bar{N}_t}$.

Moreover, the integral in (5.44) is approximated using a forward Euler approximation. For each $i \in [1, 2, \dots, \ell]$, we consider the discretization $t_{i-1} = \bar{t}_0 < \bar{t}_1 < \dots < \bar{t}_{\bar{N}_t} = t_i$ of

the time domain $[t_{i-1}, t_i]$, corresponding to the i^{th} component of the subgradient (5.44), with \tilde{N}_t uniform time steps. It follows that $\tilde{t}_n := n \times \Delta \tilde{t}$ for all $n = 0, 1, \dots, \tilde{N}_t$, and $\Delta \tilde{t} := \frac{t_i - t_{i-1}}{\tilde{N}_t}$. The numerical approximation of the i^{th} component of the ℓ -dimensional subgradient (5.44) is denoted by $(\tilde{\mathfrak{D}}\Theta(\lambda^\ell))_i$ and given by:

$$\left(\tilde{\mathfrak{D}}\Theta(\lambda^\ell)\right)_i = \frac{1}{M_{\text{SG}}} \sum_{m=1}^{M_{\text{SG}}} \sum_{n=0}^{\tilde{N}_t} \left(\mathcal{I}_{\{\phi_{\text{out}}(\tilde{t}_n, \tilde{\mathbf{X}}^{\tilde{N}_t}(\tilde{t}_n, \omega^{(m)})), \phi(\tilde{t}_n, \tilde{\mathbf{X}}^{\tilde{N}_t}(\tilde{t}_n, \omega^{(m)}))\} \geq \phi_{\text{th}}\}} - \epsilon \right) \Delta \tilde{t}, \quad (5.52)$$

where $\omega^{(m)}$ denotes the m^{th} iid realization of the random variables required to generate sample paths of $\tilde{\mathbf{X}}^{\tilde{N}_t}$. The time discretization for subgradient computation and the Lagrange multiplier function λ^ℓ generally need not coincide. In this case, Brownian bridge interpolation is necessary to evaluate the optimally controlled paths $\tilde{\mathbf{X}}^{\tilde{N}_t}$ at the time discretization points $\{\tilde{t}_n\}_{n=0}^{\tilde{N}_t}$ for each $i \in [1, \dots, \ell]$. The same M_{SG} realisations of the optimally controlled path $\tilde{\mathbf{X}}^{\tilde{N}_t}$ are used to estimate all ℓ components of $\tilde{\mathfrak{D}}\Theta(\lambda^\ell)$. Algorithm 14 in Appendix B.8 presents the corresponding pseudo-algorithm. (5.52) is a MC approximation of (5.44). It is subject to noisy statistical error controlled in a probabilistic sense by the choice of M_{SG} using the CLT (see Proposition 2.3.5). Thus, we only have access to noisy subgradients of the dual function. Hence, we apply the SSM to solve the convex, non-smooth optimization problem 5.2.4.

Remark 5.3.2 (Subgradient estimation error). *The error in the subgradient estimation consists of two parts: a statistical MC error from approximating the expected value in (5.44) using a sample average with M_{SG} samples and a discretization error from approximating the integral in the expectation in (5.44) using a forward Euler summation with \tilde{N}_t steps. From the CLT, the MC approximation error is $\mathcal{O}\left(M_{\text{SG}}^{-\frac{1}{2}}\right)$. The discretization error for the first-order forward Euler summation is $\mathcal{O}\left(\tilde{N}_t^{-1}\right)$, yielding a total error in the subgradient estimation of $\mathcal{O}\left(M_{\text{SG}}^{-\frac{1}{2}} + \tilde{N}_t^{-1}\right)$. When solving the overall optimization problem up to a relative tolerance of TOL with respect to the subgradient, we must estimate the subgradients up to an absolute tolerance of ϵTOL . This requirement implies that $M_{\text{SG}} = \mathcal{O}\left(\epsilon^{-2}\text{TOL}^{-2}\right)$ and $\tilde{N}_t = \mathcal{O}\left(\epsilon^{-1}\text{TOL}^{-1}\right)$, yielding a total computational work of $\mathcal{O}\left(\epsilon^{-3}\text{TOL}^{-3}\right)$ for each subgradient estimation.*

5.3.3 Dual optimization

After approximating the dual function $\Theta(\lambda^\ell)$ and its subgradient $\mathfrak{D}\Theta(\lambda^\ell)$ for a given Lagrange multiplier function $\lambda^\ell \in \Lambda^\ell$, we must solve the non-smooth convex dual optimization problem in Problem 5.2.4. We solve this in two stages: (i) by constructing an approximation λ^ℓ , of λ (as in (5.40)), and (ii) by numerically optimizing the amplitudes \mathbf{Y}^ℓ of the approximated function λ^ℓ . Algorithm 5 illustrates the overall procedure.

Lagrange multiplier refinement

We construct a finite-dimensional piecewise constant approximation, λ^ℓ , of the Lagrange multiplier function, λ , as defined in (5.43). Increasing the refinement level ℓ enhances the approximation quality of λ by λ^ℓ , but also increases the optimization dimension (also ℓ). This work determines a sufficient ℓ that controls the violation of the relaxed chance

Algorithm 5: Numerical dual optimization procedure

Input: TOL, TOL_{init}, k_{\max} , \bar{N}_{iter} , N_{iter} , $Y_1^1 = 1$,
 $\beta_F, M_{\text{SG}}, N_t, N_1, N_2, N_3, \bar{N}_t, \tilde{N}_t, \mathbf{X}(0)$

Output: Optimal controls ϕ^* , optimal Lagrange multiplier function $\lambda^\ell(t)$

Construct $\lambda^1(t)$ with Y_1^1 using (5.40);

Obtain \tilde{Y}_1^1 using initialization Algorithm 15 with inputs TOL_{init}, $Y_1^1 = 1$,
 $\beta_F, M_{\text{SG}}, N_t, N_1, N_2, N_3, \bar{N}_t, \mathbf{X}(0)$;

Construct $\lambda^1(t)$ with \tilde{Y}_1^1 using (5.40);

Obtain \hat{Y}_1^1 with the LMBM routine [7] with starting point \tilde{Y}_1^1 , number of iterations
 N_{iter} and parameters specified in Appendix D.3;

Construct $\lambda^1(t)$ with \hat{Y}_1^1 using (5.40);

while $\ell < \bar{N}_t$ **do**

$\ell \leftarrow 2\ell$;

Compute $\bar{\Theta}(\lambda^\ell) = v(0, \mathbf{X}(0))$ by solving (5.36) using Algorithm 13 with $\lambda^\ell(t)$
and parameters N_t, N_1, N_2, N_3 ;

Estimate $\bar{\mathfrak{D}}\Theta(\lambda^\ell)$ using Algorithm 14 with parameters $\bar{N}_t, M_{\text{SG}}, \tilde{N}_t$;

$k = 0$;

while $k < k_{\max}$ and $\|\bar{\mathfrak{D}}\Theta(\lambda^\ell)\| > \text{TOL}\epsilon$ **do**

$$\begin{cases} \mathbf{Y}^\ell \leftarrow \mathbf{Y}^\ell + C_{\text{SSM}} \frac{\bar{\mathfrak{D}}\Theta(\lambda^\ell)}{\|\bar{\mathfrak{D}}\Theta(\lambda^\ell)\|}, & k \leq \bar{N}_{\text{iter}} \\ \mathbf{Y}^\ell \leftarrow \mathbf{Y}^\ell + \frac{C_{\text{SSM}}}{k+1} \frac{\bar{\mathfrak{D}}\Theta(\lambda^\ell)}{\|\bar{\mathfrak{D}}\Theta(\lambda^\ell)\|}, & k > \bar{N}_{\text{iter}} \end{cases};$$

Construct $\lambda^\ell(t)$ with \mathbf{Y}^ℓ using (5.40);

Compute and store $\omega^{(k)} := \bar{\Theta}(\lambda^\ell) = v(0, \mathbf{X}(0))$ by solving (5.36) using
Algorithm 13 with $\lambda^\ell(t)$ and parameters N_t, N_1, N_2, N_3 ;

Estimate $\bar{\mathfrak{D}}\Theta(\lambda^\ell)$ using Algorithm 14 with parameters $\bar{N}_t, M_{\text{SG}}, \tilde{N}_t$;

$k \leftarrow k + 1$;

end

Save \mathbf{Y}^ℓ corresponding to $\max\{\omega^{(1)}, \dots, \omega^{(k)}\}$;

Construct $\lambda^{2\ell}(t)$ with \mathbf{Y}^ℓ using (5.40);

end

constraint (5.30). We start with $\ell = 1$ and keep uniformly doubling it, until the violation of the constraint (5.30) is sufficiently controlled.

Numerical optimization

We solve the optimization of $\Theta(\lambda^\ell)$ for $\lambda^\ell \in \Lambda^\ell$ with respect to the amplitudes $\mathbf{Y}^\ell \in (\mathbb{R}^+)^{\ell}$ as defined in Problem 5.2.4. For a given evaluation point λ^ℓ , we have access to the dual function value (approximated in Section 5.3.1) and a noisy subgradient (approximated in (5.3.2)). Since $\Theta(\lambda^\ell)$ is a non-smooth function of \mathbf{Y}^ℓ , the SSM [104] is employed to solve the non-smooth optimization problem. We first define the subgradient, which is an extension of the gradient for continuous, non-differentiable functions.

Definition 5.3.1 (Subdifferential of a concave function [132]). *Let the function $\Theta : \mathbb{R}^\ell \rightarrow \mathbb{R}$ be a concave, continuous function that is non-differentiable. The subdifferential $\bar{\partial}\Theta(\mathbf{Y}^\ell)$ at $\mathbf{Y}^\ell \in \mathbb{R}^\ell$ is the set of all vectors $\boldsymbol{\zeta} \in \mathbb{R}^\ell$ such that*

$$\Theta(\mathbf{y}) \leq \Theta(\mathbf{Y}^\ell) + \langle \boldsymbol{\zeta}, \mathbf{y} - \mathbf{Y}^\ell \rangle, \quad \forall \mathbf{y} \in \mathbb{R}^\ell.$$

Each vector $\boldsymbol{\zeta} \in \bar{\partial}\Theta(\mathbf{Y}^\ell)$ is called a subgradient of Θ at \mathbf{Y}^ℓ .

Next, we characterize the global maximum of a concave, non-smooth function.

Proposition 5.3.1 (Global maximum of a concave, non-smooth function [132]). *Let the function $\Theta : \mathbb{R}^\ell \rightarrow \mathbb{R}$ be a concave, continuous function that is non-differentiable. Then the following statements are equivalent.*

1. Function Θ attains its global maximum value at $\mathbf{Y}^\ell \in \mathbb{R}^\ell$.
2. $\mathbf{0} \in \bar{\partial}\Theta(\mathbf{Y}^\ell)$.
3. The directional derivative $\Theta'(\mathbf{Y}^\ell, \mathbf{d}) \leq 0$ for all $\mathbf{d} \in \mathbb{R}^\ell$.

Proposition 5.3.1 implies that any dual optimization algorithm would need to find an optimal point \mathbf{Y}^* such that a subgradient of the dual function Θ evaluated at \mathbf{Y}^* is the zero vector $\mathbf{0}$. Moreover, we only have access to stochastic estimates of the subgradient (5.52). A popular technique to solve such problems is the SSM. The SSM is a subgradient-based search algorithm using a dual function evaluation and noisy evaluation of an arbitrary subgradient at each evaluation point. The SSM is seen as an extension of the stochastic gradient descent method for continuous, non-differentiable functions. For a given refinement level ℓ of the Lagrange multiplier, the SSM has the following update rule:

$$\mathbf{Y}^\ell[k+1] = \mathbf{Y}^\ell[k] + \gamma[k]\bar{\mathcal{D}}\Theta(\lambda^\ell[k]), \quad (5.53)$$

where $\mathbf{Y}^\ell[k]$ represents the k^{th} -iterate of \mathbf{Y}^ℓ , $\gamma[k]$ denotes a suitable stepsize, $\lambda^\ell[k]$ denotes the piece-wise constant Lagrange multiplier function (5.40) defined using the amplitudes $\mathbf{Y}^\ell[k]$, and $\bar{\mathcal{D}}\Theta(\lambda^\ell[k])$ denotes the subgradient (5.52) of the dual function Θ at $\lambda^\ell[k]$. At each refinement level ℓ , the SSM runs until the subgradient norm $\|\bar{\mathcal{D}}\Theta(\lambda^\ell[k])\|$ is below a prescribed relative tolerance TOL or the number of iterations k exceeds a threshold k_{\max} . The algorithm output at each ℓ is the optimal amplitude vector \mathbf{Y}^ℓ . However, this optimal value is not necessarily reached at the final iteration of the SSM because the objective function to be maximised does not increase at every step of the subgradient method [104]. Hence, we track the highest value of the dual function $\Theta(\lambda^\ell[k])$ and store its corresponding \mathbf{Y}^ℓ as the optimal value.

The choices for the starting point and step-size are crucial for quick convergence of the **SSM**, especially in higher dimensions. This work uses a nonsummable diminishing step-size of $\mathcal{O}\left(\frac{1}{k+1}\right)$, where k denotes the iteration number. The convergence of the **SSM** with this step-size is proven in [104]. The associated constant, denoted by C_{SSM} in Algorithm 5, is tuned once such that the **SSM** step taken in $(\mathbb{R}^+)^{\ell}$ is $\mathcal{O}\left(\mathbf{Y}^{\ell}\right)$. C_{SSM} need not be tuned again for a different set of parameters with the same problem structure. For enhanced performance, the step-size at each iteration is adjusted by dividing the current subgradient by its norm. Furthermore, the **SSM** runs for the first \bar{N}_{iter} iterations using a constant step-size before reverting to a diminishing step size. The Lagrange multiplier function $\lambda^{\ell}(t)$ constructed using the computed optimal amplitudes \mathbf{Y}^{ℓ} is used as the starting point of the **SSM** at refinement level 2ℓ . However, we must still choose a good starting point at level $\ell = 1$. We devise an initialization algorithm and complement it with a non-smooth deterministic optimiser.

Remark 5.3.3 (*SSM convergence rate*). *A convergence rate of $\mathcal{O}\left(k^{-\frac{1}{2}}\right)$ for the **SSM** has been proven for locally Lipschitz convex functions with at most quadratic growth [133]. When solving the optimization problem up to a relative tolerance of TOL with respect to the dual function value $\Theta(\lambda)$, the **SSM** requires an expected number of $\mathcal{O}\left(|\Theta(\lambda)|^{-2}\text{TOL}^{-2}\right)$ iterations to converge under the assumptions in [133].*

Remark 5.3.4 (*Primal feasible solution*). *Note that Algorithm 5 uses the subgradient as a stopping criterion, instead of a duality gap. This is because the output of Algorithm 5 is already a primal feasible solution, at least up to given relative tolerance TOL . That is, the optimal controls produce a solution that minimises the primal cost while only violating the chance constraint (5.30) by lesser than a small value of ϵTOL .*

LMBM-boosted initialization

Level $\ell = 1$ implies that the Lagrange multiplier function $\lambda^1(t)$ is constant in time (5.40), with the constant denoted by Y_1^1 . First, we devise a heuristic initialization algorithm starting with the arbitrary starting point $Y_1^1 = 1$ and estimating the dual function value and a subgradient at that point. Then, if the subgradient at $Y_1^1 = 1$ is positive, Y_1^1 is continually increased by a factor $\beta_F > 0$ until the subgradient becomes negative or its norm reaches a prescribed relative tolerance TOL_{init} . Conversely, if the subgradient at $Y_1^1 = 1$ is negative, Y_1^1 is continually decreased by the factor β_F , until the subgradient becomes positive or its norm reaches relative tolerance TOL_{init} . The final obtained point is stored as \tilde{Y}_1^1 . Algorithm 15 in Appendix B.9 presents the pseudo-algorithm of the devised initialization procedure. Figure 5.3 depicts Algorithm 15 visually for further clarity.

To obtain an even better starting point for the **SSM**, we employ \tilde{Y}_1^1 obtained from the heuristic Algorithm 15 as a starting point for a deterministic non-smooth optimization routine called the **LMBM** [103, 134]. Bundle methods are more robust than a simple subgradient-based method for non-smooth optimization because they approximate the entire subdifferential of the objective function, enhancing its convergence speed [135]. However, the **LMBM** requires access to deterministic estimates of the objective function value and its subgradient at each evaluation point. Although we only have access to noisy **MC** estimates of subgradients (5.52), we can still employ the **LMBM** routine to reach a good starting point for the **SSM**, especially in the regime where changes in the objective function value considerably outweigh the noise in subgradient estimates. This work runs the **LMBM** for a fixed number of iterations N_{iter} . Table D.3 in Appendix D.6 specifies the parameters required to run the **LMBM** routine [7]. The output of this routine \hat{Y}_1^1 is

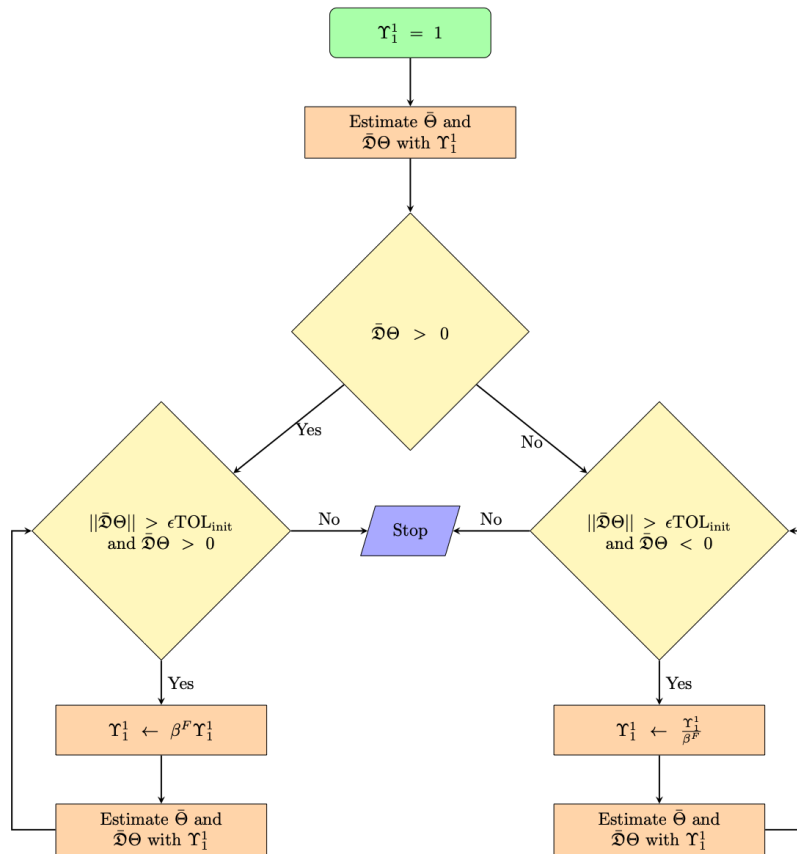


Figure 5.3: Schematic illustration of the proposed initialization Algorithm 15.

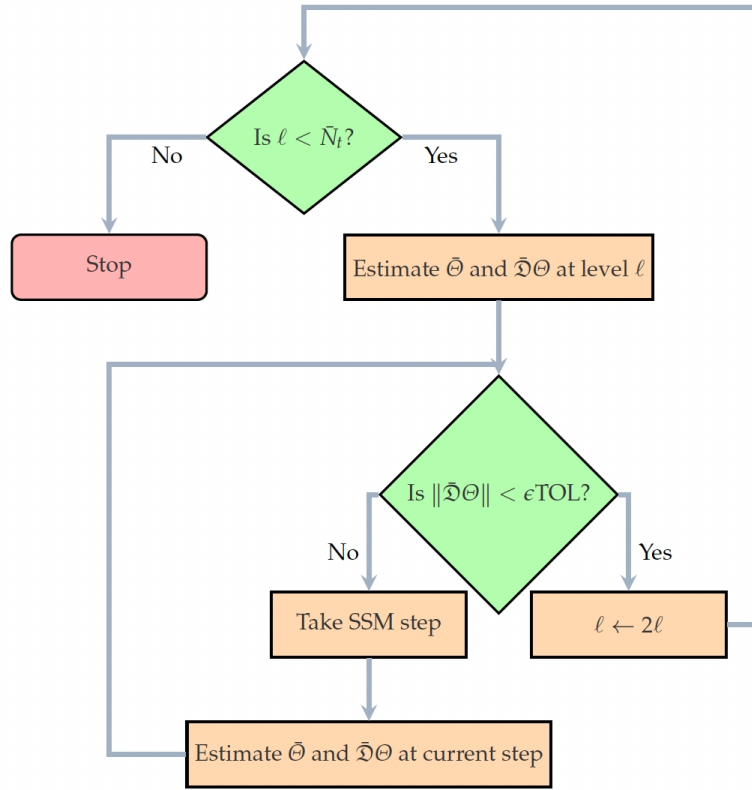


Figure 5.4: Schematic illustration of the proposed dual optimization Algorithm 5.

applied as a starting point for the *SSM* Algorithm 5. We put all these elements together to formulate the final dual optimization Algorithm 5, visually depicted in Figure 5.4.

Remark 5.3.5 (Scale of λ). *With respect to the parameter ϵ that determines the confidence level of satisfying the chance constraint (5.30), the optimal primal cost $J_{0,\mathbf{x}(0)}(\boldsymbol{\phi})$ associated with Problem 5.2.4 is a $\mathcal{O}(1)$ term, while the optimal subgradient would be $\mathcal{O}(\epsilon)$ (ideally 0). Equation (5.43) then implies that the optimal Lagrange multiplier λ is $\mathcal{O}(\frac{1}{\epsilon})$. This yields a good guess for an initial point Y_1^1 for the *LMBM*-boosted initialization procedure.*

5.4 Numerical experiments and results

This section presents a model example of a cellular base station powered by the German power grid. Here, we describe the system and all dynamics driving the operation of the base station, and provide the results of applying the proposed numerical approach to solve the optimal power procurement problem for the system.

5.4.1 Description of model cellular base station

Figure 5.2 schematically illustrates the considered base station model. Table D.1 provides the descriptions and numerical values of all coefficients used to describe the model.

The daily mobile user traffic profile primarily drives the power demand of a cellular base station (5.5). The description of the daily traffic profile consists of two facets: (i) the number of people connected to the network $N_u(t)$, and (ii) the physical distribution

of users around the base station $\rho_z(t)$. The numerical experiment applies the following sinusoidal profile to model $N_u(t)$:

$$N_u(t) = \max \left[\underline{N}_u, \frac{\bar{N}_u}{2^q} \left(1 + \sin \left(\frac{\pi t}{6} + \pi \right) \right)^q \right]. \quad (5.54)$$

This model and the stochastic versions of it have been empirically demonstrated to approximate practical user patterns closely for calibrated values of the smoothness parameter q [136, 137]. The parameter q determines the rate of increase or decrease in $N_u(t)$ during the day. A higher value of q indicates a steeper increase or decrease rate. In our numerical experiment, we fix $q = 3$. This model also considers two peak hours (about 09.00 to 10.00 and 20.00 to 21.00) and few connected users at night (00.00 to 05.00). Moreover, $\underline{N}_u \in \mathbb{N}$ and $\bar{N}_u \in \mathbb{N}$ denote the minimum and maximum number of users, respectively, that are connected to the base station at any time. This work sets $\underline{N}_u = 100$ and $\bar{N}_u = 2000$. These values were selected using the data on the average population densities of German cities [138] and the geographical area served by an average cellular base station in Germany [139]. Figure 5.5 depicts the model for $N_u(t)$.






Figure 5.5: Typical daily cellular user traffic profile, described by (5.54).

The user distribution $\rho_z(t)$ crucially determines the estimation of the outage proportion ϕ_{out} (5.10) at given time t . Some simple two-dimensional distributions have been employed to model the mobile user distribution in urban zones efficiently. For example, the uniform distribution can be employed to model the mobile user distribution in suburban areas very well, and the Gaussian distribution is a good model for mobile user distributions in industrial zones [140]. Figure 5.6 visualizes this model. Analytical expressions can be derived for the corresponding $\phi_{\text{out}}(t)$ for these distributions. Appendix D.2 provides the closed-form expressions. For the numerical experiments, we apply the symmetric two-dimensional Gaussian distribution centered on the site of the base station \mathbf{x}_{BS} with a diagonal covariance matrix with variance $\sigma_u^2 = 300^2 \text{m}^2$ along both directions.

We apply the SDE in (5.7) to model the instantaneous Nakagami fading channel, whose invariant distribution is a shifted-gamma distribution with the shape parameter μ , scale parameter θ , and shift $\underline{\zeta}$. We set $\mu = 3$ and $\theta = 1$ so that a quick, smooth convergence occurs to the invariant distribution and $\mathbb{E}[\underline{\zeta}(t)^{-2}]$ is bounded for all $t \in [0, T]$. We set the shift $\underline{\zeta}$, corresponding to the value of ζ that can still achieve an outage proportion of $\frac{\phi_{\text{th}}}{2}$ in the worst-case scenario when $N_u(t) = \bar{N}_u$. Moreover, we set the initial distribution



(a) Uniform distribution in suburban areas.




(b) Gaussian distribution in industrial zones.

Figure 5.6: Visualizing the spatial distribution of mobile users in urban zones using simple analytical two-dimensional distributions.


μ_0^ζ the same as the invariant distribution of process ζ . Figure 5.7 visualizes the sample paths and the invariant distribution of the SDE in (5.7).

We set $K_{\text{net}}(t) = 0.01 \text{ €/h}$ per person for $t \in [0, T]$. The price was taken from the Vodafone 5G prices in Germany [141]. The constant price assumption is generally valid because regulators set these prices and do not move within a day.

We apply the SDE in (5.13) to model the instantaneous normalized generated renewable power. This work uses the 2023 German wind power forecast and production data with a 15-minute frequency from the operator 50Hertz [142] to calibrate the SDE in (5.13). The forecast data are employed to construct $p(t)$ and $\dot{p}(t)$, and the discrepancy between the forecast and true production data is used to calibrate the parameters α and θ_0 in (5.13). Apart from these parameters, we also calibrate an additional parameter ζ . The discrepancy between the forecast and production at time $t = 0$ is not typically zero in real-world applications. ζ is defined as the time before $t = 0$, for which the forecast error can be set



(a) Plot of the sample mean and 95% confidence intervals of sample paths of the SDE in (5.7).



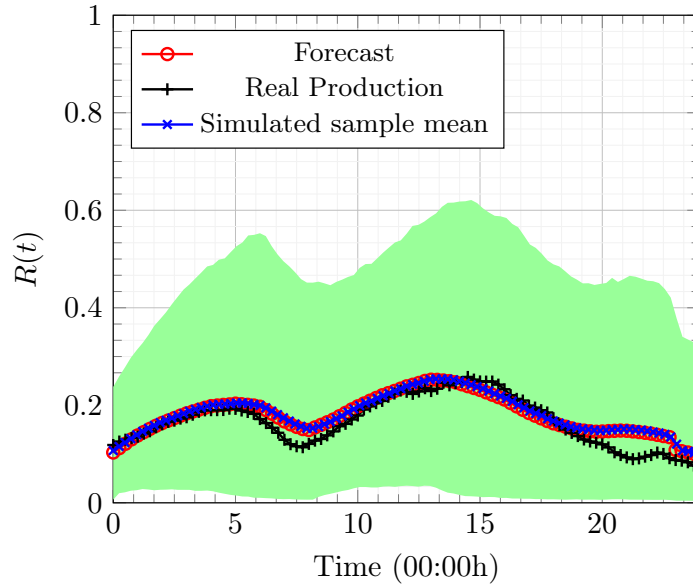
(b) Visualization of the invariant distribution of sample paths of the SDE in (5.7), verifying that the invariant distribution of the SDE in (5.7) is the shifted-gamma distribution with shape parameter $\mu = 3$ and scale parameter $\theta = 1$.

Figure 5.7: Results of numerical simulation of the SDE in (5.7) governing the Nakagami wireless channel dynamics.


to zero. This approach ensures that uncertainty also exists in renewable power production at $t = 0$. This leads to the initial distribution μ_0^r of the normalized wind power being a normal distribution with mean $p(-\zeta) + \dot{p}(-\zeta)\zeta$ and variance $2\alpha\theta_0 p(-\zeta)(1 - p(-\zeta))\zeta$. We follow the calibration procedure detailed in [125]. The mean-reversion parameter $\theta(t)$ in (5.13) is defined as follows:

$$\theta(t) := \max \left(\theta_0, \frac{\alpha\theta_0 + |\dot{p}(t)|}{\min(p(t), 1 - p(t))} \right).$$

The calibrated values for the above data are $\alpha = 0.34$, $\theta_0 = 2.3948$, and $\zeta = 0.054$. Given the deterministic forecasts of wind power from the same operator in 2024, we can generate sample paths of the SDE in (5.13) with parameters α, θ_0, ζ calibrated from 2023 data. Figure 5.8 illustrates the results for two sample days in 2024.



(a) 12.03.2024



(b) 09.04.2024

Figure 5.8: Normalized wind power forecast and real production data in Germany in the region operated by *50Hertz* in 2024. Mean path and 95% confidence intervals of the SDE in (5.13) calibrated from 2023 data.

We apply (5.14) to model the instantaneous normalized battery charge. As given in (5.14), the maximum charging capacity \underline{P}_A and the maximum discharge capacity \bar{P}_A are defined as functions of the stored battery charge. This function is typically called the characteristic curve of the battery. Our numerical experiment employs a simple characteristic curve illustrated in Figure 5.9. The characteristic curve imposes (5.15),

implying that the battery cannot discharge power when it has zero charge and can no longer charge itself when fully charged. This characteristic curve ensures that the constraints (5.27) and (5.28) are linear in the state \mathbf{X} , implying that the admissible set of controls $\mathcal{A}(t, \mathbf{X}(t))$ is a convex set for all $t \in [0, T]$ and all realizations of $\mathbf{X}(t)$.




Figure 5.9: Simple characteristic curve of a battery, described by Eqs. (5.14) and (5.15).

Forecasts of buying and selling prices, $K_b(t)$ and $K_s(t)$, are constructed from publicly available day-ahead spot prices of grid power. We set $K_b(t) = K_s(t)$, implying that the buying and selling prices of power are the same. This assumption is valid because the market commonly ensures that this holds in real life. The numerical experiment employs the German day-ahead spot-price data from 2024 available in [6]. However, these data are generally available only at a one-hour frequency. Hence, we smoothed and extended the data to domain $[0, T]$ by fitting a polynomial curve to them and using the polynomial function as the input $K_b(t)$. Figure 5.10 illustrates an example of the spot-price data.




Figure 5.10: Day-ahead spot-price for grid power in Germany on 07.09.2023. Data are from [6].

5.4.2 Numerical results

We present the results of solving the chance-constrained SOC problem for the model cellular base station (Section 5.4.1) using the numerical approach in Section 5.3.




Figure 5.11: Numerical verification of the error convergence rate of the PDE discretization with respect to N_x of the proposed upwind FDM scheme in Section 5.3.1 for the HJB PDE (5.36) with given inputs in Table D.1 and Section 5.4.1.

Value function of the dual problem 5.2.4

The numerical solution of the HJB PDE in (5.36) for a given Lagrange multiplier function $\lambda^\ell(t)$ (5.40) is obtained using the scheme in Section 5.3.1. We apply a uniform rectangular spatial grid with $N_1 = N_2 = N_3 = N_x = 10$, where N_t is set as the lowest value that satisfies the CFL condition (5.51). In this case, we obtain $N_t = 800$. The presence of second-order derivatives in the HJB PDE in (5.36) yields $N_t = \mathcal{O}(N_x^2)$ for a numerically stable solution, explaining the high value obtained for N_t . These discretization parameter choices ensure that the relative numerical PDE discretization error is below 1%, as seen in Figure 5.11. The relative errors in Figure 5.11 were estimated using a reference solution with $N_x = 20$ and $N_t = 3200$. For each N_x , the parameter N_t was chosen to satisfy the CFL condition (5.51). Additionally, we set $\bar{\chi}$ as the upper 95% quantile of the invariant distribution of process $\zeta(t)$.

The optimization problem in (5.49) must be solved at each point in the grid τ . For the simple cases of uniform or Gaussian cellular user distribution, we obtain closed-form expressions for ϕ_{out} (see Appendix D.2). The resulting ϕ_{out} ensures that (5.49) is a convex optimization problem and can be solved analytically, when the path loss exponent $\eta \geq 2$. The parameter η generally takes values between 2 and 4 in wireless networks [143], ensuring the validity of this assumption. In this numerical experiment, we set $\eta = 2$ corresponding to free space path loss [143]. Appendix D.3 presents additional details. Figure 5.12 presents three slices of the numerical value function v with the optimal Lagrange multiplier function $\lambda^\ell(t)$, each with two fixed state variables. The numerical dual function approximation $\bar{\Theta}(\lambda^\ell)$ is obtained by evaluating $v(0, \mathbf{X}(0))$.




Figure 5.12: Slices of the numerical value function v on the axes of time and state variables. For all slices on each component of \mathbf{x} , the rest of the state variables are fixed at 0.

Obtaining reference costs for comparison

This work compares the values for the dual cost above with reference values obtained from two related, easy-to-solve **SOC** problems. This comparison was conducted to ensure that the output of the numerical algorithm makes sense with respect to these reference values. First, we solve the same optimization problem as in Problem 5.2.1, but with the chance constraint (5.30) replaced with a deterministic **a.s.** constraint, solving the following problem.

Problem 5.4.1 (Primal problem with the **a.s.** QoS constraint). *Given the initial data $\mathbf{X}(0) = [A_0, R_0 \sim \mu_0^R, \xi_0 \sim \mu_0^\xi]$, and deterministic forecasts for the daily cellular user traffic profile $N_u(t)$, energy spot prices $K_b(t), K_s(t)$, renewable energy forecast $p(t)$, and mobile network usage price $K_{net}(t)$ for $0 \leq t \leq T$, we solve the following:*

$$\boldsymbol{\phi}^* = \arg \min_{\{\boldsymbol{\phi} \in \mathcal{A}\}} J_{0, \mathbf{X}(0)}(\boldsymbol{\phi}) = w(\mathcal{C}_1 - \mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + (1 - w)\mathcal{C}_2, \quad (5.55)$$

where the minimization is subject to dynamics (5.20), and constraints (5.21)-(5.28) and the following **a.s.** constraint (instead of the chance constraint in (5.30)).

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = \phi_{\text{th}}, \quad \forall t \in [0, T]. \quad (5.56)$$

Problem 5.4.1 is a standard continuous-time SOC problem; hence the dynamic programming principle can be used to derive and numerically solve the associated HJB PDE, as explained in Section 2.4. Appendix D.4 details the solution to Problem 5.4.1. The value function for this problem is expected to be higher than the dual cost of Problem 5.2.4 because the network operator would need to buy more power to keep ϕ_{out} constant at all times, whereas the operator is allowed to relax the constraint probabilistically in Problem 5.2.4 as described by the chance constraint (5.30). Next, we solve the same optimization problem as Problem 5.2.1, but without the QoS constraint (5.30).

Problem 5.4.2 (Primal problem with no QoS constraint). *Given the initial data $\mathbf{X}(0) = [A_0, R_0 \sim \mu_0^R, \xi_0 \sim \mu_0^\xi]$, and deterministic forecasts for the daily cellular user traffic profile $N_u(t)$, energy spot prices $K_b(t), K_s(t)$, renewable energy forecast $p(t)$, and mobile network usage price $K_{\text{net}}(t)$ for $0 \leq t \leq T$, we solve the following:*

$$\boldsymbol{\phi}^* = \arg \min_{\{\boldsymbol{\phi} \in \mathcal{A}\}} J_{0, \mathbf{X}(0)}(\boldsymbol{\phi}) = w(\mathcal{C}_1 - \mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + (1-w)\mathcal{C}_2, \quad (5.57)$$


where the minimization is subject to dynamics (5.20), and constraints (5.21)-(5.28).

Problem 5.4.2 is also a standard continuous-time SOC problem; hence the dynamic programming principle can be applied to derive and numerically solve the associated HJB PDE. Appendix D.5 details the solution to Problem 5.4.2. Without the Lagrangian relaxation, the value function associated with this problem is the same as that of the relaxed Problem 5.2.2. Thus, this problem can be solved using the numerical scheme in Section 5.3.1, setting $\lambda(t) = 0$. The value function for this problem is expected to be less than the dual cost of Problem 5.2.4 because the network operator only maximises revenue without ensuring that QoS is achieved. The HJB PDEs corresponding to Problems 5.4.1 and 5.4.2 are numerically solved only once because no associated dual problem needs to be solved.


Dual optimization results

We run Algorithm 5 with parameters set to values in Appendix D.6. Initially, we set $\phi_{\text{th}} = 10^{-3}$ and $\epsilon = 0.1$, implying that the network operator must ensure that the proportion of users in outage ϕ_{out} should be less than 10^{-3} with a probability of 90% for all $t \in [0, T]$. Figure 5.13 numerically verifies the concavity of the numerical dual function $\bar{\Theta}(\lambda^\ell)$ for $\ell = 1$. The dual costs in Figure 5.13a are estimated by numerically solving the HJB equation (5.36) using the upwind FDM scheme in Algorithm B.7 for λ^1 constructed with different values of Y_1^1 . Figure 5.13a shows a slight increase in the dual function values at lower values of Y_1^1 before decreasing, indicating that a unique function maximum exists. The 1D subgradients in Figure 5.13b are estimated using the EM MC Algorithm B.8 with inputs $\ell = 1$, $\bar{N}_1 = \bar{N}_t = 2^6$, and $M_{\text{SG}} = 10^3$. Figure 5.13b confirms a decreasing subgradient starting from a positive subgradient value, further verifying the dual function's concave nature. Moreover, the subgradients converge to the value $-\epsilon$ for high values of Y_1^1 , implying that the average $\mathbb{P}[\phi_{\text{out}}(t, \mathbf{X}(t), \boldsymbol{\phi}^*(t, \mathbf{X}(t))) \geq \phi_{\text{th}}]$ over the time interval $[0, T]$ converges to 0 in the limit of high Lagrange multiplier values.

Figure 5.14 presents the results of the initialization Algorithm 15 and the LMBM routine. Figure 5.14a reveals the dual cost versus Y_1^1 at each iteration of Algorithm 15. The dual cost increases with the number of iterations, and it is within the two reference values attained by solving Problems 5.4.1 and 5.4.2. Figure 5.14b plots the corresponding estimated subgradient values. The initialization algorithm reaches the prescribed initial



(a) Dual function



(b) Subgradient

Figure 5.13: Numerical verification of the dual function's concavity. Dual cost and subgradient values computed with λ^1 constructed from different values of γ_1^1 for $\ell = 1$. Both plots numerically confirm that the finite-dimensional dual maximization Problem 5.2.4 is a convex, non-smooth problem.

tolerance in three iterations. Figure 5.14 reveals that the LMBM offers a more optimal point than the output of Algorithm 15, providing an excellent starting point for the dual optimization in Algorithm 5.

Figure 5.15 illustrates the evolution of the dual function values and its corresponding subgradient at each iteration of the SSM for a fixed ℓ . Figure 5.15a verifies that not every step of the SSM is taken towards an increasing direction of the dual function. This behaviour illustrates the difficulty of solving the non-smooth optimization in Problem 5.2.4. Due to the satisfactory choice of starting values and step sizes (see Section 5.3), Algorithm 5 converges for each ℓ within k_{\max} iterations with respect to the subgradient norm.

Figure 5.16 plots the evolving optimal dual function $\bar{\Theta}(\lambda^\ell)$ and its corresponding subgradient $\bar{\mathcal{D}}\Theta(\lambda^\ell)$ constructed using the optimal amplitudes \mathbf{Y}^ℓ at each level ℓ of


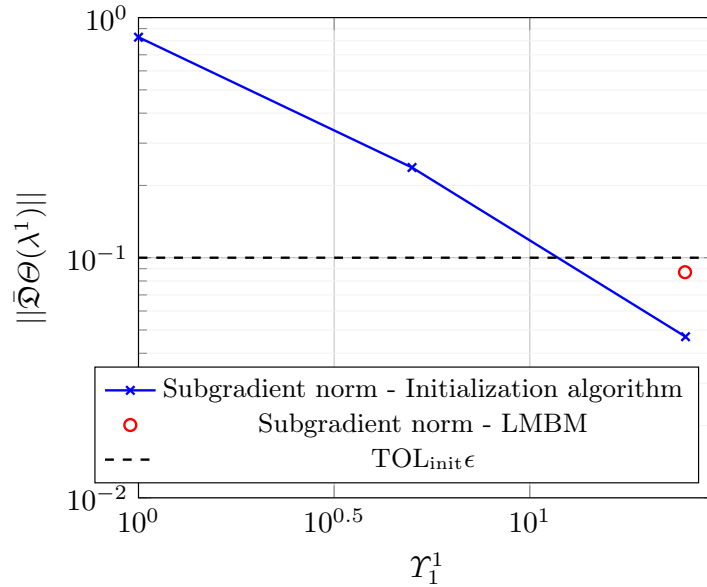
(a) Dual cost $\bar{\Theta}(\lambda^1)$ (b) Subgradient norm $\|\bar{\Delta}\Theta(\lambda^1)\|$.

Figure 5.14: Dual cost and subgradient norm computed with λ^1 constructed from Υ_1^1 attained at each iteration of Algorithm 15, along with the corresponding quantities using the final output of LMBM routine [7]. The initialization routine stops within three iterations for the considered problem. The advanced, non-smooth optimization LMBM routine provides an even more optimal point for the upcoming SSM.

Algorithm 5. Table 5.1 lists the optimal dual cost at each stage of Algorithm 5. Table 5.1 presents minor increases in the dual function evaluations as the refinement level ℓ increases, implying that we successively obtain more optimal solutions as the Lagrange multiplier function $\lambda^\ell(t)$ becomes finer. We normalize the subgradient vector norm by the number of dimensions for a fair comparison between the norm of vectors of varying dimensions.

Table 5.1 also highlights the advantage of applying a probabilistic constraint to achieve the QoS in cellular networks, instead of doing it a.s. for all $t \in [0, T]$ (as in Problem 5.4.1).



(a) Dual cost $\bar{\Theta}(\lambda^2)$ (b) Subgradient norm $\|\bar{\Delta}\Theta(\lambda^2)\|$.

Figure 5.15: Dual cost and subgradient norm computed with λ^2 constructed from \mathcal{Y}^2 at each SSM iteration of Algorithm 5 for $\ell = 2$. The SSM at the given refinement level converges within 20 iterations, suggesting that a good starting point was used.

The optimal dual cost achieved by Algorithm 5 is much closer to the expected optimal cost of the unconstrained Problem 5.4.2 and around four times less expensive than satisfying the QoS constraint (5.56) a.s. (Problem 5.4.1). In rare occasions when an unreasonable surge in demand ($N_u(t)$) occurs, or a very low wireless fading channel ($\zeta(t)$) exists due to environmental conditions or low availability of renewable power ($P_R(t)$), the proposed numerical approach allows the network operator to relax the stringent QoS constraint (5.56) to save substantial expenditure while still serving as many customers as possible. This approach contrasts the solution to Problem 5.4.1, which enforces (5.56) even in such periods, significantly increasing operating expenditure.

Figure 5.17 presents the evolution of the optimal Lagrange multiplier function as the



(a) Dual cost $\bar{\Theta}(\lambda^\ell)$ (b) Subgradient norm $\|\bar{\Delta}\Theta(\lambda^\ell)\|$.

Figure 5.16: Dual cost and subgradient norm computed with λ^ℓ constructed from the optimal \mathbf{Y}^ℓ obtained from the *SSM* at each refinement level ℓ of Algorithm 5. The *SSM* converges within k_{\max} iterations at each refinement level ℓ , indicating the efficiency of the *LMBM*-boosted initialization procedure and a good choice of step sizes for the *SSM*.

refinement level ℓ increases in Algorithm 5. For the given example, Algorithm 5 reaches prescribed *TOL* at $\ell = 8$, and does not refine the Lagrange multiplier function further. High-demand periods (see Figure 5.5) correspond to high Lagrange multiplier function values. This behaviour is expected because a high λ^ℓ forces the mobile operator to satisfy the constraint in (5.30) in times of high demand, although it significantly increases costs. The oscillation frequency of the base station power demand dictates the oscillations in the Lagrange multiplier function.

Description	ℓ	Optimal cost (€)
Problem 5.4.1		-173.57
Problem 5.4.2		-667.83
Initialization Algorithm 15	1	-661.52
LMBM	1	-663.62
Dual optimization Algorithm 5	2	-661.8
Dual optimization Algorithm 5	4	-661.59
Dual optimization Algorithm 5	8	-660.51

Table 5.1: Comparing the optimal dual costs achieved in various stages of Algorithm 5 with those of Problems 5.4.1 and 5.4.2. The LMBM procedure already gives a decent solution to the dual problem, as there is not much difference between the LMBM solution and the final dual cost obtained. As the Lagrange multiplier becomes finer, we successively obtain more optimal solutions.




Figure 5.17: Evolution of the optimal Lagrange multiplier function $\lambda^\ell(t)$ at each refinement level ℓ in Algorithm 5. The spikes in the Lagrange multiplier function coincide with periods of high cellular traffic.

Optimal power procurement policy

With the optimal Lagrange multiplier function λ^ℓ from Algorithm 5, the dual cost is -660.51 € (see Table 5.1). Figure 5.18 illustrates the optimal power procurement policy for the model cellular base station for the next 24 hours using the optimal solution, including the (a) net power consumption of the base station $C_{\text{scaI}}N_u(t)P_{\text{tx}}(t) + C_{\text{offset}}$ (refer (5.5)), (b) power drawn from the battery to run the base station $P_A(t)$, (c) power bought from the grid $P_F(t)$, (d) power sold back to the grid $P_S(t)$, and (e) optimal charge stored in battery $A(t)$. Figure 5.18a indicates that the base station consumes high power during periods of high demand (refer Figure 5.5). Figures 5.18b, 5.18c, and 5.18d reveal that the battery provides most of the power to run the base station throughout the day. The clean, renewable energy stored in the battery powers a large proportion of base station operations so that the mobile operator does not need to buy expensive grid power. This approach reduces costs and the carbon footprint of the base station. The operator must only buy

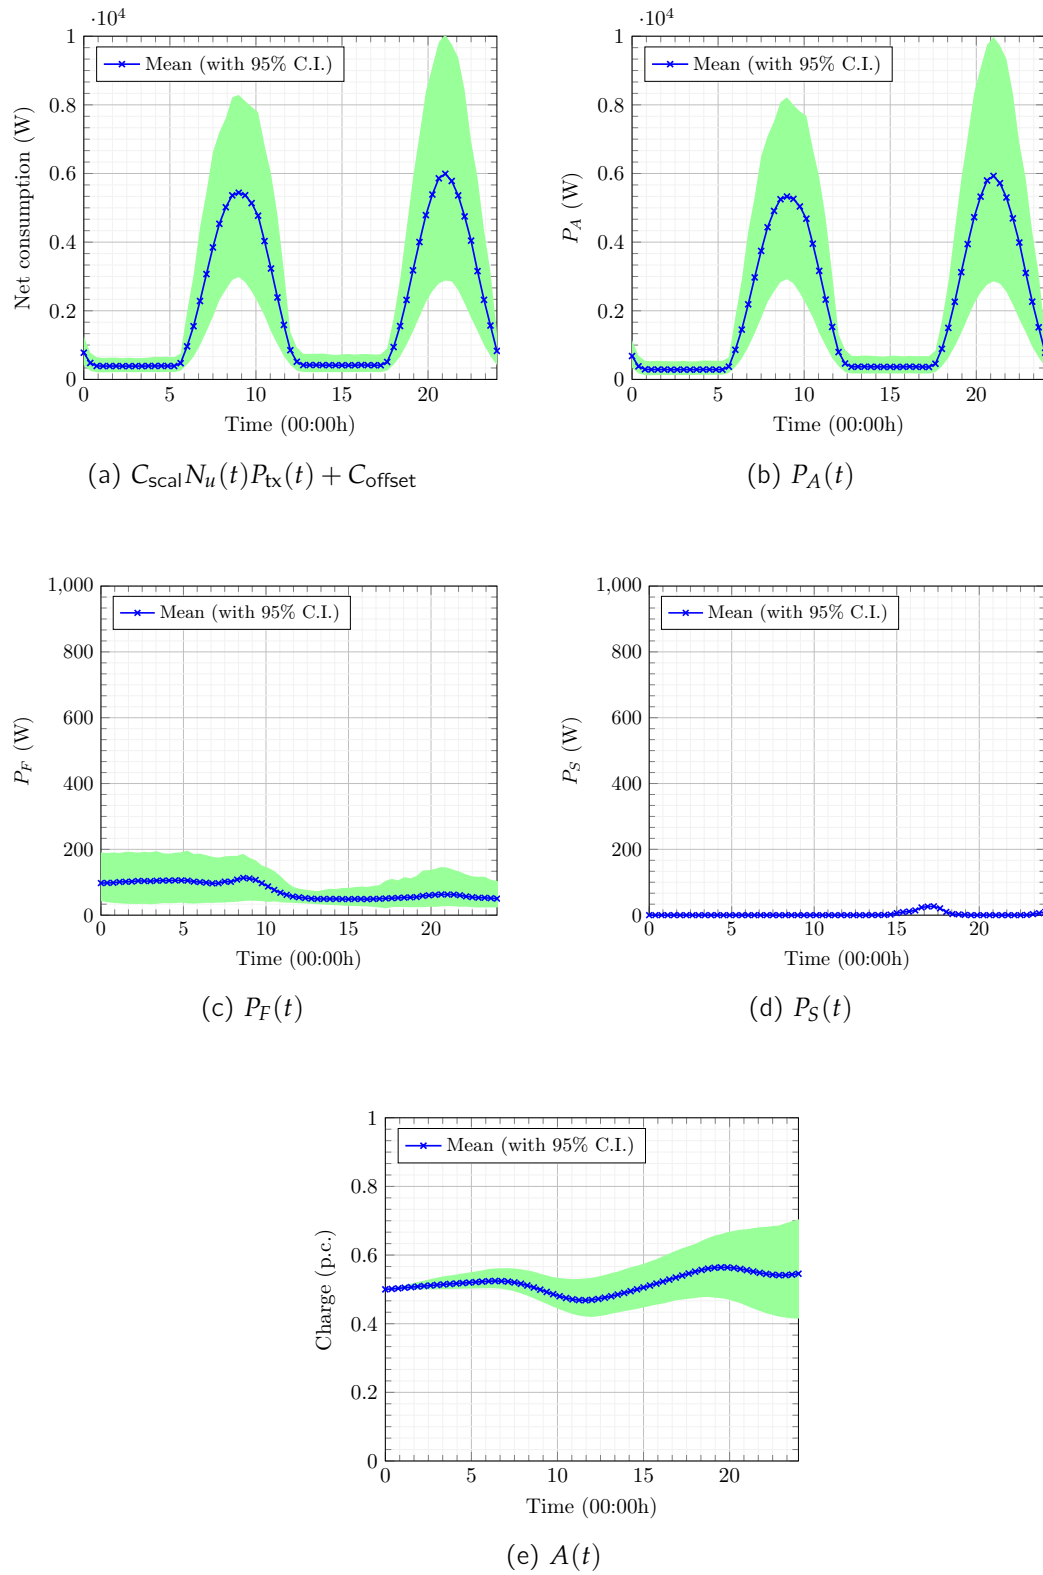




Figure 5.18: Optimal power procurement policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5. These curves are the practical outputs received by the base station operator to help make decisions regarding power procurement in the next 24 hours.

some power from the grid to satisfy high demand during peak hours (Figure 5.18c). The base station sells a negligible amount of power back to the grid (Figure 5.18d). The optimal policy also advises the operator to store incoming renewable power and charge the battery during low-demand hours before using that power to serve customers during high-demand hours (Figure 5.18e). Figure 5.18e also indicates that the battery has some stored energy at the end of the day that could be used the next day. This observation establishes the effect of implementing the running horizon framework (see Section 5.2.1), although we solve the problem for a short-term horizon.



(a) Solution to Problem 5.4.2



(b) Solution to Problem 5.2.4.

Figure 5.19: MC estimates of the violation in constraint $\mathbb{P}[\phi_{\text{out}}(t) > \phi_{\text{th}}] - \epsilon$ for each $t \in [0, T]$ using optimal controls from solving (a) Problem 5.4.2 and (b) Problem 5.2.4 using Algorithm 5. The unconstrained Problem 5.4.2 provides a solution that does not satisfy the QoS at all. The solution to the dual Problem 5.2.4 satisfies the probabilistic QoS constraint, time-averaged on the Lagrange multiplier discretization.

We estimate the time-pointwise violation of the constraint (5.30) using MC and compare it with the final subgradient vector $\tilde{\mathcal{D}}\Theta(\lambda^\ell)$ from Algorithm 5 to validate the optimality

of the solution further. We run approximate sample numerical paths of the optimally controlled dynamics, according to (2.7), with the optimal policy above. Figure 5.19b plots this information for a discrete set of points in $t \in [0, T]$. We also compare this with the optimal policy attained from solving the unconstrained problem in Problem 5.4.2 in Figure 5.19a. Figure 5.19a reveals that the user outage proportion ϕ_{out} for the optimal policy with no constraints on the QoS is consistently greater than the threshold ϕ_{th} . In contrast, the proposed optimal solution controls the chance constraint (5.30) in an integral (averaged) sense (refer (5.44)), even though there are some oscillations around zero in the pointwise violation of the constraint (5.30). Figure 5.19b illustrates this behavior, with the time-averaged violation in the chance constraint (5.30) close to zero up to the prescribed relative tolerance TOL.

Sensitivity analysis

We investigate the performance of Algorithm 5 under certain scenarios that could be of interest to network operators, which are listed below.

- Scenario A: No incoming renewable power for a day ($P_R(t) = 0$ for all $t \in [0, T]$).
- Scenario B: The wireless fading channel is substantially low due to extreme weather ($\underline{\xi}$ is halved).
- Scenario C: High weight is assigned to minimizing operating expenditure ($w = 0.99$).
- Scenario D: The price users pay to connect to the network substantially reduces ($K_{\text{net}} = 0.001$ €/h per person).
- Scenario E: The QoS chance constraint (5.30) must be satisfied with low confidence ($\epsilon = 0.2$).

All parameters and dynamics are set as in Table D.1 and Section 5.4.1, except for the prescribed parameter change to simulate each scenario. For each scenario, we solve three problems: (i) dual optimization (Problem 5.2.4), (ii) Problem 5.4.1 with a.s. QoS constraint (5.56), and (iii) Problem 5.4.2 with no QoS constraint. Table 5.2 reports the optimal costs in each case for all five scenarios.

	Dual Algorithm 5	Problem 5.4.2	Problem 5.4.1
Scenario A	−441.8	−564.9	1941.9
Scenario B	−649.28	−660.27	4160.4
Scenario C	−1334.7	−1338.8	1980.5
Scenario D	−80.52	−92.69	1209.5
Scenario E	−664.03	−667.83	−173.57

Table 5.2: Comparing the optimal dual costs (in €) achieved with the dual optimization Algorithm 5 for each scenario with the optimal primal costs of the corresponding unconstrained and a.s. QoS constrained problems. In all scenarios, the optimal cost of the chance-constrained problem lies in between the unconstrained and a.s. constrained problems.

In all scenarios, the optimal dual cost lies between the primal costs of the corresponding unconstrained and a.s. QoS constrained problems. Moreover, the optimal dual cost is closer to the primal cost of the unconstrained problem in all cases because it is costly to satisfy the QoS constraint a.s. at all times. The optimal dual cost in Scenario A is

33% higher than the optimal dual cost in Table 5.1 because there is no "free" incoming renewable energy for the operator to run the base station, increasing the expenditure. The optimal dual cost of Scenario C is 102% higher than the optimal dual cost in Table 5.1, directly implying the high weight assigned to financial profits, with no regard for its carbon footprint. The optimal cost of Scenario E confirms the intuition that it is less expensive to satisfy the QoS with low confidence (80%), compared to 90% confidence. To quantify differences in the optimal procurement policy for each scenario, we produced plots analogous to Figure 5.18, and approximated the area below the mean curve for each plot to obtain the expected energy balance. The optimal control curves for each scenario are demonstrated in Appendix D.8. Table 5.3 lists these values. The base scenario in Table 5.3 refers to the original problem described in Section 5.4.1.

Expected energy	Consumed	Battery	Bought	Sold
Base scenario	4199.6	4079.9	119.63	49.64
Scenario A	4040	3295.7	744.36	29.94
Scenario B	4584.3	4440.6	143.68	43.5
Scenario C	4402.8	2347.5	2055.2	5363.6
Scenario D	3883.3	3811.6	71.64	50.99
Scenario E	3870.8	3763.3	107.48	52.35

Table 5.3: Comparison of the expected energy balance (in Wh) using the optimal power procurement policy from the optimal value function v obtained from the dual optimization Algorithm 5 for each scenario.

In Scenario A, the operator buys six times more energy than the base scenario. The lack of incoming renewable power forces the network operator to buy more power from the grid to satisfy the QoS constraint (5.30). In Scenario C, the operator buys almost half the energy consumed by the base station from the grid and sells a considerable amount of energy stored in the battery. The optimal policy is to make as much financial profit as possible because a high weight is assigned to minimizing financial costs. The base station consumes 9% more energy in Scenario B than the base scenario because the energy demand of satisfying QoS during periods of low fading channel (extreme weather) is higher. An 8% decrease occurs in the energy consumed by the base station in Scenario E compared with the base scenario, because less energy is required to satisfy the QoS constraint (5.30) with lower confidence.

We also perform scenario simulations over probability distributions of both model and algorithmic parameters in Table D.1 and Table D.2, respectively. Table D.4 in Appendix D.7 lists the various probability distributions used to randomize the simulations. Algorithm 5 is run independently for 50 iid samples of the various parameters with no additional tuning between runs. Figure 5.20 illustrates the optimal dual cost and subgradient achieved in each case. Figure 5.20b reveals that Algorithm 5 achieves the prescribed subgradient tolerance in all 50 cases, demonstrating the robustness of the proposed numerical procedure. These results emphasize the practicality of applying the proposed numerical approach as a decision-making tool for stakeholders in the wireless communications industry.

5.5 Conclusions and outlook

In Chapter 5, we used SOC to solve the optimal power procurement problem for green cellular wireless networks powered by a hybrid energy system and equipped with an energy storage capacity under uncertainty in renewable energy and wireless fading channels. The




Figure 5.20: Optimal dual cost and subgradient norms obtained in 50 independent runs of Algorithm 5 with randomized model and algorithmic parameters as specified in Table D.4. Figure 5.20b demonstrates that Algorithm 5 achieves the prescribed subgradient norm tolerance in all 50 random scenarios, having a wide range of corresponding optimal costs as seen in Figure 5.20a.

QoS of the wireless network was posed as a probabilistic constraint with a high confidence, to address uncertainties efficiently in the system. We developed a novel Lagrangian relaxation procedure in continuous-time to relax the above probabilistic constraint, yielding a relaxed Markovian sub-problem that can be solved by the dynamic programming principle. We developed a robust numerical procedure to solve the corresponding dual problem based on a finite-dimensional approximation of the Lagrange multiplier function, and a SSM algorithm driven by the numerical solution of the associated HJB PDE. An LMBM-boosted initialization procedure and an adaptive refinement of the Lagrange multiplier function enhanced algorithm performance. The practical application of the numerical procedure was demonstrated on a model cellular base station driven by the German power system and daily cellular traffic demand. The results illustrate the viability and efficiency of the proposed approach for real-world SOC problems with a probabilistic constraint. Extensions

of this work could include the following.

1. The cellular network model described in Section 5.2.1 can be improved by including stochasticity in other dynamics, such as energy spot prices $K_b(t), K_s(t)$, and the cellular user traffic profile $N_u(t)$. Moreover, a base station on–off switching mechanism [137] can also be added to the current modelling framework. These additions would result in a more realistic algorithm that is robust to random real-life fluctuations but would increase the dimension of the associated HJB PDE. Neural network-based [144, 59, 145] or dimension-reduction-based [146, 61] approaches can be investigated to solve high-dimensional HJBs.
2. Modern-day cellular networks aim to achieve the QoS constraint with extremely high confidence, corresponding to $\epsilon \ll 10^{-3}$. The proposed numerical approach in Section 5.3 cannot achieve this, because the naïve MC method is employed to estimate the subgradient (see Section 5.3.2). One way to mitigate this problem could be to numerically solve the Fokker–Planck PDE (see Proposition 2.3.2) for the joint distribution of the optimally controlled SDE to get a deterministic estimate of the subgradient. Alternatively, IS could be applied on the path space of the optimally controlled SDE (see Section 3.2) to reduce the relative statistical error of the MC subgradient estimates.
3. A detailed modelling framework for mobile user distributions can significantly enhance the applicability of the cellular network model presented in this Chapter. As discussed in Chapter 4, MV-SDEs provide a powerful framework for representing systems of interacting users. The simulation techniques and rare event probability estimation methods developed in Chapter 4 can be integrated into the modelling and solution framework outlined here. This integration would yield a novel, sophisticated, and practically relevant cellular network model.

Appendix A

Proofs

A.1 Proof of Lemma 3.2.1

The proof of (3.24) proceeds by establishing both inequalities " \geq " and " \leq " for the value function using the principle of optimality. By constructing suitably chosen admissible controls over subintervals and applying the tower property of conditional expectation, one shows that the value function v must satisfy (3.24). First, consider " \geq ": Let $\zeta^* \in \mathcal{A}$ be the optimal control policy for $s \in (t, T)$ that minimises the second moment. From the value function definition (3.19), we obtain

$$\begin{aligned}
 v(t, \mathbf{x}) &= \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta^*}(T)) > K\}} \exp \left\{ - \int_t^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \\
 &\quad \left. \left. - 2 \int_t^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \mid \mathbf{x}^{\zeta^*}(t) = \mathbf{x} \right] \\
 &= \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta^*}(T)) > K\}} \exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \\
 &\quad \left. \left. - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
 &\quad \left. \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \\
 &\quad \left. \left. - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \mid \mathbf{x}^{\zeta^*}(t) = \mathbf{x} \right] \\
 &= \mathbb{E} \left[\mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta^*}(T)) > K\}} \exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \right. \\
 &\quad \left. \left. - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
 &\quad \left. \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
 &\quad \left. \mid \mathbf{x}^{\zeta^*}(t) = \mathbf{x}, \mathcal{F}_{t+\delta} \right] \mid \mathbf{x}^{\zeta^*}(t) = \mathbf{x} \right]. \tag{A.1}
 \end{aligned}$$

where $\mathcal{F}_{t+\delta}$ is the natural filtration of the Wiener process at time $t + \delta$. Considering Markovianity of the process \mathbf{X}^{ζ^*} , (A.1) can be expressed as

$$\begin{aligned}
v(t, \mathbf{x}) &= \mathbb{E} \left[\mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{X}^{\zeta^*}(T)) > K\}} \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{X}^{\zeta^*}(s))\|^2 ds \right. \right. \right. \\
&\quad \left. \left. - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{X}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \mid \mathbf{X}^{\zeta^*}(t+\delta) \right] \\
&\quad \exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{X}^{\zeta^*}(s))\|^2 ds \right. \\
&\quad \left. - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{X}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \mid \mathbf{X}^{\zeta^*}(t) = \mathbf{x} \Big] \\
&= \mathbb{E} \left[\exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{X}^{\zeta^*}(s))\|^2 ds - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{X}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. J_{t+\delta, \mathbf{X}^{\zeta^*}(t+\delta)}(\zeta^*) \mid \mathbf{X}^{\zeta^*}(t) = \mathbf{x} \right]. \tag{A.2}
\end{aligned}$$

As ζ^* may not be an optimal control from time $t + \delta$ to T given $\mathbf{X}^{\zeta^*}(t + \delta)$, based on the definition of the value function (3.19), we get

$$J_{t+\delta, \mathbf{X}^{\zeta^*}(t+\delta)}(\zeta^*) \geq v(t + \delta, \mathbf{X}^{\zeta^*}(t + \delta)). \tag{A.3}$$

By substituting (A.3) into (A.2), we get

$$\begin{aligned}
v(t, \mathbf{x}) &\geq \mathbb{E} \left[\exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{X}^{\zeta^*}(s))\|^2 ds - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{X}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. v(t + \delta, \mathbf{X}^{\zeta^*}(t + \delta)) \mid \mathbf{X}^{\zeta^*}(t) = \mathbf{x} \right]. \tag{A.4}
\end{aligned}$$

Taking the minimum over all controls in $[t, t + \delta]$, we obtain

$$\begin{aligned}
v(t, \mathbf{x}) &\geq \min_{\zeta^*: [t, t+\delta] \rightarrow \mathbb{R}^d} \mathbb{E} \left[\exp \left\{ - \int_t^{t+\delta} \|\zeta^*(s, \mathbf{X}^{\zeta^*}(s))\|^2 ds \right. \right. \\
&\quad \left. \left. - 2 \int_t^{t+\delta} \langle \zeta^*(s, \mathbf{X}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. v(t + \delta, \mathbf{X}^{\zeta^*}(t + \delta)) \mid \mathbf{X}^{\zeta^*}(t) = \mathbf{x} \right]. \tag{A.5}
\end{aligned}$$

Next, let us consider the second inequality " \leq ". Let $\zeta^+ \in \mathcal{A}$ be some arbitrary control from t to $t + \delta$. Given $\mathbf{X}^{\zeta^*}(t + \delta)$, let $\zeta^* \in \mathcal{A}$ be the optimal control from $t + \delta$ to T , and define a new control $\zeta' = (\zeta^+, \zeta^*) \in \mathcal{A}$ over $[t, T]$. According to (3.19),

$$\begin{aligned}
v(t, \mathbf{x}) &\leq J_{t, \mathbf{x}}(\zeta') \\
&= \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta'}(T)) > K\}} \exp \left\{ - \int_t^T \|\zeta'(s, \mathbf{x}^{\zeta'}(s))\|^2 ds \right. \right. \\
&\quad \left. \left. - 2 \int_t^T \langle \zeta'(s, \mathbf{x}^{\zeta'}(s)), d\mathbf{W}(s) \rangle \right\} \middle| \mathbf{x}^{\zeta'}(t) = \mathbf{x} \right] \\
&= \mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta'}(T)) > K\}} \exp \left\{ - \int_t^{t+\delta} \|\zeta^+(s, \mathbf{x}^{\zeta^+}(s))\|^2 ds \right. \right. \\
&\quad \left. \left. - 2 \int_t^{t+\delta} \langle \zeta^+(s, \mathbf{x}^{\zeta^+}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \\
&\quad \left. \left. - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \middle| \mathbf{x}^{\zeta'}(t) = \mathbf{x} \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta'}(T)) > K\}} \exp \left\{ - \int_t^{t+\delta} \|\zeta^+(s, \mathbf{x}^{\zeta^+}(s))\|^2 ds \right. \right. \right. \\
&\quad \left. \left. - 2 \int_t^{t+\delta} \langle \zeta^+(s, \mathbf{x}^{\zeta^+}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\
&\quad \left. \middle| \mathbf{x}^{\zeta^*}(t) = \mathbf{x}, \mathcal{F}_{t+\delta} \right] \middle| \mathbf{x}^{\zeta'}(t) = \mathbf{x} \right], \tag{A.6}
\end{aligned}$$

and we can express (A.6) as

$$\begin{aligned}
v(t, \mathbf{x}) &\leq \mathbb{E} \left[\mathbb{E} \left[\mathcal{I}_{\{g(\mathbf{x}^{\zeta^*}(T)) > K\}} \exp \left\{ - \int_{t+\delta}^T \|\zeta^*(s, \mathbf{x}^{\zeta^*}(s))\|^2 ds \right. \right. \right. \\
&\quad \left. \left. - 2 \int_{t+\delta}^T \langle \zeta^*(s, \mathbf{x}^{\zeta^*}(s)), d\mathbf{W}(s) \rangle \right\} \middle| \mathbf{x}^{\zeta^*}(t+\delta) \right] \\
&\quad \exp \left\{ - \int_t^{t+\delta} \|\zeta^+(s, \mathbf{x}^{\zeta^+}(s))\|^2 ds \right. \\
&\quad \left. - 2 \int_t^{t+\delta} \langle \zeta^+(s, \mathbf{x}^{\zeta^+}(s)), d\mathbf{W}(s) \rangle \right\} \middle| \mathbf{x}^{\zeta'}(t) = \mathbf{x} \right]. \tag{A.7}
\end{aligned}$$

Considering optimality of control ζ^* in $[t+\delta, T]$, we can express (A.7) as

$$v(t, \mathbf{x}) \leq \mathbb{E} \left[\exp \left\{ - \int_t^{t+\delta} \|\zeta^+(s, \mathbf{x}^{\zeta^+}(s))\|^2 ds - 2 \int_t^{t+\delta} \langle \zeta^+(s, \mathbf{x}^{\zeta^+}(s)), d\mathbf{W}(s) \rangle \right\} \right]$$

$$v(t + \delta, \mathbf{X}^{\zeta^*}(t + \delta)) \quad \Bigg| \quad \mathbf{X}^{\zeta'}(t) = \mathbf{x} \Bigg].$$

By taking the minimum over all controls $\zeta^+ \in \mathcal{A}$ over $[t, t + \delta]$, we get

$$\begin{aligned} v(t, \mathbf{x}) \leq \min_{\zeta^+ : [t, t + \delta] \rightarrow \mathbb{R}^d} \mathbb{E} \left[\exp \left\{ - \int_t^{t + \delta} \|\zeta^+(s, \mathbf{X}^{\zeta^+}(s))\|^2 ds \right. \right. \\ \left. \left. - 2 \int_t^{t + \delta} \langle \zeta^+(s, \mathbf{X}^{\zeta^+}(s)), d\mathbf{W}(s) \rangle \right\} \right. \\ \left. v(t + \delta, \mathbf{X}^{\zeta^*}(t + \delta)) \quad \Bigg| \quad \mathbf{X}^{\zeta'}(t) = \mathbf{x} \right]. \end{aligned} \quad (\text{A.8})$$

Equations (A.5) and (A.8) prove the equality (3.24).

A.2 Proof of Theorem 3.2.1

Theorem 3.2.1 is derived by applying the dynamic programming Lemma 3.2.1 over an infinitesimal time interval and using the continuity of the value function and cost functions. From Lemma 3.2.1, the value function defined in (3.19) satisfies the following equation:

$$\begin{aligned} v(t, \mathbf{x}) = \min_{\zeta \in \mathcal{A}} \mathbb{E} \left[\exp \left\{ - \int_t^{t + \delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right\} \exp \left\{ - 2 \int_t^{t + \delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right\} \right. \\ \left. v(t + \delta, \mathbf{X}^\zeta(t + \delta)) \quad \Bigg| \quad \mathbf{X}^\zeta(t) = \mathbf{x} \right]. \end{aligned} \quad (\text{A.9})$$

Since we have a multiplicative cost structure, we use Taylor–Maclaurin expansions of the exponential function and bound the remainder terms to derive the HJB equation. First, we have the following Taylor–Maclaurin expansions for $\exp(x)$ for $\delta \rightarrow 0$:

$$\begin{aligned} \exp \left\{ - \int_t^{t + \delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right\} = 1 - \int_t^{t + \delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \\ + \frac{1}{2} \left(\int_t^{t + \tau} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^2, \end{aligned} \quad (\text{A.10})$$

$$\begin{aligned} \exp \left\{ - 2 \int_t^{t + \delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right\} = 1 - 2 \int_t^{t + \delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \\ + 2 \left(\int_t^{t + \delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 - \frac{4}{3} \left(\int_t^{t + \tau} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^3, \end{aligned} \quad (\text{A.11})$$

for some $0 < \tau < \delta$. Next, we write down Itô's formula for $v(t + \delta, \mathbf{X}^\zeta(t + \delta))$,

$$\begin{aligned} v(t + \delta, \mathbf{X}^\zeta(t + \delta)) = v(t, \mathbf{X}^\zeta(t)) + \int_t^{t + \delta} \left(\frac{\partial v}{\partial s}(s, \mathbf{X}^\zeta(s)) \right. \\ \left. + \langle \mathbf{a}(s, \mathbf{X}^\zeta(s)) + \mathbf{b}(s, \mathbf{X}^\zeta(s))\zeta(s, \mathbf{X}^\zeta(s)), \nabla v(s, \mathbf{X}^\zeta(s)) \rangle \right) ds \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{2} (\mathbf{b}(s, \mathbf{X}^\zeta(s)) \mathbf{b}(s, \mathbf{X}^\zeta(s))^{\text{tr}} : \nabla^2 v) \, ds \\
& + \int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle. \tag{A.12}
\end{aligned}$$

Substituting (A.10), (A.11), and (A.12) in (A.9), we get

$$\begin{aligned}
v(t, \mathbf{x}) = \min_{\zeta \in \mathcal{A}} & \left\{ v(t, \mathbf{x}) + v(t, \mathbf{x}) \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \right. \\
& + \mathbb{E} \left[\int_t^{t+\delta} \left(\frac{\partial v}{\partial s}(s, \mathbf{X}^\zeta(s)) \right. \right. \\
& \quad + \langle \mathbf{a}(s, \mathbf{X}^\zeta(s)) + \mathbf{b}(s, \mathbf{X}^\zeta(s)) \zeta(s, \mathbf{X}^\zeta(s)), \nabla v(s, \mathbf{X}^\zeta(s)) \rangle \\
& \quad \left. \left. + \frac{1}{2} (\mathbf{b}(s, \mathbf{X}^\zeta(s)) \mathbf{b}(s, \mathbf{X}^\zeta(s))^{\text{tr}} : \nabla^2 v(s, \mathbf{X}^\zeta(s))) \, ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \right. \\
& \left. - 2 \mathbb{E} \left[\int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \zeta(s, \mathbf{X}^\zeta(s)), \nabla v(s, \mathbf{X}^\zeta(s)) \rangle ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] + \mathcal{R} \right\}. \tag{A.13}
\end{aligned}$$

Here \mathcal{R} is the residual term. By repeatedly applying Hölder's inequality, Itô isometry and Burkholder-Davis-Gundy inequality (to bound higher moments of martingales), one can prove that \mathcal{R} is bounded and $\mathcal{O}(\delta^{\frac{3}{2}})$. We illustrate this on some leading terms in \mathcal{R} .

$$\begin{aligned}
\text{i. } & 2 \mathbb{E} \left[\left(\int_t^{t+\delta} \left(\frac{\partial v}{\partial s} + \langle \mathbf{a} + \mathbf{b} \zeta, \nabla v \rangle + \frac{1}{2} (\mathbf{b} \mathbf{b}^{\text{tr}} : \nabla^2 v) \right) (s, \mathbf{X}^\zeta(s)) ds \right) \right. \\
& \quad \left. \left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right) \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\
& \leq 2 \mathbb{E} \left[\left(\int_t^{t+\delta} \left(\frac{\partial v}{\partial s} + \langle \mathbf{a} + \mathbf{b} \zeta, \nabla v \rangle + \frac{1}{2} (\mathbf{b} \mathbf{b}^{\text{tr}} : \nabla^2 v) \right) (s, \mathbf{X}^\zeta(s)) ds \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq 2 \delta^{\frac{1}{2}} \mathbb{E} \left[\int_t^{t+\delta} \left| \left(\frac{\partial v}{\partial s} + \langle \mathbf{a} + \mathbf{b} \zeta, \nabla v \rangle + \frac{1}{2} (\mathbf{b} \mathbf{b}^{\text{tr}} : \nabla^2 v) \right) (s, \mathbf{X}^\zeta(s)) \right|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq \delta^{\frac{3}{2}} \mathbb{E} \left[\sup_{t < s < t+\delta} \left| \left(\frac{\partial v}{\partial s} + \langle \mathbf{a} + \mathbf{b} \zeta, \nabla v \rangle + \frac{1}{2} (\mathbf{b} \mathbf{b}^{\text{tr}} : \nabla^2 v) \right) (s, \mathbf{X}^\zeta(s)) \right|^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\sup_{t < s < t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}}.
\end{aligned}$$

$$\begin{aligned}
\text{ii. } & 2\mathbb{E} \left[\left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 \right. \\
& \quad \left. \left(\int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right) \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\
& \leq 2\mathbb{E} \left[\left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^4 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\left(\int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq 2C_{\text{BDG}} \mathbb{E} \left[\left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\int_t^{t+\delta} \|\mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq 2\delta^{\frac{1}{2}} C_{\text{BDG}} \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^4 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\int_t^{t+\delta} \|\mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq 2\delta^{\frac{3}{2}} C_{\text{BDG}} \mathbb{E} \left[\sup_{t < s < t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^4 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \quad \mathbb{E} \left[\sup_{t < s < t+\delta} \|\mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s))\|^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}}.
\end{aligned}$$

$$\begin{aligned}
\text{iii. } & \frac{4}{3} \mathbb{E} \left[v(t, \mathbf{X}^\zeta(t)) \left(\int_t^{t+\tau} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^3 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\
& \leq \frac{4}{3} v(t, \mathbf{x}) C_{\text{BDG}} \mathbb{E} \left[\left(\int_t^{t+\tau} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^{\frac{3}{2}} \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\
& \leq \frac{4}{3} v(t, \mathbf{x}) C_{\text{BDG}} \mathbb{E} \left[\left(\int_t^{t+\tau} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^3 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq \frac{4}{3} v(t, \mathbf{x}) C_{\text{BDG}} \mathbb{E} \left[\left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^3 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\
& \leq \frac{4}{3} \delta v(t, \mathbf{x}) C_{\text{BDG}} \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^6 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}}
\end{aligned}$$

$$\leq \frac{4}{3} \delta^{\frac{3}{2}} v(t, \mathbf{x}) C_{\text{BDG}} \mathbb{E} \left[\sup_{t < s < t + \delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^6 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}}.$$

$$\begin{aligned} \text{iv. } & \mathbb{E} \left[\left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right) \right. \\ & \quad \left. \left(\int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right) \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\ & \leq \mathbb{E} \left[\left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \quad \mathbb{E} \left[\left(\int_t^{t+\delta} \langle \mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \leq \delta^{\frac{1}{2}} \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^4 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \quad \mathbb{E} \left[\int_t^{t+\delta} \|\mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \leq \delta^{\frac{3}{2}} \mathbb{E} \left[\sup_{t < s < t + \delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^4 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \quad \mathbb{E} \left[\sup_{t < s < t + \delta} \|\mathbf{b}(s, \mathbf{X}^\zeta(s)) \nabla v(s, \mathbf{X}^\zeta(s))\|^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}}. \end{aligned}$$

$$\begin{aligned} \text{v. } & 2\mathbb{E} \left[v(t, \mathbf{X}^\zeta(t)) \left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right) \right. \\ & \quad \left. \left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right) \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right] \\ & \leq 2v(t, \mathbf{x}) \mathbb{E} \left[\left(\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \quad \mathbb{E} \left[\left(\int_t^{t+\delta} \langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle \right)^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \leq 2v(t, \mathbf{x}) \delta^{\frac{1}{2}} \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^4 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \\ & \quad \mathbb{E} \left[\int_t^{t+\delta} \|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds \mid \mathbf{X}^\zeta(t) = \mathbf{x} \right]^{\frac{1}{2}} \end{aligned}$$

$$\leq 2v(t, \mathbf{x})\delta^{\frac{3}{2}}\mathbb{E}\left[\sup_{t < s < t+\delta}\|\zeta(s, \mathbf{X}^\zeta(s))\|^4 \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right]^{\frac{1}{2}}$$

$$\mathbb{E}\left[\sup_{t < s < t+\delta}\|\zeta(s, \mathbf{X}^\zeta(s))\|^2 \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right]^{\frac{1}{2}}.$$

Taking the limit $\delta \rightarrow 0$ of the dynamic programming lemma (A.9),

$$0 = \lim_{\delta \rightarrow 0} \min_{\zeta: [t, t+\delta] \rightarrow \mathbb{R}^d} \frac{1}{\delta} \mathbb{E}\left[\exp\left\{-\int_t^{t+\delta}\|\zeta(s, \mathbf{X}^\zeta(s))\|^2 ds\right\}\right. \\ \left.\exp\left\{-2\int_t^{t+\delta}\langle \zeta(s, \mathbf{X}^\zeta(s)), d\mathbf{W}(s) \rangle\right\} v(t+\delta, \mathbf{X}^\zeta(t+\delta))\right. \\ \left.- v(t, \mathbf{x}) \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right]. \quad (\text{A.14})$$

Inserting (A.13) to (A.14),

$$0 = \lim_{\delta \rightarrow 0} \min_{\zeta: [t, t+\delta] \rightarrow \mathbb{R}^d} \frac{1}{\delta} \left[v(t, \mathbf{x}) + v(t, \mathbf{x}) \mathbb{E}\left[\int_t^{t+\delta}\|\zeta(s, \mathbf{X}^\zeta(s))\|^2 dt \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right] \right. \\ \left. + \mathbb{E}\left[\int_t^{t+\delta}\left(\frac{\partial v}{\partial s} + \langle \mathbf{a} + \mathbf{b}\zeta, \nabla v \rangle + \frac{1}{2}(\mathbf{b}\mathbf{b}^{\text{tr}}) : \nabla^2 v\right)(s, \mathbf{X}^\zeta(s)) ds \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right] \right. \\ \left. - 2\mathbb{E}\left[\int_t^{t+\delta}\langle \mathbf{b}(s, \mathbf{X}^\zeta(s))\zeta(s, \mathbf{X}^\zeta(s)), \nabla v(s, \mathbf{X}^\zeta(s)) \rangle ds \mid \mathbf{X}^\zeta(t) = \mathbf{x}\right] \right. \\ \left. + \mathcal{O}(\delta^{\frac{3}{2}}) - v(t, \mathbf{x}) \right]$$

Given the regularity of v required to apply Itô formula and the regularity of $\mathbf{a}, \mathbf{b}, \zeta$ required for the existence and uniqueness of a solution to (3.17), we can apply the limit $\delta \rightarrow 0$ and condition on $\mathbf{X}^\zeta(t) = \mathbf{x}$ to get,

$$\min_{\alpha \in \mathbf{U}} \left\{ \frac{\partial v}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a}(t, \mathbf{x}), \nabla v(t, \mathbf{x}) \rangle + \frac{1}{2}(\mathbf{b}(t, \mathbf{x})\mathbf{b}(t, \mathbf{x})^{\text{tr}}) : \nabla^2 v(t, \mathbf{x}) \right. \\ \left. - \langle \mathbf{b}(t, \mathbf{x})\alpha, \nabla v(t, \mathbf{x}) \rangle + \|\alpha\|^2 v(t, \mathbf{x}) \right\} = 0 \\ \Rightarrow \frac{\partial v}{\partial t}(t, \mathbf{x}) + \langle \mathbf{a}(t, \mathbf{x}), \nabla v(t, \mathbf{x}) \rangle + \frac{1}{2}(\mathbf{b}(t, \mathbf{x})\mathbf{b}(t, \mathbf{x})^{\text{tr}}) : \nabla^2 v(t, \mathbf{x}) \\ + \min_{\alpha \in \mathbf{U}} \left\{ \|\alpha\|^2 v(t, \mathbf{x}) - \langle \mathbf{b}(t, \mathbf{x})\alpha, \nabla v(t, \mathbf{x}) \rangle \right\} = 0. \quad (\text{A.15})$$

For fixed $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$, the Markov control function ζ appears in the quantity to be minimised through $\alpha \in \mathbf{U}$. That is, the optimal Markov control function ζ^* is defined

by the map between each $(t, \mathbf{x}) \in [0, T] \times \mathbb{R}^d$ and $\mathbf{a} \in \mathcal{U}$. Neglecting the trivial solution $v(t, \mathbf{x}) = 0$, we obtain the minimiser for (A.15) as in (3.27). By substituting optimal control function ζ^* in (A.15), we get (3.26), which solves for value function v .

A.3 Proof of Theorem 4.5.1

To establish the optimal computational complexity of the DLMC estimator, we formulated a constrained optimization problem (4.44) that minimises total computational work subject to the prescribed bias and variance constraints. This problem is solved using the Lagrangian multiplier method, yielding the optimal parameters. We have the corresponding Lagrangian,

$$\begin{aligned} \mathcal{L} = & M_1 NP^2 + M_1 M_2 NP + \lambda_1 \left(\frac{C_1}{P} + \frac{C_2}{N} - \theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]| \right) \\ & + \lambda_2 \left(C_v^2 \left(\frac{C_3}{PM_1} + \frac{C_4}{M_1 M_2} \right) - (1 - \theta)^2 \text{TOL}_r^2 |\mathbb{E}[g(\mathbf{X}(T))]|^2 \right), \end{aligned} \quad (\text{A.16})$$

where $\lambda_1, \lambda_2 \in \mathbb{R}^+$ are Lagrange multipliers. Hence, we obtain optimality conditions for (A.16) as follows:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial M_1} = 0 & \implies NP^2 + M_2 NP = \lambda_2 C_v^2 \left(\frac{C_3}{PM_1^2} + \frac{C_4}{M_2 M_1^2} \right), \\ \frac{\partial \mathcal{L}}{\partial M_2} = 0 & \implies M_1 NP = \frac{\lambda_2 C_4 C_v^2}{M_1 M_2^2}, \\ \frac{\partial \mathcal{L}}{\partial N} = 0 & \implies M_1 P^2 + M_1 M_2 P = \frac{\lambda_1 C_2}{N^2}, \\ \frac{\partial \mathcal{L}}{\partial P} = 0 & \implies 2M_1 NP + M_1 M_2 N = \frac{\lambda_1 C_1}{P^2} + \frac{\lambda_2 C_v^2 C_3}{M_1 P^2}, \\ \frac{\partial \mathcal{L}}{\partial \lambda_1} = 0 & \implies \frac{C_1}{P} + \frac{C_2}{N} = \theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|, \\ \frac{\partial \mathcal{L}}{\partial \lambda_2} = 0 & \implies C_v^2 \left(\frac{C_3}{PM_1} + \frac{C_4}{M_1 M_2} \right) = (1 - \theta)^2 \text{TOL}_r^2 |\mathbb{E}[g(\mathbf{X}(T))]|^2. \end{aligned}$$

By solving the above equations for $P, N, M_1, M_2, \lambda_1$, and λ_2 , we get

$$\begin{aligned} P &= \frac{2C_1}{\theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|} = \mathcal{O}(\text{TOL}_r^{-1}), \\ N &= \frac{2C_2}{\theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|} = \mathcal{O}(\text{TOL}_r^{-1}), \\ M_1 &= \frac{\theta C_v^2 (C_3 + \sqrt{C_3 C_4})}{2C_1 (1 - \theta)^2 \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|} = \mathcal{O}(\text{TOL}_r^{-1}), \\ M_2 &= \frac{2C_1 \sqrt{C_4}}{\sqrt{C_3} \theta \text{TOL}_r |\mathbb{E}[g(\mathbf{X}(T))]|} = \mathcal{O}(\text{TOL}_r^{-1}), \end{aligned} \quad (\text{A.17})$$

By substituting optimal parameters (A.17) into (4.43), we get

$$\mathcal{W} = M_1 NP^2 + M_1 M_2 NP = \frac{4C_1 C_2 C_v^2 (\sqrt{C_3} + \sqrt{C_4})^2}{\theta^2 (1 - \theta)^2 |\mathbb{E}[g(\mathbf{X}(T))]|^4 \text{TOL}_r^4} = \mathcal{O}(\text{TOL}_r^{-4}). \quad (\text{A.18})$$

A.4 Proof of Theorem 4.6.1

We derive the complexity rate of the **MLDLMC** estimator under Assumptions 4.6.1 and 4.6.2. The following analysis reveals how the **MLDLMC** estimator achieves improved asymptotic complexity compared to the **DLMC** estimator. For a given level $L \in \mathbb{N}$, (4.69) provides the optimal number of samples to satisfy the variance constraint in (4.61) for the proposed **MLDLMC** estimator (4.54). We bound the estimator's computational cost as follows:

$$\begin{aligned}
\mathcal{W}[\mathcal{A}_{\text{MLDLMC}}] &\lesssim \sum_{\ell=0}^L (M_{1,\ell} P_\ell^2 N_\ell + M_{1,\ell} M_{2,\ell} P_\ell N_\ell) \\
&\lesssim \sum_{\ell=0}^L \left((\mathcal{M}_{1,\ell} + 1) P_\ell^2 N_\ell + (\mathcal{M}_{1,\ell} + 1) \left(\frac{\tilde{\mathcal{M}}_\ell}{\lceil \mathcal{M}_{1,\ell} \rceil} + 1 \right) P_\ell N_\ell \right) \\
&\leq \underbrace{\sum_{\ell=0}^L (\mathcal{M}_{1,\ell} P_\ell^2 N_\ell + \tilde{\mathcal{M}}_\ell P_\ell N_\ell)}_{=: W_1 \text{ (objective function of (4.67))}} \\
&\quad + \underbrace{\sum_{\ell=0}^L (P_\ell^2 N_\ell + P_\ell N_\ell)}_{=: W_2 \text{ (cost of one realization per level)}} + \underbrace{\sum_{\ell=0}^L \mathcal{M}_{1,\ell} P_\ell N_\ell}_{=: W_3} + \underbrace{\sum_{\ell=0}^L \frac{\tilde{\mathcal{M}}_\ell}{\lceil \mathcal{M}_{1,\ell} \rceil} P_\ell N_\ell}_{=: W_4}.
\end{aligned} \tag{A.19}$$

Because $P_\ell > 1$, W_3 is always dominated by the first term in W_1 (i.e., $\sum_{\ell=0}^L \mathcal{M}_{1,\ell} P_\ell^2 N_\ell$). We analyse each term individually. By substituting (4.68) in W_1 , we have that

$$W_1 = \frac{C_v^2}{(1-\theta)^2 \text{TOL}_r^2 |\mathbb{E}[G]|^2} \left(\sum_{\ell=0}^L \sqrt{P_\ell N_\ell} (\sqrt{V_{1,\ell} P_\ell} + \sqrt{V_{2,\ell}}) \right)^2. \tag{A.20}$$

By selecting level L that satisfies (4.58) and using Assumption 4.6.1,

$$\frac{|\mathbb{E}[G - G_L]|}{|\mathbb{E}[G]|} \leq C_b 2^{-\tilde{\alpha} L} \approx \theta \text{TOL}_r, \tag{A.21}$$

and hence,

$$L = \left\lceil \frac{1}{\tilde{\alpha}} \log \left(\frac{C_b}{\theta |\mathbb{E}[G]|} \text{TOL}_r^{-1} \right) \right\rceil. \tag{A.22}$$

Using the hierarchies (4.52) and Assumption 4.6.2 in (A.20),

$$W_1 \lesssim \text{TOL}_r^{-2} \left(\sum_{\ell=0}^L 2^{\frac{3-\tilde{w}}{2}\ell} + 2^{\frac{2-\tilde{s}}{2}\ell} \right)^2. \tag{A.23}$$

The summation in (A.23) has two terms; thus, we obtain the following two cases:

- **Case 1:** $\tilde{s} + 1 \geq \tilde{w}$, i.e., the first term dominates. W_1 can be expressed as

$$W_1 \lesssim \text{TOL}_r^{-2} \left(\sum_{\ell=0}^L 2^{\frac{3-\tilde{w}}{2}\ell} \right)^2, \tag{A.24}$$

and simplified using (A.22) and the sum of a geometric series,

$$W_1 \lesssim \begin{cases} \text{TOL}_r^{-2} & , \tilde{w} > 3, \\ \text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2 & , \tilde{w} = 3, \\ \text{TOL}_r^{-2 - (\frac{3-\tilde{w}}{\alpha})} & , \tilde{w} < 3, \end{cases} \quad (\text{A.25})$$

which can be expressed more compactly:

$$W_1 \lesssim \text{TOL}_r^{-2 - \max(0, \frac{3-\tilde{w}}{\alpha})} (\log \text{TOL}_r^{-1})^2 \mathcal{J}_1, \quad (\text{A.26})$$

where

$$\mathcal{J}_1 := \begin{cases} 1, & \text{if } \tilde{w} = 3, \\ 0, & \text{otherwise.} \end{cases}$$

- **Case 2:** $\tilde{s} + 1 < \tilde{w}$, i.e., the second term dominates. W_1 can be expressed as

$$W_1 \lesssim \text{TOL}_r^{-2} \left(\sum_{\ell=0}^L 2^{\frac{2-\tilde{s}}{2}\ell} \right)^2, \quad (\text{A.27})$$

and simplify using the sum of a geometric series,

$$W_1 \lesssim \begin{cases} \text{TOL}_r^{-2} & , \tilde{s} > 2, \\ \text{TOL}_r^{-2} (\log \text{TOL}_r^{-1})^2 & , \tilde{s} = 2, \\ \text{TOL}_r^{-2 - (\frac{2-\tilde{s}}{\alpha})} & , \tilde{s} < 2, \end{cases} \quad (\text{A.28})$$

which can be expressed more compactly:

$$W_1 \lesssim \text{TOL}_r^{-2 - \max(0, \frac{2-\tilde{s}}{\alpha})} (\log \text{TOL}_r^{-1})^2 \mathcal{J}_2, \quad (\text{A.29})$$

where

$$\mathcal{J}_2 := \begin{cases} 1, & \text{if } \tilde{s} = 2, \\ 0, & \text{otherwise.} \end{cases}$$

Using (4.52), for W_2 ,

$$W_2 \lesssim \text{TOL}_r^{\frac{3}{\alpha}}. \quad (\text{A.30})$$

Next, we examine W_4 :

$$\begin{aligned} W_4 &= \sum_{\ell=0}^L \frac{\tilde{\mathcal{M}}_\ell}{\lceil \mathcal{M}_{1,\ell} \rceil} P_\ell N_\ell \\ &\leq \sum_{\ell=0}^L \frac{\tilde{\mathcal{M}}_\ell}{\max\{1, \mathcal{M}_{1,\ell}\}} P_\ell N_\ell \\ &= \sum_{\ell=0}^L \min\{\tilde{\mathcal{M}}_\ell P_\ell N_\ell, \frac{\tilde{\mathcal{M}}_\ell}{\mathcal{M}_{1,\ell}} P_\ell N_\ell\} \\ &\leq \sum_{\ell=0}^L \tilde{\mathcal{M}}_\ell P_\ell N_\ell. \end{aligned} \quad (\text{A.31})$$

W_4 has the same or lower order than that of W_1 . Next, we must ensure that W_1 dominates W_2 as $\text{TOL}_r \rightarrow 0$ for the proposed **MIDL**MC method to be feasible. Comparing (A.30) to W_1 for the cases, the following condition ensures W_1 is the dominant term:

$$\tilde{\alpha} \geq \frac{1}{2} \min(\tilde{w}, 1 + \tilde{s}, 3). \quad (\text{A.32})$$

Thus, (A.32), (A.26), and (A.29) complete the proof.

A.5 Proof of Theorem 4.7.1

We derive the complexity rate of the **MIDL**MC estimator under Assumptions 4.7.1 and 4.7.2. The following analysis reveals how the **MIDL**MC estimator achieves improved asymptotic complexity compared to the **MLDL**MC estimator. We start by explicitly writing the optimal index set (4.115) for a given level $v \in \mathbb{R}^+$:

$$\mathcal{I}(v) = \left\{ \boldsymbol{\alpha} \in \mathbb{N}^2 : \frac{\exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle}{\exp\langle \bar{\mathbf{g}}, \boldsymbol{\alpha} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \boldsymbol{\alpha} \rangle} \geq v \right\}. \quad (\text{A.33})$$

We make use of the following optimality result to prove Theorem 4.7.1.

Lemma A.5.1 (Optimal index sets). *The set $\mathcal{I}(v)$ in (A.33) is optimal in the sense that any other set $\tilde{\mathcal{I}}$, such that $\tilde{B}(\mathcal{I}(v)) = \tilde{B}(\tilde{\mathcal{I}})$, is a set with more work (i.e., $\tilde{W}_1(\mathcal{I}(v)) < \tilde{W}_1(\tilde{\mathcal{I}})$).*

Proof. For any $\boldsymbol{\alpha} \in \mathcal{I}(v)$ and $\hat{\boldsymbol{\alpha}} \notin \mathcal{I}(v)$,

$$\frac{\exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle}{\exp\langle \bar{\mathbf{g}}, \boldsymbol{\alpha} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \boldsymbol{\alpha} \rangle} \geq v, \quad (\text{A.34})$$

$$\frac{\exp\langle -\boldsymbol{\rho}, \hat{\boldsymbol{\alpha}} \rangle}{\exp\langle \bar{\mathbf{g}}, \hat{\boldsymbol{\alpha}} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\boldsymbol{\alpha}} \rangle} < v. \quad (\text{A.35})$$

We divide \mathbb{N}^2 into the following disjoint sets:

$$\begin{aligned} \mathcal{J}_1 &:= \mathcal{I}(v) \cap \tilde{\mathcal{I}}^c, & \mathcal{J}_2 &:= \mathcal{I}(v) \cap \tilde{\mathcal{I}}, \\ \mathcal{J}_3 &:= \mathcal{I}(v)^c \cap \tilde{\mathcal{I}}, & \mathcal{J}_4 &:= \mathcal{I}(v)^c \cap \tilde{\mathcal{I}}^c, \end{aligned}$$

where $\mathcal{I}(v)^c$ is the complement of the set $\mathcal{I}(v)$. By construction of $\tilde{\mathcal{I}}$, we obtain

$$\begin{aligned} \tilde{B}(\mathcal{I}(v)) &= \tilde{B}(\tilde{\mathcal{I}}) \\ \sum_{\boldsymbol{\alpha} \in \mathcal{J}_3 \cup \mathcal{J}_4} \exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle &= \sum_{\boldsymbol{\alpha} \in \mathcal{J}_1 \cup \mathcal{J}_4} \exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle \\ \sum_{\boldsymbol{\alpha} \in \mathcal{J}_3} \exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle &= \sum_{\boldsymbol{\alpha} \in \mathcal{J}_1} \exp\langle -\boldsymbol{\rho}, \boldsymbol{\alpha} \rangle. \end{aligned} \quad (\text{A.36})$$

In addition, \mathcal{J}_1 and \mathcal{J}_3 cannot be empty for sets satisfying $\tilde{B}(\mathcal{I}(v)) = \tilde{B}(\tilde{\mathcal{I}})$, unless $\mathcal{I}(v) = \tilde{\mathcal{I}}$. We consider the difference between the work of the two sets:

$$\begin{aligned} \tilde{W}_1(\mathcal{I}(v)) - \tilde{W}_1(\tilde{\mathcal{I}}) &= \sum_{\boldsymbol{\alpha} \in \mathcal{J}_1 \cup \mathcal{J}_2} \exp\langle \bar{\mathbf{g}}, \hat{\boldsymbol{\alpha}} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\boldsymbol{\alpha}} \rangle \\ &\quad - \sum_{\boldsymbol{\alpha} \in \mathcal{J}_2 \cup \mathcal{J}_3} \exp\langle \bar{\mathbf{g}}, \hat{\boldsymbol{\alpha}} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\boldsymbol{\alpha}} \rangle \\ &= \sum_{\boldsymbol{\alpha} \in \mathcal{J}_1} \exp\langle \bar{\mathbf{g}}, \hat{\boldsymbol{\alpha}} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\boldsymbol{\alpha}} \rangle \end{aligned} \quad (\text{A.37})$$

$$- \sum_{\alpha \in \mathcal{J}_3} \exp\langle \bar{\mathbf{g}}, \hat{\alpha} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\alpha} \rangle.$$

From (A.34) and (A.35),

$$\forall \alpha \in \mathcal{J}_1, \quad \exp\langle \bar{\mathbf{g}}, \hat{\alpha} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\alpha} \rangle \leq \frac{1}{v} \exp\langle -\rho, \alpha \rangle \quad \text{and}, \quad (\text{A.38})$$

$$\forall \alpha \in \mathcal{J}_3, \quad -\exp\langle \bar{\mathbf{g}}, \hat{\alpha} \rangle - \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\alpha} \rangle < -\frac{1}{v} \exp\langle -\rho, \alpha \rangle. \quad (\text{A.39})$$

By inserting (A.38) and (A.39) into (A.37), we obtain

$$\begin{aligned} \tilde{W}_1(\mathcal{I}(v)) - \tilde{W}_1(\bar{\mathcal{I}}) &\leq \frac{1}{v} \underbrace{\left(\sum_{\alpha \in \mathcal{J}_1} \exp\langle -\rho, \alpha \rangle - \sum_{\alpha \in \mathcal{J}_3} \exp\langle -\rho, \alpha \rangle \right)}_{=0 \text{ from (A.36)}} \\ &\Rightarrow \tilde{W}_1(\mathcal{I}(v)) \leq \tilde{W}_1(\bar{\mathcal{I}}). \end{aligned}$$

□

The index set $\mathcal{I}(v)$ is optimal in the sense of Lemma A.5.1; thus, we bound $\tilde{W}_1(\mathcal{I}(v))$ by the optimal work performed by an index set that can be analysed according to the classical multi-index setting [24]. From the definition of the optimal index set (4.115), two obvious candidates exist for such index sets:

$$\mathcal{I}_{\bar{\delta}}(\bar{L}) := \{ \alpha \in \mathbb{N}^2 : \langle \bar{\delta}, \alpha \rangle \leq \bar{L} \}, \quad (\text{A.40})$$

$$\mathcal{I}_{\bar{\bar{\delta}}}(\bar{L}) := \{ \alpha \in \mathbb{N}^2 : \langle \bar{\bar{\delta}}, \alpha \rangle \leq \bar{L} \}. \quad (\text{A.41})$$

Next, we separately analyse the optimal work of the two index sets.

Set 1: $\mathcal{I}_{\bar{\delta}}$

Following [24], a set threshold $\bar{L}_1 \in \mathbb{R}^+$ exists such that the following bias constraint is satisfied:

$$\lim_{\text{TOL}_r \downarrow 0} \frac{\tilde{B}(\mathcal{I}_{\bar{\delta}}(\bar{L}_1))}{\frac{(1-\theta)\text{TOL}_r \mathbb{E}[G]}{Q_B}} \leq 1. \quad (\text{A.42})$$

For \bar{L}_1 that satisfies (A.42), we estimate the work of the corresponding index set:

$$\begin{aligned} \tilde{W}_1(\mathcal{I}_{\bar{\delta}}(\bar{L}_1)) &= \sum_{\{\alpha \in \mathbb{N}^2 : \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_1\}} (\exp\langle \bar{\mathbf{g}}, \hat{\alpha} \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\alpha} \rangle) \\ &= \underbrace{\sum_{\{\alpha \in \mathbb{N}^2 : \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_1\}} \exp\langle \bar{\mathbf{g}}, \hat{\alpha} \rangle}_{=:\tilde{W}_{11}} + \underbrace{\sum_{\{\alpha \in \mathbb{N}^2 : \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_1\}} \exp\langle \bar{\bar{\mathbf{g}}}, \hat{\alpha} \rangle}_{=:\tilde{W}_{12}}. \end{aligned} \quad (\text{A.43})$$

From [24], we know that, as $\text{TOL}_r \rightarrow 0$,

$$\tilde{W}_{11} \lesssim \begin{cases} \left(\log \text{TOL}_r^{-1} \right)^{d_1}, & \chi_{11} \leq 0 \\ \text{TOL}_r^{-\frac{\chi_{11}}{\eta_1}} \left(\log \text{TOL}_r^{-1} \right)^{j_{11}}, & \chi_{11} > 0 \end{cases}, \quad (\text{A.44})$$

and

$$\tilde{W}_{12} \lesssim \begin{cases} \left(\log \text{TOL}_r^{-1}\right)^{d_2}, & \chi_{12} \leq 0 \\ \text{TOL}_r^{-\frac{\chi_{12}}{\eta_1}} \left(\log \text{TOL}_r^{-1}\right)^{j_{12}}, & \chi_{12} > 0 \end{cases}, \quad (\text{A.45})$$

with the following additional notation:

$$\begin{aligned} j_{11} &:= (e_1 - 1) \left(1 + \frac{\chi_{11}}{\eta_1}\right), \\ j_{12} &:= (\aleph_1 - 1) + (e_1 - 1) \frac{\chi_{11}}{\eta_1}. \end{aligned}$$

Set 2: $\mathcal{I}_{\bar{\delta}}$

Following [24], a set threshold $\bar{L}_2 \in \mathbb{R}^+$ exists such that the bias constraint (A.42) is satisfied. The work of the corresponding index set is written as follows:

$$\begin{aligned} \tilde{W}_1(\mathcal{I}_{\bar{\delta}}(\bar{L}_2)) &= \sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_2\}} (\exp\langle \bar{\mathbf{g}}, \alpha \rangle + \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle) \\ &= \underbrace{\sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_2\}} \exp\langle \bar{\mathbf{g}}, \alpha \rangle}_{=:\tilde{W}_{21}} + \underbrace{\sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}_2\}} \exp\langle \bar{\bar{\mathbf{g}}}, \alpha \rangle}_{=:\tilde{W}_{22}}. \end{aligned} \quad (\text{A.46})$$

From [24], we know that, as $\text{TOL}_r \rightarrow 0$, asymptotically

$$\tilde{W}_{21} \lesssim \begin{cases} \left(\log \text{TOL}_r^{-1}\right)^{d_1}, & \chi_{21} \leq 0 \\ \text{TOL}_r^{-\frac{\chi_{21}}{\eta_2}} \left(\log \text{TOL}_r^{-1}\right)^{j_{21}}, & \chi_{21} > 0 \end{cases}, \quad (\text{A.47})$$

and

$$\tilde{W}_{22} \lesssim \begin{cases} \left(\log \text{TOL}_r^{-1}\right)^{d_2}, & \chi_{22} \leq 0 \\ \text{TOL}_r^{-\frac{\chi_{22}}{\eta_2}} \left(\log \text{TOL}_r^{-1}\right)^{j_{22}}, & \chi_{22} > 0 \end{cases}, \quad (\text{A.48})$$

with the following additional constants:

$$\begin{aligned} j_{21} &:= (\aleph_2 - 1) + (e_2 - 1) \frac{\chi_{21}}{\eta_2}, \\ j_{22} &:= (e_2 - 1) \left(1 + \frac{\chi_{22}}{\eta_2}\right). \end{aligned}$$

From Lemma A.5.1, we can bound the optimal work of the index set (4.115) as follows:

$$\tilde{W}_1(\mathcal{I}(v_1)) \leq \min(\tilde{W}_1(\mathcal{I}_{\bar{\delta}}(\bar{L}_1)), \tilde{W}_1(\mathcal{I}_{\bar{\delta}}(\bar{L}_2))). \quad (\text{A.49})$$

By inserting (A.44) and (A.45) into (A.43) and (A.47), inserting (A.48) into (A.46), and combining all into (A.49), we obtain the following:

$$\tilde{W}_1(\mathcal{I}(v_1)) \lesssim \text{TOL}_r^{-\zeta} \left(\log \text{TOL}_r^{-1}\right)^\varrho. \quad (\text{A.50})$$

We also analyse $\tilde{W}_2(\mathcal{I}(v_1))$ because we must ensure that this work does not dominate the first term in (4.101). Following the same analysis of the work as for \tilde{W}_1 , we have

$$\begin{aligned} \tilde{W}_2(\mathcal{I}(v_1)) &= \sum_{\{\alpha \in \mathbb{N}^2: \exp\langle \bar{\delta}, \alpha \rangle + \exp\langle \bar{\delta}, \alpha \rangle \leq L\}} \exp\langle \lambda, \alpha \rangle \\ &\leq \min \left(\sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}\}} \exp\langle \lambda, \alpha \rangle, \sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}\}} \exp\langle \lambda, \alpha \rangle \right). \end{aligned} \quad (\text{A.51})$$

From [24], we know that

$$\sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}\}} \exp\langle \lambda, \alpha \rangle \lesssim \text{TOL}_r^{-\frac{\Gamma_1}{\eta_1}} \left(\log \text{TOL}_r^{-1} \right)^{m_1}, \quad (\text{A.52})$$

$$\sum_{\{\alpha \in \mathbb{N}^2: \langle \bar{\delta}, \alpha \rangle \leq \bar{L}\}} \exp\langle \lambda, \alpha \rangle \lesssim \text{TOL}_r^{-\frac{\Gamma_2}{\eta_2}} \left(\log \text{TOL}_r^{-1} \right)^{m_2}. \quad (\text{A.53})$$

where we define the following constants:

$$\begin{aligned} \Gamma_1 &:= 2C_{\bar{\delta}} \log(2) \min \left(\frac{2}{2 - w_1 + 2b_1}, \frac{1}{1 - w_2 + 2b_2} \right), \\ \Gamma_2 &:= 2C_{\bar{\delta}} \log(2) \min \left(\frac{2}{1 - s_1 + 2b_1}, \frac{1}{1 - s_2 + 2b_2} \right), \\ g_1 &:= \begin{cases} 2, & \frac{2}{2 - w_1 + 2b_1} = \frac{1}{1 - w_2 + 2b_2} \\ 1, & \text{otherwise} \end{cases}, \\ g_2 &:= \begin{cases} 2, & \frac{2}{1 - s_1 + 2b_1} = \frac{1}{1 - s_2 + 2b_2} \\ 1, & \text{otherwise} \end{cases}, \\ m_1 &:= (g_1 - 1) + (e_1 - 1) \frac{\Gamma_1}{\eta_1}, \\ m_2 &:= (g_2 - 1) + (e_2 - 1) \frac{\Gamma_2}{\eta_2}. \end{aligned}$$

By inserting (A.52) and (A.53) into (A.51), we attain

$$\tilde{W}_2(\mathcal{I}(v_1)) \lesssim \text{TOL}_r^{-2\Psi} \left(\log \text{TOL}_r^{-1} \right)^{\varkappa}, \quad (\text{A.54})$$

where

$$\varkappa := \begin{cases} (g_1 - 1) + (e_1 - 1) \frac{\Gamma_1}{\eta_1}, & \Psi = \frac{2}{2 - w_1 + 2b_1} \text{ or } \frac{1}{1 - w_2 + 2b_2} \\ (g_2 - 1) + (e_2 - 1) \frac{\Gamma_2}{\eta_2}, & \Psi = \frac{2}{1 - s_1 + 2b_1} \text{ or } \frac{1}{1 - s_2 + 2b_2} \end{cases}.$$

By inserting (A.50) and (A.54) into (4.101), we obtain

$$\mathcal{W}[\mathcal{A}_{\text{MIDL MC}}(\mathcal{I}(v_1))] \lesssim \text{TOL}_r^{-2-2\zeta} \left(\log \text{TOL}_r^{-1} \right)^{2\varrho} + \text{TOL}_r^{-2\Psi} \left(\log \text{TOL}_r^{-1} \right)^{\varkappa}. \quad (\text{A.55})$$

To ensure that the first term in (A.55) dominates the second term, we must satisfy the condition (4.120), which concludes the proof.

Appendix B

Algorithms

B.1 Estimating the bias at level ℓ for adaptive DLMC

Algorithm 6: Estimating $\mathbb{E}[\Delta G_\ell]$ with an antithetic sampler

Inputs: ℓ, M_1, M_2, ζ ;

for $m_1 = 1, \dots, M_1$ **do**

Generate realization of random variables $\omega_{1:P_\ell}^{(m_1)}$;

Generate $\mu^{P_\ell|N_\ell}(\omega_{1:P_\ell}^{(m_1)})$ realization with P_ℓ -particle system and N_ℓ time steps using (4.16);

From $\mu^{P_\ell|N_\ell}(\omega_{1:P_\ell}^{(m_1)})$, generate $\mu^{P_\ell-1|N_\ell-1}(\omega_{1:\frac{P_\ell}{2}}^{(m_1)})$ and $\mu^{P_\ell-1|N_\ell-1}(\omega_{\frac{P_\ell}{2}:P_\ell}^{(m_1)})$;

for $m_2 = 1, \dots, M_2$ **do**

Generate realization of random variables $\bar{\omega}^{(m_2)}$;

Given $\mu^{P_\ell|N_\ell}(\omega_{1:P_\ell}^{(m_1)})$ and ζ , generate a sample path of (4.20) using N_ℓ time steps;

Compute $G_\ell(\omega_{1:P_\ell}^{(m_1)} \times \bar{\omega}^{(m_2)})$;

Given $\mu^{P_\ell-1|N_\ell-1}(\omega_{1:\frac{P_\ell}{2}}^{(m_1)})$ and ζ , generate a sample path of (4.20) using $N_{\ell-1}$ time steps;

Compute $G_{\ell-1}(\omega_{1:\frac{P_\ell}{2}}^{(m_1)} \times \bar{\omega}^{(m_2)})$;

Given $\mu^{P_\ell-1|N_\ell-1}(\omega_{\frac{P_\ell}{2}:P_\ell}^{(m_1)})$ and ζ , generate a sample path of (4.20) using $N_{\ell-1}$ time steps;

Compute $G_{\ell-1}(\omega_{\frac{P_\ell}{2}:P_\ell}^{(m_1)} \times \bar{\omega}^{(m_2)})$;

end

end

Approximate $\mathbb{E}[G_\ell - G_{\ell-1}]$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} G_\ell(\omega_{1:P_\ell}^{(m_1)} \times \bar{\omega}^{(m_2)}) - \frac{1}{2} \left(G_{\ell-1}(\omega_{1:\frac{P_\ell}{2}}^{(m_1)}, \bar{\omega}^{(m_2)}) + G_{\ell-1}(\omega_{\frac{P_\ell}{2}:P_\ell}^{(m_1)}, \bar{\omega}^{(m_2)}) \right)$;

B.2 Estimating $V_{1,L}$ and $V_{2,L}$ for adaptive DLMC

Algorithm 7: Estimating $V_{1,L}$ and $V_{2,L}$ for adaptive DLMC

Inputs: $P_L, N_L, M_1, M_2, \zeta$;

for $m_1 = 1, \dots, M_1$ **do**

Generate realization of random variables $\omega_{1:P_L}^{(m_1)}$;

Generate $\mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)})$ realization with P_L -particle system and N_L time steps using (4.16);

for $m_2 = 1, \dots, M_2$ **do**

Generate realization of random variables $\bar{\omega}^{(m_2)}$;

Given $\mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)})$ and ζ , generate a sample path of (4.20) using N_L time steps;

Compute $g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T))(\omega_{1:P_L}^{(m_1)} \times \bar{\omega}^{(m_2)})$;

Compute $\mathbb{L}^{P_L|N_L}(\omega_{1:P_L}^{(m_1)} \times \bar{\omega}^{(m_2)})$ using (4.32);

end

Approximate $\mathbb{E} \left[g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L} \mid \mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)}) \right]$ by

$$\frac{1}{M_2} \sum_{m_2=1}^{M_2} g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L}(\omega_{1:P_L}^{(m_1)} \times \bar{\omega}^{(m_2)});$$

Approximate $\text{Var} \left[g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L} \mid \mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)}) \right]$ by sample variance

$$\text{of } \left\{ g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L}(\omega_{1:P_L}^{(m_1)} \times \bar{\omega}^{(m_2)}) \right\}_{m_2=1}^{M_2};$$

end

Approximate $V_{1,L}$ by sample variance of

$$\left\{ \mathbb{E} \left[g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L} \mid \mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)}) \right] \right\}_{m_1=1}^{M_1};$$

Approximate $V_{2,L}$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \text{Var} \left[g(\bar{\mathbf{X}}_{\zeta}^{P_L|N_L}(T)) \mathbb{L}^{P_L|N_L} \mid \mu^{P_L|N_L}(\omega_{1:P_L}^{(m_1)}) \right]$;

B.3 Estimating level differences for the **MLDLMC** estimator

Algorithm 8: IS scheme to estimate $\mathbb{E} [\Delta G_\ell]$ using the antithetic sampler (4.56)

Inputs: ℓ, M_1, M_2, ζ ;

for $m_1 = 1, \dots, M_1$ **do**

Generate $\mu_\ell(\omega_{1:P_\ell}^{(\ell, m_1)})$ using (4.28);

for $a = 1, 2$ **do**

Generate $\mu_{\ell-1}^{(a)}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)})$ using (4.28);

end

for $m_2 = 1, \dots, M_2$ **do**

Given $\mu_\ell(\omega_{1:P_\ell}^{(\ell, m_1)})$ and ζ , generate sample path of (4.20) at level ℓ with $\bar{\omega}^{(\ell, m_2)}$;

Compute $G_\ell(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

for $a = 1, 2$ **do**

Given $\mu_{\ell-1}^{(a)}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)})$ and ζ , generate sample path of (4.20) at level $\ell - 1$ with $\bar{\omega}^{(\ell, m_2)}$;

Compute $G_{\ell-1}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

end

$\hat{G}_{\ell-1}(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)}) = \frac{1}{2} \sum_{a=1}^2 G_{\ell-1}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

end

$\Delta G_\ell^{(m_1, m_2)} = (G_\ell - \hat{G}_{\ell-1})(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

end

Approximate $\mathbb{E} [G_\ell - G_{\ell-1}]$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} \Delta G_\ell^{(m_1, m_2)}$;

B.4 Estimating variances for the adaptive **MLDLMC** algorithm

Algorithm 9: Estimating $V_{1,\ell}$ and $V_{2,\ell}$ for adaptive **MLDLMC**

Inputs: ℓ, M_1, M_2, ζ ;

for $m_1 = 1, \dots, M_1$ **do**

Generate $\mu_\ell(\omega_{1:P_\ell}^{(\ell, m_1)})$ using (4.28);

for $a = 1, 2$ **do**

Generate $\mu_{\ell-1}^{(a)}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)})$ using (4.28);

end

for $m_2 = 1, \dots, M_2$ **do**

Given $\mu_\ell(\omega_{1:P_\ell}^{(\ell, m_1)})$ and ζ , generate sample path of (4.18) at level ℓ with $\bar{\omega}^{(\ell, m_2)}$;

Compute $G_\ell(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

for $a = 1, 2$ **do**

Given $\mu_{\ell-1}^{(a)}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)})$ and ζ , generate sample path of (4.18) at level $\ell - 1$ with $\bar{\omega}^{(\ell, m_2)}$;

Compute $G_{\ell-1}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

end

$\hat{G}_{\ell-1}(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)}) = \frac{1}{2} \sum_{a=1}^2 G_{\ell-1}(\omega_{((a-1)P_{\ell-1}+1):aP_{\ell-1}}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

end

$\Delta G_\ell^{(m_1, m_2)} = (G_\ell - \hat{G}_{\ell-1})(\omega_{1:P_\ell}^{(\ell, m_1)} \times \bar{\omega}^{(\ell, m_2)})$;

Approximate $\mathbb{E} [\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}(\omega_{1:P_\ell}^{(\ell, m_1)})]$ by $\frac{1}{M_2} \sum_{m_2=1}^{M_2} \Delta G_\ell^{(m_1, m_2)}$;

Approximate $\text{Var} [\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}(\omega_{1:P_\ell}^{(\ell, m_1)})]$ by sample variance of $\left\{ \Delta G_\ell^{(m_1, m_2)} \right\}_{m_2=1}^{M_2}$;

end

Approximate $V_{1,\ell}$ by sample variance of $\left\{ \mathbb{E} [\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}(\omega_{1:P_\ell}^{(\ell, m_1)})] \right\}_{m_1=1}^{M_1}$;

Approximate $V_{2,\ell}$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \text{Var} [\Delta G_\ell \mid \{\mu_\ell, \mu_{\ell-1}\}(\omega_{1:P_\ell}^{(\ell, m_1)})]$.

B.5 Estimating first-order mixed differences for the MIDLMC estimator

Algorithm 10: IS scheme to estimate $\mathbb{E} [\Delta \mathcal{G}_\alpha]$ using the antithetic sampler (4.83)

Inputs: $\alpha, M_1, M_2, \zeta;$

for $m_1 = 1, \dots, M_1$ **do**

Generate realization of random variables $\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)};$

Using $\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}$, generate realization of empirical law μ at all required discretizations as per (4.83) by generating sample paths of particle system (4.9);

for $m_2 = 1, \dots, M_2$ **do**

Generate realization of random variables $\bar{\omega}^{(\alpha, m_2)};$

Using $\bar{\omega}^{(\alpha, m_2)}$, generate sample path of (4.20) at all required discretizations as per (4.83) with control $\zeta;$

Compute $\Delta \mathcal{G}_\alpha(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)})$ using (4.83);

end

end

Approximate $\mathbb{E} [\Delta \mathcal{G}_\alpha]$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \frac{1}{M_2} \sum_{m_2=1}^{M_2} \Delta \mathcal{G}_\alpha(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)});$

B.6 Estimating variances for the adaptive MIDLMC algorithm

Algorithm 11: Estimating $\{V_{1,\alpha}, V_{2,\alpha}\}$ for adaptive MIDLMC

Inputs: $\alpha, M_1, M_2, \zeta;$

for $m_1 = 1, \dots, M_1$ **do**

Generate realization of random variables $\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)};$

Using $\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}$, generate realization of empirical law μ at all required discretizations as per (4.83) by generating sample paths of particle system (4.9);

for $m_2 = 1, \dots, M_2$ **do**

Generate realization of random variables $\bar{\omega}^{(\alpha, m_2)};$

Using $\bar{\omega}^{(\alpha, m_2)}$, generate sample path of (4.20) at all required discretizations as per (4.83) with control $\zeta;$

Compute $\Delta \mathcal{G}_\alpha(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)})$ using (4.83);

end

Approximate $\mathbb{E} [\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}]$ by $\frac{1}{M_2} \sum_{m_2=1}^{M_2} \Delta \mathcal{G}_\alpha(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)});$

Approximate $\text{Var} [\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}]$ by sample variance of

$$\left\{ \Delta \mathcal{G}_\alpha(\omega_{1:P_{\alpha_1}}^{(\alpha, m_1)} \times \bar{\omega}^{(\alpha, m_2)}) \right\}_{m_2=1}^{M_2};$$

end

Approximate $V_{1,\alpha}$ by sample variance of $\left\{ \mathbb{E} [\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}] \right\}_{m_1=1}^{M_1};$

Approximate $V_{2,\alpha}$ by $\frac{1}{M_1} \sum_{m_1=1}^{M_1} \text{Var} [\Delta \mathcal{G}_\alpha \mid \omega_{1:P_{\alpha_1}}^{(\alpha, m_1)}].$

Algorithm 12: Sample variance extrapolation for adaptive MIDLMC

Inputs: $\mathcal{I}, \{\bar{M}_1, \bar{M}_2\}, (w_1, w_2), (s_1, s_2)$;
 Estimate $\{V_{1,\alpha}, V_{2,\alpha}\}$ for $\alpha \in \{0, 1, 2\} \times \{0, 1, 2\}$ using Algorithm 11 with $\{\bar{M}_1, \bar{M}_2\}$;
for $\alpha \in \mathcal{I}$ **do**
 if $\alpha_1 = 0, 1$ **then**
 $V_{1,\alpha} = \max\left(\frac{V_{1,\alpha-(0,1)}}{2^{w_2}}, \frac{V_{1,\alpha-(0,2)}}{2^{2w_2}}\right)$; $V_{2,\alpha} = \max\left(\frac{V_{2,\alpha-(0,1)}}{2^{s_2}}, \frac{V_{2,\alpha-(0,2)}}{2^{2s_2}}\right)$;
 else if $\alpha_2 = 0, 1$ **then**
 $V_{1,\alpha} = \max\left(\frac{V_{1,\alpha-(1,0)}}{2^{w_1}}, \frac{V_{1,\alpha-(2,0)}}{2^{2w_1}}\right)$; $V_{2,\alpha} = \max\left(\frac{V_{2,\alpha-(1,0)}}{2^{s_1}}, \frac{V_{2,\alpha-(2,0)}}{2^{2s_1}}\right)$;
 else
 $V_{1,\alpha} = \max\left(\frac{V_{1,\alpha-(0,1)}}{2^{w_2}}, \frac{V_{1,\alpha-(1,0)}}{2^{w_1}}\right)$; $V_{2,\alpha} = \max\left(\frac{V_{2,\alpha-(0,1)}}{2^{s_2}}, \frac{V_{2,\alpha-(1,0)}}{2^{s_1}}\right)$;
 end
end

B.7 Upwind FDM for HJB PDE (5.36)**Algorithm 13:** Numerical HJB PDE solver

Input: $\lambda(t), N_t, N_1, N_2, N_3$;
Initialization: $v_{(i,j,k)}^{N_t} = -wP_K \bar{A} a_i$, $\Delta_a^\pm v|_{(i,j,k)}^{N_t} = -wP_K \bar{A}$ at grid $\tau_a \times \tau_r \times \tau_\chi$;
for $n = N_t, \dots, 2$ **do**
 for $i = 1, \dots, N_1$ **do**
 for $j = 1, \dots, N_2$ **do**
 for $k = 1, \dots, N_3$ **do**
 Using $\Delta_a^\pm v|_{(i,j,k)}^n$, compute optimal controls $(\phi^*)_{(i,j,k)}^n$ by solving (5.49);
 Using $v_{(i,j,k)}^n$ and $(\phi^*)_{(i,j,k)}^n$, compute $v_{(i,j,k)}^{n-1}$ with update rule (5.50);
 end
 end
 end
 end
 Using $v_{(i,j,k)}^{n-1}$, compute and save gradients $\Delta_a^\pm v|_{(i,j,k)}^{n-1}$ using (5.46) at grid $\tau_a \times \tau_r \times \tau_\chi$;
end
Output: v at grid τ .

B.8 MC with EM scheme for subgradient (5.52) estimation

Algorithm 14: Numerical subgradient estimation

Input: $\ell, \bar{N}_t, \tilde{N}_t, M_{\text{SG}}$;

Initialization: $\bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_0) = \mathbf{X}(0)$;

for $m = 1, \dots, M_{\text{SG}}$ **do**

 Generate m^{th} iid realization of ω ;

for $n = 0, \dots, \bar{N}_t - 1$ **do**

 Estimate $\Delta_a^\pm v(\bar{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n))$ by linearly interpolating from $\Delta_a^\pm v$ values computed in Algorithm 13 at grid τ ;

 Using $\bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n)$ and $\Delta_a^\pm v(\bar{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n))$, compute optimal controls $\boldsymbol{\phi}^*(\bar{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n))$ by solving (5.49);

 Using $\omega^{(m)}$ and $\boldsymbol{\phi}^*(\bar{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_n))$, compute $\bar{\mathbf{X}}^{\bar{N}_t}(\bar{t}_{n+1})$ using EM scheme (2.7);

end

for $i = 1, \dots, \ell$ **do**

 Construct uniform time grid $\tilde{\tau} : t_{i-1} = \tilde{t}_0 < \tilde{t}_1 < \dots < \tilde{t}_{\tilde{N}_t} = t_i$ with $\Delta \tilde{t} = \frac{t_i - t_{i-1}}{\tilde{N}_t}$;

for $n = 0, \dots, \tilde{N}_t$ **do**

 Interpolate $\bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n)$ from $\bar{\mathbf{X}}^{\bar{N}_t}$ using Brownian bridge interpolation;

 Estimate $\Delta_a^\pm v(\tilde{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n))$ by linearly interpolating from $\Delta_a^\pm v$ values computed in Algorithm 13 at grid τ ;

 Using $\bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n)$ and $\Delta_a^\pm v(\tilde{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n))$, compute optimal controls $\boldsymbol{\phi}^*(\tilde{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n))$ by solving (5.49);

 Compute $\zeta^{(m)}(n) := \left(\mathcal{I}_{\{\phi_{\text{out}}(\tilde{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n), \boldsymbol{\phi}(\tilde{t}_n, \bar{\mathbf{X}}^{\bar{N}_t}(\tilde{t}_n))) \geq \phi_{\text{th}}\}} - \epsilon \right) \Delta \tilde{t}$;

end

$\varphi^{(m)}(i) := \sum_{n=0}^{\tilde{N}_t} \zeta^{(m)}(n)$;

end

end

$(\bar{\mathcal{D}}\Theta(\lambda^\ell))_i = \frac{1}{M_{\text{SG}}} \sum_{m=1}^{M_{\text{SG}}} \varphi^{(m)}(i)$ for each $i = 1, \dots, \ell$;

Output: $\bar{\mathcal{D}}\Theta(\lambda^\ell)$.

B.9 Initialization algorithm for the SSM

Algorithm 15: Initialization

Input: $Y_1^1 = 1$, TOL_{init} , β_F , M_{SG} , N_t , N_1 , N_2 , N_3 , \tilde{N}_t , $\mathbf{X}(0)$

Output: \tilde{Y}_1^1

Construct $\lambda^1(t)$ with Y_1^1 using (5.40);

Compute $\tilde{\Theta}(\lambda^1) = v(0, \mathbf{X}(0))$ by solving (5.36) using Algorithm 13 with $\lambda^1(t)$ and parameters N_t, N_1, N_2, N_3 ;

Estimate $\tilde{\mathcal{D}}\Theta(\lambda^1)$ using Algorithm 14 with $\ell = 1, \tilde{N}_t = \tilde{N}_t$ and parameters \tilde{N}_t, M_{SG} ;

if $\tilde{\mathcal{D}}\Theta(\lambda^1) > 0$ **then**

while $|\tilde{\mathcal{D}}\Theta(\lambda^1)| > TOL_{init}\epsilon$ **or** $\tilde{\mathcal{D}}\Theta(\lambda^1) > 0$ **do**

$Y_1^1 \leftarrow \beta_F Y_1^1$;

 Construct $\lambda^1(t)$ with Y_1^1 using (5.40);

 Compute $\tilde{\Theta}(\lambda^1) = v(0, \mathbf{X}(0))$ by solving (5.36) using Algorithm 13 with $\lambda^1(t)$ and parameters N_t, N_1, N_2, N_3 ;

 Estimate $\tilde{\mathcal{D}}\Theta(\lambda^1)$ using Algorithm 14 with $\ell = 1, \tilde{N}_t = \tilde{N}_t$ and parameters \tilde{N}_t, M_{SG} ;

end

else

while $|\tilde{\mathcal{D}}\Theta(\lambda^1)| > TOL_{init}\epsilon$ **or** $\tilde{\mathcal{D}}\Theta(\lambda^1) < 0$ **do**

$Y_1^1 \leftarrow \frac{Y_1^1}{\beta_F}$;

 Construct $\lambda^1(t)$ with Y_1^1 using (5.40);

 Compute $\tilde{\Theta}(\lambda^1) = v(0, \mathbf{X}(0))$ by solving (5.36) using Algorithm 13 with $\lambda^1(t)$ and parameters N_t, N_1, N_2, N_3 ;

 Estimate $\tilde{\mathcal{D}}\Theta(\lambda^1)$ using Algorithm 14 with $\ell = 1, \tilde{N}_t = \tilde{N}_t$ and parameters \tilde{N}_t, M_{SG} ;

end

end

$\tilde{Y}_1^1 = Y_1^1$;

Construct $\lambda^1(t)$ with \tilde{Y}_1^1 using (5.40);

Appendix C

Numerical solvers for control PDEs

C.1 FDM solver for SOC-based IS control PDE

For convenience, we consider here only the 1D version of (4.24)

$$\begin{cases} \frac{\partial u}{\partial t}(t, x) + a \left(x, \frac{1}{P} \sum_{j=1}^P \kappa_1(x, X_j^P(t)) \right) \frac{\partial u}{\partial x}(t, x) \\ + \frac{1}{2} b^2 \left(x, \frac{1}{P} \sum_{j=1}^P \kappa_2(x, X_j^P(t)) \right) \frac{\partial^2 u}{\partial x^2}(t, x) = 0, & (t, x) \in [0, T] \times \mathbb{R} \\ u(T, x) = |g(x)|, & x \in \mathbb{R} \end{cases} \quad (\text{C.1})$$

with optimal control

$$\zeta(t, x) = b \left(x, \frac{1}{P} \sum_{j=1}^P \kappa_2(x, X_j^P(t)) \right) \frac{\partial \log u}{\partial x}(t, x). \quad (\text{C.2})$$

We aim to estimate the control ζ given $\{X_p^P(t)\}_{p=1}^P \sim \mu_t^P$. We define the second-order linear differential operator \mathcal{G} acting on $C^{1,2}([0, T] \times \mathbb{R})$ class of functions by

$$\mathcal{G} := -a \left(x, \frac{1}{P} \sum_{j=1}^P \kappa_1(x, X_j^P(t)) \right) \frac{\partial}{\partial x} - \frac{b^2}{2} \left(x, \frac{1}{P} \sum_{j=1}^P \kappa_2(x, X_j^P(t)) \right) \frac{\partial^2}{\partial x^2}.$$

Then, (C.1) can be expressed as

$$\begin{cases} \frac{\partial u}{\partial t}(t, x) = \mathcal{G}u(t, x) & , (t, x) \in [0, T] \times \mathbb{R} \\ u(T, x) = |g(x)| & , x \in \mathbb{R}. \end{cases}$$

To discretize the PDE using FDM, both the space- and time-domains must be discretized. This is done by introducing the following grids:

$$0 = t_0 < t_1 < \dots < t_{N_t} = T, \quad \text{and} \quad (\text{C.3})$$

$$-x_b = x_0 < x_1 < \dots < x_{N_x} = x_b. \quad (\text{C.4})$$

Note that for space dimension x , we approximate the domain \mathbb{R} by an appropriately chosen $[-x_b, x_b]$. In addition, we use uniform grids in both dimensions.

$$t_n = n\Delta t, \quad n = 0, \dots, N_t \quad \text{and} \quad (\text{C.5})$$

$$x_i = -x_b + i\Delta x, \quad i = 0, \dots, N_x. \quad (\text{C.6})$$

Let $u_i^n := u(t_n, x_i) := u(n\Delta t, -x_b + i\Delta x)$. Then, discretization is performed in x by using central differences:

$$\left. \frac{\partial u}{\partial x} \right|_i^n = \frac{u_{i+1}^n - u_{i-1}^n}{2\Delta x} + \mathcal{O}(\Delta x^2). \quad (\text{C.7})$$

$$\left. \frac{\partial^2 u}{\partial x^2} \right|_i^n = \frac{u_{i+1}^n - 2u_i^n + u_{i-1}^n}{\Delta x^2} + \mathcal{O}(\Delta x^2). \quad (\text{C.8})$$

Therefore, the operator \mathcal{G} applied to u at time t_n and state x_i can be expressed as

$$\begin{aligned} \mathcal{G}u \Big|_i^n &= -a \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_1(x_i, X_j^P(t_n)) \right) \left. \frac{\partial u}{\partial x} \right|_i^n - \frac{b^2}{2} \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right) \left. \frac{\partial^2 u}{\partial x^2} \right|_i^n \\ &= -a \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_1(x_i, X_j^P(t_n)) \right) \frac{u_{i+1}^n - u_{i-1}^n}{2\Delta x} \\ &\quad - \frac{b^2}{2} \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right) \frac{u_{i+1}^n - 2u_i^n + u_{i-1}^n}{\Delta x^2} + \mathcal{O}(\Delta x^2). \end{aligned} \quad (\text{C.9})$$

By grouping the coefficients together, we get

$$\begin{aligned} \mathcal{G}u \Big|_i^n &= \underbrace{\left(\frac{a \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_1(x_i, X_j^P(t_n)) \right)}{2\Delta x} - \frac{b^2 \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right)}{2\Delta x^2} \right)}_{=d_{i,i}^n} u_{i-1}^n \\ &\quad - \underbrace{\left(\frac{a \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_1(x_i, X_j^P(t_n)) \right)}{2\Delta x} + \frac{b^2 \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right)}{2\Delta x^2} \right)}_{=d_{i,i}^n} u_{i+1}^n \\ &\quad + \underbrace{\left(\frac{b^2 \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right)}{\Delta x^2} \right)}_{=d_i^n} u_i^n + \mathcal{O}(\Delta x^2) \\ &= d_{i,i}^n u_{i-1}^n + d_i^n u_i^n + d_{i,i}^n u_{i+1}^n + \mathcal{O}(\Delta x^2). \end{aligned} \quad (\text{C.10})$$

We formulate a linear system for each time step by using (C.10). Let $\mathbf{u}^n := [u_0^n, \dots, u_{N_x}^n]^{\text{tr}}$; consequently, the finite difference can be expressed as

$$(\mathcal{G}u \Big|_i^n)_{i=0}^{N_x} = \mathbf{G}^n \mathbf{u}^n + \mathcal{O}(\Delta x^2), \quad (\text{C.11})$$

where $\mathbf{G}^n \in \mathbb{R}^{N_x \times N_x}$ is the discretization operator at time t_n ,

$$\mathbf{G}^n = \begin{pmatrix} d_0^n & d_{u,0}^n & 0 & 0 & \dots & 0 \\ d_{l,1}^n & d_1^n & d_{u,1}^n & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & d_{l,N_x-1}^n & d_{N_x-1}^n & d_{u,N_x-1}^n \\ 0 & 0 & \dots & 0 & d_{l,N_x}^n & d_{N_x}^n \end{pmatrix}. \quad (\text{C.12})$$

We use the generic θ -family of time-discretization schemes,

$$\frac{u_i^n - u_i^{n-1}}{\Delta t} = \theta \mathbf{g}_i^n \mathbf{u}_i^n + (1 - \theta) \mathbf{g}_i^{n-1} \mathbf{u}_i^{n-1}, \quad (\text{C.13})$$

where $\mathbf{g}_i^n \in \mathbb{R}^{1 \times N_x}$ is the i^{th} row of \mathbf{G}^n . Hence, the time-step update in matrix-vector form can be expressed as

$$\mathbf{u}^{n-1} = \left(\mathbf{I} + (1 - \theta) \Delta t \mathbf{G}^{n-1} \right)^{-1} \left(\mathbf{I} - \theta \Delta t \mathbf{G}^n \right) \mathbf{u}^n. \quad (\text{C.14})$$

Different θ values correspond to different time-stepping schemes. We used $\theta = 0.5$, corresponding to the Crank–Nicholson scheme, with convergence $\mathcal{O}(\Delta t^2)$. In addition, the terminal time condition at $t = T$ is denoted as

$$\mathbf{u}^{N_t} = (g(x_i))_{i=0}^{N_x} = (g(-x_b + i\Delta x))_{i=0}^{N_x}. \quad (\text{C.15})$$

However, (C.1) does not specify the boundary conditions. To circumvent this problem, we used nonreflective boundary conditions at $x = -x_b$ and $x = x_b$. These conditions emphasize that the actual domain is not bounded at the boundaries, i.e., $u_{N_x}^n$ and u_0^n are extrapolated from known interior values:

$$u_0^n \approx u_2^n + \underbrace{\frac{x_2 - x_0}{x_2 - x_1}}_{= \frac{2\Delta x}{\Delta x} = 2} (u_1^n - u_2^n) = 2u_1^n - u_2^n. \quad (\text{C.16})$$

$$u_{N_x}^n \approx u_{N_x-2}^n + \underbrace{\frac{x_{N_x} - x_{N_x-2}}{x_{N_x-1} - x_{N_x-2}}}_{= \frac{2\Delta x}{\Delta x} = 2} (u_{N_x-1}^n - u_{N_x-2}^n) = 2u_{N_x-1}^n - u_{N_x-2}^n. \quad (\text{C.17})$$

We then prescribe (C.16) and (C.17) as Dirichlet boundary conditions. This also modifies certain \mathbf{G}^n elements:

$$\begin{aligned} \mathcal{G}u|_1^n &= d_{l,1}^n \underbrace{u_0^n}_{\approx 2u_1^n - u_2^n} + d_1^n u_1^n + d_{u,1}^n u_2^n \\ &\approx (2d_{l,1}^n + d_1^n) u_1^n + (d_{u,1}^n - d_{l,1}^n) u_2^n \\ \implies \mathbf{g}_1^n &= (0, \underbrace{2d_{l,1}^n + d_1^n}_{= \tilde{d}_1^n}, \underbrace{d_{u,1}^n - d_{l,1}^n}_{= \tilde{d}_{u,1}^n}, 0, \dots, 0). \end{aligned} \quad (\text{C.18})$$

$$\mathcal{G}u|_{N_x-1}^n = d_{l,N_x-1}^n u_{N_x-2}^n + d_{N_x-1}^n u_{N_x-1}^n + d_{u,N_x-1}^n \underbrace{u_{N_x}^n}_{\approx 2u_{N_x-1}^n - u_{N_x-2}^n}$$

$$\begin{aligned} &\approx (d_{l,N_x-1}^n - d_{u,N_x-1}^n)u_{N_x-2}^n + (d_{N_x-1}^n + 2d_{u,N_x-1}^n)u_{N_x-1}^n \\ \implies \mathbf{g}_{N_x-1}^n &= (0, \dots, 0, \underbrace{d_{l,N_x-1}^n - d_{u,N_x-1}^n}_{=\tilde{d}_{l,N_x-1}^n}, \underbrace{d_{N_x-1}^n + 2d_{u,N_x-1}^n}_{=\tilde{d}_{N_x-1}^n}, 0). \end{aligned} \quad (\text{C.19})$$

Hence, the modified discretization matrix, \mathbf{G}^n , now includes the boundary conditions:

$$\mathbf{G}^n = \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 \\ 0 & \tilde{d}_1^n & \tilde{d}_{u,1}^n & 0 & \dots & 0 \\ 0 & d_{l,2}^n & d_2^n & d_{u,2}^n & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \tilde{d}_{l,N_x-1}^n & \tilde{d}_{N_x-1}^n & 0 \\ 0 & 0 & \dots & 0 & 0 & 1 \end{pmatrix}. \quad (\text{C.20})$$

As the boundary conditions are fixed by Dirichlet conditions, the computation needs to be performed only in the inner $\mathbb{R}^{(N_x-2) \times (N_x-2)}$ sub-matrix. After obtaining the numerical solution for u , we apply central differences to (C.2) to obtain numerical solutions to the optimal control:

$$\zeta(t_n, x_i) \approx b \left(x_i, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_i, X_j^P(t_n)) \right) \frac{\log u_{i+1}^n - \log u_{i-1}^n}{2\Delta x}. \quad (\text{C.21})$$

The boundary points use one-sided finite difference formulas using an interior point:

$$\zeta(t_n, x_0) \approx b \left(x_0, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_0, X_j^P(t_n)) \right) \frac{\log u_1^n - \log u_0^n}{\Delta x}. \quad (\text{C.22})$$

$$\zeta(t_n, x_{N_x}) \approx b \left(x_{N_x}, \frac{1}{P} \sum_{j=1}^P \kappa_2(x_{N_x}, X_j^P(t_n)) \right) \frac{\log u_{N_x}^n - \log u_{N_x-1}^n}{\Delta x}. \quad (\text{C.23})$$

Therefore, we obtain optimal control at grid points (C.3) and (C.4). We extend this control to the entire domain $[0, T) \times \mathbb{R}$ using linear interpolation in both time and space. As the PDE (C.1) belongs to the advection-diffusion type, the selection of an appropriate value for $\frac{\Delta t}{\Delta x}$ is crucial for ensuring stable numerical solutions according to the well-known CFL condition. For our example (4.12), we select $\frac{\Delta t}{\Delta x} = 0.2$.

C.2 Upwind FDM solver for the HJB PDE

For a given Lagrange multiplier function λ , we rewrite the HJB PDE (5.36) as follows:

$$\begin{cases} \frac{\partial v}{\partial t}(t, \mathbf{x}) + \min_{\phi \in \mathcal{A}(t, \mathbf{x})} \left(\langle \mathbf{a}(t, \mathbf{x}, \phi), \nabla v(t, \mathbf{x}) \rangle + g(t, \mathbf{x}, \phi) \right. \\ \left. + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \phi) > \phi_{\text{th}}\}} - \epsilon \right) + \frac{1}{2} (\mathbf{b}(t, \mathbf{x}) \mathbf{b}(t, \mathbf{x})^{\text{tr}}) : \nabla^2 v(t, \mathbf{x}) = 0 \right. \\ \left. v(T, \mathbf{x}) = -w P_K \bar{A} a, \quad \forall \mathbf{x} \in \mathbf{S}, \right. \end{cases} \quad (\text{C.24})$$

Using the finite differences introduced in (5.47), the discretized version of the PDE (C.24)

is written as follows.

$$\left\{ \begin{array}{l} \Delta_t^- v(t, \mathbf{x}) + a_1^+(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \Delta_a^+ v(t, \mathbf{x}) + a_2^+(t, \mathbf{x}) \Delta_r^+ v(t, \mathbf{x}) \\ + a_3^+(t, \mathbf{x}) \Delta_\chi^+ v(t, \mathbf{x}) - a_1^-(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \Delta_a^- v(t, \mathbf{x}) \\ - a_2^-(t, \mathbf{x}) \Delta_r^- v(t, \mathbf{x}) - a_3^-(t, \mathbf{x}) \Delta_\chi^- v(t, \mathbf{x}) \\ + g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) > \phi_{\text{th}}\}} - \epsilon \right) \\ + \frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \Delta_r^2 v(t, \mathbf{x}) + \frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \Delta_\chi^2 v(t, \mathbf{x}) = 0 \\ v(t_{N_i}, \mathbf{x}) = -w P_K \bar{A} a. \end{array} \right. \quad (\text{C.25})$$

In (C.25), we denote that the controls $\boldsymbol{\phi}(t, \mathbf{x})$ depend on the value function derivative $\Delta_a^\pm v(t, \mathbf{x})$, emphasizing the non-linearity of the PDE. Substituting the definitions of these finite differences, we get the following time-update scheme.

$$\begin{aligned} v(t - \Delta t, \mathbf{x}) = & v(t, \mathbf{x}) \left(1 - b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} - b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta \chi^2} \right. \\ & - \frac{\Delta t}{\Delta a} |a_1(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x})))| \\ & \left. - \frac{\Delta t}{\Delta r} |a_2(t, \mathbf{x})| - \frac{\Delta t}{\Delta \chi} |a_3(t, \mathbf{x})| \right) \\ & + v(t, [a + \Delta a, r, \chi]) \frac{\Delta t}{\Delta a} a_1^+(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \\ & + v(t, [a - \Delta a, r, \chi]) \frac{\Delta t}{\Delta a} a_1^-(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \\ & + v(t, [a, r + \Delta r, \chi]) \left(\frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} + a_2^+(t, \mathbf{x}) \frac{\Delta t}{\Delta r} \right) \\ & + v(t, [a, r - \Delta r, \chi]) \left(\frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} + a_2^-(t, \mathbf{x}) \frac{\Delta t}{\Delta r} \right) \\ & + v(t, [a, r, \chi + \Delta \chi]) \left(\frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta \chi^2} + a_3^+(t, \mathbf{x}) \frac{\Delta t}{\Delta \chi} \right) \\ & + v(t, [a, r, \chi - \Delta \chi]) \left(\frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta \chi^2} + a_3^-(t, \mathbf{x}) \frac{\Delta t}{\Delta \chi} \right) \\ & + \Delta t \left(g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \right. \\ & \left. + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) > \phi_{\text{th}}\}} - \epsilon \right) \right). \end{aligned} \quad (\text{C.26})$$

In (C.26), we see that the values of $v(t, \mathbf{x})$ and $\Delta_a^\pm v(t, \mathbf{x})$ are required to compute the value function at time step $t - \Delta t$. From the final time condition, we know that $v(T, \mathbf{x}) = -w P_K \bar{A} a$ and $\Delta_a^\pm v(T, \mathbf{x}) = -w P_K \bar{A}$. Next, we show and establish conditions under which the upwind FDM scheme is a convergent numerical scheme.

C.2.1 Convergence analysis

First, we present the consistency of the this numerical scheme. For this, the scheme's truncation error ϵ_{trunc} is written as follows.

$$\begin{aligned} \epsilon_{\text{trunc}}(t, \mathbf{x}) = & \Delta_t^- v(t, \mathbf{x}) + a_1^+(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \Delta_a^+ v(t, \mathbf{x}) + a_2^+(t, \mathbf{x}) \Delta_r^+ v(t, \mathbf{x}) \\ & + a_3^+(t, \mathbf{x}) \Delta_\chi^+ v(t, \mathbf{x}) - a_1^-(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \Delta_a^- v(t, \mathbf{x}) \\ & - a_2^-(t, \mathbf{x}) \Delta_r^- v(t, \mathbf{x}) - a_3^-(t, \mathbf{x}) \Delta_\chi^- v(t, \mathbf{x}) + \frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \Delta_r^2 v(t, \mathbf{x}) \\ & + \frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \Delta_\chi^2 v(t, \mathbf{x}) + g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \\ & + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) > \phi_{\text{th}}\}} - \epsilon \right) \\ & - \left(\frac{\partial v}{\partial t}(t, \mathbf{x}) + a_1 \left(t, \mathbf{x}, \boldsymbol{\phi} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right) \right) \frac{\partial v}{\partial a}(t, \mathbf{x}) + a_2(t, \mathbf{x}) \frac{\partial v}{\partial r}(t, \mathbf{x}) \\ & + a_3(t, \mathbf{x}) \frac{\partial v}{\partial \chi}(t, \mathbf{x}) + \frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) + \frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \frac{\partial^2 v}{\partial \chi^2}(t, \mathbf{x}) \\ & + g \left(t, \mathbf{x}, \boldsymbol{\phi} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right) + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x})) > \phi_{\text{th}}\}} - \epsilon \right). \end{aligned}$$

For the subsequent analysis, we use the following Taylor series expansions of the sufficiently smooth functions a_1, g . For notational convenience, we represent some function $f \left(t, \mathbf{x}, \boldsymbol{\phi} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right) := f \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right)$.

$$\begin{aligned} a_1 \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) + \mathcal{O}(\Delta a) \right) & \approx a_1 \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) + \mathcal{O}(\Delta a) \frac{\partial a_1}{\partial \left(\frac{\partial v}{\partial a} \right)} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \\ & + \mathcal{O}(\Delta a^2) \frac{\partial^2 a_1}{\partial \left(\frac{\partial a_1}{\partial a} \right)^2} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right), \\ g \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) + \mathcal{O}(\Delta a) \right) & \approx g \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) + \mathcal{O}(\Delta a) \frac{\partial g}{\partial \left(\frac{\partial v}{\partial a} \right)} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \\ & + \mathcal{O}(\Delta a^2) \frac{\partial^2 g}{\partial \left(\frac{\partial a_1}{\partial a} \right)^2} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right). \end{aligned}$$

We cannot use such an expansion for the discontinuous indicator function in the running cost. The standard numerical treatment for dealing with the indicator function is to generate a sequence of mollifiers [147] and then prove consistency in the limit of the mollifying factor being zero. This detail is not the main focus of the thesis, and hence is omitted here. Inserting the above expansions in (C.26), the truncation error is re-written as:

$$\epsilon_{\text{trunc}}(t, \mathbf{x}) = \mathcal{O}(\Delta t) + \mathcal{O}(\Delta a) \left(a_1 \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) + \frac{\partial v}{\partial a}(t, \mathbf{x}) \frac{\partial a_1}{\partial \left(\frac{\partial v}{\partial a} \right)} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right)$$

$$\begin{aligned}
& + \frac{\partial g}{\partial \left(\frac{\partial v}{\partial a}\right)} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \Big) + \mathcal{O}(\Delta r)a_2(t, \mathbf{x}) + \mathcal{O}(\Delta \chi)a_3(t, \mathbf{x}) \\
& + \mathcal{O}(\Delta a^2) \left(\frac{\partial a_1}{\partial \left(\frac{\partial v}{\partial a}\right)} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) + \frac{\partial v}{\partial a}(t, \mathbf{x}) \frac{\partial^2 a_1}{\partial \left(\frac{\partial v}{\partial a}\right)^2} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right. \\
& \quad \left. + \frac{\partial^2 g}{\partial \left(\frac{\partial v}{\partial a}\right)^2} \left(t, \mathbf{x}, \frac{\partial v}{\partial a}(t, \mathbf{x}) \right) \right) + \mathcal{O}(\Delta r^2)b_{2,2}(t, \mathbf{x})^2 \\
& + \mathcal{O}(\Delta \chi^2)b_{3,3}(t, \mathbf{x})^2.
\end{aligned}$$

Hence, the upwind FDM scheme is consistent if and only if $\epsilon_{\text{trunc}} \rightarrow 0$ in the limit $\Delta t, \Delta a, \Delta r, \Delta \chi \rightarrow 0$. Hence, we require the following conditions for a consistent numerical scheme.

1. The value function v is smooth enough to allow existence of a unique solution to (C.24).
2. Drift function \mathbf{a} and the squared diffusion function \mathbf{b}^2 are bounded and Lipschitz with respect to \mathbf{x} and $\boldsymbol{\phi}$.
3. The mapping of the derivatives of the value function $\frac{\partial v}{\partial a}$ to the optimal controls $\boldsymbol{\phi}$ is C^1 with bounded derivatives up to first order.

Second, we establish conditions required for the upwind FDM scheme to be monotone. The numerical scheme (C.26) is monotone if and only if

$$\frac{\partial v(t - \Delta t, \mathbf{x})}{\partial v(t, [a \pm \Delta a, r \pm \Delta r, \chi \pm \Delta \chi])} \geq 0. \quad (\text{C.27})$$

(C.27) implies that $v(t - \Delta t, \mathbf{x})$ is a non-decreasing monotone function of $v(t, [a \pm \Delta a, r \pm \Delta r, \chi \pm \Delta \chi])$. Applying (C.27) to the time-update scheme (C.26), we get the condition (5.51) required for a monotone scheme. Lastly, we check for conditions that ensure stability of the numerical scheme. In our analysis, we show numerical stability in the L^∞ -norm. Using the triangle inequality on (C.26),

$$\begin{aligned}
|v(t - \Delta t, \mathbf{x})| & \leq |v(t, \mathbf{x})| \left| 1 - b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} - b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta \chi^2} \right. \\
& \quad \left. - \frac{\Delta t}{\Delta a} |a_1(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x})))| \right. \\
& \quad \left. - \frac{\Delta t}{\Delta r} |a_2(t, \mathbf{x})| - \frac{\Delta t}{\Delta \chi} |a_3(t, \mathbf{x})| \right| \\
& + |v(t, [a + \Delta a, r, \chi])| \frac{\Delta t}{\Delta a} a_1^+(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \\
& + |v(t, [a - \Delta a, r, \chi])| \frac{\Delta t}{\Delta a} a_1^-(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \\
& + |v(t, [a, r + \Delta r, \chi])| \left(\frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} + a_2^+(t, \mathbf{x}) \frac{\Delta t}{\Delta r} \right) \\
& + |v(t, [a, r - \Delta r, \chi])| \left(\frac{1}{2} b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} + a_2^-(t, \mathbf{x}) \frac{\Delta t}{\Delta r} \right)
\end{aligned}$$

$$\begin{aligned}
& + |v(t, [a, r, \chi + \Delta\chi])| \left(\frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta\chi^2} + a_3^+(t, \mathbf{x}) \frac{\Delta t}{\Delta\chi} \right) \\
& + |v(t, [a, r, \chi - \Delta\chi])| \left(\frac{1}{2} b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta\chi^2} + a_3^-(t, \mathbf{x}) \frac{\Delta t}{\Delta\chi} \right) \\
& + \Delta t \left| g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \right. \\
& \quad \left. + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) > \phi_{\text{th}}\}} - \epsilon \right) \right|.
\end{aligned}$$

We further bound this using the supremum of the numerical solution v at time t over all \mathbf{x} .

$$\begin{aligned}
|v(t - \Delta t, \mathbf{x})| & \leq \left(\sup_{\mathbf{x} \in \mathbb{S}} |v(t, \mathbf{x})| \right) \left(\left| 1 - b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} - b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta\chi^2} \right. \right. \\
& \quad - \frac{\Delta t}{\Delta a} |a_1(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x})))| \\
& \quad - \frac{\Delta t}{\Delta r} |a_2(t, \mathbf{x})| - \frac{\Delta t}{\Delta\chi} |a_3(t, \mathbf{x})| \left. \right| \\
& \quad + \frac{\Delta t}{\Delta a} |a_1(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x})))| \\
& \quad + b_{2,2}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta r^2} + |a_2(t, \mathbf{x})| \frac{\Delta t}{\Delta r} \\
& \quad + b_{3,3}(t, \mathbf{x})^2 \frac{\Delta t}{\Delta\chi^2} + |a_3(t, \mathbf{x})| \frac{\Delta t}{\Delta\chi} \left. \right) \\
& \quad + \Delta t \sup_{\mathbf{x} \in \mathbb{S}} \left| g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) \right. \\
& \quad \left. + \lambda(t) \left(\mathcal{I}_{\{\phi_{\text{out}}(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x}))) > \phi_{\text{th}}\}} - \epsilon \right) \right|.
\end{aligned}$$

The numerical scheme achieves the stability condition [148]

$$\|v(t - \Delta t, \mathbf{x})\|_\infty \leq \|v(t, \mathbf{x})\|_\infty + \Delta t \|g(t, \mathbf{x}, \boldsymbol{\phi}(t, \mathbf{x}, \Delta_a^\pm v(t, \mathbf{x})))\|_\infty,$$

under the following conditions.

1. The running cost function g is bounded.
2. CFL condition (5.51) holds.

Here, $\|v(t, \mathbf{x})\|_\infty := \sup_{\mathbf{x} \in \mathbb{S}} |v(t, \mathbf{x})|$ indicates the L^∞ -norm of the value function. Consistency, monotonicity, and stability implies convergence of the numerical scheme for the nonlinear HJB equation (C.24) [149, 150, 131].

C.2.2 Boundary conditions

In general, boundary conditions must be imposed on the HJB PDE (5.36) to obtain a unique numerical solution. However, boundary conditions at the boundaries $a = 0, 1$, $r = 0, 1$, and $\chi = 0$ are unnecessary. To understand this, we make the following observations.

- When $a = 0$, the constraints in (5.15) ensure $a_1 \geq 0$.
- When $a = 1$, the constraints in (5.15) ensure $a_1 \leq 0$.
- When $r = 0$, the construction of the SDE (5.13) in [125] ensures that $a_2 \geq 0$ and $b_{2,2} = 0$.
- When $r = 1$, the construction of the SDE (5.13) in [125] ensures that $a_2 \leq 0$ and $b_{2,2} = 0$.
- When $\chi = \underline{\xi}$, the drift and diffusion functions in SDE (5.7) ensures that $a_3 \geq 0$ and $b_{3,3} = 0$.

At the boundary $\chi = \bar{\chi}$, we impose a nonreflective boundary condition to indicate that the actual domain is not bounded at $\chi = \bar{\chi}$ (like in Appendix C.1). This leads to the following boundary condition at $\chi = \bar{\chi}$:

$$v(t, [a, r, \bar{\chi} + \Delta\chi]) = 2v(t, [a, r, \bar{\chi}]) - v(t, [a, r, \bar{\chi} - \Delta\chi]).$$

This leads to the following update rule at $k = N_3$ for all n, i, j :

$$\begin{aligned} v_{(i,j,N_3)}^{n-1} = & v_{(i,j,N_3)}^n \left(1 - \left(b_{2,2} |_{(i,j,N_3)}^n \right)^2 \frac{\Delta t}{\Delta r^2} - \frac{\Delta t}{\Delta a} \left| a_1 |_{(i,j,N_3)}^n \right| \right. \\ & \left. - \frac{\Delta t}{\Delta r} \left| a_2 |_{(i,j,N_3)}^n \right| - \frac{\Delta t}{\Delta \chi} \left| a_3 |_{(i,j,N_3)}^n \right| \right) \\ & + v_{(i+1,j,N_3)}^n \frac{\Delta t}{\Delta a} a_1^+ |_{(i,j,N_3)}^n + v_{(i-1,j,N_3)}^n \frac{\Delta t}{\Delta a} a_1^- |_{(i,j,N_3)}^n \\ & + v_{(i,j+1,N_3)}^n \left(\left(b_{2,2} |_{(i,j,N_3)}^n \right)^2 \frac{\Delta t}{2\Delta r^2} + a_2^+ |_{(i,j,N_3)}^n \frac{\Delta t}{\Delta r} \right) \\ & + v_{(i,j-1,N_3)}^n \left(\left(b_{2,2} |_{(i,j,N_3)}^n \right)^2 \frac{\Delta t}{2\Delta r^2} + a_2^- |_{(i,j,N_3)}^n \frac{\Delta t}{\Delta r} \right) \\ & - v_{(i,j,N_3-1)}^n a_3^- |_{(i,j,N_3)}^n \frac{\Delta t}{\Delta \chi} + \Delta t g^\lambda |_{(i,j,N_3)}^n. \end{aligned} \quad (\text{C.28})$$

This leads to a unique numerical solution for the HJB PDE (5.36) using an upwind FDM scheme.

Appendix D

Details of numerical problem described in Chapter 5

D.1 Base station model parameters

	Unit	Description	Value
C_{scal}		Base station power-loss scaling factor	7.84
C_{offset}	Watt (W)	Base station offset power	71.5
\bar{P}_{tx}	Watt (W)	Maximum base station transmission limit	5×10^3
\mathbf{x}_{BS}		Base station location	[0,0]
κ		Path loss constant	1
η		Path loss exponent	2
SNR_{th}	Decibel (dB)	Signal-to-noise ratio threshold	15
σ_0	Watt (W)	Ambient transmission noise	3.16×10^{-8}
\bar{P}_R	Watt (W)	Max. renewable power-production capacity	10^4
\bar{A}	Watt-hour (Wh)	Maximum battery charge capacity	10^4
\underline{P}_A	Watt (W)	Maximum battery charge capacity	7.5×10^3
\bar{P}_A	Watt (W)	Maximum battery discharge capacity	3×10^4
C_1	€/Wh	Pollutant emission Coefficient 1	4×10^{-4}
C_2	€/W ² h	Pollutant emission Coefficient 2	10^{-4}
P_K	€/Wh	Fictitious cost per unit battery charge	0.0064
w		Pareto parameter	0.5

Table D.1: Parameters and coefficients for modeling the cellular base station described in Figure 5.2.

D.2 Outage proportion for simple cellular user distributions

Equation (5.10) demonstrates that the outage proportion for a given distribution of cellular users $\rho_z(t)$ at time t can be written as follows:

$$\phi_{\text{out}}(t, \zeta(t), P_{\text{tx}}(t)) = \int_{\mathbb{R}^2} \mathcal{I} \left\{ \|z - \mathbf{x}_{\text{BS}}\| > \left(\frac{P_{\text{tx}}(t)\zeta(t)\kappa}{\sigma_0 10^{-\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{1}{\eta}} \right\} \rho_z(t) \mathrm{d}z.$$

First, we consider the mobile user distribution to be a two-dimensional uniform dis-

tribution in the geographical domain Π served by the base station. The domain area is denoted as A_Π . Then,

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = 1 - \frac{\pi}{A_\Pi} \left(\frac{P_{\text{tx}}(t) \xi(t) \kappa}{\sigma_0 10^{\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{2}{\eta}}, \quad \text{for } \rho_z(t) = \mathfrak{U}(\Pi). \quad (\text{D.1})$$

Here $\mathfrak{U}(\Pi)$ denotes the described uniform user distribution in the region Π . Next, we consider the case in which the mobile user distribution is a two-dimensional Gaussian distribution centered at the site of the base station \mathbf{x}_{BS} with a diagonal covariance matrix with variance σ_u^2 along both directions:

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = \exp \left(-\frac{1}{2\sigma_u^2} \left(\frac{P_{\text{tx}}(t) \xi(t) \kappa}{\sigma_0 10^{\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{2}{\eta}} \right), \quad \text{for } \rho_z(t) = \mathcal{N}(\mathbf{x}_{\text{BS}}, \sigma_u^2 \mathbb{I}_2), \quad (\text{D.2})$$

where \mathbb{I}_2 denotes the two-dimensional identity matrix. Next, we consider the case in which the mobile user distribution is a two-dimensional Gaussian distribution centered at $\boldsymbol{\mu} \in \mathbb{R}^2$ (not at \mathbf{x}_{BS}) and a diagonal covariance matrix with variance σ_u^2 along both directions:

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = \mathfrak{Q}_1 \left(\|\mathbf{x}_{\text{BS}} - \boldsymbol{\mu}\|, \left(\frac{P_{\text{tx}}(t) \xi(t) \kappa}{\sigma_u \sigma_0 10^{\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{1}{\eta}} \right), \quad \text{for } \rho_z(t) = \mathcal{N}(\boldsymbol{\mu}, \sigma_u^2 \mathbb{I}_2), \quad (\text{D.3})$$

where $\mathfrak{Q}_\beta(c_1, c_2)$ denotes the generalized Marcum Q-function of order β with non-centrality parameter c_1 and threshold parameter c_2 [151]. Next, we consider the case in which the mobile user distribution is a two-dimensional Gaussian distribution centered at the base station \mathbf{x}_{BS} with a diagonal covariance matrix with variance σ_1^2 along the x -direction and σ_2^2 along the y -direction.

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = \mathbb{P} \left[\tilde{\chi}^2([\sigma_1^2, \sigma_2^2], [1, 1], [0, 0]) \geq \left(\frac{P_{\text{tx}}(t) \xi(t) \kappa}{\sigma_0 10^{\frac{\text{SNR}_{\text{th}}}{10}}} \right)^{\frac{2}{\eta}} \right], \quad (\text{D.4})$$

for $\rho_z(t) = \mathcal{N} \left(\mathbf{x}_{\text{BS}}, \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \right)$ where $\tilde{\chi}^2(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$ is a two-dimensional generalized chi-squared random variable with weight \mathbf{c}_1 , degrees of freedom \mathbf{c}_2 , and shift \mathbf{c}_3 . There is no analytical expression for this probability, but numerical implementation is available in Python and Matlab up to arbitrary precision.

D.3 Analytical solution of the Hamiltonian (5.37) for simple cellular user distributions

For each point $(t, \mathbf{x}) \in [0, T] \times \mathcal{S}$, we must solve the optimization problem in (5.37). For uniform or Gaussian cellular user distributions with closed-form expressions for ϕ_{out} (see Appendix D.2), this simplifies to solving the following constrained convex optimization problem.

$$\left\{ \begin{array}{l} \min_{P_F, P_S, P_{tx}} a_1 P_F + a_2 P_F^2 + b_1 P_S + c_1 P_{tx} + c_2 \exp\{-c_3 P_{tx}\} \\ \quad + \lambda(t) \mathcal{I}_{\{\phi_{out}(t, \xi, P_{tx}) > \phi_{th}\}} \\ \text{s.t. } P_{tx} - d_3 \leq 0 \\ \quad - P_F + P_S + d_1 P_{tx} + d_2 \leq 0 \\ \quad P_F - P_S - d_1 P_{tx} + d_4 \leq 0 \\ \quad - P_F \leq 0 \\ \quad - P_{tx} \leq d_5 \\ \quad - P_S \leq 0 \\ \quad P_F - d_1 P_{tx} - C_{offset} \leq 0 \end{array} \right. \quad (D.5)$$

where

$$\begin{aligned} a_1 &:= \frac{1}{A} \frac{\partial v}{\partial a}(t, \mathbf{x}) + w K_B(t) + (1-w) C_1, \\ a_2 &:= (1-w) C_2, \\ b_1 &:= -\frac{1}{A} \frac{\partial v}{\partial a}(t, \mathbf{x}) - w K_B(t), \\ c_1 &:= \begin{cases} -\frac{C_{scal} N_u(t)}{A} \frac{\partial v}{\partial a}(t, \mathbf{x}) - \frac{w K_{net}(t) \pi \chi \kappa}{A_{II} \sigma_0 10^{-\frac{SNR_{th}}{10}}} & , \text{if } \rho_z(t) = \mathcal{U}(II), \\ -\frac{C_{scal} N_u(t)}{A} \frac{\partial v}{\partial a}(t, \mathbf{x}) & , \text{if } \rho_z(t) = \mathcal{N}(\mathbf{x}_{BS}, \sigma_u^2 \mathbb{I}_2). \end{cases} \\ c_2 &:= \begin{cases} 0 & , \text{if } \rho_z(t) = \mathcal{U}(II), \\ w K_{net}(t) N_u(t) & , \text{if } \rho_z(t) = \mathcal{N}(\mathbf{x}_{BS}, \sigma_u^2 \mathbb{I}_2). \end{cases} \\ c_3 &:= \begin{cases} 0 & , \text{if } \rho_z(t) = \mathcal{U}(II), \\ \frac{\chi \kappa}{2 \sigma_u^2 \sigma_0 10^{-\frac{SNR_{th}}{10}}} & , \text{if } \rho_z(t) = \mathcal{N}(\mathbf{x}_{BS}, \sigma_u^2 \mathbb{I}_2). \end{cases} \\ d_1 &:= C_{scal} N_u(t), \\ d_2 &:= -\bar{P}_A(a) - r \bar{P}_R + C_{offset}, \\ d_3 &:= \frac{\bar{P}_{tx}}{N_u(t)}, \\ d_4 &:= \underline{P}_A(a) + r \bar{P}_R - C_{offset}, \\ d_5 &:= \begin{cases} \frac{(1-\phi_{th}) A_{II} \sigma_0 10^{-\frac{SNR_{th}}{10}}}{\pi \chi \kappa} & , \text{if } \rho_z(t) = \mathcal{U}(II), \\ -\frac{2 \sigma_u^2 \sigma_0 10^{-\frac{SNR_{th}}{10}} \log \phi_{th}}{\xi \kappa} & , \text{if } \rho_z(t) = \mathcal{N}(\mathbf{x}_{BS}, \sigma_u^2 \mathbb{I}_2). \end{cases} \end{aligned}$$

The discontinuous indicator function in the objective function poses a difficulty while deriving the analytical solution to (D.5). Decomposing (D.5) simplifies this into two

constrained convex optimization subproblems, which are

$$\left\{ \begin{array}{l} \mathcal{H}_1 := \min_{P_F, P_S, P_{\text{tx}}} a_1 P_F + a_2 P_F^2 + b_1 P_S + c_1 P_{\text{tx}} + c_2 \exp\{-c_3 P_{\text{tx}}\} \\ \text{s.t. } -P_{\text{tx}} \leq -d_5 \\ P_{\text{tx}} - d_3 \leq 0 \\ -P_F + P_S + d_1 P_{\text{tx}} + d_2 \leq 0 \\ P_F - P_S - d_1 P_{\text{tx}} + d_4 \leq 0 \\ -P_F \leq 0 \\ -P_S \leq 0 \\ P_F - d_1 P_{\text{tx}} - C_{\text{offset}} \leq 0. \end{array} \right.$$

and

$$\left\{ \begin{array}{l} \mathcal{H}_2 := \min_{P_F, P_S, P_{\text{tx}}} a_1 P_F + a_2 P_F^2 + b_1 P_S + c_1 P_{\text{tx}} + c_2 \exp\{-c_3 P_{\text{tx}}\} + \lambda(t) \\ \text{s.t. } P_{\text{tx}} \leq d_5 \\ -P_F + P_S + d_1 P_{\text{tx}} + d_2 \leq 0 \\ P_F - P_S - d_1 P_{\text{tx}} + d_4 \leq 0 \\ -P_F \leq 0 \\ -P_{\text{tx}} \leq 0 \\ -P_S \leq 0 \\ P_F - d_1 P_{\text{tx}} - C_{\text{offset}} \leq 0. \end{array} \right.$$

The optimal controls for (D.5) are given as

$$\arg \min_{P_F, P_S, P_{\text{tx}}} (\mathcal{H}_1, \mathcal{H}_2).$$

Analytical solutions to these subproblems can be derived by solving the corresponding Karush–Kuhn–Tucker conditions for constrained optimization problems [130].

D.4 Solution to Problem 5.4.1

We gather all the elements of the SOC Problem 5.4.1.

1. **Time horizon:** $[0, T]$ where 0 denotes the start of the current day and T denotes the end of next day.
2. **(Controlled) state process:** The stochastic process $\mathbf{X} : 0, T \times \Omega \rightarrow \mathbb{R}^3$ follows the 3-dimensional Itô SDE (5.20). The state space is $\mathbb{S} = [0, 1] \times [0, 1] \times [\underline{\xi}, \infty)$.
3. **Control process:** The control process is denoted by the \mathcal{F}_t -adapted Markov process $\Phi = [P_A, P_F, P_S, P_{\text{tx}}]$. That is, $\Phi(t) = \phi(t, \mathbf{X}(t))$, where the policy is denoted by ϕ .
4. **Admissible controls:** The policy ϕ must satisfy the following equality a.s. constraints arising from the power balance equation (5.5) of the base station, and the a.s. QoS constraint, respectively.

$$P_A(t) + P_F(t) - C_{\text{scal}} N_u(t) P_{\text{tx}}(t) - C_{\text{offset}} = 0, \quad 0 \leq t \leq T,$$

$$\phi_{\text{out}}(t, \xi(t), P_{\text{tx}}(t)) = \phi_{\text{th}}, \quad 0 \leq t \leq T.$$

The policy ϕ must also satisfy the following inequality a.s. constraints arising from the capacity constraints of the transmission and energy systems introduced before.

$$\begin{aligned} -P_A(t) &\leq 0, \\ -P_F(t) &\leq 0, \\ -P_{\text{tx}}(t) &\leq 0 \\ -P_S(t) &\leq 0, \\ \bar{P}_R R(t) - P_A(t) - P_S(t) - \underline{P}_A(A(t)) &\leq 0, \\ P_A(t) + P_S(t) - \bar{P}_R R(t) - \bar{P}_A(A(t)) &\leq 0. \end{aligned}$$

Note that we have removed the constraint (5.25) that defines a maximum limit for the transmitted power to ensure that the SOC Problem 5.4.1 is always feasible. Let $\mathbb{A}_1(t, \mathbf{X}(t))$ denote the set of controls that satisfy the above constraints at time t for given state $\mathbf{X}(t)$. For the battery characteristic functions \underline{P}_A and \bar{P}_A defined in Figure 5.9, and the simple user distributions introduced in Section D.2, $\mathbb{A}_1(t, \mathbf{X}(t))$ is a convex set for given $t \in [0, T]$ and state $\mathbf{X}(t)$. Subsequently, \mathcal{A}_1 denotes the set of all admissible policies ϕ satisfying the constraints for all $t \in [0, T]$ and for all realizations of \mathbf{X} .

$$\mathcal{A}_1 := \left\{ \phi : \phi \text{ is } \mathcal{F}_t\text{-adapted and Markovian, } \phi(t, \mathbf{X}(t, \omega)) \in \mathbb{A}_1(t, \mathbf{X}(t, \omega)), \right. \\ \left. \forall \omega \in \Omega, \quad \forall t \in [0, T] \right\}.$$

5. **Cost function:** The net cost to be minimised is written as

$$\begin{aligned} \mathcal{U} &= w(\mathcal{C}_1 - \mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + (1-w)\mathcal{C}_2 \\ &= \mathbb{E} \left[\int_0^T g(t, \mathbf{X}(t), \phi(t, \mathbf{X}(t))) dt + h(\mathbf{X}(T)) \right]. \end{aligned}$$

The objective is to minimise the following cost functional.

$$J_{t, \mathbf{x}}(\phi) = \mathbb{E} \left[\int_t^T g(s, \mathbf{X}(s), \phi(s, \mathbf{X}(s))) ds + h(\mathbf{X}(T)) \mid \mathbf{X}(t) = \mathbf{x} \right].$$

6. **Value function:** Let $v : [0, T] \times \mathbb{S} \rightarrow \mathbb{R}$ denote the value function that describes the minimal possible cost of the system.

$$v(t, \mathbf{x}) = \min_{\phi \in \mathcal{A}_1} J_{t, \mathbf{x}}(\phi).$$

For this standard SOC problem, the dynamic programming principle can be applied and the resulting HJB PDE follows.

$$\frac{\partial v}{\partial t}(t, \mathbf{x}) + (\dot{p}(t) - \theta(t)(r - p(t))) \frac{\partial v}{\partial r}(t, \mathbf{x}) - \theta(\chi - \underline{\xi} - \mu) \frac{\partial v}{\partial \chi}(t, \mathbf{x})$$

$$\begin{aligned}
& + \alpha \theta_0 r (1-r) \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) + \theta (\chi - \underline{\xi}) \frac{\partial^2 v}{\partial \chi^2}(t, \mathbf{x}) - w K_{\text{net}}(t) N_u(t) (1 - \phi_{\text{th}}) \\
& + \min_{\boldsymbol{\phi} \in \mathcal{A}_1(t, \mathbf{x})} \left[\frac{\bar{P}_R r - P_S - P_A}{\bar{A}} \frac{\partial v}{\partial a}(t, \mathbf{x}) + w (K_b(t) P_F - K_S(t) P_S) \right. \\
& \quad \left. + (1-w) (C_1 P_F + C_2 P_F^2) \right] = 0 \\
v(T, \mathbf{x}) & = -w P_K \bar{A} a.
\end{aligned}$$

The above PDE is solved numerically using the same upwind FDM scheme described in Section 5.3.1 with an analytical solution to the Hamiltonian.

D.5 Solution to Problem 5.4.2

We gather all the elements of the SOC Problem 5.4.2.

1. **Time horizon:** $[0, T]$ where 0 denotes the start of the current day and T denotes the end of next day.
2. **(Controlled) state process:** The stochastic process $\mathbf{X} : 0, T \times \Omega \rightarrow \mathbb{R}^3$ follows the 3-dimensional Itô SDE (5.20). The state space is $\mathcal{S} = [0, 1] \times [0, 1] \times [\underline{\xi}, \infty)$.
3. **Control process:** The control process is denoted by the \mathcal{F}_t -adapted Markov process $\boldsymbol{\Phi} = [P_A, P_F, P_S, P_{\text{tx}}]$. That is, $\boldsymbol{\Phi}(t) = \boldsymbol{\phi}(t, \mathbf{X}(t))$, where the policy is denoted by $\boldsymbol{\phi}$.
4. **Admissible controls:** The policy $\boldsymbol{\phi}$ must satisfy the following equality a.s. constraint arising from the power balance equation (5.5) of the base station.

$$P_A(t) + P_F(t) - C_{\text{scal}} N_u(t) P_{\text{tx}}(t) - C_{\text{offset}} = 0, \quad 0 \leq t \leq T.$$

The policy $\boldsymbol{\phi}$ must also satisfy the following inequality a.s. constraints arising from the capacity constraints of the transmission and energy systems introduced before.

$$\begin{aligned}
& -P_A(t) \leq 0, \\
& -P_F(t) \leq 0, \\
& -P_{\text{tx}}(t) \leq 0, \\
& P_{\text{tx}}(t) \leq \frac{\bar{P}_{\text{tx}}}{N_u(t)}, \\
& -P_S(t) \leq 0, \\
& \bar{P}_R R(t) - P_A(t) - P_S(t) - \underline{P}_A(A(t)) \leq 0, \\
& P_A(t) + P_S(t) - \bar{P}_R R(t) - \bar{P}_A(A(t)) \leq 0.
\end{aligned}$$

The admissible set \mathcal{A} of policies is the same as that of Problem 5.2.2 defined in (5.29).

5. **Cost function:** The net cost to be minimised is written as

$$\begin{aligned}
\mathcal{U} & = w (C_1 - \mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + (1-w) C_2 \\
& = \mathbb{E} \left[\int_0^T g(t, \mathbf{X}(t), \boldsymbol{\phi}(t, \mathbf{X}(t))) dt + h(\mathbf{X}(T)) \right].
\end{aligned}$$

The objective is to minimise the following cost functional.

$$J_{t,\mathbf{x}}(\boldsymbol{\phi}) = \mathbb{E} \left[\int_t^T g(s, \mathbf{X}(s), \boldsymbol{\phi}(s, \mathbf{X}(s))) ds + h(\mathbf{X}(T)) \middle| \mathbf{X}(t) = \mathbf{x} \right].$$

6. **Value function:** Let $v : [0, T] \times \mathbf{S} \rightarrow \mathbb{R}$ denote the value function that describes the minimal possible cost of the system.

$$v(t, \mathbf{x}) = \min_{\boldsymbol{\phi} \in \mathcal{A}} J_{t,\mathbf{x}}(\boldsymbol{\phi}).$$

For this standard **SOC** problem, the dynamic programming principle can be applied and the resulting **HJB PDE** follows.

$$\begin{aligned} & \frac{\partial v}{\partial t}(t, \mathbf{x}) + (\dot{p}(t) - \theta(t)(r - p(t))) \frac{\partial v}{\partial r}(t, \mathbf{x}) - \theta(\chi - \underline{\zeta} - \mu) \frac{\partial v}{\partial \chi}(t, \mathbf{x}) \\ & + \alpha \theta_0 r(1 - r) \frac{\partial^2 v}{\partial r^2}(t, \mathbf{x}) + \theta(\chi - \underline{\zeta}) \frac{\partial^2 v}{\partial \chi^2}(t, \mathbf{x}) \\ & + \min_{\boldsymbol{\phi} \in \mathcal{A}(t, \mathbf{x})} \left[\frac{\bar{P}_R r - P_S - P_A}{\bar{A}} \frac{\partial v}{\partial a}(t, \mathbf{x}) + w \left(K_b(t) P_F - K_S(t) P_S \right. \right. \\ & \quad \left. \left. - K_{\text{net}}(t) N_u(t) (1 - \phi_{\text{out}}(t, \chi, P_{\text{tx}})) \right) + (1 - w) (C_1 P_F + C_2 P_F^2) \right] = 0 \\ & v(T, \mathbf{x}) = -w P_K \bar{A} a. \end{aligned}$$

The above **PDE** is solved numerically using the same upwind **FDM** scheme described in Section 5.3.1 with an analytical solution to the Hamiltonian.

D.6 Simulation parameters in Section 5.4

Parameter	Description	Value
ϕ_{th}	Mobile user outage proportion threshold	10^{-3}
ϵ	Confidence level of violating the constraint (5.30)	0.1
A_0	Normalized battery charge level at $t = 0$	0.5
$\mu_0^{\tilde{\epsilon}}$	Initial distribution of the wireless fading channel	$\gamma(3,1)$
N_1	Discretization of (5.36) in the a domain	10
N_2	Discretization of (5.36) in the r domain	10
N_3	Discretization of (5.36) in the χ domain	10
N_t	Discretization of (5.36) in the t domain	800
TOL	Prescribed relative tolerance for Algorithm 5	0.1
TOL _{init}	Prescribed relative tolerance for Algorithm 15	1
k_{max}	Prescribed maximum iterations in Algorithm 5	50
\tilde{N}_{iter}	Initial number of SSM iterations with a constant step-size	10
N_{iter}	Prescribed number of LMBM iterations	50
β_F	Factor of increase/decrease in Algorithm 15	5
M_{SG}	Number of sample paths in Algorithm 14	10^3
\tilde{N}_t	Time discretization parameter in Algorithm 14	64
\tilde{N}_t	Time discretization parameter in Algorithm 14	$\frac{\tilde{N}_t}{\ell}$

Table D.2: Simulation parameters to run Algorithm 5 and produce numerical results as described in Section 5.4.

Parameter	Description	Value
RPAR(1)	Tolerance for changes in the function value	0.1
RPAR(2)	Second tolerance for changes in the function value	-1 (ignored)
RPAR(3)	Minimum acceptable function value	0
RPAR(4)	Tolerance for the first termination parameter	10^{-2}
RPAR(5)	Tolerance for the second termination parameter	10^{-2}
RPAR(6)	Distance measure parameter	0.5
RPAR(7)	Line search parameter	0.2
RPAR(8)	Maximum step size	10
IPAR(1)	Exponent for distance measure	2
IPAR(2)	Maximum iterations	50
IPAR(3)	Maximum function evaluations	100
IPAR(4)	Maximum iterations with change smaller than RPAR(1)	5
IPAR(5)	Printout specification	-1
IPAR(6)	Selection of method	0 (LMBM)
IPAR(7)	Selection of scaling strategy	0

Table D.3: Parameters to run the LMBM routine [7] in Algorithm 5 to produce numerical results described in Section 5.4. See [7] for more details.

D.7 Sensitivity analysis settings in Section 5.4

Parameter	Distribution
\bar{P}_{tx}	$\log_{10}(\bar{P}_{\text{tx}}) \sim \mathcal{U}[3, 4]$
σ_0	$\sigma_0 \sim \mathcal{N}(3.1623 \times 10^{-8}, 10^{-16})$
\bar{P}_R	$\log_{10}(\bar{P}_R) \sim \mathcal{U}[2, 4]$
\bar{A}	$\log_{10}(\bar{A}) \sim \mathcal{U}[3, 4]$
w	$w \sim \mathcal{U}[0, 1]$
ϕ_{th}	$\log_{10}(\phi_{\text{th}}) \sim \mathcal{U}[-2, -4]$
A_0	$A_0 \sim \mathcal{U}[0, 1]$
K_{net}	$\log_{10}(K_{\text{net}}) \sim \mathcal{U}[-1, 1]$

Table D.4: Distributions of model and algorithmic parameters used to run randomized multi-scenario simulations of Algorithm 5.

D.8 Detailed results from multi-scenario simulations in Section 5.4

In this Section, we provide plots detailing the results of applying Algorithm 5 to the five scenarios introduced in the sensitivity analysis in Section 5.4. Figure D.1 reveals the 1D subgradient norm versus Y_1^1 at each iteration of Algorithm 15 for each scenario. In Scenarios A, B, and D, the subgradient norm goes below the prescribed initial relative tolerance TOL_{init} . In Scenarios C and E, Algorithm 15 stops as the subgradient changes sign, even though its absolute value is above the prescribed tolerance.

Figures D.2 and D.3 plots the evolving optimal dual function $\bar{\Theta}(\lambda^\ell)$ and its corresponding subgradient $\bar{\mathcal{D}}\Theta(\lambda^\ell)$ constructed using the optimal amplitudes \mathbf{Y}^ℓ at each level ℓ of Algorithm 5. In all scenarios, Algorithm 5 converges to the prescribed relative tolerance TOL with respect to the subgradient norm at each refinement level ℓ . As expected, the dual costs given by Algorithm 5 always lie above the primal optimal costs of the corresponding unconstrained Problem 5.4.2.

Figures D.4 to D.8 illustrate the optimal power procurement policy for the model cellular base stations under the various scenarios for the next 24 hours using the optimal solution from Algorithm 5. Table 5.3 presents the areas under the mean curves in Figures D.4 to D.8 to compare the optimal policies among different scenarios. Figure D.6 demonstrates that a large amount of grid power is bought in Scenarios A and C, either due to no incoming renewable power or due to low weightage given to minimizing the base station's carbon footprint. Figure D.7 reveals that it is not optimal to sell much power from the battery even though the battery has some stored charge, highlighting the effect of implementing the running horizon framework.

Figure D.9 presents the evolution of the optimal Lagrange multiplier functions as the refinement level ℓ increases in Algorithm 5. Figure D.9 reveals that when the required confidence $1 - \epsilon$ of satisfying the chance constraint (5.30) decreases, the Lagrange multiplier also decreases as we expect from Remark 5.3.5. Figure D.10 visualizes the time-pointwise violation of the constraint (5.30) using MC and compares it with the final subgradient vector $\bar{\mathcal{D}}\Theta(\lambda^\ell)$ from Algorithm 5 for different scenarios. Figure D.10 indicates that the proposed optimal solution controls the chance constraint (5.30) in an integral sense (see (5.44)) for all scenarios, although there are some oscillations around zero in the time-pointwise constraint violation values. Figure D.10e also indicates that the oscillations

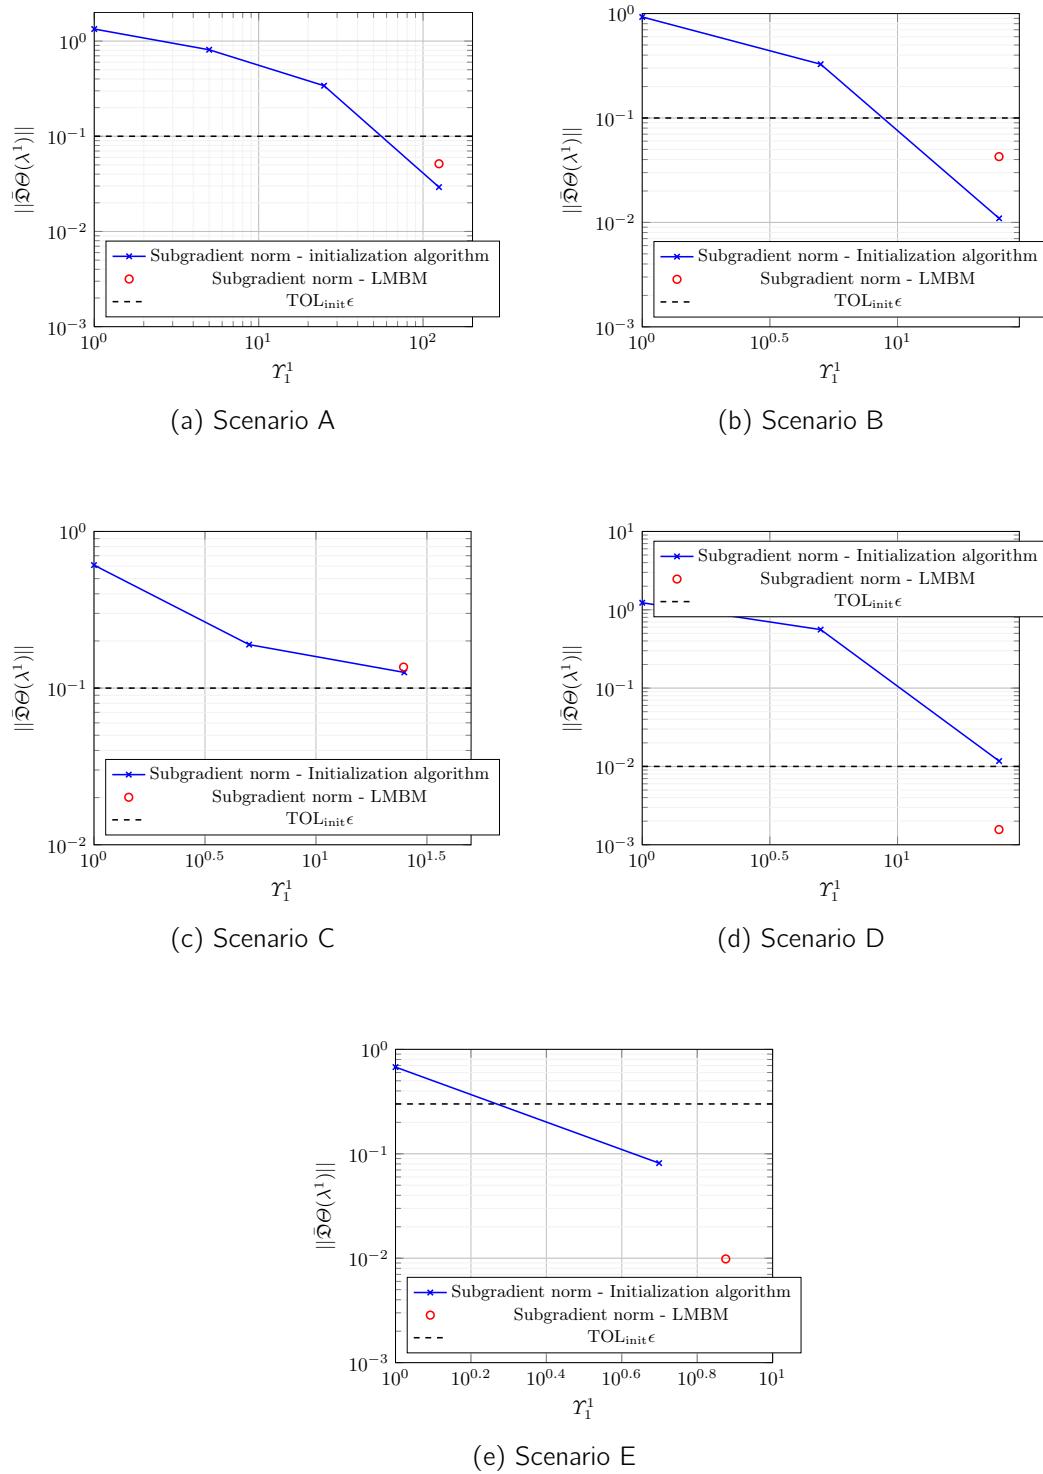




Figure D.1: Subgradient norm computed with λ^1 constructed from Υ_1^1 attained at each iteration of Algorithm 15 for different scenarios, along with the corresponding quantities using the final output of the LMBM routine [7].


increase when the required confidence $1 - \epsilon$ of satisfying the chance constraint (5.30) reduces.




(a) Scenario A




(b) Scenario B



(c) Scenario C



(d) Scenario D



(e) Scenario E

Figure D.2: Dual cost computed with λ^ℓ constructed from the optimal \mathbf{Y}^ℓ obtained from the SSM at each refinement level ℓ of Algorithm 5 for different scenarios.

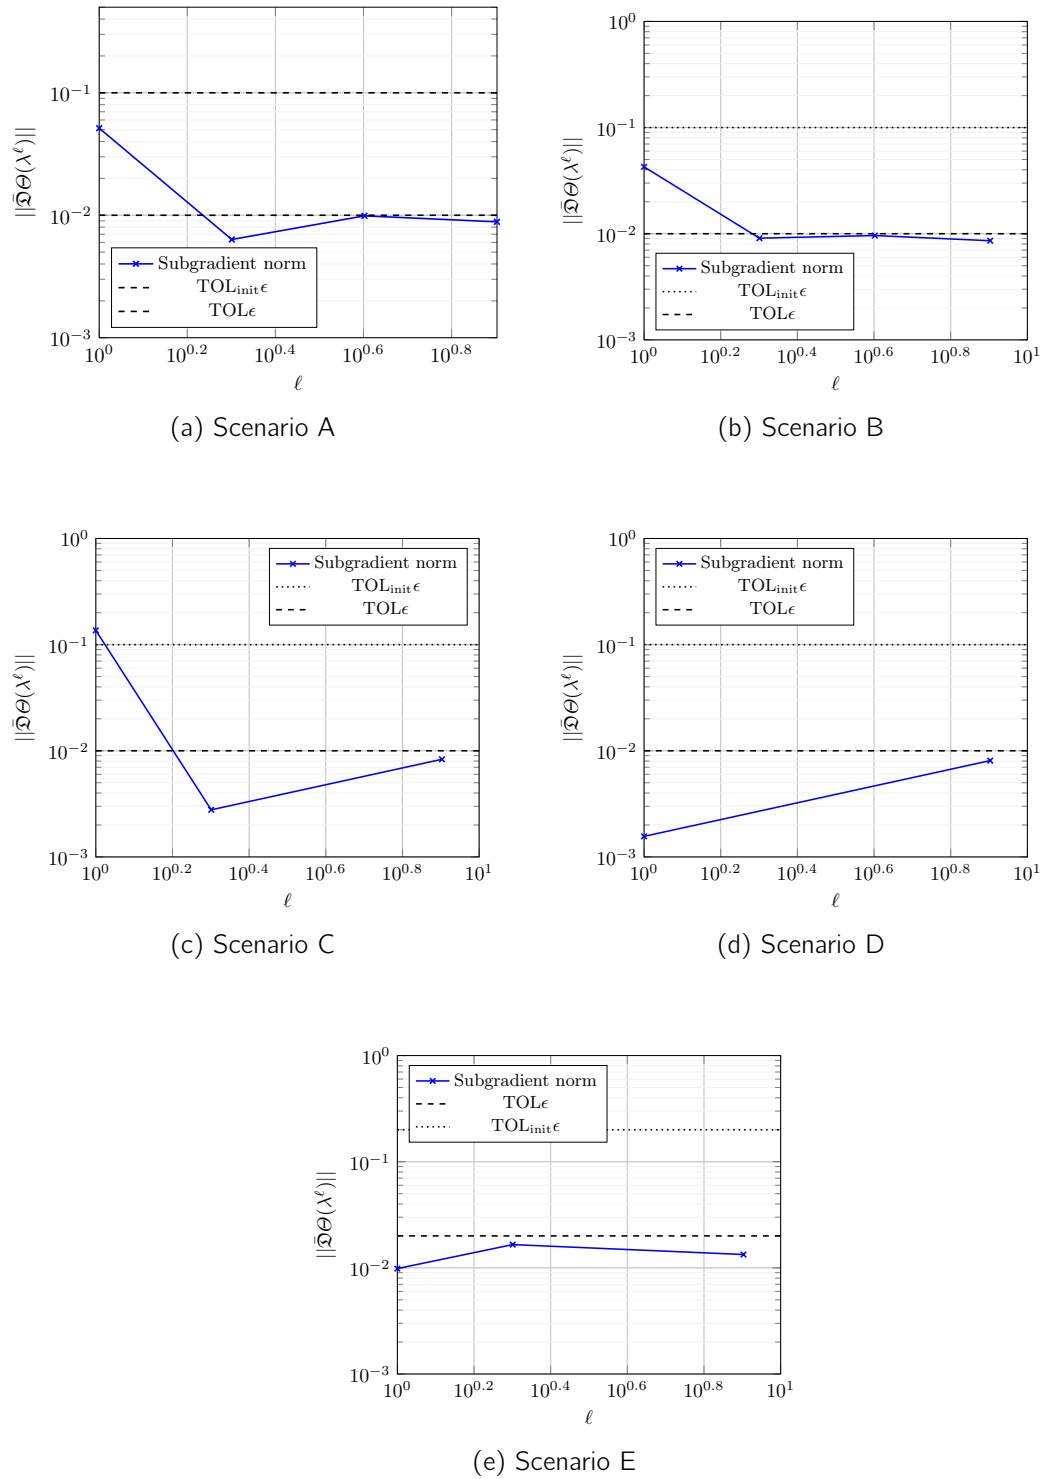


Figure D.3: Subgradient norm computed with λ^ℓ constructed from the optimal \mathbf{Y}^ℓ obtained from the **SSM** at each refinement level ℓ of Algorithm 5 for different scenarios.

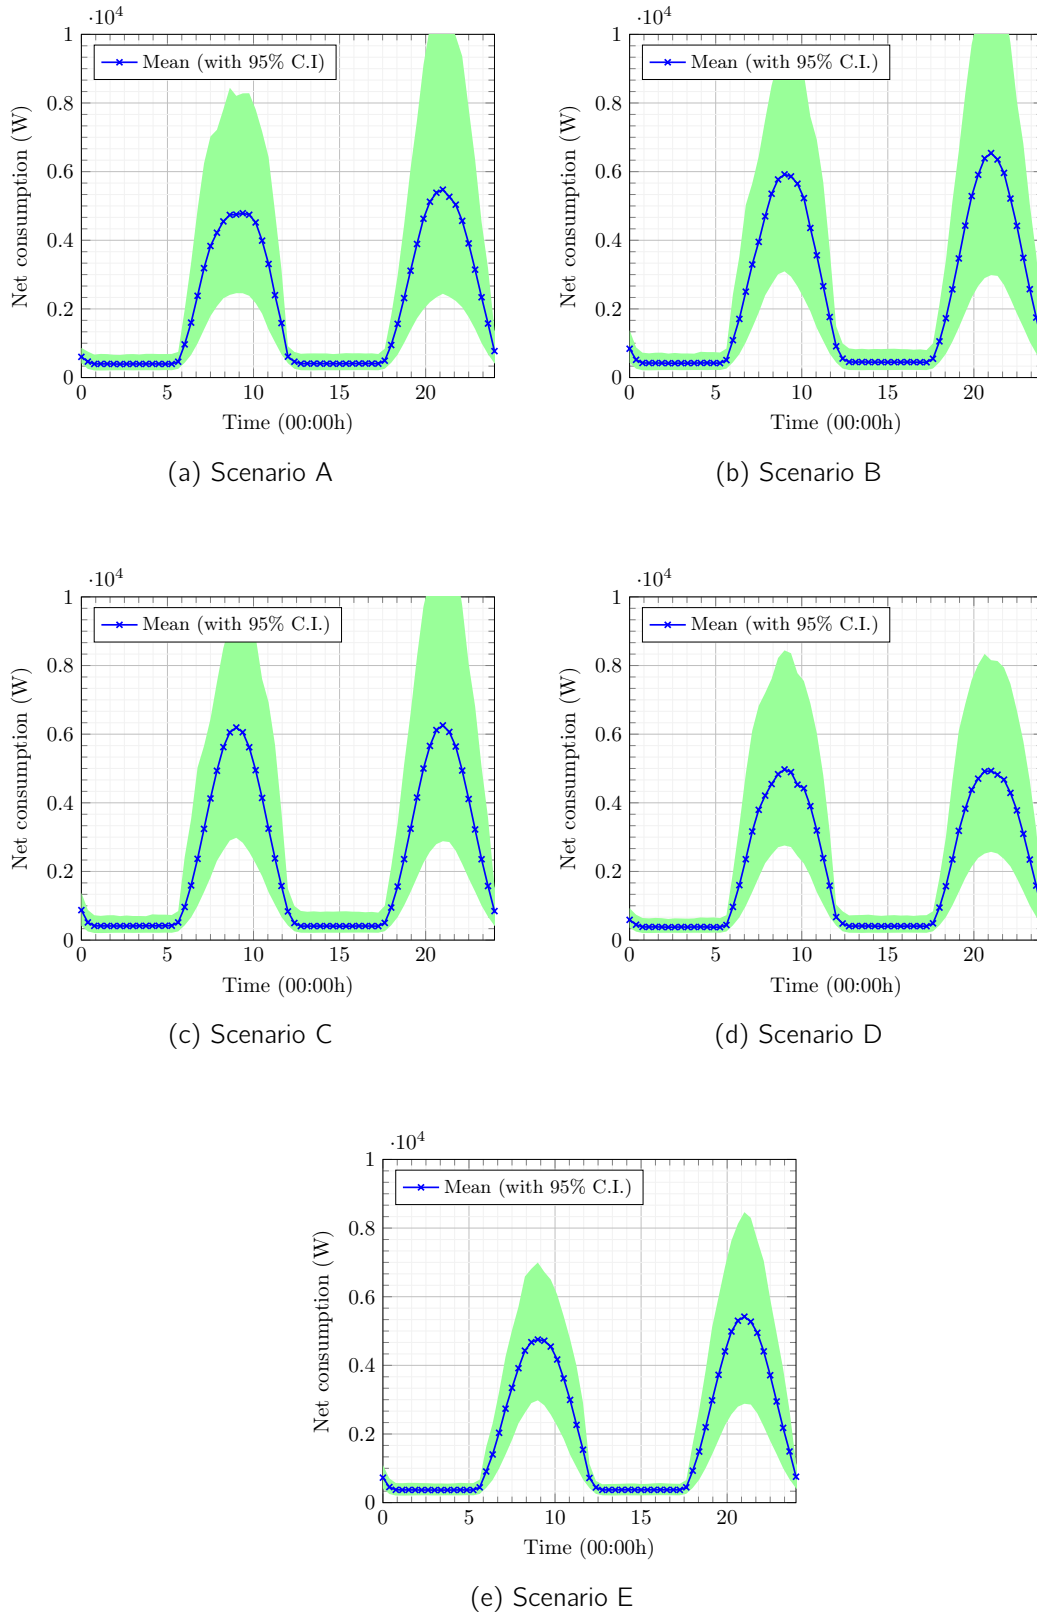


Figure D.4: Optimal base station power consumption ($C_{\text{scal}}N_u(t)P_{\text{tx}}(t) + C_{\text{offset}}$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.

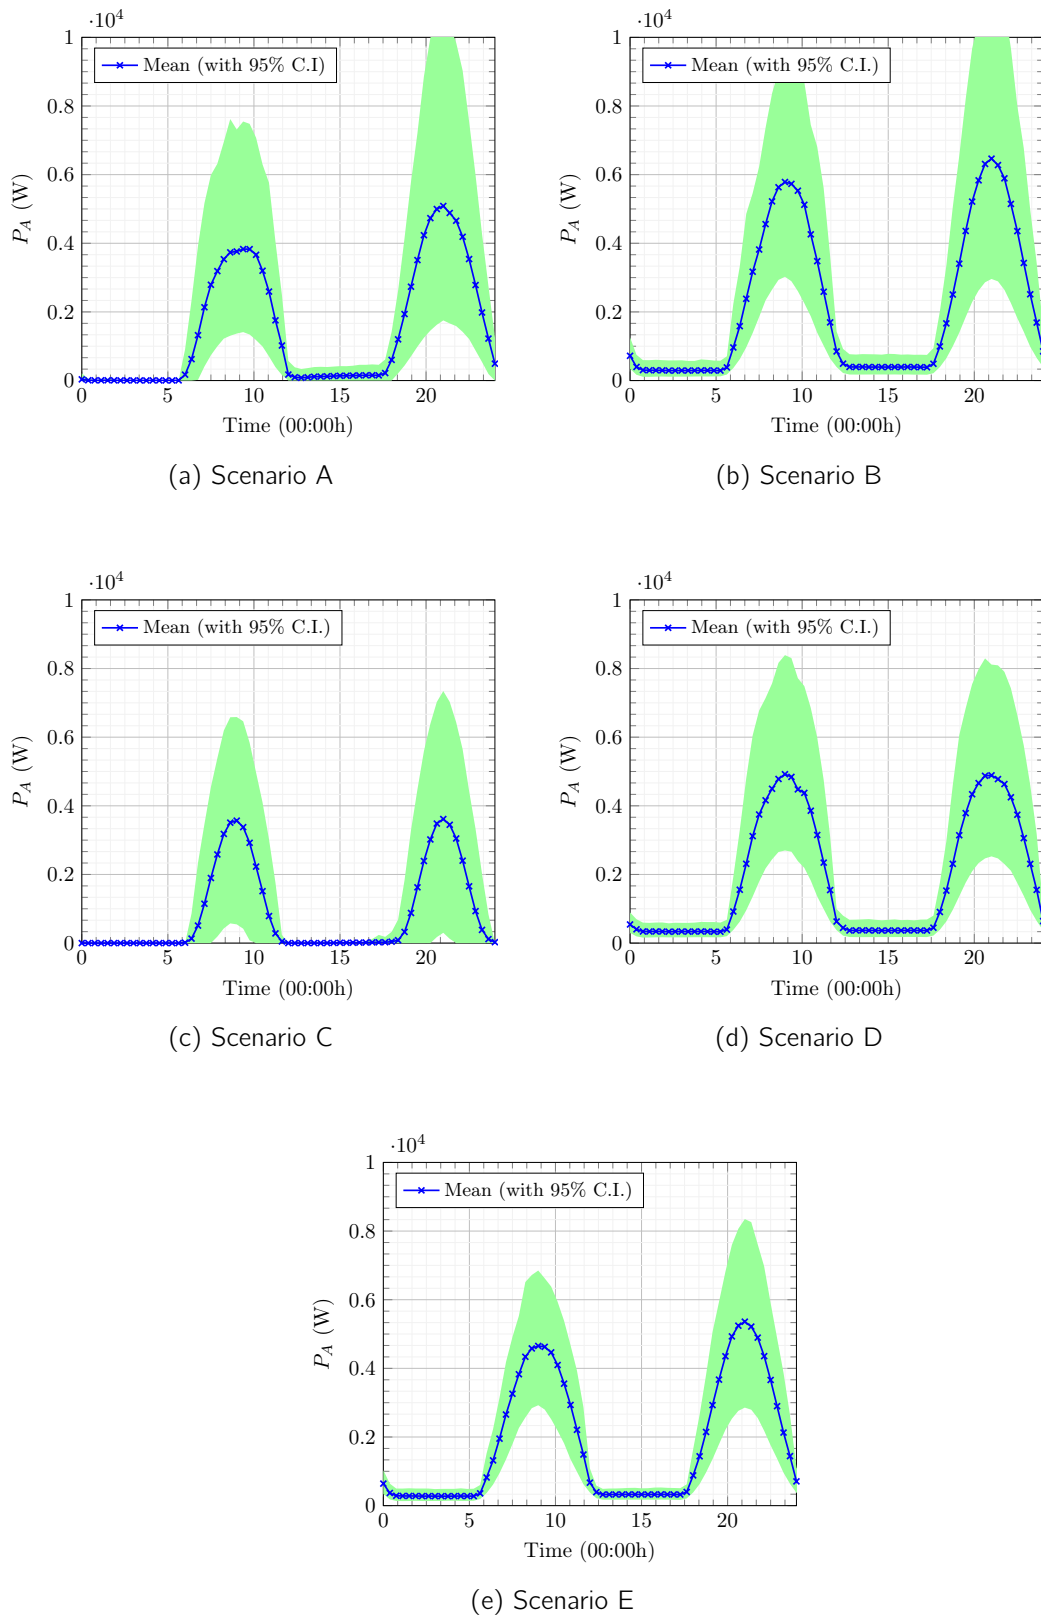





Figure D.5: Optimal battery power usage policy ($P_A(t)$) for transmission, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.




(a) Scenario A




(b) Scenario B



(c) Scenario C



(d) Scenario D



(e) Scenario E

Figure D.6: Optimal grid power purchase ($P_F(t)$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.

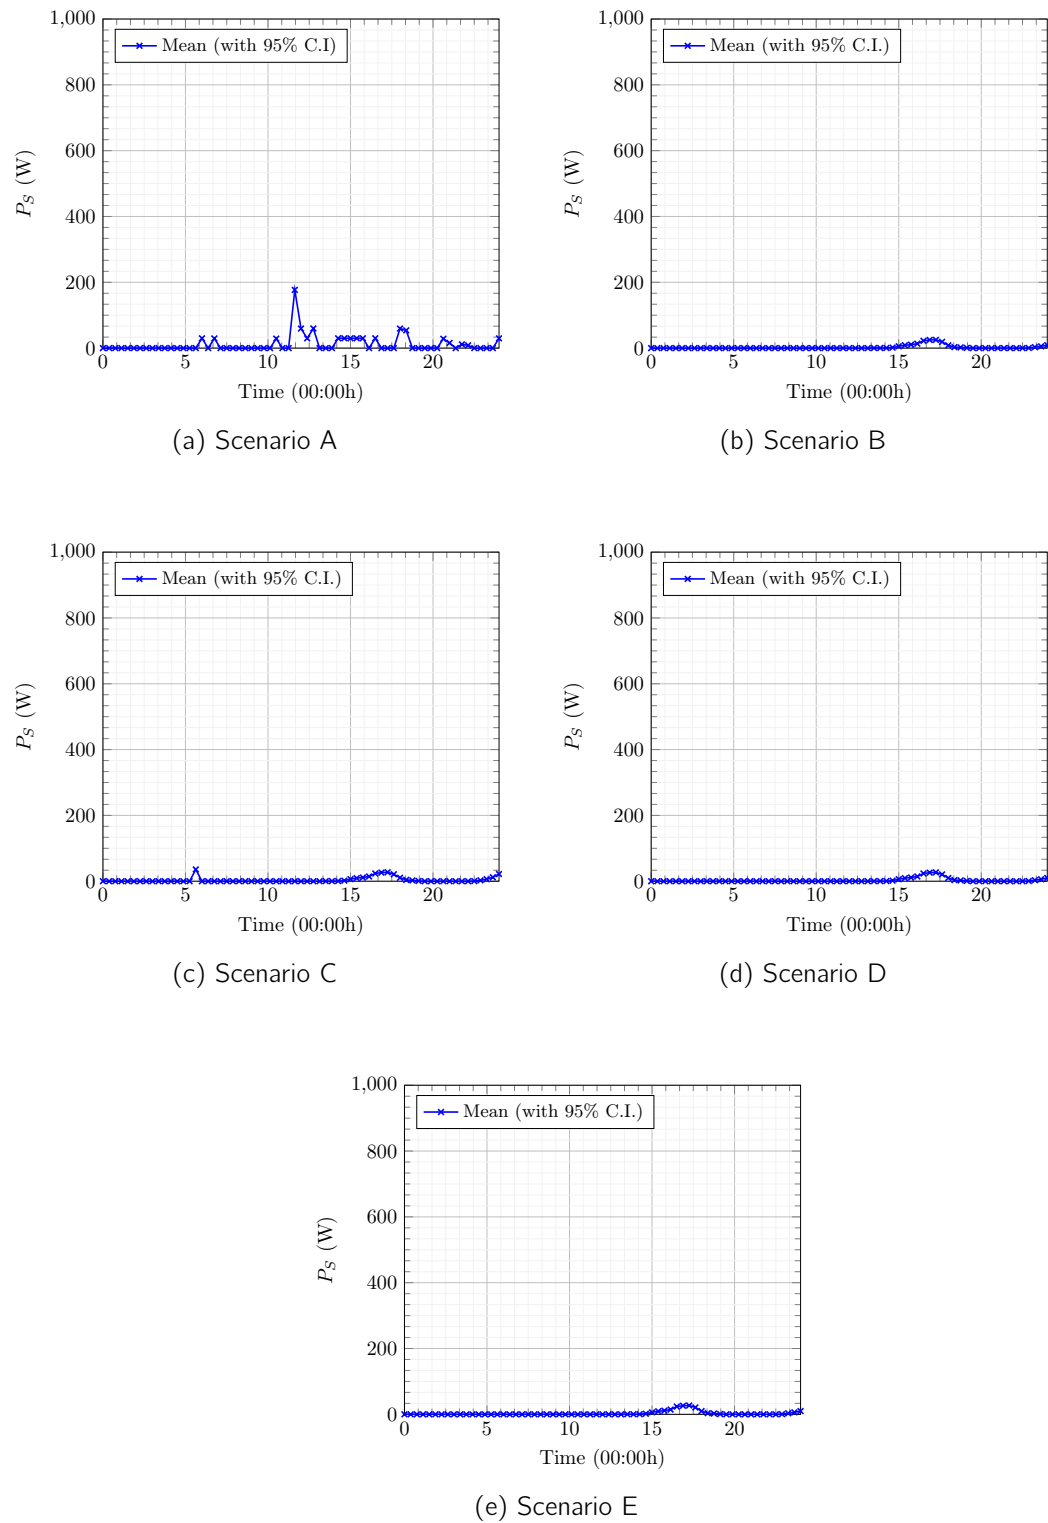





Figure D.7: Optimal power selling ($P_S(t)$) policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.




(a) Scenario A




(b) Scenario B



(c) Scenario C



(d) Scenario D



(e) Scenario E

Figure D.8: Optimal battery charge retention $A(t)$ policy, with 95% confidence intervals, using the optimal value function v from Algorithm 5 for different scenarios.

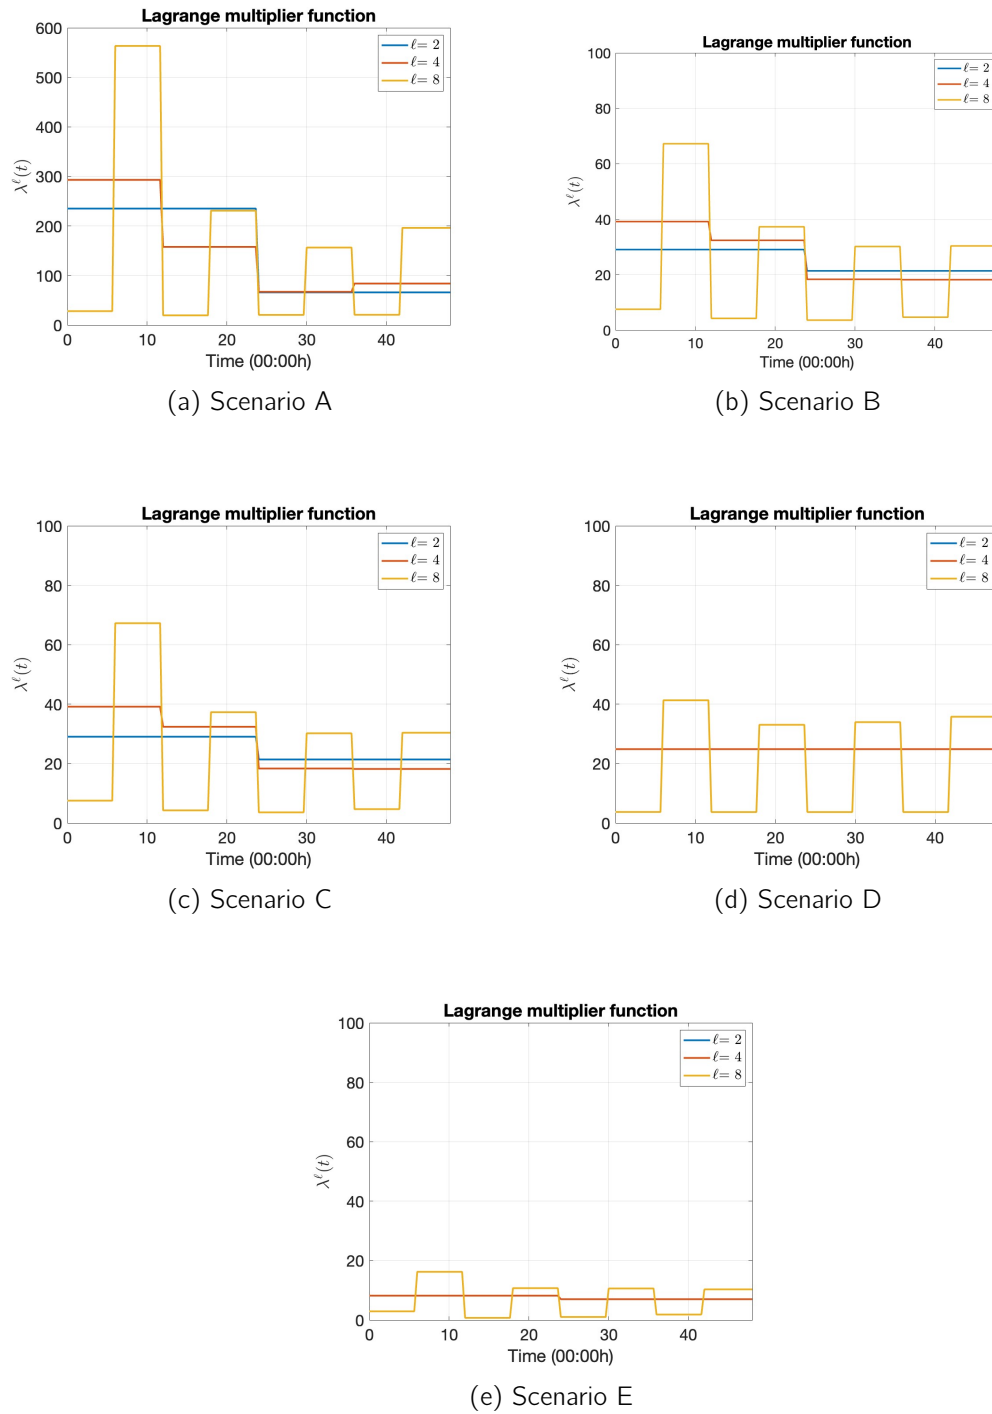


Figure D.9: Evolution of the optimal Lagrange multiplier function $\lambda^\ell(t)$ at each refinement level ℓ in Algorithm 5 for different scenarios.

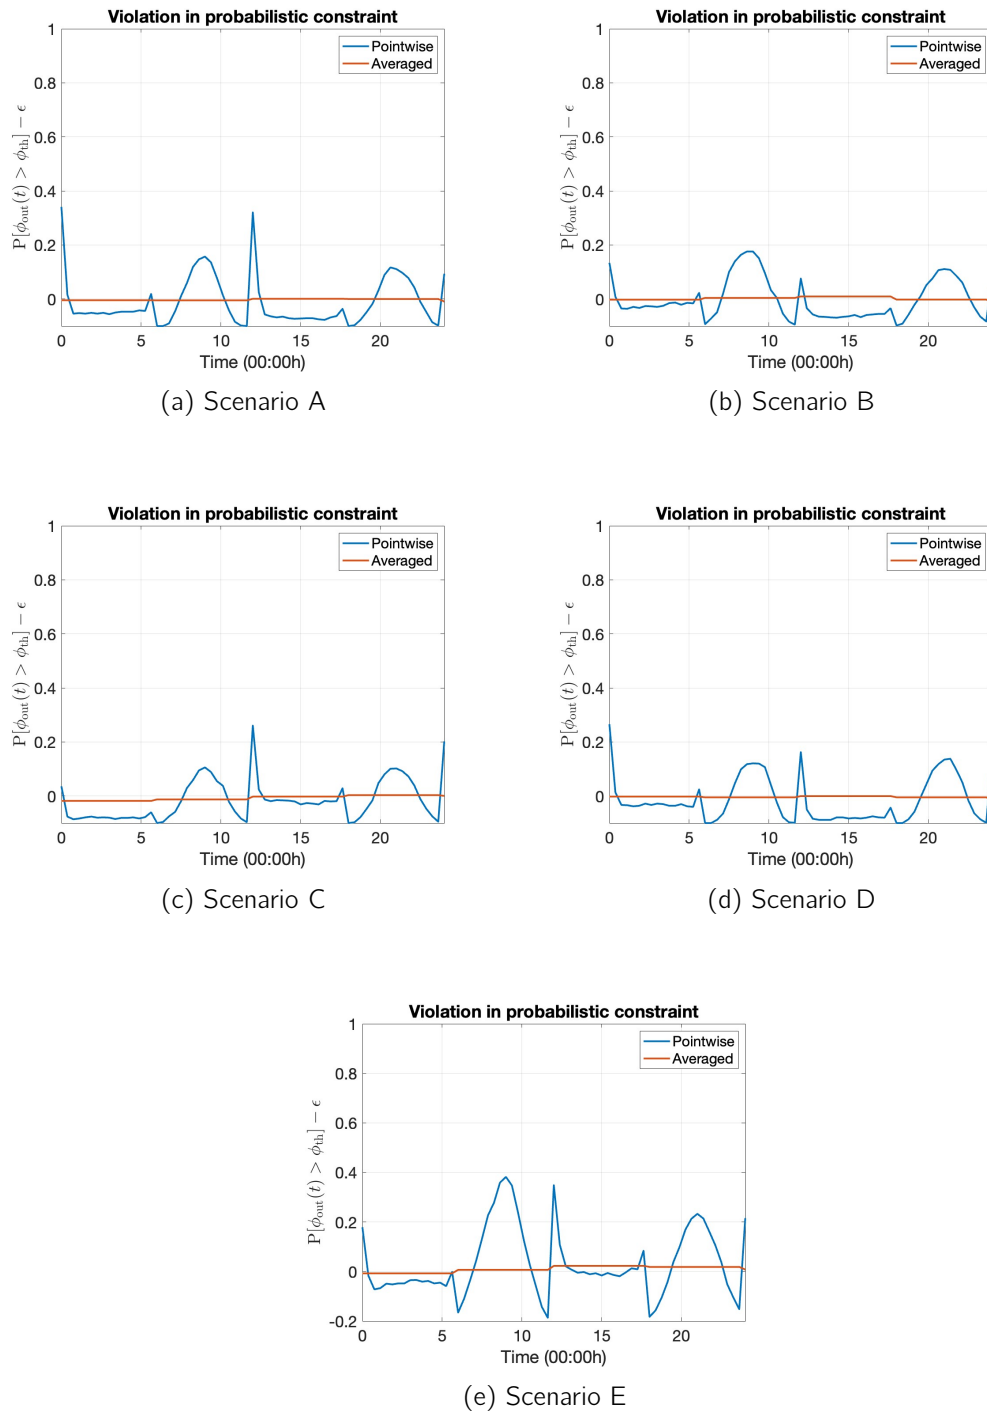


Figure D.10: MC estimates of the violation in constraint $\mathbb{P}[\phi_{\text{out}}(t) > \phi_{\text{th}}] - \epsilon$ for each $t \in [0, T]$ using optimal controls obtained from Algorithm 5 for different scenarios.

Bibliography

- [1] N. Ben Rached, A.-L. Haji-Ali, S. M. Subbiah Pillai, and R. Tempone, "Double-loop importance sampling for McKean–Vlasov stochastic differential equation," *Statistics and Computing*, vol. 34, no. 6, p. 197, 2024.
- [2] N. Ben Rached, A.-L. Haji-Ali, S. M. Subbiah Pillai, and R. Tempone, "Multilevel importance sampling for rare events associated with the McKean–Vlasov equation," *Statistics and Computing*, vol. 35, no. 1, p. 1, 2024.
- [3] N. Ben Rached, S. M. Subbiah Pillai, and R. Tempone, "Optimal power procurement for green cellular wireless networks under uncertainty and chance constraints," *Entropy*, vol. 27, no. 3, 2025.
- [4] D. Helbing, P. Molnár, I. J. Farkas, and K. Bolay, "Self-organizing pedestrian movement," *Environment and Planning B: Planning and Design*, vol. 28, pp. 361–383, 2025/04/08 2001.
- [5] N. B. Rached, H. Ghazzai, A. Kadri, and M. S. Alouini, "Energy management optimization for cellular networks under renewable energy generation uncertainty," *IEEE Transactions on Green Communications and Networking*, vol. 1, no. 2, pp. 158–166, 2017.
- [6] SMARD, "Market data download center," 2025. Accessed on 19 January 2025.
- [7] N. Karmitsa, "LMBM–FORTRAN subroutines for large-scale nonsmooth minimization: User's manual," *TUCS Techn. Rep*, vol. 77, p. 856, 2007.
- [8] H. J. Kushner, "Numerical methods for stochastic control problems in continuous time," *SIAM Journal on Control and Optimization*, vol. 28, pp. 999–1048, 2024/12/16 1990.
- [9] A. Agarwal, S. de Marco, E. Gobet, and G. Liu, "Rare event simulation related to financial risks: efficient estimation and sensitivity analysis," Dec. 2017. working paper or preprint.
- [10] C. Le Priol, J. M. Monteiro, and F. Bouchet, "Using rare event algorithms to understand the statistics and dynamics of extreme heatwave seasons in South Asia," *Environmental Research: Climate*, vol. 3, no. 4, p. 045016, 2024.
- [11] J. Morio and M. Balesdent, *Estimation of rare event probabilities in complex aerospace and other systems : A practical approach*. Woodhead Publishing in mechanical engineering; no. 720; number 720, Amsterdam: Woodhead Publishing is an imprint of Elsevier, 2016.

- [12] D. P. Kroese, T. Taimre, and Z. I. Botev, *Handbook of Monte Carlo methods*. John Wiley & Sons, 2013.
- [13] N. J. Newton, "Variance reduction for simulated diffusions," *SIAM journal on applied mathematics*, vol. 54, no. 6, pp. 1780–1805, 1994.
- [14] R. Funke and A. Y. Shevlyakov, "On a generalization of a formula of Clark," *Theory Random Processes*, vol. 7, pp. 93–96, 1977.
- [15] C. Hartmann, L. Richter, C. Schütte, and W. Zhang, "Variational characterization of free energy: Theory and algorithms," *Entropy*, vol. 19, no. 11, 2017.
- [16] J.-D. Deuschel and D. W. Stroock, *Large deviations*, vol. 342. American Mathematical Soc., 2001.
- [17] H. P. McKean, "A class of Markov processes associated with nonlinear parabolic equations," *Proceedings of the National Academy of Sciences*, vol. 56, pp. 1907–1911, 2025/03/20 1966.
- [18] A. L. Haji Ali, "Pedestrian flow in the mean field limit," 2012.
- [19] R. Erban and J. Haskovec, "From individual to collective behaviour of coupled velocity jump processes: a locust example," *arXiv preprint arXiv:1104.2584*, 2011.
- [20] N. Bush, B. M. Hambly, H. Haworth, L. Jin, and C. Reisinger, "Stochastic evolution equations in portfolio credit modelling," *SIAM Journal on Financial Mathematics*, vol. 2, pp. 627–664, 2025/03/20 2011.
- [21] G. dos Reis, G. Smith, and P. Tankov, "Importance sampling for McKean-Vlasov SDEs," *Applied Mathematics and Computation*, vol. 453, p. 128078, 2023.
- [22] Z. W. Bezemek and M. Heldman, "Importance sampling for the empirical measure of weakly interacting diffusions," *Applied Mathematics & Optimization*, vol. 89, no. 1, p. 7, 2023.
- [23] M. B. Giles, "Multilevel Monte Carlo path simulation," *Operations Research*, vol. 56, pp. 607–617, 2025/03/24 2008.
- [24] A.-L. Haji-Ali, F. Nobile, and R. Tempone, "Multi-index Monte Carlo: when sparsity meets sampling," *Numerische Mathematik*, vol. 132, no. 4, pp. 767–806, 2016.
- [25] A. Kebaier and J. Lelong, "Coupling importance sampling and multilevel Monte Carlo using sample average approximation," *Methodology and Computing in Applied Probability*, vol. 20, no. 2, pp. 611–641, 2018.
- [26] M. Ben Alaya, K. , Hajji, , and A. Kebaier, "Adaptive importance sampling for multilevel Monte Carlo Euler method," *Stochastics*, vol. 95, pp. 303–327, 02 2023.
- [27] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, and R. Spigler, "The Kuramoto model: A simple paradigm for synchronization phenomena," *Reviews of Modern Physics*, vol. 77, pp. 137–185, 04 2005.
- [28] A. Charnes, W. W. Cooper, and G. H. Symonds, "Cost horizons and certainty equivalents: An approach to stochastic programming of heating oil," *Management Science*, vol. 4, pp. 235–263, 2024/12/05 1958.

- [29] A. Charnes and W. W. Cooper, "Chance-constrained programming," *Management Science*, vol. 6, pp. 73–79, 2024/12/05 1959.
- [30] N. H. Agnew, R. A. Agnew, J. Rasmussen, and K. R. Smith, "An application of change constrained programming to portfolio selection in a casualty insurance firm," *Management Science*, vol. 15, pp. B512–B520, 2025/04/09/ 1969.
- [31] Y. Dvorkin, "A chance-constrained stochastic electricity market," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 2993–3003, 2020.
- [32] C. Zhao, J. Wang, R. Zhang, D. Niyato, G. Sun, H. Du, D. I. Kim, and A. Jamalipour, "Generative AI-enabled wireless communications for robust low-altitude economy networking," *arXiv preprint arXiv:2502.18118*, 2025.
- [33] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, pp. 969–996, 2024/12/05 2006.
- [34] D. Bertsekas, *Dynamic programming and optimal control: Volume I*, vol. 4. Athena scientific, 2012.
- [35] K. W. Ding, M. H. Wang, and N. J. Huang, "Distributionally robust chance constrained problem under interval distribution information," *Optimization Letters*, vol. 12, no. 6, pp. 1315–1328, 2018.
- [36] X. Chang, Y. Xu, W. Gu, H. Sun, M. Y. Chow, and Z. Yi, "Accelerated distributed hybrid stochastic/robust energy management of smart grids," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 8, pp. 5335–5347, 2021.
- [37] A. Zhou, M. Yang, T. Wu, and L. Yang, "Distributionally robust energy management for islanded microgrids with variable moment information: An MISOCP approach," *IEEE Transactions on Smart Grid*, vol. 14, no. 5, pp. 3668–3680, 2023.
- [38] P. Du, H. Lei, I. S. Ansari, J. Du, and X. Chu, "Distributionally robust optimization based chance-constrained energy management for hybrid energy powered cellular networks," *Digital Communications and Networks*, vol. 9, no. 3, pp. 797–808, 2023.
- [39] X. Liu, H. Li, and H. Wang, "Probability constrained robust multicast beamforming in cognitive radio network," in *2013 8th International Conference on Communications and Networking in China (CHINACOM)*, pp. 708–712, 2013.
- [40] Z. Liu, Z. Liu, Y. Xie, K. Y. Chan, Y. Yuan, and Y. Yang, "Power allocation in D2D enabled cellular network with probability constraints: A robust Stackelberg game approach," *Ad Hoc Networks*, vol. 133, p. 102891, 2022.
- [41] L. Andrieu, G. Cohen, and F. J. Vázquez-Abad, "Gradient-based simulation optimization under probability constraints," *European Journal of Operational Research*, vol. 212, no. 2, pp. 345–351, 2011.
- [42] L. Pfeiffer, X. Tan, and Y.-L. Zhou, "Duality and approximation of stochastic optimal control problems under expectation constraints," *SIAM Journal on Control and Optimization*, vol. 59, pp. 3231–3260, 2024/12/16 2021.
- [43] C. Ben Hammouda, E. Rezvanova, E. von Schwerin, and R. Tempone, "Lagrangian relaxation for continuous-time optimal control of coupled hydrothermal power systems including storage capacity and a cascade of hydropower systems with time delays,"

- Optimal Control Applications and Methods*, vol. 45, pp. 2279–2311, 2024/12/02 2024.
- [44] J. Jacod and P. Protter, *Probability essentials*. Springer Science & Business Media, 2012.
- [45] R. Durrett, *Probability: theory and examples*, vol. 49. Cambridge university press, 2019.
- [46] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Springer, 1992.
- [47] B. Oksendal, *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [48] D. Talay and L. Tubaro, “Expansion of the global error for numerical schemes solving stochastic differential equations,” *Stochastic Analysis and Applications*, vol. 8, pp. 483–509, 01 1990.
- [49] M. Lefebvre, *Applied probability and statistics*. Springer Science & Business Media, 2007.
- [50] N. Touzi, “Stochastic control and application to finance,” *Scuola Normale Superiore, Pisa. Special Research Semester on Financial Mathematics*, 2002.
- [51] K. Ross, “Stochastic control in continuous time,” *Lecture Notes on Continuous Time Stochastic Control*, pp. P33–P37, 2008.
- [52] H. Pham, *Continuous-time stochastic control and optimization with financial applications*, vol. 61. Springer Science & Business Media, 2009.
- [53] R. V. Gamkrelidze, “Discovery of the maximum principle,” *Journal of Dynamical and Control Systems*, vol. 5, no. 4, pp. 437–451, 1999.
- [54] W. H. Fleming and H. M. Soner, *Controlled Markov processes and viscosity solutions*, vol. 25. Springer Science & Business Media, 2006.
- [55] C. Bayer, H. Hoel, E. von Schwerin, and R. Tempone, “On nonasymptotic optimal stopping criteria in Monte Carlo simulations,” *SIAM Journal on Scientific Computing*, vol. 36, pp. A869–A885, 2025/06/06 2014.
- [56] C. Ben Hammouda, N. Ben Rached, R. Tempone, and S. Wiechert, “Learning-based importance sampling via stochastic optimal control for stochastic reaction networks,” *Statistics and Computing*, vol. 33, no. 3, p. 58, 2023.
- [57] E. Ben Amar, N. Ben Rached, A.-L. Haji-Ali, and R. Tempone, “State-dependent importance sampling for estimating expectations of functionals of sums of independent random variables,” *Statistics and Computing*, vol. 33, no. 2, p. 40, 2023.
- [58] A. Melnikov, *Discrete time stochastic analysis: further results and applications*, pp. 65–80. Cham: Springer Nature Switzerland, 2023.
- [59] N. Nüsken and L. Richter, “Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space,” *Partial Differential Equations and Applications*, vol. 2, no. 4, p. 48, 2021.

- [60] C. Hartmann, O. Kebiri, L. Neureither, and L. Richter, "Variational approach to rare event simulation using least-squares regression," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, p. 063107, 2/20/2025 2019.
- [61] C. Hartmann, C. Schütte, and W. Zhang, "Model reduction algorithms for optimal control and importance sampling of diffusions," *Nonlinearity*, vol. 29, no. 8, p. 2298, 2016.
- [62] 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 Journal on Scientific Computing*, vol. 36, pp. A2654–A2672, 2025/02/20 2014.
- [63] H. P. Awad, P. W. Glynn, and R. Y. Rubinstein, "Zero-variance importance sampling estimators for Markov process expectations," *Mathematics of Operations Research*, vol. 38, pp. 358–388, 2025/03/21 2012.
- [64] P. Hinds and M. V. Tretyakov, "Neural variance reduction for stochastic differential equations," *arXiv preprint arXiv:2209.12885*, 2022.
- [65] A.-S. Sznitman, "Topics in propagation of chaos," *Ecole d'étéde probabilités de Saint-Flour XIX—1989*, vol. 1464, pp. 165–251, 1991.
- [66] D. M. Forrester, "Arrays of coupled chemical oscillators," *Scientific Reports*, vol. 5, no. 1, p. 16994, 2015.
- [67] U. Dobramysl, S. Rüdiger, and R. Erban, "Particle-based multiscale modeling of calcium puff dynamics," *Multiscale Modeling & Simulation*, vol. 14, pp. 997–1016, 2025/03/20 2016.
- [68] D. Crisan, , and J. Xiong, "Approximate McKean–Vlasov representations for a class of SPDEs," *Stochastics*, vol. 82, pp. 53–68, 02 2010.
- [69] Y. Mishura and A. Veretennikov, "Existence and uniqueness theorems for solutions of McKean–Vlasov stochastic equations," *Theory of Probability and Mathematical Statistics*, vol. 103, pp. 59–101, 2020.
- [70] W. R. Hammersley, D. Šiška, and Ł. Szpruch, "Weak existence and uniqueness for McKean–Vlasov SDEs with common noise," 2021.
- [71] C. Kumar, Neelima, C. Reisinger, and W. Stockinger, "Well-posedness and tamed schemes for McKean–Vlasov equations with common noise," *The Annals of Applied Probability*, vol. 32, pp. 3283–3330, 10 2022.
- [72] W. Hong, S. Hu, and W. Liu, "McKean–Vlasov SDE and SPDE with locally monotone coefficients," *The Annals of Applied Probability*, vol. 34, pp. 2136–2189, 4 2024.
- [73] C. Cuchiero and J. Möller, "Polynomial McKean–Vlasov SDEs," *arXiv preprint arXiv:2502.19203*, 2025.
- [74] L.-P. Chaintron and A. Diez, "Propagation of chaos: a review of models, methods and applications. II. Applications," *arXiv preprint arXiv:2106.14812*, 2021.
- [75] R. Buckdahn, J. Li, S. Peng, and C. Rainer, "Mean-field stochastic differential equations and associated PDEs," 2017.

- [76] D. Crisan and E. McMurray, "Smoothing properties of McKean–Vlasov SDEs," *Probability Theory and Related Fields*, vol. 171, no. 1, pp. 97–148, 2018.
- [77] M. Bossy and D. Talay, "Convergence rate for the approximation of the limit law of weakly interacting particles: application to the Burgers equation," *The Annals of Applied Probability*, vol. 6, no. 3, pp. 818–861, 1996.
- [78] C. Graham, T. G. Kurtz, S. Méléard, P. E. Protter, M. Pulvirenti, and D. Talay, "Asymptotic behaviour of some interacting particle systems; McKean–Vlasov and Boltzmann models," *Probabilistic Models for Nonlinear Partial Differential Equations: Lectures given at the 1st Session of the Centro Internazionale Matematico Estivo (CIME) held in Montecatini Terme, Italy, May 22–30, 1995*, pp. 42–95, 1996.
- [79] M. Bossy and D. Talay, "A stochastic particle method for the McKean–Vlasov and the Burgers equation," *Mathematics of computation*, vol. 66, no. 217, pp. 157–192, 1997.
- [80] D. Cumin and C. P. Unsworth, "Generalising the Kuramoto model for the study of neuronal synchronisation in the brain," *Physica D: Nonlinear Phenomena*, vol. 226, no. 2, pp. 181–196, 2007.
- [81] S. Ogawa, "Monte Carlo simulation of nonlinear diffusion processes," *Japan Journal of Industrial and Applied Mathematics*, vol. 9, no. 1, pp. 25–33, 1992.
- [82] A.-L. Haji-Ali and R. Tempone, "Multilevel and multi-index Monte Carlo methods for the McKean–Vlasov equation," *Statistics and Computing*, vol. 28, no. 4, pp. 923–935, 2018.
- [83] Y. Li, X. Mao, Q. Song, F. Wu, and G. Yin, "Strong convergence of Euler–Maruyama schemes for McKean–Vlasov stochastic differential equations under local Lipschitz conditions of state variables," *IMA Journal of Numerical Analysis*, vol. 43, pp. 1001–1035, 3/21/2025 2023.
- [84] V. N. Kolokoltsov, , and M. Troeva, "On mean field games with common noise and McKean–Vlasov SPDEs," *Stochastic Analysis and Applications*, vol. 37, pp. 522–549, 07 2019.
- [85] A.-L. Haji-Ali, H. Hoel, and R. Tempone, "Weak convergence analysis in the particle limit of the McKean–Vlasov equations using stochastic flows of particle systems," *arXiv preprint arXiv:2101.00886*, 2021.
- [86] V. Lemaire and G. Pagès, "Multilevel Richardson–Romberg extrapolation," 2017.
- [87] M. B. Giles, "Multilevel Monte Carlo methods," *Acta Numerica*, vol. 24, pp. 259–328, 2015.
- [88] M. S. Rosin, L. F. Ricketson, A. M. Dimits, R. E. Caflisch, and B. I. Cohen, "Multilevel Monte Carlo simulation of Coulomb collisions," *Journal of Computational Physics*, vol. 274, pp. 140–157, 2014.
- [89] Ł. Szpruch and A. Tse, "Antithetic multilevel particle system sampling method for McKean–Vlasov SDEs," *arXiv preprint arXiv:1903.07063*, 2019.
- [90] W. Fang and M. B. Giles, "Multilevel Monte Carlo method for ergodic SDEs without contractivity," *Journal of Mathematical Analysis and Applications*, vol. 476, no. 1, pp. 149–176, 2019.

- [91] C. Ben Hammouda, N. Ben Rached, and R. Tempone, "Importance sampling for a robust and efficient multilevel Monte Carlo estimator for stochastic reaction networks," *Statistics and Computing*, vol. 30, no. 6, pp. 1665–1689, 2020.
- [92] R. B. Ash and C. A. Doléans-Dade, *Probability and measure theory*. Academic press, 2000.
- [93] N. Collier, A.-L. Haji-Ali, F. Nobile, E. von Schwerin, and R. Tempone, "A continuation multilevel Monte Carlo algorithm," *BIT Numerical Mathematics*, vol. 55, no. 2, pp. 399–432, 2015.
- [94] A.-L. Haji-Ali, F. Nobile, E. von Schwerin, and R. Tempone, "Optimization of mesh hierarchies in multilevel Monte Carlo samplers," *Stochastics and Partial Differential Equations Analysis and Computations*, vol. 4, no. 1, pp. 76–112, 2016.
- [95] M. B. Giles, D. J. Higham, and X. Mao, "Analysing multi-level Monte Carlo for options with non-globally Lipschitz payoff," *Finance and Stochastics*, vol. 13, no. 3, pp. 403–413, 2009.
- [96] I. Gyöngy, "Mimicking the one-dimensional marginal distributions of processes having an Ito differential," *Probability Theory and Related Fields*, vol. 71, no. 4, pp. 501–516, 1986.
- [97] A.-L. Haji-Ali, J. Spence, and A. L. Teckentrup, "Adaptive multilevel Monte Carlo for probabilities," *SIAM Journal on Numerical Analysis*, vol. 60, pp. 2125–2149, 2025/03/26 2022.
- [98] C. Bayer, C. Ben Hammouda, and R. Tempone, "Multilevel Monte Carlo with numerical smoothing for robust and efficient computation of probabilities and densities," *SIAM Journal on Scientific Computing*, vol. 46, pp. A1514–A1548, 2025/03/26 2024.
- [99] C.-H. Rhee and P. W. Glynn, "Unbiased estimation with square root convergence for SDE models," *Operations Research*, vol. 63, pp. 1026–1043, 2025/06/06 2015.
- [100] L. Pfeiffer, *Sensitivity analysis for optimal control problems. Stochastic optimal control with a probability constraint*. Theses, Ecole Polytechnique X, Nov 2013.
- [101] B. Bouchard, R. Elie, and C. Imbert, "Optimal control under stochastic target constraints," *SIAM Journal on Control and Optimization*, vol. 48, pp. 3501–3531, 2024/12/05 2010.
- [102] B. Bouchard, R. Elie, and N. Touzi, "Stochastic target problems with controlled loss," *SIAM Journal on Control and Optimization*, vol. 48, pp. 3123–3150, 2024/12/05 2009.
- [103] M. Haarala, K. Miettinen, and M. M. Mäkelä, "New limited memory bundle method for large-scale nonsmooth optimization," *Optimization Methods and Software*, vol. 19, no. 6, pp. 673–692, 2004.
- [104] S. Boyd and A. Mutapcic, "Stochastic subgradient methods," *Lecture Notes for EE364b, Stanford University*, vol. 97, 2008.
- [105] S. Ma and D. Sun, "Chance constrained robust beamforming in cognitive radio networks," *IEEE Communications Letters*, vol. 17, no. 1, pp. 67–70, 2013.

- [106] International Telecommunication Union, "Measuring digital development: Facts and figures 2024," tech. rep., International Telecommunication Union Development Sector, Geneva, Switzerland, 2024.
- [107] International Energy Agency, "Tracking clean energy progress 2023," tech. rep., International Energy Agency, Paris, France, 2023.
- [108] W. Strielkowski, M. Dvořák, P. Rovný, E. Tarkhanova, and N. Baburina, "5G wireless networks in the future renewable energy systems," *Frontiers in Energy Research*, vol. 9, p. 714803, 2021.
- [109] GSMA Intelligence, "Going green: benchmarking the energy efficiency of mobile networks," tech. rep., GSM Association, London, United Kingdom, 2023.
- [110] T. Han and N. Ansari, "Powering mobile networks with green energy," *IEEE Wireless Communications*, vol. 21, no. 1, pp. 90–96, 2014.
- [111] A. Saleem, X. Zhang, Y. Xu, U. A. Albalawi, and O. S. Younes, "A critical review on channel modeling: Implementations, challenges and applications," *Electronics*, vol. 12, no. 9, 2023.
- [112] M. Klöppel, A. Gabash, A. Geletu, and P. Li, "Chance constrained optimal power flow with non-Gaussian distributed uncertain wind power generation," in *2013 12th International Conference on Environment and Electrical Engineering*, pp. 265–270, 2013.
- [113] H. Wu, M. Shahidehpour, Z. Li, and W. Tian, "Chance-constrained day-ahead scheduling in stochastic power system operation," *IEEE Transactions on Power Systems*, vol. 29, no. 4, pp. 1583–1591, 2014.
- [114] Y. Huang, L. Wang, W. Guo, Q. Kang, and Q. Wu, "Chance constrained optimization in a home energy management system," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 252–260, 2018.
- [115] A. Azgin and M. Krunz, "Scheduling in wireless cellular networks under probabilistic channel information," in *Proceedings. 12th International Conference on Computer Communications and Networks (IEEE Cat. No.03EX712)*, pp. 89–94, 2003.
- [116] M. J. Farooq, H. Ghazzai, and A. Kadri, "A stochastic geometry-based demand response management framework for cellular networks powered by smart grid," in *2016 IEEE Wireless Communications and Networking Conference*, pp. 1–6, 2016.
- [117] G. Claßen, *Optimisation under Data Uncertainty in Wireless Communication Networks*. PhD thesis, RWTH Aachen University, Aachen, 2015.
- [118] U. Challita, Z. Dawy, G. Turkiyyah, and J. Naoum-Sawaya, "A chance constrained approach for LTE cellular network planning under uncertainty," *Computer Communications*, vol. 73, pp. 34–45, 2016.
- [119] D. An, Q. Yang, W. Yu, X. Yang, X. Fu, and W. Zhao, "Sto2Auc: A stochastic optimal bidding strategy for microgrids," *IEEE Internet of Things Journal*, vol. 4, no. 6, pp. 2260–2274, 2017.
- [120] A. Bhattacharya, J. P. Kharoufeh, and B. Zeng, "Managing energy storage in microgrids: A multistage stochastic programming approach," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 483–496, 2018.

- [121] M. NAKAGAMI, *The m-Distribution—A General Formula of Intensity Distribution of Rapid Fading*, pp. 3–36. Pergamon, 1960.
- [122] A. Abdi, K. Wills, H. A. Barger, M. . S. Alouini, and M. Kaveh, “Comparison of the level crossing rate and average fade duration of Rayleigh, Rice and Nakagami fading models with mobile channel data,” in *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference (Cat. No.00CH37152)*, vol. 4, pp. 1850–1857 vol.4, 2000.
- [123] K. Pearson and O. M. F. E. Henrici, “X. Contributions to the mathematical theory of evolution.—II. Skew variation in homogeneous material,” *Philosophical Transactions of the Royal Society of London. (A.)*, vol. 186, pp. 343–414, 2025/01/30 1997.
- [124] J. L. FORMAN and M. SØRENSEN, “The Pearson diffusions: A class of statistically tractable diffusion processes,” *Scandinavian Journal of Statistics*, vol. 35, pp. 438–465, 2025/03/27 2008.
- [125] R. Caballero, A. Kebaier, M. Scavino, and R. Tempone, “Quantifying uncertainty with a derivative tracking SDE model and application to wind power forecast data,” *Statistics and Computing*, vol. 31, no. 5, p. 64, 2021.
- [126] J. Brucker, W. G. Bessler, and R. Gasper, “Grey-box modelling of lithium-ion batteries using neural ordinary differential equations,” *Energy Informatics*, vol. 4, no. 3, p. 15, 2021.
- [127] S. Tamilselvi, S. Gunasundari, N. Karuppiah, A. Razak RK, S. Madhusudan, V. M. Nagarajan, T. Sathish, M. Z. M. Shamim, C. A. Saleel, and A. Afzal, “A review on battery modelling techniques,” *Sustainability*, vol. 13, no. 18, 2021.
- [128] J. S. Edge, S. O’Kane, R. Prosser, N. D. Kirkaldy, A. N. Patel, A. Hales, A. Ghosh, W. Ai, J. Chen, J. Yang, S. Li, M.-C. Pang, L. Bravo Diaz, A. Tomaszewska, M. W. Marzook, K. N. Radhakrishnan, H. Wang, Y. Patel, B. Wu, and G. J. Offer, “Lithium ion battery degradation: what you need to know,” *Physical Chemistry Chemical Physics*, vol. 23, no. 14, pp. 8200–8221, 2021.
- [129] M. J. Farooq, H. Ghazzai, and A. Kadri, “Optimized energy procurement for cellular networks powered by smart grid based on stochastic geometry,” in *2015 IEEE Globecom Workshops (GC Wkshps)*, pp. 1–6, 2015.
- [130] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge: Cambridge University Press, 2004.
- [131] B. Sun and B. Z. Guo, “Convergence of an upwind finite-difference scheme for Hamilton–Jacobi–Bellman equation in optimal control,” *IEEE Transactions on Automatic Control*, vol. 60, no. 11, pp. 3012–3017, 2015.
- [132] M. Stein, “Dual problem algorithms,” *Seminar for Optimization in Communications and Signal Processing, Sarntal*, 2009.
- [133] B. Grimmer, “Convergence rates for deterministic and stochastic subgradient methods without Lipschitz continuity,” *SIAM Journal on Optimization*, vol. 29, pp. 1350–1365, 2025/02/03 2019.

- [134] N. Karmitsa, A. Bagirov, and M. M. Mäkelä, "Comparing different nonsmooth minimization methods and software," *Optimization Methods and Software*, vol. 27, pp. 131–153, 02 2012.
- [135] M. Mäkelä, "Survey of bundle methods for nonsmooth optimization," *Optimization Methods and Software*, vol. 17, pp. 1–29, 01 2002.
- [136] M. A. Marsan and M. Meo, "Energy efficient management of two cellular access networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 4, pp. 69–73, 2010.
- [137] N. B. Rached, H. Ghazzai, A. Kadri, and M. S. Alouini, "A time-varied probabilistic ON/OFF switching algorithm for cellular networks," *IEEE Communications Letters*, vol. 22, no. 3, pp. 634–637, 2018.
- [138] S. B. (Destatis), "Official population statistics of Germany," 2025. Accessed on 19 January 2025.
- [139] Statista, "Number of 5G base stations in Germany (2023)," 2023. Accessed on 19 January 2025.
- [140] S. Nguyen and R. Akl, "Approximating user distributions in WCDMA networks using 2-D Gaussian," Citeseer, 2005.
- [141] V. Germany, "Vodafone Germany tariffs," 2025. Accessed on 19 January 2025.
- [142] 50Hertz, "Wind power production grid feed-in data," 2025. Accessed on 19 January 2025.
- [143] T. S. Rappaport, *Wireless communications. Principles and practice*, Upper Saddle River, NJ: Prentice Hall, 1996.
- [144] H. Pham, X. Warin, and M. Germain, "Neural networks-based backward scheme for fully nonlinear PDEs," *SN Partial Differential Equations and Applications*, vol. 2, no. 1, p. 16, 2021.
- [145] A. Capponi and C.-A. Lehalle, eds., *NeuralNetworks-Based Algorithms for Stochastic Control and PDEs in Finance*, pp. 426–452. Cambridge: Cambridge University Press, 2023.
- [146] S. Steck and K. Urban, "A reduced basis method for the Hamilton-Jacobi-Bellman equation with application to the European Union emission trading scheme," *arXiv preprint arXiv:1503.07560*, 2015.
- [147] K. O. Friedrichs, "The identity of weak and strong extensions of differential operators," *Transactions of the American Mathematical Society*, vol. 55, pp. 132–151, 2025/04/04/ 1944.
- [148] J. C. Strikwerda, "Finite difference methods for the Stokes and Navier–Stokes equations," *SIAM Journal on Scientific and Statistical Computing*, vol. 5, pp. 56–68, 2025/03/31 1984.
- [149] P. E. Souganidis, "Approximation schemes for viscosity solutions of Hamilton-Jacobi equations," *Journal of Differential Equations*, vol. 59, no. 1, pp. 1–43, 1985.

-
- [150] G. Barles and P. E. Souganidis, "Convergence of approximation schemes for fully nonlinear second order equations," *Asymptotic Analysis*, vol. 4, no. 3, pp. 271–283, 1991.
- [151] J. Marcum, "A statistical theory of target detection by pulsed radar," *IRE Transactions on Information Theory*, vol. 6, no. 2, pp. 59–267, 1960.