



Stochastic Particle Systems and Optimization

Branching Processes, Mean Field Games and Impulse Control

Dem Fachbereich IV der Universität Trier zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften (Dr. rer. nat.)
eingereichte Dissertation.

vorgelegt von

Daniel Hoffmann

Datum der Disputation:

28.10.2020

Berichterstatter:

Univ.-Prof. Dr. Frank Thomas Seifried (Universität Trier)

Univ.-Prof. Dr. Volker Schulz (Universität Trier)

Univ.-Prof. Dr. Sören Christensen (Christian-Albrechts-Universität zu Kiel)

Daniel Hoffmann
Fachbereich IV – Mathematik
Universität Trier
54286 Trier

Email: daniel.hoffmann.mathematik@gmail.com

Version: 2020-11-04, 2035

Für meine Eltern

Petra und Herbert Hoffmann

Abstract

This thesis addresses three different topics from the fields of mathematical finance, applied probability and stochastic optimal control. Correspondingly, it is subdivided into three independent main chapters each of which approaches a mathematical problem with a suitable notion of a stochastic particle system.

In Chapter 1, we extend the branching diffusion Monte Carlo method of Henry-Labordère et. al. [HOT⁺19] to the case of parabolic PDEs with mixed local-nonlocal analytic nonlinearities. We investigate branching diffusion representations of classical solutions, and we provide sufficient conditions under which the branching diffusion representation solves the PDE in the viscosity sense. Our theoretical setup directly leads to a Monte Carlo algorithm, whose applicability is showcased in two stylized high-dimensional examples. As our main application, we demonstrate how our methodology can be used to value financial positions with defaultable, systemically important counterparties.

In Chapter 2, we formulate and analyze a mathematical framework for continuous-time mean field games with finitely many states and common noise, including a rigorous probabilistic construction of the state process. The key insight is that we can circumvent the master equation and reduce the mean field equilibrium to a system of forward-backward systems of (random) ordinary differential equations by conditioning on common noise events. We state and prove a corresponding existence theorem, and we illustrate our results in three stylized application examples. In the absence of common noise, our setup reduces to that of Gomes, Mohr and Souza [GMS13] and Cecchin and Fischer [CF20].

In Chapter 3, we present a heuristic approach to tackle stochastic impulse control problems in discrete time. Based on the work of Bensoussan [Ben08] we reformulate the classical Bellman equation of stochastic optimal control in terms of a discrete-time QVI, and we prove a corresponding verification theorem. Taking the resulting optimal impulse control as a starting point, we devise a self-learning algorithm that estimates the continuation and intervention region of such a problem. Its key features are that it explores the state space of the underlying problem by itself and successively learns the behavior of the optimally controlled state process. For illustration, we apply our algorithm to a classical example problem, and we give an outlook on open questions to be addressed in future research.

Finally, several parts of the appendix complement the main part of this thesis: In Appendix A we provide some auxiliary technical results, and we give a brief outline of the theory of Carathéodory solutions of ODEs; in Appendix B we sketch an extension of the branching diffusion methodology of Section 1.1 to nonlocal Dirichlet problems; Appendix C shows the relevant pricing PDEs for Section 1.4.3 in a synopsis; Appendix D contains the proof of the existence result for a mean field equilibrium (see Theorem 2.3.6), relevant assumptions and auxiliary results; and Appendix E provides the master equation for the mean field games considered in Chapter 2.

Zusammenfassung

Die vorliegende Arbeit behandelt drei verschiedene Themen aus den Bereichen Finanzmathematik, angewandte Wahrscheinlichkeitstheorie und stochastische Optimalsteuerung. Sie ist folglich in drei unabhängige Kapitel unterteilt, von welchen jedes sich einem mathematischen Problem mit einem geeigneten stochastischen Partikelsystem nähert.

In Kapitel 1 erweitern wir die Branching Diffusion-Monte-Carlo-Methode von Henry-Labordère et. al. [HOT⁺19] auf den Fall parabolischer partieller Differentialgleichungen (PDEs) mit gemischt lokal-nichtlokalen analytischen Nichtlinearitäten. Dabei untersuchen wir die Darstellung klassischer Lösungen mittels Branching Diffusions und geben hinreichende Bedingungen dafür an, dass eine ebensolche Darstellung die vorliegende PDE im Viskositätssinne löst. Unsere theoretischen Überlegungen führen unmittelbar zu einem Monte-Carlo-Algorithmus, dessen Anwendbarkeit wir in zwei stilisierten hochdimensionalen Beispielen demonstrieren. Unsere Hauptanwendung zeigt auf, wie unser Verfahren zur Bewertung von Derivatepositionen mit ausfallgefährdeten systemrelevanten Kontrahenten verwendet werden kann.

In Kapitel 2 formulieren und analysieren wir einen mathematischen Rahmen für Mean Field Games in stetiger Zeit mit endlich vielen Zuständen und Common Noise – einschließlich einer rigorosen probabilistischen Konstruktion des Zustandsprozesses. Die wesentliche Erkenntnis ist, dass wir die Master Equation umgehen und, durch Bedingen auf Common Noise-Ereignisse, das Mean Field-Gleichgewicht auf ein Vorwärts-Rückwärtssystem von (zufälligen) gewöhnlichen Differentialgleichungen reduzieren können. Wir formulieren und beweisen ein zugehöriges Existenzresultat und illustrieren unsere Ergebnisse in drei stilisierten Anwendungsbeispielen. Ohne Common Noise reduziert sich unser Modell auf das von Gomes, Mohr und Souza [GMS13] und Cecchin und Fischer [CF20].

In Kapitel 3 stellen wir einen heuristischen Zugang zu stochastischen Impulskontrollproblemen in diskreter Zeit vor. Basierend auf der Arbeit von Bensoussan [Ben08] formulieren wir die klassische Bellman-Gleichung der stochastischen Optimalsteuerung in die Form einer quasi-variationellen Ungleichung in diskreter Zeit um und beweisen einen zugehörigen Verifikationssatz. Ausgehend von der resultierenden optimalen Impulskontrolle entwickeln wir sodann einen selbstlernenden Algorithmus, der die (Nicht-)Eingreifregion eines solchen Problems approximiert. Seine wesentlichen Eigenschaften sind, dass er den Zustandsraum des zugrundeliegenden Problems selbst erkundet und das Verhalten des optimal gesteuerten Zustandsprozesses sukzessive lernt. Zur Veranschaulichung wenden wir unseren Algorithmus auf ein klassisches Beispielproblem an und geben einen Ausblick auf offene, in Zukunft zu behandelnde Forschungsfragen.

Schließlich ergänzen die Teile des Anhangs den Hauptteil dieser Arbeit: In Anhang A behandeln wir einige technische Hilfsresultate und geben einen kurzen Abriss der Theorie der Carathéodory-Lösungen gewöhnlicher Differentialgleichungen; in Anhang B skizzieren wir eine Erweiterung der Branching Diffusion-Methodik aus Abschnitt 1.1 auf nichtlokale Dirichlet-Probleme; Anhang C stellt die relevanten Bewertungsgleichungen für Abschnitt 1.4.3 gegenüber; Anhang D beinhaltet den Beweis des Existenzresultats für ein Mean Field-Gleichgewicht (siehe Theorem 2.3.6), zugehörige Annahmen sowie Hilfsresultate; Anhang E behandelt die Master Equation für die Mean Field Games aus Kapitel 2.

Publications

The structure and style of this thesis have their origins in three independent research projects that I have been pursuing in recent years. In the course of these projects, two research papers have been written and thus, upon submission of this thesis, essential parts of two of its three main chapters and related parts of the appendix have already been made publicly available:

1. Essential parts of Chapter 1 have been submitted for publication in a scientific journal; the corresponding preprint [BHS20a] is available at <https://ssrn.com/abstract=3451280>.
2. Essential parts of Chapter 2 and Appendix D have been submitted for publication in a scientific journal; the corresponding preprint [BHS20b] is available at <https://ssrn.com/abstract=3458336>.

Each of Chapter 1 (together with Appendix B/C) and Chapter 2 (together with Appendix D/E) thus consists of a revised and partly extended version of the respective aforementioned original research article. ◇

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor Professor Dr. Frank Thomas Seifried for his constant support and his sustained interest in my work. His excellent guidance and our countless inspiring discussions of ideas have greatly contributed to the completion of this thesis. I am very grateful that he shared his passion for mathematics and his wealth of experience with me.

I wish to thank Professor Dr. Volker Schulz for the advice he gave me on numerical issues related to my work as well as for agreeing to both co-supervise and review this thesis.

Furthermore, I am very grateful to Professor Dr. Sören Christensen for accepting to act as an external referee for this thesis.

I owe my sincere thanks to Professor Dr. Christoph Belak for his helpful suggestions and many insightful discussions on various topics related to my work; moreover, I would like to thank him for our fruitful collaboration together with Professor Seifried that has led to two research papers.

In the last more than three years of my entire PhD studies I have been a member of the DFG Research Training Group 2126 on “Algorithmic Optimization” (ALOP). I gratefully acknowledge the provided educational and financial support. I have enjoyed the pleasant working atmosphere at the Department of Mathematics and hence wish to thank all its members as well as my colleagues – with special thanks to the secretaries Martina Shaw, Monika Thieme-Trapp and Doris Karpa.

Schließlich möchte ich meiner Familie von Herzen danken: Ein ganz besonderer Dank geht an meine Eltern, Petra und Herbert, die mir jederzeit und in allen Lebenslagen mit Rat und Tat zur Seite gestanden und immer an mich geglaubt haben. Ohne ihre unermüdliche Fürsorge wäre mein bisheriger Weg nicht möglich gewesen.

Ich danke meinem Onkel Horst für sein stets offenes Ohr und seinen konstruktiven persönlichen Rat. Vermutlich hätte ich mich ohne seine Ermutigung niemals näher mit Mathematik beschäftigt. Und ich danke Christine, die mich während meiner Zeit als Doktorand stets unterstützt und geduldig Verständnis für meine Arbeit aufgebracht hat.

Trier und Idenheim, im Juni 2020

Daniel Hoffmann

Someone's sitting in the shade today because someone planted a tree a long time ago.

WARREN BUFFETT

Contents

Abstract	A
Zusammenfassung	C
Publications	E
Acknowledgements	G
1 Branching Diffusions with Jumps and Nonlocal Nonlinear PDEs	1
1.1 Branching Diffusion Representations for Nonlocal PDEs	3
1.1.1 Preliminaries	3
1.1.2 Branching Mechanism	4
1.1.3 Branching Diffusion Dynamics	5
1.1.4 Branching Diffusion Representation	8
1.2 Viscosity Solutions and Branching Diffusion Representations	13
1.2.1 Viscosity Solutions of Nonlocal Nonlinear PDEs	14
1.2.2 Sufficient Conditions for Uniform Integrability of $\{\Psi^{t,x}\}$	17
1.3 Monte Carlo Simulation: A High-Dimensional Example	20
1.4 Valuation with Systemically Important Counterparties	23
1.4.1 Valuation with Systemic Risk	23
1.4.2 Branching Diffusion with Jumps Approach	26
1.4.3 Numerical Illustration	28
1.4.4 Another High-Dimensional Example	32
Conclusion and Outlook	35
2 Continuous-Time Finite-State Mean Field Games with Common Noise	37
2.1 Mean Field Model	39
2.1.1 Probabilistic Setting and Common Noise	40
2.1.2 Optimization Problem	41
2.1.3 State Dynamics	42

2.2	Solution of the Optimization Problem	46
2.3	Equilibrium	50
2.3.1	Aggregation	51
2.3.2	Mean Field Equilibrium System	56
2.4	Applications	57
2.4.1	A Decentralized Agricultural Production Model	59
2.4.2	An SIR Model with Random One-Shot Vaccination	65
2.4.3	Evacuation of a Room with Randomly Opening Doors	73
	Conclusion and Outlook	84
3	A Heuristic Approach to Discrete-Time Impulse Control	85
3.1	Theoretical Setup	87
3.1.1	State Space and Dynamics	87
3.1.2	Stochastic Control Problem and Value Function	88
3.2	Reformulation of the Dynamic Programming Equation	91
3.2.1	The Bellman Principle and a Classical Verification Theorem	91
3.2.2	A Verification QVI in Discrete Time	92
3.2.3	Excursion: Existence of a Measurable Maximizer	98
3.3	A Self-Learning Monte Carlo Algorithm	100
3.3.1	Theoretical Aspects	100
3.3.2	Implementation: Smart Voronoi Tessellations	103
3.4	Numerical Results	109
3.4.1	Impulse Control of Brownian Motion towards Zero	109
3.4.1.1	A One-Dimensional Example without Proportional Costs	110
3.4.1.2	A One-Dimensional Example with Fixed and Proportional Costs	112
3.4.1.3	Some Examples in Higher Dimensions	113
3.4.2	A Problem with Disconnected Continuation Region	117
3.4.3	A Short Remark on the Numerical Results	118
	Conclusion and Outlook	118
A	Some Auxiliary Results	123
A.1	Modifications and Integrals	123
A.2	Counting Processes	124
A.3	Carathéodory Solutions of ODEs	126
B	Branching Diffusions and Nonlocal Dirichlet Problems	129
C	Counterparties and Pricing PDEs	133

D Existence of Mean Field Equilibria	135
E The Master Equation	147
List of Figures	I
List of Tables	III
References	V

Chapter 1

Branching Diffusions with Jumps and Nonlocal Nonlinear PDEs

The present chapter and Appendix B/C are a revised and partly extended version of the article [BHS20a] that has already been made publicly available as a preprint and submitted for publication in a scientific journal. ◇

The objective of this chapter is to derive probabilistic representations of solutions of a certain class of nonlinear parabolic partial differential equations with nonlocal terms in the nonlinearity. The representation is based on a branching diffusion mechanism with jumps at branching times and makes it possible to compute solutions by direct (non-nested) Monte Carlo simulation, leading to a numerical algorithm that does not suffer from the curse of dimensionality.

The class of partial differential equations under consideration takes the form

$$\partial_t u(t, x) + \mathcal{A}[u](t, x) + \int_{\Xi} f(t, x, \xi, \mathcal{J}[u](t, x, \xi)) \gamma(d\xi) = 0, \quad (\text{PDE})$$

where \mathcal{A} denotes the infinitesimal generator of an Itô diffusion, i.e. a (possibly degenerate) linear partial differential operator of second order; \mathcal{J} is a nonlocal operator; and the nonlinearity f is analytic in the jump terms.

In recent years, there has been significant progress in the realm of probabilistic representations of partial differential equations with analytic nonlinearities acting on zeroth- and first-order derivatives. We refer in particular to [RRM10], [Hen12a] and [HTT14] for the zeroth-order case and [HOT⁺19] for the first-order case. We also refer to [BTWZ17] for an extension of the branching diffusion approach to the case of locally analytic nonlinearities, [BTW19] for the case of Lipschitz nonlinearities, [HT18] and [War18] for higher-order partial differential equations, and [BCL⁺15] and [AC18] for an extension to elliptic equations. Furthermore, beyond these recent articles, connections between branching diffusion processes and partial differential equations have already been established, e.g., in [Sko64], [Wat65] and [McK75].

Our main contribution is to extend the branching diffusion approach to the case of nonlocal terms inside the nonlinearity. This extension is achieved by introducing jump marks in the branching

diffusion underlying the probabilistic representation result: We consider a branching diffusion similar to the one introduced in Henry-Labordère et. al. [HOT⁺19] with the additional feature that, at each branching time, a subset of offspring particles may jump away from their parent's position. We refer to the resulting object as a *branching diffusion with jumps*.¹

Our main motivation behind the derivation of a probabilistic representation is to open up the possibility to apply numerical algorithms for nonlocal nonlinear PDEs in high dimensions via Monte Carlo simulation. The effectiveness and efficiency of such algorithms is showcased in an example of a nonlocal nonlinear PDE, which we solve in dimensions up to 100. In addition, we show how our Monte Carlo methodology can be used in the pricing of (equivalently, the computation of credit valuation adjustments for) financial positions where the counterparty is a systemically important financial institution whose default causes a devaluation of the underlying. This jump-at-default model represents a particularly realistic setup for wrong-way risk; see, e.g., [PS13], [ML15a], [BP18] or [BPP18]. Our Monte Carlo approach complements existing methods by making it possible to price systemic defaultable positions for settings where the underlying dynamics are not tractable by PDE methods; the latter are considered in or developed by, e.g., [DH96], [HL99], or [KL16]; for probabilistic representation results we refer, e.g., to [CG15].

The remainder of the chapter is organized as follows: Section 1.1 provides the stochastic construction of the branching diffusion with jumps and derives a probabilistic representation of classical solutions of (PDE) in terms of it. In Section 1.2 we conversely establish sufficient conditions for the branching diffusion representation to yield a viscosity solution of (PDE). Section 1.3 illustrates how the branching diffusion representation allows for efficient simulation of solutions via the Monte Carlo method in a stylized high-dimensional example. Finally, in Section 1.4 we showcase our methodology in the context of pricing with a systemically important, defaultable counterparty.

Three parts of the appendix complement the exposition of this chapter: Appendix A provides some auxiliary technical results; in Appendix B we sketch how the representation result in Section 1.1 (see Theorem 1.1.6) can be transferred to the case of certain nonlocal Dirichlet problems; and Appendix C provides the relevant pricing PDEs for the considered scenarios in Section 1.4.3.

¹This is not to be confused with the straightforward notion of branching jump-diffusions, where jump terms appear in the infinitesimal generator \mathcal{A} ; see also the discussion below Theorem 1.1.6.

1.1 Branching Diffusion Representations for Nonlocal PDEs

Throughout this chapter, we fix a time horizon $T > 0$, a non-empty set $\mathcal{I} \subseteq \mathbb{N}_0^m$ of multi-indices, and a jump distribution γ on an abstract measurable space (Ξ, \mathfrak{B}) . The goal of this section is to provide a stochastic representation of a classical solution $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ of a nonlocal partial differential equation of the form

$$\begin{aligned} \partial_t u(t, x) + \mathcal{A}[u](t, x) + \int_{\Xi} f(t, x, \xi, \mathcal{J}[u](t, x, \xi)) \gamma(d\xi) &= 0, & (t, x) \in [0, T] \times \mathbb{R}^d, & \text{(PDE)} \\ u(T, x) &= g(x), & x \in \mathbb{R}^d, & \text{(TC)} \end{aligned}$$

where

$$\mathcal{A}[u](t, x) \triangleq \mu(t, x)^\top \nabla_x u(t, x) + \frac{1}{2} \operatorname{tr} \left[\sigma(t, x) \sigma(t, x)^\top \nabla_x^2 u(t, x) \right]$$

is the infinitesimal generator of a diffusion process; the nonlinearity

$$f : [0, T] \times \mathbb{R}^d \times \Xi \times \mathbb{R}^m \rightarrow \mathbb{R}, \quad f(t, x, \xi, y) \triangleq \sum_{i \in \mathcal{I}} c_i(t, x, \xi) y^i$$

is (multivariate) analytic² in y with measurable coefficients $c_i : [0, T] \times \mathbb{R}^d \times \Xi \rightarrow \mathbb{R}$ for $i \in \mathcal{I}$; and the jump operator \mathcal{J} is given by

$$\mathcal{J}_\ell[u](t, x, \xi) \triangleq u(t, \Gamma_\ell(t, x, \xi)) \quad \text{for } \ell \in [1:m],$$

where the jump maps $\Gamma_\ell : [0, T] \times \mathbb{R}^d \times \Xi \rightarrow \mathbb{R}^d$, $\ell \in [1:m]$, are measurable. Finally, the function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is assumed to be measurable as well.

1.1.1 Preliminaries

For $n \in \mathbb{N}$, we denote by $\mathbf{N}_n \triangleq \bigcup_{\nu=1}^n \mathbb{N}^\nu$ the set of all \mathbb{N} -words with length at most n and by $\mathbf{N} \triangleq \bigcup_{n \in \mathbb{N}} \mathbf{N}_n$ the set of finite \mathbb{N} -words. We work on a probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ that supports the following random variables, all of which are taken to be mutually independent:

- A family $\{W^{(k)}\}_{k \in \mathbf{N}}$ of independent \mathbb{R}^d -valued Brownian motions that serve as driving noise for the emerging diffusion processes with infinitesimal generator \mathcal{A} .
- A family $\{\Delta^{(k)}\}_{k \in \mathbf{N}}$ of i.i.d. \mathbb{R}^d -valued random variables with distribution γ .
- A family $\{\tau^{(k)}\}_{k \in \mathbf{N}}$ of i.i.d. $(0, \infty)$ -valued random variables serving as lifetimes of the particles underlying the branching mechanism. We assume³ that the distribution of $\tau^{(k)}$, $k \in \mathbf{N}$, admits a continuous density $\rho : (0, \infty) \rightarrow [0, \infty)$ that is strictly positive on $(0, T]$ and such that

$$F(T) \triangleq \int_T^\infty \rho(s) ds > 0. \tag{1.1}$$

²We use standard multi-index notation and write $y^i \triangleq \prod_{\ell=1}^m y_\ell^{i_\ell}$ for $y \in \mathbb{R}^m$ and $|i| \triangleq \sum_{\ell=1}^m i_\ell$, $i \in \mathbb{N}_0^m$. The case of a multivariate polynomial obtains if \mathcal{I} is finite.

³The stated conditions are standard given the literature; see, e.g., Assumption 3.1(i) in [HOT⁺19].

- A family $\{I^{(k)}\}_{k \in \mathbf{N}}$ of i.i.d. \mathcal{I} -valued random variables modeling the number of offspring of each particle as well as marks for their initial positions. We assume³ that

$$p_i \triangleq \mathbb{P}(I^{(k)} = i) > 0 \quad \text{for all } i \in \mathcal{I}, k \in \mathbf{N}, \quad \text{and} \quad \sum_{i \in \mathcal{I}} |i| p_i < \infty.$$

In the following, we describe the branching mechanism and the spatial dynamics separately to finally obtain the branching diffusion representation of (PDE).

1.1.2 Branching Mechanism

We fix an initial time $t \in [0, T]$. The branching mechanism is defined by recursion on successive generations, for each $\omega \in \Omega$.⁴

Given a particle of generation $n \in \mathbb{N}$ labeled $k = (k_1, \dots, k_n) \in \mathbb{N}^n$, we denote its parent particle by $k- \triangleq (k_1, \dots, k_{n-1}) \in \mathbb{N}^{n-1}$. The branching time of particle k is given by $T_t^{(k)} \triangleq (T_t^{(k-)} + \tau^{(k)}) \wedge T$; on the event $\{T_t^{(k)} < T\}$ the particle k is removed at time $T_t^{(k)}$ and branches into $|I^{(k)}|$ descendants, which are labeled by $(k_1, \dots, k_n, k_{n+1}) \in \mathbb{N}^{n+1}$ for $k_{n+1} \in [1 : |I^{(k)}|]$. We attach the jump mark $J^{(k, k_{n+1})} \triangleq 1$ to the first $I_1^{(k)}$ of these offspring particles, i.e. if $k_{n+1} \in [1 : I_1^{(k)}]$, the mark $J^{(k, k_{n+1})} \triangleq 2$ to the following $I_2^{(k)}$ offspring particles, etc., so each offspring particle (k, k_{n+1}) carries a mark $J^{(k, k_{n+1})} \in [1 : m]$. This iteration is well-defined and uniquely determines the branching dynamics if we assume that the mechanism starts with a single particle with label (1) of generation 1 at time t and

$$(1)- \triangleq () \triangleq \emptyset \quad \text{and} \quad T_t^\emptyset \triangleq t.⁵$$

In the following, we only refer to $k \in \mathbf{N}$ as a *particle* if either $k = (1)$ or if $k-$ is a particle and $k_n \in [1 : |I^{(k-)}|]$. Figure 1.1 below visualizes this branching mechanism.

We denote the random set of all particles of generation $n \in \mathbb{N}$ alive at time $s \in [t, T]$ by

$$\mathcal{K}_t^n(s) \triangleq \begin{cases} \{k \in \mathbb{N}^n : k \text{ is a particle and } T_t^{(k-)} \leq s < T_t^{(k)}\} & \text{if } s \in [t, T), \\ \{k \in \mathbb{N}^n : k \text{ is a particle and } T_t^{(k)} = T\} & \text{if } s = T. \end{cases}$$

The set of all particles of generation n alive before or at time s is given by

$$\bar{\mathcal{K}}_t^n(s) \triangleq \bigcup_{r \in [t, s]} \mathcal{K}_t^n(r).$$

Finally, the set of all particles alive at time s and the set of all particles alive before or at time s are defined as

$$\mathcal{K}_t(s) \triangleq \bigcup_{n \in \mathbb{N}} \mathcal{K}_t^n(s) \quad \text{and} \quad \bar{\mathcal{K}}_t(s) \triangleq \bigcup_{n \in \mathbb{N}} \bar{\mathcal{K}}_t^n(s), \quad \text{respectively.}$$

⁴Note that the following construction of the branching mechanism is due to §2.2 in [HOT⁺19].

⁵For the sake of completeness, one could formally also define $I^\emptyset \triangleq e_1$ to indicate that the *non-particle* \emptyset has exactly one descendant, to wit, the initial particle (1); here, e_1 denotes the first unit vector in \mathbb{R}^m . However, note that as there is no need for particle (1) to carry a jump mark (see below), setting it to 1 is an arbitrary choice.




Figure 1.1: Illustration of the branching mechanism (without jump marks).

For ease of notation, we subsequently write

$$\mathcal{K}_t^n \triangleq \mathcal{K}_t^n(T), \quad \mathcal{K}_t \triangleq \mathcal{K}_t(T), \quad \bar{\mathcal{K}}_t^n \triangleq \bar{\mathcal{K}}_t^n(T), \quad \bar{\mathcal{K}}_t \triangleq \bar{\mathcal{K}}_t(T). \quad (1.2)$$

As in Proposition 2.4 in [HOT⁺19], the total number of particles is almost surely finite, i.e.

$$\#\bar{\mathcal{K}}_t < \infty \quad \mathbb{P}\text{-a.s. for all } t \in [0, T]; \quad (1.3)$$

see also Theorem IV.1.1 in [AN72] and Chapter VI §§12f. in [Har63].

1.1.3 Branching Diffusion Dynamics

The next step is to specify the dynamics of the individual particles. We first impose some standard regularity assumptions on the coefficient functions in the infinitesimal generator \mathcal{A} :

Assumption 1.1.1. The functions

$$\mu : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d \quad \text{and} \quad \sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$$

are measurable and satisfy the following Lipschitz and linear growth conditions: There exists a constant $L > 0$ such that

$$\begin{aligned} \|\mu(t, x) - \mu(t, y)\| + \|\sigma(t, x) - \sigma(t, y)\| &\leq L \cdot \|x - y\|, & t \in [0, T], \quad x, y \in \mathbb{R}^d, \\ \|\mu(t, x)\|^2 + \|\sigma(t, x)\|^2 &\leq L^2 \cdot (1 + \|x\|^2), & t \in [0, T], \quad x \in \mathbb{R}^d. \end{aligned} \quad (1.4) \quad \diamond$$

Under this assumption, classical results such as Theorem 5.2.5/9 in [KS98] or Theorem 3.21 in

[PR14] imply that the stochastic differential equation

$$\begin{aligned}\bar{X}_s^{t,x} &= x, & s \in [0, t], \\ d\bar{X}_s^{t,x} &= \mu(s, \bar{X}_s^{t,x})ds + \sigma(s, \bar{X}_s^{t,x})d\bar{W}_s, & s \in [t, T],\end{aligned}\tag{1.5}$$

admits a pathwise unique strong solution for each starting configuration $(t, x) \in [0, T] \times \mathbb{R}^d$. Here, \bar{W} is an \mathbb{R}^d -valued Brownian motion on $(\Omega, \mathfrak{A}, \mathbb{P})$ that is independent of all other random variables that occurred so far. The natural filtration of $\bar{X}^{t,x}$ augmented by all \mathbb{P} -nullsets is denoted by $\bar{\mathfrak{F}}^{t,x} = \{\bar{\mathfrak{F}}_s^{t,x}\}_{s \in [0, T]}$. An application of the results in Chapter 3.7 of [PR14] yields the following:

Lemma 1.1.2 (Properties of the Diffusion). The random field $\{\bar{X}_s^{t,x}\}_{s, t \in [0, T], x \in \mathbb{R}^d}$ defined via (1.5) can be chosen such that it satisfies the following conditions:

(i) *Continuity with respect to initial data:* The map

$$\bar{X} : [0, T] \times \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d, \quad (t, x, s) \mapsto \bar{X}_s^{t,x},$$

is almost surely continuous.

(ii) *Flow property:* For all $(t, x) \in [0, T] \times \mathbb{R}^d$ and any $[t, T]$ -valued $\bar{\mathfrak{F}}^{t,x}$ -stopping time τ ,

$$\bar{X}_{\tau+s}^{\tau, \bar{X}_\tau^{t,x}} \mathbf{1}_{\{\tau+s \leq T\}} = \bar{X}_{\tau+s}^{t,x} \mathbf{1}_{\{\tau+s \leq T\}}, \quad s \in [0, \infty).$$

(iii) *Strong Markov property:* For all $(t, x) \in [0, T] \times \mathbb{R}^d$ and $s \in [0, \infty)$, for any $[t, T]$ -valued $\bar{\mathfrak{F}}^{t,x}$ -stopping time τ , and for every bounded measurable function $h : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, we have

$$\mathbb{E} \left[h(\tau + s, \bar{X}_{\tau+s}^{t,x}) \mathbf{1}_{\{\tau+s \leq T\}} \middle| \bar{\mathfrak{F}}_\tau^{t,x} \right] = \mathbb{E} \left[h(\tau + s, \bar{X}_{\tau+s}^{t,x}) \mathbf{1}_{\{\tau+s \leq T\}} \middle| (\tau, \bar{X}_\tau^{t,x}) \right]. \quad \diamond$$

The branching diffusion is constructed by attaching to each particle in the branching mechanism a diffusion with the same dynamics as \bar{X} , but with a different driving noise and a suitable initial condition. To make this precise, we fix $x \in \mathbb{R}^d$ and define for each $k \in \bar{\mathcal{K}}_t$ with $k = (1, k_2, \dots, k_n) \in \mathbb{N}^n$ an associated diffusion $X^{(k)} = X^{k,t,x} = \{X_s^{k,t,x}\}_{s \in [T_t^{(k-)}, T_t^{(k)}]}$ as the unique strong solution of

$$X_{T_t^{(k-)}}^{(k)} = \Gamma_{J^{(k)}}(T_t^{(k-)}, X_{T_t^{(k-)}}^{(k-)}, \Delta^{(k-)}), \tag{1.6}$$

$$dX_s^{(k)} = \mu(s, X_s^{(k)})ds + \sigma(s, X_s^{(k)})dW_s^{(k)}, \quad s \in [T_t^{(k-)}, T_t^{(k)}]. \tag{1.7}$$

Note that replacing (1.6) with $X_t^{(1)} = x$ in case of $k = (1)$ renders this iteration well-defined. It follows that $X^{(k)}$ has the same dynamics as \bar{X} , but with the different, independent, driving noise $W^{(k)}$. The lifetime of $X^{(k)}$ coincides with the lifetime of the particle k . Moreover, in case of $k \neq (1)$, the initial value of $X^{(k)}$ is the terminal value of the diffusion $X^{(k-)}$ associated with its parent particle $k-$, plus an additional jump whose size is given by the jump map $\Gamma_{J^{(k)}}$ corresponding to its mark $J^{(k)}$ and the jump parameter $\Delta^{(k-)}$.

The resulting family $\{X^{(k)}\}_{k \in \bar{\mathcal{K}}_t}$ is referred to as a *branching diffusion with jumps*.⁶

Figure 1.2 below visualizes a sample path of a branching Brownian motion with jumps on $[0, 1]$. For later reference, we encode all information available up to generation $n \in \mathbb{N}$ by setting

$$\mathfrak{F}^n \triangleq \sigma(W^{(k)}, \tau^{(k)}, \Delta^{(k)}, I^{(k)} : k \in \mathbf{N}_n).$$

For notational convenience, we furthermore write $\mathfrak{F}^0 \triangleq \{\emptyset, \Omega\}$ for the trivial σ -algebra. Finally, for $n \in \mathbb{N}_0$, we enlarge these σ -algebras by the branching time information of one future generation, i.e. we set

$$\mathfrak{G}^n \triangleq \mathfrak{F}^n \vee \sigma(\tau^{(k)} : k \in \mathbf{N}_{n+1}).$$

For $n \in \mathbb{N}$ and $k \in \bar{\mathcal{K}}_t^n$, we observe that

$$\text{conditional on } \mathfrak{G}^{n-1}, \text{ the laws of } X^{(k)} \text{ and } \bar{X}_{T_t^{(k-)}, X_{T_t^{(k-)}}}^{(k)} \text{ on } [T_t^{(k-)}, T_t^{(k)}] \text{ are identical.} \quad (1.8)$$

We stress that $X^{(k)}$ and $\bar{X}_{T_t^{(k-)}, X_{T_t^{(k-)}}}^{(k)}$ do not coincide pathwise since the dynamics of $X^{(k)}$ are driven by $W^{(k)}$, while those of $\bar{X}_{T_t^{(k-)}, X_{T_t^{(k-)}}}^{(k)}$ are driven by \bar{W} . Replacing the driving Brownian motion with a new, independent one for each offspring particle – without changing its distribution – will be the key step in the branching representation below.




Figure 1.2: Sample path of a one-dimensional branching Brownian motion with jumps on $[0, 1]$.

⁶For a similar construction of a *branching diffusion* (albeit without jumps) we refer to §2.3 in [HOT⁺19]; see also, e.g., §2.2 in [HTT14].

1.1.4 Branching Diffusion Representation

We now address the branching diffusion representation of classical solutions of (PDE). To begin with, we specify suitable boundedness assumptions on the coefficient functions c_i and the terminal condition g .

Assumption 1.1.3. The functions $c_i : [0, T] \times \mathbb{R}^d \times \Xi \rightarrow \mathbb{R}$, $i \in \mathcal{I}$, and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ in (PDE) and (TC) are bounded and measurable. \diamond

In order for the possibly infinite series in the nonlinearity of (PDE) to be defined unambiguously, we subsequently agree on the following definition of classical solutions.

Definition 1.1.4 (Classical Solution). Under Assumption 1.1.3, a continuous function

$$u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}, \quad (t, x) \mapsto u(t, x)$$

is said to be a *classical solution* of (PDE) with terminal condition (TC) if

(i) $u \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$;

(ii) for each $(t, x) \in [0, T] \times \mathbb{R}^d$, it holds that

$$\sum_{i \in \mathcal{I}} \int_{\Xi} |c_i(t, x, \xi) \mathcal{J}[u](t, x, \xi)^i| \gamma(d\xi) < +\infty;$$

(iii) and u satisfies (PDE) for each $(t, x) \in [0, T] \times \mathbb{R}^d$ and (TC) for each $x \in \mathbb{R}^d$. \diamond

We are now in a position to establish the main result of this section, which allows us to represent a classical solution of (PDE) by means of a functional of the branching diffusion. The key idea is to introduce *randomization* across subsequent generations and subsequently exploit the conditional independence structure.⁷

Remark 1.1.5 (Randomization). We fix a particle $k \in \overline{\mathcal{K}}_t^n$ of generation $n \in \mathbb{N}$. By a slight abuse of notation, we drop any indices pertaining to the initial position (t, x) and write

$$\overline{X} \triangleq \overline{X}^{T^{(k-)}, X_{T^{(k-)}}^{(k)}} \quad \text{and} \quad (\Delta, \tau, I) \triangleq (\Delta^{(k)}, \tau^{(k)}, I^{(k)}).$$

Under suitable regularity and integrability assumptions, any classical solution u of (PDE) admits a Feynman-Kač representation⁸ of the form

$$\begin{aligned} u(T^{(k-)}, X_{T^{(k-)}}^{(k)}) &= \mathbb{E} \left[g(\overline{X}_T) + \int_{T^{(k-)}}^T f(r, \overline{X}_r, \Delta, \mathcal{J}[u](r, \overline{X}_r, \Delta)) dr \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[g(\overline{X}_T) + \int_{T^{(k-)}}^T \sum_{i \in \mathcal{I}} c_i(r, \overline{X}_r, \Delta) \mathcal{J}[u](r, \overline{X}_r, \Delta)^i dr \middle| \mathfrak{F}^{n-1} \right]. \end{aligned} \quad (1.9)$$

⁷These ideas, to be made rigorous in Remark 1.1.5 and the proof of Theorem 1.1.6 below, are well-known in the literature on branching diffusion representations; see, e.g., [HOT⁺19] and the references therein.

⁸As demonstrated in its proof, this holds in particular under the conditions of Theorem 1.1.6.

The key idea underlying the branching diffusion representation is to represent the right-hand side recursively in terms of the branching diffusion $X^{(k)}$, thus eliminating the integral, sum and nonlinearity within the conditional expectation. More precisely, we claim that

$$\begin{aligned} & u(T^{(k-)}, X_{T^{(k-)}}^{(k)}) \\ &= \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T-T^{(k-)})} + \mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)^I \middle| \mathfrak{F}^{n-1} \right]. \end{aligned} \quad (1.10)$$

Let us start by considering the first summand in (1.10). Since $T^{(k)} = (T^{(k-)} + \tau) \wedge T$ and \bar{X} and $X^{(k)}$ have the same conditional distribution given \mathfrak{G}^{n-1} , see (1.8), by using the tower property of conditional expectation, we have

$$\begin{aligned} \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T-T^{(k-)})} \middle| \mathfrak{F}^{n-1} \right] &= \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T-T^{(k-)})} \middle| \mathfrak{G}^{n-1} \right] \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{\{\tau \geq T-T^{(k-)}\}} \frac{g(\bar{X}_T)}{F(T-T^{(k-)})} \middle| \mathfrak{G}^{n-1} \right] \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\mathbb{1}_{\{\tau \geq T-T^{(k-)}\}} \frac{g(\bar{X}_T)}{F(T-T^{(k-)})} \middle| \mathfrak{F}^{n-1} \right]. \end{aligned}$$

But then, since τ is independent of \mathfrak{F}^{n-1} and \bar{X} , we can simply integrate with respect to the density of τ and use the definition of F in (1.1) to obtain

$$\mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T-T^{(k-)})} \middle| \mathfrak{F}^{n-1} \right] = \mathbb{E} \left[\frac{g(\bar{X}_T)}{F(T-T^{(k-)})} F(T-T^{(k-)}) \middle| \mathfrak{F}^{n-1} \right] = \mathbb{E} [g(\bar{X}_T) \middle| \mathfrak{F}^{n-1}]$$

as in (1.9). The second term in (1.10) is slightly more involved, but can be handled similarly: First, we use the conditional identity in law given \mathfrak{G}^{n-1} of \bar{X} and $X^{(k)}$ in (1.8) and the tower property of conditional expectation to obtain

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)^I \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)^I \middle| \mathfrak{G}^{n-1} \right] \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, \bar{X}_{T^{(k)}})}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, \bar{X}_{T^{(k)}}, \Delta)^I \middle| \mathfrak{G}^{n-1} \right] \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, \bar{X}_{T^{(k)}})}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, \bar{X}_{T^{(k)}}, \Delta)^I \middle| \mathfrak{F}^{n-1} \right]. \end{aligned}$$

Next, independence of I and τ from all other objects involved allows us to integrate with respect to the associated probability mass function and density, and we obtain

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}<T\}} \frac{c_I(T^{(k)}, \bar{X}_{T^{(k)}})}{\rho(T^{(k)}-T^{(k-)})p_I} \mathcal{J}[u](T^{(k)}, \bar{X}_{T^{(k)}}, \Delta)^I \middle| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\int_{T^{(k-)}}^T \sum_{i \in I} c_i(r, \bar{X}_r, \Delta) \mathcal{J}[u](r, \bar{X}_r, \Delta)^i dr \middle| \mathfrak{F}^{n-1} \right]. \end{aligned}$$

Combining the above equations, we have therefore argued that

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{\{T^{(k)} < T\}} \frac{c_I(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)}{\rho(T^{(k)} - T^{(k-)})} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta) \Big| \mathfrak{F}^{n-1} \right] \\ &= \mathbb{E} \left[\int_{T^{(k-)}}^T \sum_{i \in \mathcal{I}} c_i(r, \bar{X}_r, \Delta) \mathcal{J}[u](r, \bar{X}_r, \Delta)^i dr \Big| \mathfrak{F}^{n-1} \right] \end{aligned}$$

and thus (1.10) holds. The key advantage of (1.10) is that the nonlinearity is represented via

$$\mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta)^I = \prod_{\ell=1}^m u(T^{(k)}, \Gamma_\ell(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta))^{I_\ell},$$

where it is possible to iterate over all generations of particles, using conditional independence across generations: Indeed, the terms $\Gamma_\ell(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})$ appearing as arguments in the function u correspond to the initial positions of the particles of generation $n+1$; see (1.6). This leads to the branching diffusion representation made rigorous in Theorem 1.1.6 below. \diamond

We next state the first main result of this chapter: A branching diffusion representation of classical solutions of (PDE).⁹

Theorem 1.1.6 (Branching Representation of Classical Solutions). Suppose Assumptions 1.1.1 and 1.1.3 hold, let $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a classical solution of (PDE) satisfying (TC) and fix $(t, x) \in [0, T] \times \mathbb{R}^d$. For each $n \in \mathbb{N}_0$, iteratively define the random variables¹⁰ $\mathcal{G}_0^{t,x} \triangleq \mathcal{C}_0^{t,x} \triangleq 1$,

$$\mathcal{G}_n^{t,x} \triangleq \mathcal{G}_{n-1}^{t,x} \prod_{k \in \mathcal{K}^n} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \quad \text{and} \quad \mathcal{C}_n^{t,x} \triangleq \mathcal{C}_{n-1}^{t,x} \prod_{k \in \bar{\mathcal{K}}^n \setminus \mathcal{K}^n} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})}{\rho(T^{(k)} - T^{(k-)})} p_{I^{(k)}}$$

and

$$\begin{aligned} \mathcal{R}_n^{t,x} &\triangleq \prod_{k \in \mathcal{K}^{n+1}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \\ &\times \prod_{k \in \bar{\mathcal{K}}^{n+1} \setminus \mathcal{K}^{n+1}} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)}) \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}}}{\rho(T^{(k)} - T^{(k-)}) p_{I^{(k)}}}, \end{aligned}$$

and set

$$\Psi_n^{t,x} \triangleq \mathcal{G}_n^{t,x} \mathcal{C}_n^{t,x} \mathcal{R}_n^{t,x} \quad \text{and} \quad \Psi^{t,x} \triangleq \lim_{n \rightarrow \infty} \Psi_n^{t,x}. \quad (1.11)$$

Suppose that

- (i) the family $\{\Psi_n^{t,x}\}_{n \in \mathbb{N}_0}$ is uniformly integrable;

⁹Note that this result is similar to Proposition 3.4 in [HOT⁺19]; it partially extends the latter to the present setup of a PDE with a *nonlocal* analytic nonlinearity.

¹⁰For notational convenience, we suppress the dependence on (t, x) on the right-hand sides of these definitions but highlight here that $X^{(k)} = X^{k,t,x}$ and $T^{(k)} = T_t^{(k)}$ as well as $\bar{\mathcal{K}}^n = \bar{\mathcal{K}}_t^n$, $\mathcal{K}^n = \mathcal{K}_t^n$, $\bar{\mathcal{K}} = \bar{\mathcal{K}}_t$ and $\mathcal{K} = \mathcal{K}_t$.

(ii) for every $(s, y) \in [t, T] \times \mathbb{R}^d$, it holds that

$$\sum_{i \in \mathcal{I}} \mathbb{E} \left[\int_s^T \int_{\Xi} |c_i(r, \bar{X}_r^{s,y}, \xi) \mathcal{J}[u](r, \bar{X}_r^{s,y}, \xi)| \gamma(d\xi) dr \right] < +\infty;$$

(iii) for any $(s, y) \in [t, T] \times \mathbb{R}^d$, the local martingale

$$M^{s,y} \triangleq \int_s^\cdot \nabla_x u(r, \bar{X}_r^{s,y})^\top \sigma(r, \bar{X}_r^{s,y}) d\bar{W}_r$$

is a martingale.

Then $\Psi^{t,x}$ is integrable and u admits the branching diffusion representation

$$u(t, x) = \mathbb{E}[\Psi^{t,x}]. \quad \diamond$$

Before we turn to the proof, note that unwinding the definitions we have $\Psi^{t,x} = \mathcal{G}^{t,x} \mathcal{C}^{t,x}$ where

$$\mathcal{G}^{t,x} = \prod_{k \in \mathcal{K}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \quad \text{and} \quad \mathcal{C}^{t,x} = \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})}{\rho(T^{(k)} - T^{(k-)}) p_{I^{(k)}}}, \quad (1.11)$$

so the branching diffusion representation in Theorem 1.1.6 can be written more explicitly as

$$u(t, x) = \mathbb{E} \left[\prod_{k \in \mathcal{K}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})}{\rho(T^{(k)} - T^{(k-)}) p_{I^{(k)}}} \right]. \quad (1.12)$$

Proof of Theorem 1.1.6. Since $\lim_{n \rightarrow \infty} \Psi_n^{t,x} = \Psi^{t,x}$ \mathbb{P} -a.s. by (1.3) and as $\{\Psi_n^{t,x}\}_{n \in \mathbb{N}_0}$ is uniformly integrable by (i), it follows from Vitali's convergence theorem that $\Psi^{t,x}$ is integrable and $\lim_{n \rightarrow \infty} \mathbb{E}[\Psi_n^{t,x}] = \mathbb{E}[\Psi^{t,x}]$. Hence, to prove the result, it suffices to show that

$$u(t, x) = \mathbb{E}[\Psi_n^{t,x}] \quad \text{for each } n \in \mathbb{N}_0. \quad (1.13)$$

Step 1: Feynman-Kač representation and randomization. Using Itô's lemma and the fact that u solves (PDE) subject to (TC), we have

$$\begin{aligned} u(s, y) &= u(T, \bar{X}_T^{s,y}) - \int_s^T \left[\partial_t u(r, \bar{X}_r^{s,y}) + \mathcal{A}[u](r, \bar{X}_r^{s,y}) \right] dr - M_T^{s,y} \\ &= g(\bar{X}_T^{s,y}) + \int_s^T \int_{\Xi} f(r, \bar{X}_r^{s,y}, \xi, \mathcal{J}[u](r, \bar{X}_r^{s,y}, \xi)) \gamma(d\xi) dr - M_T^{s,y} \end{aligned}$$

for any $(s, y) \in [t, T] \times \mathbb{R}^d$. Fix a particle $k \in \bar{\mathcal{K}}^n$ of generation $n \in \mathbb{N}$, write $\bar{X}^{(k)} \triangleq \bar{X}^{T^{(k-)}, X_{T^{(k-)}}^{(k)}}$, and note that $(T^{(k-)}, X_{T^{(k-)}}^{(k)})$ is \mathfrak{F}^{n-1} -measurable. Choosing $(s, y) = (T^{(k-)}, X_{T^{(k-)}}^{(k)})$ in the Itô representation above, taking conditional expectations and using (ii) and (iii) we obtain

$$u(T^{(k-)}, X_{T^{(k-)}}^{(k)}) = \mathbb{E} \left[g(\bar{X}_T^{(k)}) + \int_{T^{(k-)}}^T \int_{\Xi} f(r, \bar{X}_r^{(k)}, \xi, \mathcal{J}[u](r, \bar{X}_r^{(k)}, \xi)) \gamma(d\xi) dr \middle| \mathfrak{F}^{n-1} \right].$$

Since $\Delta^{(k)}$ has distribution γ , is independent of \mathfrak{F}^{n-1} and $\bar{X}^{(k)}$ and $\Delta^{(k)}$ are independent, it follows that

$$u(T^{(k-)}, X_{T^{(k-)}}^{(k)}) = \mathbb{E} \left[g(\bar{X}_T^{(k)}) + \int_{T^{(k-)}}^T f(r, \bar{X}_r^{(k)}, \Delta^{(k)}, \mathcal{J}[u](r, \bar{X}_r^{(k)}, \Delta^{(k)})) dr \middle| \mathfrak{F}^{n-1} \right],$$

and by the argument in Remark 1.1.5 this can be further rewritten as

$$\begin{aligned} u(T^{(k-)}, X_{T^{(k-)}}^{(k)}) &= \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \right. \\ &\quad \left. + \mathbb{1}_{\{T^{(k)}<T\}} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})}{\rho(T^{(k)} - T^{(k-)})p_{I^{(k)}}} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}} \middle| \mathfrak{F}^{n-1} \right]. \end{aligned} \quad (1.14)$$

Step 2: Induction. We establish (1.13) by induction on n . For $n = 0$, let $k = (1)$ be the only particle of generation 1, recall that \mathfrak{F}^0 is trivial, and note that (1.14) rewrites as

$$\begin{aligned} u(t, x) &= \mathbb{E} \left[\mathbb{1}_{\{T^{(k)}=T\}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \right. \\ &\quad \left. + \mathbb{1}_{\{T^{(k)}<T\}} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)}) \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}}}{\rho(T^{(k)} - T^{(k-)})p_{I^{(k)}}} \right] \\ &= \mathbb{E}[\mathcal{R}_0^{t,x}] = \mathbb{E}[\Psi_0^{t,x}]. \end{aligned}$$

Now, let $n \in \mathbb{N}$ and suppose that the claim is true for $n - 1$, i.e.

$$u(t, x) = \mathbb{E}[\Psi_{n-1}^{t,x}] = \mathbb{E}[\mathcal{G}_{n-1}^{t,x} \mathcal{C}_{n-1}^{t,x} \mathcal{R}_{n-1}^{t,x}]. \quad (1.15)$$

Let $k \in \bar{\mathcal{K}}^n$ be an arbitrary particle of generation n . On the event $\{k \in \bar{\mathcal{K}}^n \setminus \mathcal{K}^n\} = \{T^{(k)} < T\}$, the particle k branches into $|I^{(k)}|$ offspring particles (k, k_{n+1}) , $k_{n+1} \in [1 : |I^{(k)}|]$, of which the first $I_1^{(k)}$ have mark 1, i.e. jump to $\Gamma_1(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})$, the next $I_2^{(k)}$ have mark 2, i.e. jump to $\Gamma_2(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})$, and so forth. Thus, on the event $\{k \in \bar{\mathcal{K}}^n \setminus \mathcal{K}^n\}$, we have

$$\begin{aligned} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}} &= \prod_{\ell=1}^m u(T^{(k)}, \Gamma_\ell(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)}))^{I_\ell^{(k)}} \\ &= \prod_{k_{n+1}=1}^{|I^{(k)}|} u(T^{(k)}, X_{T^{(k)}}^{(k, k_{n+1})}) = \prod_{\bar{k} \in \bar{\mathcal{K}}^{n+1}, \bar{k}-=k} u(T^{(\bar{k}-)}, X_{T^{(\bar{k}-)}}^{(\bar{k})}) \\ &= \prod_{\bar{k} \in \bar{\mathcal{K}}^{n+1}, \bar{k}-=k} \mathbb{E} \left[\mathbb{1}_{\{T^{(\bar{k})}=T\}} \frac{g(X_T^{(\bar{k})})}{F(T - T^{(\bar{k}-)})} \right. \\ &\quad \left. + \mathbb{1}_{\{T^{(\bar{k})}<T\}} \frac{c_{I^{(\bar{k})}}(T^{(\bar{k})}, X_{T^{(\bar{k})}}^{(\bar{k})}, \Delta^{(\bar{k})}) \mathcal{J}[u](T^{(\bar{k})}, X_{T^{(\bar{k})}}^{(\bar{k})}, \Delta^{(\bar{k})})^{I^{(\bar{k})}}}{\rho(T^{(\bar{k})} - T^{(\bar{k}-)})p_{I^{(\bar{k})}}} \middle| \mathfrak{F}^n \right], \end{aligned}$$

where the final identity is due to (1.14). Thus using \mathfrak{F}^n -conditional independence of individual

offspring particles, we have

$$\begin{aligned}
& \prod_{k \in \bar{\mathcal{K}}^n \setminus \mathcal{K}^n} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}} \\
&= \mathbb{E} \left[\prod_{\bar{k} \in \bar{\mathcal{K}}^{n+1}} \left(\mathbb{1}_{\{T^{(\bar{k})} = T\}} \frac{g(X_{T^{(\bar{k})}}^{(\bar{k})})}{F(T - T^{(\bar{k}-)})} \right. \right. \\
&\quad \left. \left. + \mathbb{1}_{\{T^{(\bar{k})} < T\}} \frac{c_{I^{(\bar{k})}}(T^{(\bar{k})}, X_{T^{(\bar{k})}}^{(\bar{k})}, \Delta^{(\bar{k})}) \mathcal{J}[u](T^{(\bar{k})}, X_{T^{(\bar{k})}}^{(\bar{k})}, \Delta^{(\bar{k})})^{I^{(\bar{k})}}}{\rho(T^{(\bar{k})} - T^{(\bar{k}-)}) p_{I^{(\bar{k})}}}} \right) \Big| \mathfrak{F}^n \right] \\
&= \mathbb{E}[\mathcal{R}_n^{t,x} | \mathfrak{F}^n]. \tag{1.16}
\end{aligned}$$

But then, using (1.16) and the definition of $\mathcal{R}_{n-1}^{t,x}$, we obtain

$$\mathcal{R}_{n-1}^{t,x} = \frac{\mathcal{G}_n^{t,x} \mathcal{C}_n^{t,x}}{\mathcal{G}_{n-1}^{t,x} \mathcal{C}_{n-1}^{t,x}} \prod_{k \in \bar{\mathcal{K}}^n \setminus \mathcal{K}^n} \mathcal{J}[u](T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})^{I^{(k)}} = \frac{\mathcal{G}_n^{t,x} \mathcal{C}_n^{t,x}}{\mathcal{G}_{n-1}^{t,x} \mathcal{C}_{n-1}^{t,x}} \mathbb{E}[\mathcal{R}_n^{t,x} | \mathfrak{F}^n].$$

Plugging this into (1.15) yields the claim since $(\mathcal{G}_n^{t,x}, \mathcal{C}_n^{t,x})$ is \mathfrak{F}^n -measurable and thus

$$u(t, x) = \mathbb{E}[\mathcal{G}_{n-1}^{t,x} \mathcal{C}_{n-1}^{t,x} \mathcal{R}_{n-1}^{t,x}] = \mathbb{E}[\mathcal{G}_n^{t,x} \mathcal{C}_n^{t,x} \mathbb{E}[\mathcal{R}_n^{t,x} | \mathfrak{F}^n]] = \mathbb{E}[\Psi_n^{t,x}]. \quad \square$$

Remark 1.1.7. A sufficient condition for (ii) in Theorem 1.1.6 to hold is that u is bounded and $\sum_{i \in \mathcal{I}} \|c_i\|_\infty \|u\|_\infty^{|i|} < +\infty$; note that this condition simplifies to $\sum_{i \in \mathcal{I}} \|c_i\|_\infty < +\infty$ if $\|u\|_\infty \leq 1$. (iii) holds if $\nabla_x u$ and σ are bounded, or, more generally, if $\mathbb{E}[\int_s^T \|\sigma(r, \bar{X}_r^{s,y})^\top \nabla_x u(r, \bar{X}_r^{s,y})\|^2 dr] < +\infty$ for all $(s, y) \in [t, T] \times \mathbb{R}^d$. If u is bounded, similar arguments as in Section 1.2.2 can be used to obtain sufficient conditions for (i). \diamond

The branching diffusion representation in Theorem 1.1.6 can be extended in several ways: Upon combining our approach with that in [HOT⁺19], one can also treat mixed local-nonlocal analytic nonlinearities that include first-order derivatives. Moreover, the notion of branching diffusions with jumps can also be extended to that of branching *jump*-diffusions with jumps, allowing for an additional (linear) nonlocal term in the infinitesimal generator. Inspired by [AC18], it is furthermore possible to consider nonlocal Dirichlet problems; we sketch a corresponding extension of Theorem 1.1.6 in Appendix B.

1.2 Viscosity Solutions and Branching Diffusion Representations

In the previous section, our point of view was to start with a classical solution of (PDE) and derive its branching diffusion representation. This result was achieved under the assumption that a classical solution exists.

In this section, we study the converse question: Can the branching diffusion representation be used to define a solution of the PDE? It is clear that this representation does in general not yield a sufficiently regular solution to qualify as a classical solution, hence we subsequently work with the weaker concept of viscosity solutions.

The main result of this section gives sufficient conditions under which the branching diffusion representation defines a viscosity solution of (PDE). We then derive conditions under which the family $\{\Psi^{t,x}\}_{(t,x) \in [0,T] \times \mathbb{R}^d}$ is uniformly integrable, which implies one of the key assumptions needed to obtain this viscosity property.

1.2.1 Viscosity Solutions of Nonlocal Nonlinear PDEs

We first provide a definition of viscosity solutions of (PDE) appropriate to deal with the nonlocal terms in the nonlinearity.¹¹

Definition 1.2.1 (Viscosity Solution). Suppose Assumption 1.1.3 holds and let $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous function such that

$$\sum_{i \in \mathcal{I}} \int_{\Xi} |c_i(t, x, \xi) \mathcal{J}[u](t, x, \xi)^i| \gamma(d\xi) < +\infty \quad \text{for all } (t, x) \in [0, T] \times \mathbb{R}^d.$$

We say that

1. u is a *viscosity subsolution* of (PDE) if for all $(t, x) \in [0, T] \times \mathbb{R}^d$ and all test functions $\varphi \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$ with $\varphi(t, x) = u(t, x)$ and $\varphi \geq u$ we have

$$-\partial_t \varphi(t, x) - \mathcal{A}[\varphi](t, x) - \int_{\Xi} f(t, x, \xi, \mathcal{J}[u](t, x, \xi)) \gamma(d\xi) \leq 0;$$

2. u is a *viscosity supersolution* of (PDE) if for all $(t, x) \in [0, T] \times \mathbb{R}^d$ and all test functions $\varphi \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$ with $\varphi(t, x) = u(t, x)$ and $\varphi \leq u$ we have

$$-\partial_t \varphi(t, x) - \mathcal{A}[\varphi](t, x) - \int_{\Xi} f(t, x, \xi, \mathcal{J}[u](t, x, \xi)) \gamma(d\xi) \geq 0;$$

3. u is a *viscosity solution* of (PDE) if it is both a viscosity sub- and supersolution. \diamond

With this definition in place, we can state our second main result.¹²

Theorem 1.2.2 (Viscosity Property of the Branching Representation). For all $(t, x) \in [0, T] \times \mathbb{R}^d$, let $\Psi^{t,x}$ be given as in (1.11) and define

$$u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}, \quad (t, x) \mapsto u(t, x) \triangleq \mathbb{E}[\Psi^{t,x}]. \quad (1.17)$$

In addition to Assumptions 1.1.1 and 1.1.3, assume that the SDE coefficients μ and σ as well as the PDE coefficients c_i , $i \in \mathcal{I}$, Γ_ℓ , $\ell \in [1:m]$, and g are continuous.¹³ Moreover, suppose that the following conditions are satisfied for all $(t, x) \in [0, T] \times \mathbb{R}^d$:

- (i) There exists $\varepsilon > 0$ such that the family $\{\Psi^{s,y}\}_{(s,y) \in B_\varepsilon(t,x)}$ is uniformly integrable.

¹¹Note that the following definition is inspired by the standard definition of viscosity solutions of local PDEs; see, e.g., Definition 6.3 in [Tou13].

¹²Note that this result is consistent with the existing literature on branching diffusion representations for local parabolic PDEs; see, e.g., Theorem 4.1 in [Hen12a] or Theorem 3.5 in [HOT⁺19].

¹³As regards c_i and Γ_ℓ , we assume that for every fixed $\xi \in \Xi$ the maps $(t, x) \mapsto c_i(t, x, \xi)$ and $(t, x) \mapsto \Gamma_\ell(t, x, \xi)$ are continuous.

(ii) There exist $\delta > 0$ and a measurable function $\zeta : \mathcal{I} \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$|c_i(s, y, \xi) \mathcal{J}[u](s, y, \xi)^i| \leq \zeta(i, \xi) \quad \text{for all } (s, y) \in \overline{B_\delta(t, x)}$$

with

$$\sum_{i \in \mathcal{I}} \int_{\Xi} |\zeta(i, \xi)| \gamma(d\xi) < +\infty.$$

(iii) It holds that

$$\sum_{i \in \mathcal{I}} \mathbb{E} \left[\int_t^T \int_{\Xi} |c_i(s, \bar{X}_s^{t,x}, \xi) \mathcal{J}[u](s, \bar{X}_s^{t,x}, \xi)^i| \gamma(d\xi) ds \right] < +\infty.$$

Then u is a viscosity solution of (PDE). \diamond

Note that if u is bounded and $\sum_{i \in \mathcal{I}} \|c_i\|_\infty \|u\|_\infty^{|i|} < +\infty$ (as in Remark 1.1.7), then conditions (ii) and (iii) are satisfied; sufficient conditions for (i), which simultaneously imply boundedness of u , are given in Section 1.2.2 below. \diamond

*Proof of Theorem 1.2.2.*¹⁴ Note that the uniform integrability assumption (i) implies that $\Psi^{t,x} \in L^1(\mathbb{P})$ for every $(t, x) \in [0, T] \times \mathbb{R}^d$ and hence u is well-defined. Moreover, Lemma 1.1.2 and continuity of c_i , $i \in \mathcal{I}$, Γ_ℓ , $\ell \in [1:m]$, g , ρ and F guarantee that $(t, x) \mapsto \Psi^{t,x}$ is almost surely continuous, so u is continuous by Vitali's convergence theorem. Finally, condition (ii) guarantees that u satisfies the integrability conditions required on viscosity solutions in Definition 1.2.1.

We fix $(t, x) \in [0, T] \times \mathbb{R}^d$.

Step 1: Dynamic programming representation. As before, we drop the index (t, x) in some of the random variables and processes to simplify notation. Moreover, we set

$$(X, I, \Delta) \triangleq (X^{(1)}, I^{(1)}, \Delta^{(1)}).$$

We first observe that

$$\mathbf{1}_{\{T^{(1)}=T\}} \Psi^{t,x} = \mathbf{1}_{\{T^{(1)}=T\}} \frac{g(X_T)}{F(T-t)}$$

as well as

$$\begin{aligned} \mathbf{1}_{\{T^{(1)} < T\}} \Psi^{t,x} &= \mathbf{1}_{\{T^{(1)} < T\}} \frac{c_I(T^{(1)}, X_{T^{(1)}}), \Delta)}{\rho(T^{(1)} - t) p_I} \\ &\quad \times \prod_{k \in \mathcal{K} \setminus \{1\}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \prod_{k \in \bar{\mathcal{K}} \setminus (\mathcal{K} \cup \{1\})} \frac{c_{I^{(k)}}(T^{(k)}, X_{T^{(k)}}), \Delta^{(k)}}{\rho(T^{(k)} - T^{(k-)}) p_{I^{(k)}}} \\ &= \mathbf{1}_{\{T^{(1)} < T\}} \frac{c_I(T^{(1)}, X_{T^{(1)}}), \Delta)}{\rho(T^{(1)} - t) p_I} \prod_{k_2=1}^{|I|} \Psi^{T^{(1)}, X_{T^{(1)}}^{(1, k_2)}}. \end{aligned}$$

Using the definition of the branching diffusion and its conditional independence structure, we

¹⁴The result is proved along the lines of the proof of Theorem 3.5 in [HOT⁺19].

find that

$$\begin{aligned} \mathbb{E}\left[\prod_{k_2=1}^{|I|} \Psi^{T^{(1)}, X_{T^{(1)}}^{(1, k_2)}} \Big| \mathfrak{F}^1\right] &= \prod_{k_2=1}^{|I|} \sum_{\ell=1}^m \mathbf{1}_{\{J^{(1, k_2)} = \ell\}} \mathbb{E}\left[\Psi^{s, y}\right] \Big|_{(s, y) = (T^{(1)}, \Gamma_\ell(T^{(1)}, X_{T^{(1)}}(\Delta)))} \\ &= \mathcal{J}[u](T^{(1)}, X_{T^{(1)}}(\Delta))^I. \end{aligned}$$

Putting these equations together and using the tower property of conditional expectation, it follows as in Remark 1.1.5 that¹⁵

$$\begin{aligned} u(t, x) &= \mathbb{E}\left[\mathbb{E}[\Psi^{t, x} | \mathfrak{F}^1]\right] \\ &= \mathbb{E}\left[\mathbf{1}_{\{T^{(1)} = T\}} \frac{g(X_T)}{F(T-t)} + \mathbf{1}_{\{T^{(1)} < T\}} \frac{c_I(T^{(1)}, X_{T^{(1)}}(\Delta)) \mathcal{J}[u](T^{(1)}, X_{T^{(1)}}(\Delta))^I}{\rho(T^{(1)} - t) p_I}\right] \\ &= \mathbb{E}\left[g(\bar{X}_T) + \int_t^T \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right]. \end{aligned} \quad (1.18)$$

For $\varepsilon > 0$, we now introduce the stopping time

$$\tau_\varepsilon \triangleq \inf\{s \geq t : \|\bar{X}_s - x\| \geq \varepsilon\} \wedge (t + \varepsilon) \wedge T.$$

From (1.18), the flow property and the strong Markov property of \bar{X} as noted in Lemma 1.1.2, in combination with the conditional version of Fubini's theorem, which is applicable by (iii), it follows that u satisfies the dynamic programming representation¹⁶

$$\begin{aligned} u(t, x) &= \mathbb{E}\left[g(\bar{X}_T) + \int_t^T \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right] \\ &= \mathbb{E}\left[\mathbb{E}[g(\bar{X}_T) | \mathfrak{F}_{\tau_\varepsilon}^1] + \int_{\tau_\varepsilon}^T \int_{\Xi} \sum_{i \in \mathcal{I}} \mathbb{E}[c_i(s, \bar{X}_s, \xi) \mathcal{J}[u](s, \bar{X}_s, \xi)^i | \mathfrak{F}_{\tau_\varepsilon}^1] \gamma(d\xi) ds\right] \\ &\quad + \mathbb{E}\left[\int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right] \\ &= \mathbb{E}\left[\mathbb{E}[g(\bar{X}_T) | (\tau_\varepsilon, \bar{X}_{\tau_\varepsilon})] + \int_{\tau_\varepsilon}^T \int_{\Xi} \sum_{i \in \mathcal{I}} \mathbb{E}[c_i(s, \bar{X}_s, \xi) \mathcal{J}[u](s, \bar{X}_s, \xi)^i | (\tau_\varepsilon, \bar{X}_{\tau_\varepsilon})] \gamma(d\xi) ds\right] \\ &\quad + \mathbb{E}\left[\int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[g(\bar{X}_T^{\tau_\varepsilon}, \bar{X}_{\tau_\varepsilon}) + \int_{\tau_\varepsilon}^T \int_{\Xi} f(s, \bar{X}_s^{\tau_\varepsilon}, \bar{X}_{\tau_\varepsilon}, \xi, \mathcal{J}[u](s, \bar{X}_s^{\tau_\varepsilon}, \bar{X}_{\tau_\varepsilon}, \xi)) \gamma(d\xi) ds \Big| (\tau_\varepsilon, \bar{X}_{\tau_\varepsilon})\right]\right] \\ &\quad + \mathbb{E}\left[\int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right] \\ &= \mathbb{E}\left[u(\tau_\varepsilon, \bar{X}_{\tau_\varepsilon}) + \int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds\right]. \end{aligned} \quad (1.19)$$

Step 2: Viscosity solution property. From the dynamic programming representation (1.19), the viscosity property of u follows by standard arguments.¹⁷ To keep the presentation self-

¹⁵Note that the following identity (1.18) is in fact valid for $(t, x) \in [0, T] \times \mathbb{R}^d$.

¹⁶As mentioned above, we suppress the initial condition in the notation whenever it coincides with (t, x) ; correspondingly, we consider $\bar{X} = \bar{X}^{t, x}$ and $\mathfrak{F} = \mathfrak{F}^{t, x}$.

¹⁷see, e.g., Proposition 3.6/7 in [Sei11], Proposition 7.2/3 in [Tou13] or Theorem 5.5.8 in [Zha17].

contained, we provide a proof of the subsolution property (the supersolution property is established analogously). We fix a test function $\varphi \in C^{1,2}([0, T] \times \mathbb{R}^d)$ with $\varphi(t, x) = u(t, x)$ and $\varphi \geq u$. By the dynamic programming representation (1.19) and Itô's lemma, we find that

$$\begin{aligned} \varphi(t, x) &= u(t, x) \\ &= \mathbb{E} \left[u(\tau_\varepsilon, \bar{X}_{\tau_\varepsilon}) + \int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds \right] \\ &\leq \mathbb{E} \left[\varphi(\tau_\varepsilon, \bar{X}_{\tau_\varepsilon}) + \int_t^{\tau_\varepsilon} \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds \right] \\ &= \mathbb{E} \left[\varphi(t, x) + \int_t^{\tau_\varepsilon} \partial_t \varphi(s, \bar{X}_s) + \mathcal{A}[\varphi](s, \bar{X}_s) \right. \\ &\quad \left. + \int_{\Xi} f(s, \bar{X}_s, \xi, \mathcal{J}[u](s, \bar{X}_s, \xi)) \gamma(d\xi) ds \right]. \end{aligned}$$

For $(s, y) \in [0, T] \times \mathbb{R}^d$, we now define

$$I_\varphi(s, y) \triangleq \partial_t \varphi(s, y) + \mathcal{A}[\varphi](s, y) + \int_{\Xi} f(s, y, \xi, \mathcal{J}[u](s, y, \xi)) \gamma(d\xi)$$

to arrive at

$$\mathbb{E} \left[\int_t^{\tau_\varepsilon} I_\varphi(s, \bar{X}_s) ds \right] \geq 0.$$

From condition (ii) and dominated convergence, it follows that I is continuous. But then

$$0 \leq \mathbb{E} \left[\int_t^{\tau_\varepsilon} I_\varphi(s, \bar{X}_s) ds \right] \leq \mathbb{E}[\tau_\varepsilon - t] \max_{(s, y) \in B_\varepsilon(t, x)} I_\varphi(s, y),$$

and thus, since $\mathbb{E}[\tau_\varepsilon - t] > 0$,

$$\max_{(s, y) \in B_\varepsilon(t, x)} I_\varphi(s, y) \geq 0.$$

Letting $\varepsilon \downarrow 0$ allows us to conclude that $I_\varphi(t, x) \geq 0$, so u is a viscosity subsolution in the sense of Definition 1.2.1. \square

1.2.2 Sufficient Conditions for Uniform Integrability of $\{\Psi^{t, x}\}$

In this section, we provide a ramification of the results in [HOT⁺19] for branching diffusions with jumps to give sufficient conditions for uniform integrability of $\{\Psi^{t, x}\}$ as required in Theorem 1.2.2.

Proposition 1.2.3 (Integrability Conditions). Let $\kappa \in (1, \infty)$ and define

$$C_1 \triangleq \frac{\|g\|_\infty^\kappa}{F(T)^{\kappa-1}} \quad \text{and} \quad C_2 \triangleq \sup_{i \in \mathcal{I}, t \in (0, T]} \left(\frac{\|c_i\|_\infty}{\rho(t)p_i} \right)^{\kappa-1}.$$

Then the family $\{\Psi^{t, x}\}_{(t, x) \in [0, T] \times \mathbb{R}^d}$ is bounded in $L^\kappa(\mathbb{P})$, and in particular uniformly integrable, in either of the following two cases:

(i) It holds that

$$\frac{C_1}{F(T)} \vee C_2^{\frac{\kappa}{\kappa-1}} \leq 1.$$

- (ii) The power series $\sum_{i \in \mathcal{I}} \|c_i\|_\infty x^{|i|}$ has an infinite radius of convergence and is non-degenerate, i.e. at least one coefficient c_i , $i \in \mathcal{I}$, is non-zero,¹⁸ and the terminal time $T > 0$ is sufficiently small in that

$$T < \int_{C_1}^{\infty} \left(C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty x^{|i|} \right)^{-1} dx. \quad \diamond$$

Conceptually, essential parts of the following proof coincide with the one of Theorem 3.12 in [HOT⁺19]; we also refer to §3.3 in [AC18] for similar considerations.

Proof of Proposition 1.2.3. Fix some $(t, x) \in [0, T] \times \mathbb{R}^d$. By definition of $\Psi^{t,x}$, we have

$$|\Psi^{t,x}|^\kappa = \prod_{k \in \mathcal{K}} \frac{|g(X_T^{(k)})|^\kappa}{F(T - T^{(k-)})^\kappa} \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} \frac{|c_{I(k)}(T^{(k)}, X_{T^{(k)}}^{(k)}, \Delta^{(k)})|^\kappa}{|\rho(T^{(k)} - T^{(k-)})p_{I(k)}|^\kappa}. \quad (1.20)$$

With this, under condition (i), and since F is decreasing we immediately find that

$$\mathbb{E}[|\Psi^{t,x}|^\kappa] \leq \mathbb{E} \left[\prod_{k \in \mathcal{K}} \frac{C_1}{F(T)} \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} C_2^{\frac{\kappa}{\kappa-1}} \right] \leq \mathbb{E} \left[\prod_{k \in \bar{\mathcal{K}}} \frac{C_1}{F(T)} \vee C_2^{\frac{\kappa}{\kappa-1}} \right] \leq 1$$

and the proof is complete. Let us therefore subsequently assume that we are in case (ii); we follow the approach of Theorem 3.12 in [HOT⁺19]: The main idea is to regard an upper bound of $|\Psi^{t,x}|^\kappa$ as a branching estimator for an ODE which admits a global solution. First note that

$$\mathbb{E}[|\Psi^{t,x}|^\kappa] \leq \mathbb{E} \left[\prod_{k \in \mathcal{K}} \frac{C_1}{F(T - T_t^{(k-)})} \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} \frac{C_2 \|c_{I(k)}\|_\infty}{\rho(T^{(k)} - T^{(k-)})p_{I(k)}} \right]. \quad (1.21)$$

Now consider the following ODE to be solved backwards in time:

$$\dot{\eta}(t) + C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty \eta(t)^{|i|} = 0, \quad t \in [0, T]; \quad \eta(T) = C_1. \quad (1.22)$$

Suppose that $C_1 > 0$ and define the map¹⁹

$$\varphi : [C_1, \infty) \rightarrow [0, \infty], \quad y \mapsto \varphi(y) \triangleq \int_{C_1}^y \left(C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty x^{|i|} \right)^{-1} dx.$$

Since the power series is non-degenerate, φ is a continuous mapping that is strictly increasing where finite. Upon rearranging and integrating the ODE, note that a $[C_1, \infty)$ -valued function $\eta \in \mathcal{C}^1([0, T])$ is a solution of (1.22) if and only if

$$-\int_t^T \left(C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty \eta(s)^{|i|} \right)^{-1} \dot{\eta}(s) ds = \int_t^T 1 ds, \quad t \in [0, T]; \quad \eta(T) = C_1,$$

¹⁸Otherwise (PDE) becomes linear; for that case the branching approach becomes unnecessary.

¹⁹The following arguments are borrowed from the theory of ODEs with separated variables; see, e.g., §I.1.VII in [Wal98].

and the substitution $x \triangleq \eta(s)$ shows that this is the case if and only if

$$\varphi(\eta(t)) = \int_{C_1}^{\eta(t)} \left(C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty x^{|i|} \right)^{-1} dx = T - t, \quad t \in [0, T].$$

Since φ is strictly increasing where it is finite, the latter statement is equivalent to

$$T \in \text{range}(\varphi) \cap \mathbb{R} = \left\{ \xi \in \mathbb{R} : 0 \leq \xi \leq \int_{C_1}^{\infty} \left(C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty z^{|i|} \right)^{-1} dz \right\}.$$

Thus, if $C_1 > 0$, condition (ii) on T is both necessary and sufficient to guarantee the existence of a strictly positive solution η of ODE (1.22) on $[0, T]$. Otherwise, i.e. if $C_1 = 0$, note that (1.22) is solved by $\eta \equiv 0$. With this, we are now able to define

$$\chi_n \triangleq \prod_{k \in \bigcup_{\nu=1}^n \mathcal{K}^\nu} \frac{C_1}{F(T - T^{(k-)})} \prod_{k \in \bigcup_{\nu=1}^n \bar{\mathcal{K}}^\nu \setminus \mathcal{K}^\nu} \frac{C_2 \|c_{I(k)}\|_\infty}{\rho(T^{(k)} - T^{(k-)}) p_{I(k)}} \prod_{k \in \bar{\mathcal{K}}^{n+1}} \eta(T^{(k-)})$$

as well as

$$\chi_\infty \triangleq \lim_{n \rightarrow \infty} \chi_n = \prod_{k \in \mathcal{K}} \frac{C_1}{F(T - T^{(k-)})} \prod_{k \in \bar{\mathcal{K}} \setminus \mathcal{K}} \frac{C_2 \|c_{I(k)}\|_\infty}{\rho(T^{(k)} - T^{(k-)}) p_{I(k)}}.$$

By analogous arguments as in Remark 1.1.5, we obtain

$$\eta(t) = \eta(T) + \int_t^T C_2 \sum_{i \in \mathcal{I}} \|c_i\|_\infty \eta(s)^{|i|} ds = \mathbb{E}[\chi_1] = \dots = \mathbb{E}[\chi_n], \quad t \in [0, T]; \quad n \in \mathbb{N}.$$

But then, thanks to (1.21) and Fatou's lemma, it follows that

$$\mathbb{E}[|\Psi^{t,x}|^\kappa] \leq \mathbb{E}[\chi_\infty] \leq \liminf_{n \rightarrow \infty} \mathbb{E}[\chi_n] = \eta(t) \leq \sup_{t \in [0, T]} \eta(t) < \infty, \quad (1.23)$$

and the proof is complete. \square

Under the conditions of Proposition 1.2.3, it follows in particular that $\{\Psi^{t,x}\}_{(t,x) \in [0, T] \times \mathbb{R}^d}$ is bounded in $L^1(\mathbb{P})$, so the function $u : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, $u(t, x) \triangleq \mathbb{E}[\Psi^{t,x}]$ is in fact bounded.

Moreover, in case of $\kappa = 2$ the result provides sufficient conditions for $\Psi^{t,x}$ to have a finite variance.

1.3 Monte Carlo Simulation: A High-Dimensional Example

The representation (1.17) derived in Theorem 1.2.2 makes it possible to compute solutions of (PDE) by direct (non-nested, plain vanilla) Monte Carlo simulation.²⁰ To illustrate the effectiveness and efficiency of the branching Monte Carlo algorithm, we consider the following stylized problem in $d \geq 1$ dimensions which is inspired by the one considered in [HOT⁺19, §5.2]:

$$\begin{aligned} \partial_t u(t, x) + \frac{1}{2d^2} \Delta u(t, x) + \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}} c_i(t, x) u(t, x)^{i_1} u(t, x + \xi)^{i_2} \gamma(d\xi) &= 0 \\ u(T, x) &= \cos(1_d^\top x). \end{aligned} \quad (1.24)$$

Here, Δ denotes the Laplace operator in \mathbb{R}^d ; the time horizon is $T = 1$; γ is the discrete uniform distribution supported on $\{-(\pi/2)e_i \in \mathbb{R}^d : i \in [1:d]\}$;²¹ and the set of possible descendants is given by $\mathcal{I} \triangleq \{i \in \mathbb{N}_0^2 : |i| \leq 2\}$ with coefficients

$$\begin{aligned} c_{(0,0)}(t, x) &= (\alpha + 1/(2d)) \cos(1_d^\top x) \exp\{\alpha(T-t)\} + \cos(1_d^\top x)^2/d - 1/(2d), \\ c_{(1,0)}(t, x) &= (-1) \cdot \cos(1_d^\top x) \exp\{-\alpha(T-t)\}/d, \\ c_{(0,1)}(t, x) &= (-1) \cdot c_{(1,0)}(t, x), \\ c_{(2,0)}(t, x) &= c_{(0,2)}(t, x) = \exp\{-2\alpha(T-t)\}/(2d), \\ c_{(1,1)}(t, x) &= (-2) \cdot c_{(2,0)}(t, x), \end{aligned}$$

where $\alpha = 0.2$. It is not hard to verify that the solution is given in closed form by

$$u(t, x) = \cos(1_d^\top x) e^{\alpha(T-t)} = \cos(\sum_{i=1}^d x_i) e^{\alpha(T-t)}, \quad (t, x) \in [0, T] \times \mathbb{R}^d.$$

We refer to u as the *exact solution* and use it as a benchmark to quantify the error of the estimates produced by our algorithm. All subsequent simulation results correspond to the initial configuration $(t, x) = (0, 1_d)$, i.e. we determine $u^* \triangleq u(0, 1_d)$.

The choice of parameters for the branching diffusion is reported in Table 1.1.

Parameter	Value
Law($\tau^{(1)}$)	$\Gamma(\kappa, \theta)$ with $\kappa = 0.5$ and $\theta = 2.5$
$p_{(0,0)}$	1/3
$p_{(1,0)}, p_{(0,1)}, p_{(1,1)}$	1/6
$p_{(2,0)}, p_{(0,2)}$	1/12

Table 1.1: Parameters of the branching diffusion.²²

Our simulation study is conducted as follows: Given a spatial dimension $d \in \mathbb{N}$ and a number of Monte Carlo simulations N , a simulation run consists of computing the estimator $\hat{u}_{d,N}$ of u^* as the average of N i.i.d. copies $\{\Psi_n^{0,1_d}\}_{n \in [1:N]}$ of $\Psi^{0,1_d}$, i.e. $\hat{u}_{d,N} = \frac{1}{N} \sum_{n=1}^N \Psi_n^{0,1_d}$; its standard

²⁰see, e.g., Algorithm 3.1 in [KKK10] for the *crude* Monte Carlo algorithm.

²¹For $i \in [1:d]$ we denote by e_i the i^{th} unit vector in \mathbb{R}^d .

²²The choice of the lifetime distribution is the same as in [HOT⁺19, §5.2].

deviation is estimated via

$$\widehat{\text{std}}(\hat{u}_{d,N}) \triangleq \left(\frac{1}{N(N-1)} \sum_{n=1}^N (\Psi_n^{0,1d} - \hat{u}_{d,N})^2 \right)^{\frac{1}{2}}.$$

All numerical computations are implemented in MATLAB.²³

Table 1.2 presents the simulation results for the estimator $\hat{u}_{d,N}$ for several values of d and N .

$N \backslash d$	3	5	10	20	50	100
10^2	-1.0296 (0.567)	0.4456 (0.145)	-0.9406 (0.127)	0.5478 (0.064)	1.1470 (0.122)	0.9323 (0.107)
10^3	-1.2793 (0.142)	0.3449 (0.048)	-1.0211 (0.038)	0.4732 (0.019)	1.2351 (0.038)	1.0392 (0.034)
10^4	-1.1832 (0.060)	0.3710 (0.014)	-1.0236 (0.012)	0.5010 (0.006)	1.1707 (0.012)	1.0680 (0.011)
10^5	-1.2003 (0.019)	0.3441 (0.005)	-1.0246 (0.004)	0.4952 (0.002)	1.1837 (0.004)	1.0551 (0.003)
10^6	-1.2165 (0.008)	0.3443 (0.001)	-1.0238 (0.001)	0.4983 (0.001)	1.1797 (0.001)	1.0542 (0.001)
u^*	-1.2092	0.3465	-1.0248	0.4984	1.1786	1.0532

Table 1.2: Simulation results for $\hat{u}_{d,N}$, standard deviation $\widehat{\text{std}}(\hat{u}_{d,N})$ in brackets.

To quantify the accuracy of the Monte Carlo algorithm, we further denote the relative error compared to the exact solution by

$$\text{err}_{\text{rel}}(\hat{u}_{d,N}) \triangleq \frac{|\hat{u}_{d,N} - u^*|}{|u^*|}$$

and report our simulation results in Table 1.3.

$N \backslash d$	3	5	10	20	50	100
10^2	36.3% (42.2%)	30.5% (21.6%)	10.0% (6.7%)	10.2% (6.9%)	8.0% (6.2%)	7.5% (5.8%)
10^3	12.4% (12.0%)	10.3% (8.8%)	2.8% (1.9%)	2.9% (2.1%)	2.6% (1.9%)	2.5% (1.7%)
10^4	4.6% (3.6%)	3.3% (2.6%)	0.9% (0.7%)	0.9% (0.6%)	0.8% (0.5%)	0.8% (0.6%)
10^5	1.5% (1.2%)	1.0% (0.7%)	0.3% (0.2%)	0.3% (0.2%)	0.2% (0.2%)	0.2% (0.2%)
10^6	0.8% (2.2%)	0.3% (0.2%)	0.1% (0.1%)	0.1% (0.1%)	0.1% (0.1%)	0.1% (0.1%)

Table 1.3: Relative error $\text{err}_{\text{rel}}(\hat{u}_{d,N})$, standard deviation in brackets.²⁴

The algorithm is able to achieve high levels of accuracy even for high dimensions of $d = 50$ and $d = 100$, provided N is sufficiently large. Note that in this example the precision of the results actually increases with the dimension d ; this is in line with Proposition 1.2.3, as in the above specification the norm $\|c_i\|_\infty$ appearing in the constant C_2 decreases with d .

At the same time, the running times required to achieve these levels of accuracy are rather modest; the exact values are reported in Table 1.4.²⁵

²³We use MATLAB R2019b [Mat19]; no explicit parallelization; OS: Linux Mint 18.1 Serena.

Hardware: Intel Core i7-6700 CPU @ 3.40 GHz (4 cores) with 32 GiB RAM capacity.

²⁴Estimates and standard deviations are based on 100 mutually independent simulation runs of $\hat{u}_{d,N}$.

²⁵All stated running times are wall-clock times; each of them includes the time required by both the branching Monte Carlo algorithm itself and its initialization.

$N \backslash d$	3	5	10	20	50	100
10^2	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
10^3	0.2 (0.1)	0.2 (0.1)	0.2 (0.0)	0.2 (0.1)	0.2 (0.1)	0.2 (0.1)
10^4	1.5 (0.1)	1.5 (0.1)	1.5 (0.1)	1.5 (0.1)	1.6 (0.1)	2.0 (0.1)
10^5	14.8 (0.1)	14.9 (0.1)	15.0 (0.1)	15.1 (0.1)	15.6 (0.1)	19.8 (0.2)
10^6	146.7 (0.4)	147.4 (0.4)	148.3 (0.4)	149.2 (0.4)	163.9 (1.9)	204.8 (1.6)

Table 1.4: Running time (in seconds) to compute $\hat{u}_{d,N}$, standard deviation in brackets.²⁴

Finally, Figure 1.3 displays the relative error, the standard deviation, and the running time as functions of the number of Monte Carlo samples N . As expected from the Central Limit Theorem, the slope in the logarithmic plot of the standard deviation is approximately $-1/2$. The running times also demonstrate that there is no curse of dimensionality effect.

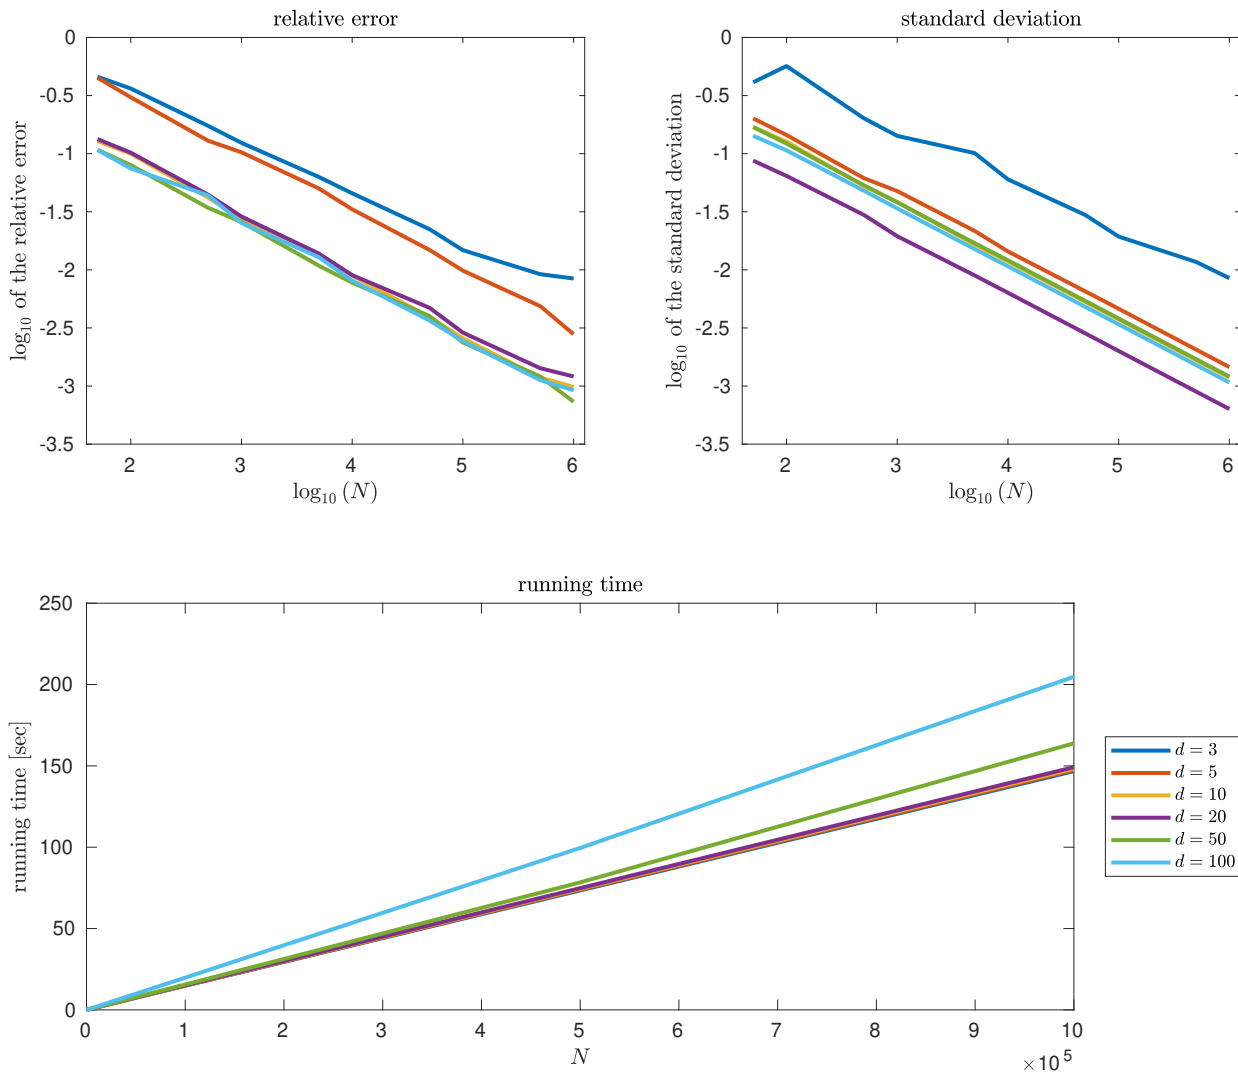


Figure 1.3: Simulation results.

1.4 Valuation with Systemically Important Counterparties

In this section we illustrate the usefulness of the theory developed in this chapter in the valuation of financial positions with systemic counterparty credit risk. Specifically, we assume that the counterparty in a given financial position is a systemically important bank (SIB);²⁶ its systemic importance is captured by jumps in the underlying risk factors, or equivalently devaluations in risky asset prices, that occur upon the SIB's default. This model setup was, e.g., proposed by [PS13]; see also [ML15a], [BP18] and [BPP18]. We wish to stress that our focus here is on situations where finite-difference or fixed point methods for the pricing PDE (as considered in or developed by, e.g., [DH96], [HL99] or [KL16]) are not applicable.

1.4.1 Valuation with Systemic Risk

We consider a financial market that is free of arbitrage; the underlying probability measure (denoted by \mathbb{P} in the preceding sections) is taken as the relevant risk-neutral pricing measure \mathbb{Q} .²⁷ The financial market consists of a locally riskless money market account $B = \{B_t\}_{t \in [0, T]}$ and $d \in \mathbb{N}$ dynamically traded risky assets with prices $X = \{X_t\}_{t \in [0, T]}$ given by an $[0, \infty)^d$ -valued semimartingale such that X/B is a local \mathbb{Q} -martingale. The financial position whose price is to be determined promises a time- T payoff $g(X_T)$ where g is a bounded continuous function of the underlyings X , provided the counterparty does not default before time T .

Crucially, the counterparty in this financial position is a defaultable, systemically important bank (SIB). This means that (i) the financial position is subject to credit risk; and (ii) if and when the SIB counterparty defaults, there is a negative impact on risky asset prices X , which simultaneously affects the value of the financial position $g(X_T)$. Assuming *fractional recovery of post-default mark-to-market value* as, e.g., in [DS99], risk-neutral pricing yields

$$\frac{V_t}{B_t} = \mathbb{E}_t^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau > T\}} \frac{g(X_T)}{B_T} + \mathbf{1}_{\{\tau \leq T\}} \frac{h(V_\tau)}{B_\tau} \right] \quad \text{on } \{\tau > t\}, \quad (1.25)$$

where V represents the value of the financial position *provided the counterparty has not defaulted*; τ is the (original) counterparty's default time; and the recovery value the investor retrieves is given by $h(V_\tau)$, where $h(v) \triangleq Rv^+ - v^-$ with a recovery rate $R \in [0, 1]$. Note that in (1.25), V_τ represents the time- τ mark-to-market price of an identical financial position with an SIB counterparty that has *not* defaulted, but is otherwise identical to the original one, immediately after the original counterparty's default. In particular, V_τ is based on *post-default* risky asset prices $X_\tau = X_{\tau-} + \Delta X_\tau$.

Specifically, we assume that default events are modeled within a classical reduced-form framework, and that the risk-neutral dynamics of X are given by an $[0, \infty)^d$ -valued jump-diffusion.²⁸ Thus

²⁶A list of banks classified as globally systemically important by the Financial Stability Board (FSB) is available at <http://www.fsb.org/work-of-the-fsb/policy-development/addressing-sifis/global-systemically-important-financial-institutions-g-sifis> (accessed: 2020-05-31, 1939).

²⁷Throughout this section, all emerging processes are assumed to be adapted with respect to a suitable filtration $\{\mathfrak{F}_t\}_{t \in [0, T]}$ satisfying the usual conditions of \mathbb{Q} -completeness and right-continuity; by $\mathbb{E}_t^{\mathbb{Q}}[\cdot] \triangleq \mathbb{E}^{\mathbb{Q}}[\cdot | \mathfrak{F}_t]$ we denote the corresponding time- t conditional expectation.

²⁸For details on *reduced-form models* we refer, e.g., to §3.3 in [BMP13] and the references therein.

the SIB counterparty's default time is given by²⁹

$$\tau \triangleq \inf\{t \in [0, T] : Y_t \neq 0\},$$

where $Y = \{Y_t\}_{t \in [0, T]}$ is a Cox process with intensity $\{\lambda(t, X_t)\}_{t \in [0, T]}$,³⁰ and the risk-neutral dynamics of B and X are given by

$$dB_t = r(t, X_t)B_t dt, \quad B_0 = 1, \quad (1.26)$$

$$dX_t = \text{diag}(X_{t-})[\mu(t, X_{t-})dt + \sigma(t, X_{t-})dW_t + \Delta_{Y_t}dY_t], \quad X_0 = x, \quad (1.27)$$

where $\{\Delta_n\}_{n \in \mathbb{N}}$ is a sequence of i.i.d. integrable $E \triangleq (-1, \infty)^d$ -valued random variables and independent of Y and W , and

$$\mu(t, x) \triangleq r(t, x)1_d - \lambda(t, x)\mathbb{E}[\Delta_1].^{31}$$

In particular, the SIB counterparty's default triggers a simultaneous devaluation in the financial position's underlyings of size Δ_1 . The corresponding pricing PDE is given by

$$\partial_t u(t, x) + \mathcal{A}^*[u](t, x) - r(t, x)u(t, x) + \lambda(t, x) \int_E \left(h \circ u(t, x + \text{diag}(\xi)x) - u(t, x) \right) \gamma(d\xi) = 0, \quad (1.28)$$

subject to the terminal condition $u(T, x) = g(x)$, where γ denotes the distribution of Δ_n , $n \in \mathbb{N}$, and the operator \mathcal{A}^* is given by

$$\mathcal{A}^*[u](t, x) \triangleq \mu(t, x)^\top \text{diag}(x) \nabla_x u(t, x) + \frac{1}{2} \text{tr} \left[\text{diag}(x) \sigma(t, x) \sigma(t, x)^\top \text{diag}(x) \nabla_x^2 u(t, x) \right].$$

To justify (1.28) formally, let $\{\tau_n\}_{n \in [1:Y_T]}$ denote the jump times of Y in $[0, T]$ in increasing order (if any); the random measure associated to the marked point process $\{(\tau_n, \Delta_n)\}_{n \in [1:Y_T]}$ is given by

$$\nu(ds, dz) \triangleq \sum_{n=1}^{Y_T} \delta_{\{(\tau_n, \Delta_n)\}}(ds, dz);$$

and the corresponding compensated random measure is defined to be

$$\bar{\nu}(ds, dz) \triangleq \nu(ds, dz) - \lambda(s, X_s) \gamma(dz) ds.^{32}$$

²⁹We stick to the common convention $\inf \emptyset \triangleq +\infty$; note that on the event $\{\tau > T\} = \{\tau = +\infty\}$ the SIB does not default in $[0, T]$.

³⁰see, e.g., Definition II.1.D1 (special case (β)) in [Bré81].

³¹We assume that the functions $r : [0, T] \times [0, \infty)^d \rightarrow \mathbb{R}$, $\lambda : [0, T] \times [0, \infty)^d \rightarrow [0, \infty)$ and $\sigma : [0, T] \times [0, \infty)^d \rightarrow \mathbb{R}^{d \times d}$ are measurable and satisfy suitable Lipschitz and linear growth conditions (in the spirit of Assumption 1.1.1) to guarantee that (1.26)-(1.27) admits a pathwise unique strong solution; for a (change-of-measure) construction of interacting Itô diffusions and point processes we refer, e.g., to [Kus99], [BS05] or [BKS13]; similar arguments also appear in Chapter 2; see Lemma 2.1.3 below. Furthermore, we assume that r and λ are bounded.

Note with the help of Lemma A.4 that X_- is a modification of X and hence in the following we may write X_t instead of X_{t-} whenever it is more convenient and does not make a difference (e.g., in dt - or dW -integrals; see Lemma A.1).

³²For details on marked point processes see, e.g., §VIII.1 in [Bré81].

Then the formal result³³ reads as follows:

Theorem 1.4.1 (Risk-Neutral Pricing). Let $u : [0, T] \times [0, \infty)^d \rightarrow \mathbb{R}$ be a classical solution of (1.28) subject to $u(T, x) = g(x)$ for $x \in [0, \infty)^d$. Moreover, suppose that

$$\begin{aligned} & \int_0^T \|\sigma(s, X_s)^\top \text{diag}(X_s) \nabla_x u(s, X_s)\|_2^2 ds \\ & + \int_0^T \int_E |h \circ u(s, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) - u(s, X_{s-})| \lambda(s, X_s) \gamma(dz) ds < +\infty \quad \mathbb{Q}\text{-a.s.}, \end{aligned}$$

and that

$$\frac{1}{B} \bullet M = \int_{(0, \cdot]} \frac{1}{B_s} dM_s$$

is a \mathbb{Q} -martingale where the local \mathbb{Q} -martingale³⁴ M is given by

$$\begin{aligned} M \triangleq & \int_0^\cdot \nabla_x u(s, X_s)^\top \text{diag}(X_s) \sigma(s, X_s) dW_s \\ & + \int_{(0, \cdot]} \int_E \left(h \circ u(s, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) - u(s, X_{s-}) \right) \bar{\nu}(ds, dz). \end{aligned}$$

Then the process $V = \{V_t\}_{t \in [0, T]}$ given by $V_t \triangleq u(t, X_t)$ for $t \in [0, T]$ satisfies (1.25). In particular, given $X_0 = x \in [0, \infty)^d$, the risk-neutral price of $g(X_T)$ issued by a defaultable SIB is given by $V_0 = u(0, x)$.

Proof. For $t \in [0, T]$ we obtain from (1.27) with Itô's formula³⁵ and (1.28)

$$\begin{aligned} V_t &= u(t, X_t) = u(0, x) + \int_0^t \left(\partial_t u(s-, X_{s-}) + \mathcal{A}^*[u](s-, X_{s-}) \right) ds \\ &+ \int_0^t \nabla_x u(s-, X_{s-})^\top \text{diag}(X_{s-}) \sigma(s-, X_{s-}) dW_s \\ &+ \int_{(0, t]} \int_E \left(u(s-, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) - u(s-, X_{s-}) \right) \nu(ds, dz) \\ &= u(0, x) + \int_0^t \left(\partial_t u(s-, X_{s-}) + \mathcal{A}^*[u](s-, X_{s-}) \right) ds + M_t \\ &+ \int_0^t \int_E \lambda(s-, X_{s-}) \left(h \circ u(s-, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) - u(s-, X_{s-}) \right) \gamma(dz) ds \\ &- \int_{(0, t]} \int_E \left(h \circ u(s-, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) - u(s-, (\mathbf{I}_{d \times d} + \text{diag}(z))X_{s-}) \right) \nu(ds, dz) \\ &= u(0, x) + \int_0^t r(s, X_s) u(s, X_s) ds + M_t - \int_{(0, t]} (h(u(s, X_s)) - u(s, X_s)) dY_s \\ &= V_0 + \int_0^t r(s, X_s) V_s ds + M_t - \int_{(0, t]} (h(V_s) - V_s) dY_s. \end{aligned}$$

³³With regard to the following proof we also refer to [BJR05] (see, e.g., the proof of Proposition 3.1 therein) for similar arguments.

³⁴For the local martingale property we refer, e.g., to Corollary VIII.C4 in [Bré81].

³⁵see, e.g., Theorem II.33 in [Pro04].

Thus, since by (1.26) it holds $1/B_t = \exp\{-\int_0^t r(s, X_s)ds\}$, thanks to Itô's product rule³⁵ we have

$$\frac{V_t}{B_t} = V_0 + (V_- \bullet \frac{1}{B})_s + \int_{(0,t]} \frac{1}{B_{s-}} dV_s = V_0 + \int_{(0,t]} \frac{1}{B_s} dM_s - \int_{(0,t]} \frac{h(V_s) - V_s}{B_s} dY_s.$$

Hence, due to the martingale assumption, it follows that the process $\frac{V}{B} + \int_{(0, \cdot]} \frac{h(V_s) - V_s}{B_s} dY_s$ is a \mathbb{Q} -martingale; since $V_T = g(X_T)$ we conclude in turn that

$$\frac{V_t}{B_t} = \mathbb{E}_t^{\mathbb{Q}} \left[\frac{g(X_T)}{B_T} + \int_{(t,T]} \frac{h(V_s) - V_s}{B_s} dY_s \right] \quad \text{for } t \in [0, T]. \quad (1.29)$$

For $t \in [0, T]$ we consider the $[t, \infty]$ -valued stopping time $\tau^t \triangleq \inf\{s \in [t, T] : \Delta Y_s \neq 0\}$; with the help of (1.29) and the tower property of conditional expectation we obtain the identity

$$\begin{aligned} \frac{V_t}{B_t} &= \mathbb{E}_t^{\mathbb{Q}} \left[\mathbb{E}_{\tau^t}^{\mathbb{Q}} \left[\frac{g(X_T)}{B_T} + \int_{(t,T]} \frac{h(V_s) - V_s}{B_s} dY_s \right] \right] \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[\mathbb{E}_{\tau^t}^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau^t > T\}} \frac{g(X_T)}{B_T} \right] \right. \\ &\quad \left. + \mathbf{1}_{\{\tau^t \leq T\}} \left(\mathbb{E}_{\tau^t}^{\mathbb{Q}} \left[\frac{g(X_T)}{B_T} + \int_{(\tau^t, T]} \frac{h(V_s) - V_s}{B_s} dY_s \right] + \frac{h(V_{\tau^t}) - V_{\tau^t}}{B_{\tau^t}} \right) \right] \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau^t > T\}} \frac{g(X_T)}{B_T} + \mathbf{1}_{\{\tau^t \leq T\}} \frac{h(V_{\tau^t})}{B_{\tau^t}} \right]. \end{aligned} \quad (1.30)$$

Since on the event $\{\tau > t\}$ we have $\tau = \tau^t$, (1.30) implies (1.25). \square

Note that the approach followed by the preceding proof is standard given the literature³⁶ as it uses Feynman-Kač type arguments to link a classical solution of (1.28) to the risk-neutral pricing formula (1.25).

Furthermore, note that, similarly as in other credit risk valuation problems with recovery of mark-to-market value, the pricing formula (1.25) and the corresponding pricing equation (1.28) are inherently implicit, with the price to be determined appearing inside the nonlinearity that represents the recovery value (see, e.g., [CG15] or [KL16] and the references therein) – with the additional complication that jumps at default imply that this term is also nonlocal.

1.4.2 Branching Diffusion with Jumps Approach

In order to solve the pricing equation (1.28), we follow §5 in [Hen12a] and approximate the recovery function h by a polynomial. Thus we assume without loss that $\|g\|_{L^\infty} \leq 1$ (since g is bounded, this can always be achieved by appropriate rescaling; see (1.28)) and approximate the recovery value function by a polynomial using $v^\pm \approx \sum_{m=0}^M \alpha_m^\pm v^m$, $v \in [-1, 1]$; see Figure 1.4 for illustration. We obtain the PDE

$$\begin{aligned} \partial_t u(t, x) + \mathcal{A}^*[u](t, x) - [r(t, x) + \lambda(t, x)]u(t, x) \\ + \int_E \sum_{m=0}^M [R\alpha_m^+ - \alpha_m^-] \lambda(t, x) u(t, x + \text{diag}(\xi)x)^m \gamma(d\xi) = 0, \end{aligned} \quad (1.31)$$

³⁶see, e.g., §4.8 in [Eth02], Chapter 6 in [Shr04] or [BJR05].

subject to the terminal condition $u(T, x) = g(x)$.




Figure 1.4: Approximation of h by a polynomial (parameters as in Tables 1.5/1.6).

Hence, setting $\mathcal{I} \triangleq \{(\ell, m) \in [0:M]^2 : \ell = 0 \text{ or } m = 0\}$ we can rewrite (1.31) as

$$\partial_t u(t, x) + \mathcal{A}^*[u](t, x) + \int_E \sum_{i \in \mathcal{I}} c_i(t, x) u(t, x)^{i_1} u(t, x + \text{diag}(\xi)x)^{i_2} \gamma(d\xi) = 0, \quad (1.32)$$

$$u(T, x) = g(x),$$

where the coefficients $c_{(\ell, m)}$, $(\ell, m) \in \mathcal{I}$, are given by

$$\begin{aligned} c_{(0, m)}(t, x) &\triangleq [R\alpha_m^+ - \alpha_m^-] \lambda(t, x) \quad \text{for } m \in [0:M] \\ c_{(1, 0)}(t, x) &\triangleq -[r(t, x) + \lambda(t, x)] \quad \text{and} \quad c_{(\ell, 0)}(t, x) \triangleq 0 \quad \text{for } \ell \in [2:M]. \end{aligned} \quad (1.33)$$

Since (1.32) is a special case of (PDE), the branching diffusion with jumps approach developed in this chapter can be applied to compute its solution via $u(0, x) = \mathbb{E}[\Psi^{0, x}]$ with

$$\begin{aligned} \Psi^{0, x} &\triangleq \prod_{k \in \mathcal{K}} \frac{g(X_T^{(k)})}{F(T - T^{(k-)})} \times \prod_{\substack{k \in \bar{\mathcal{K}} \setminus \mathcal{K}, \\ I^{(k)} = (1, 0)}} \frac{r(T^{(k)}, X_{T^{(k)}}^{(k)}) + \lambda(T^{(k)}, X_{T^{(k)}}^{(k)})}{\rho(T^{(k)} - T^{(k-)}) p_{(1, 0)}} \\ &\times \prod_{m=0}^M \prod_{\substack{k \in \bar{\mathcal{K}} \setminus \mathcal{K}, \\ I^{(k)} = (0, m)}} \frac{[R\alpha_m^+ - \alpha_m^-] \lambda(T^{(k)}, X_{T^{(k)}}^{(k)})}{\rho(T^{(k)} - T^{(k-)}) p_{(0, m)}}. \end{aligned}$$

Here $X^{(k)} \triangleq X^{k, 0, x}$, $k \in \bar{\mathcal{K}} \triangleq \bar{\mathcal{K}}_0$, is a branching diffusion with jumps as specified in Section 1.1, where the dynamics of each individual particle $k \in \bar{\mathcal{K}}$ are given by

$$dX_t^{(k)} = \text{diag}(X_{t-}^{(k)}) \left[\mu(t, X_{t-}^{(k)}) dt + \sigma(t, X_{t-}^{(k)}) dW_t \right], \quad t \in [T^{(k-)}, T^{(k)}],$$

with initial conditions

$$X_0^{(1)} = x \quad \text{and} \quad X_{T^{(k-)}}^{(k)} \triangleq \begin{cases} X_{T^{(k-)}}^{(k-)} & \text{if } I^{(k)} = (1, 0), \\ X_{T^{(k-)}}^{(k-)} + \text{diag}(\Delta^{(k-)})X_{T^{(k-)}}^{(k-)} & \text{else.} \end{cases}$$

1.4.3 Numerical Illustration

While our setup and Monte Carlo methodology allow for general diffusion dynamics and arbitrary dimensionality in the underlying risky assets (see also Section 1.4.4 below), for illustration we first use a baseline Black-Scholes-Merton model [BS73, Mer73], enhanced by the SIB's credit risk with constant default intensity and devaluations captured by the model of Kou [Kou02]. Thus we take Y as a homogeneous Poisson process with intensity $\lambda \geq 0$; the riskless rate $r \in \mathbb{R}$ is constant; and g has a single underlying ($d = 1$) with volatility $\sigma > 0$. Jumps at default are such that $-\log(1 + \Delta_1)$ is exponentially distributed with parameter $\eta \geq 0$.³⁷ We consider two financial positions representing a shifted put and shifted discount call, respectively, i.e.

$$g_{\text{short}}(X_T) \triangleq (K - X_T)^+ - L \quad \text{and} \quad g_{\text{long}}(X_T) = L - (K - X_T)^+.$$

Figure 1.5 provides an illustration of the corresponding payoffs.




Figure 1.5: Payoff profiles g_{short} and g_{long} with $K = 2$ and $L = 1$.

We compute their valuation at $(t, x) = (0, x)$. To quantify the impact of systemic risk, we compare the valuations of the financial positions g_{short} and g_{long} in three benchmark scenarios that are identical in all respects, except the choice of counterparty:

- *SIB counterparty*: The counterparty is systemically important, and their default triggers devaluations in risky asset prices (see above).

³⁷In the specification of §2 in [Kou02], this corresponds to the choice $p \triangleq 0$ and $\eta_2 \triangleq 1$.

- *Non-SIB counterparty*: The counterparty is non-systemic, but otherwise identical to the SIB; in particular, it is defaultable with the same credit risk characteristics as the SIB.
- *Default-free counterparty*: There is no counterparty credit risk.³⁸

In all three scenarios, the SIB is part of the model and will cause devaluations upon default; the scenarios thus differ only in the choice of counterparty and the resulting wrong- or right-way risk; the relevant pricing PDEs are reported in Appendix C.

Our implementation of the branching diffusion with jumps is based on the parameters specified in Table 1.6. In contrast to Section 1.3, where we deliberately employ a non-accelerated crude Monte Carlo method for performance analysis, here we exploit standard techniques for variance reduction³⁹ such as common random numbers, control variates and parallelization: Since all three pricing PDEs can be solved by means of the discussed branching methodology (see Section 1.4.2 and Appendix C), and prices can be computed explicitly in case of a default-free counterparty,³⁸ we take the branching estimator of the default-free price as a control variate. The latter is optimized by taking correlations between the estimator of the default-free price and the other two estimators into account; the corresponding covariance matrices are estimated independently from the main simulation based on 10^4 branching (jump-)diffusion samples.⁴⁰

The resulting three price estimators are simulated using common random numbers; in particular, we make use of common samples of branching ancestries, jump parameters and increments of the particles' driving noise (i.e. Brownian motions and, where applicable, jump times and jump sizes of compound Poisson processes). Note that the dynamics of each particle can be simulated exactly as each underlying (geometric) SDE admits an explicit solution suitable for this purpose. Each simulation run for the price estimators is parallelized over a pool of 44 workers using MATLAB's `parpool` and `parfor` commands; the corresponding loop simulates a total of $8.8 \cdot 10^6$ i.i.d. branching (jump-)diffusion samples using 88 simulation blocks with 10^5 samples drawn per block and 44 of these blocks running in parallel.⁴¹

The relevant model and simulation parameters are reported in Tables 1.5 and 1.6; note in particular that the expected devaluation upon the SIB's default is -50% ; see (4) in [Kou02].

Coefficient	T	r	σ	R	η	x	K	L
Value	1	0.5%	25%	40%	1	1	2	1

Table 1.5: Market coefficients.

³⁸In this scenario, the pricing PDE becomes linear (see Appendix C) and can be solved in closed form; see, e.g., [Kou02].

³⁹For details on these (and other) variance reduction principles we refer, e.g., to Chapter 4 in [Gla03] or §3.3/4 in [KKK10].

⁴⁰This approach is standard given the literature; see, e.g., §3.3.2.1 in [KKK10]; we also refer to §4.1.1/3 in [Gla03].

⁴¹We use MATLAB R2019b [Mat19]; OS: Linux Ubuntu 16.04.5 LTS.

For details on `parpool` and `parfor` we refer to the online documentation manual; see <https://de.mathworks.com/help/parallel-computing/parpool.html> and <https://de.mathworks.com/help/parallel-computing/parfor.html>, respectively (both accessed: 2020-05-31, 1940).

Hardware: $2 \times$ Intel Xeon CPU E5-2699 v4 @ 2.20 GHz (44 cores in total) with 755 GiB RAM capacity.

Parameter	Value
N	$8.8 \cdot 10^6$
$\text{Law}(\tau)$	$\Gamma(\kappa, \theta)$ with $\kappa = 0.5$ and $\theta = 2.5$
M	4
$(\alpha_0^\pm, \dots, \alpha_4^\pm)$	$(0.06, \pm 0.50, 0.82, 0.00, -0.41)$
$p(0,0)$	q_0
$p(\ell,0)$	$q_{\text{loc}} \cdot q_\ell$ for $\ell \in [1:M]$
$p(0,m)$	$(1 - q_{\text{loc}}) \cdot q_m$ for $m \in [1:M]$
q_{loc}	$\frac{ r+\lambda }{ r+\lambda + \lambda \cdot \sum_{m=0}^M R\alpha_m^+ - \alpha_m^- }$
q_m	$\frac{ c(1,0) ^2 \cdot \mathbf{1}_{\{m=1\}} + c(0,m) \cdot \sum_{m=0}^M c(0,m) }{ c(1,0) ^2 + (\sum_{m'=0}^M c(0,m'))^2}$ for $m \in [0:M]$

Table 1.6: Simulation parameters.^{22, 42}

Figures 1.6 through 1.8 display our simulation results for different specifications of the default intensity. Note that the average running time of a single simulation run is approximately 102.2 sec; for default intensities λ up to 5% it is approximately 95.7 sec.⁴³

Figure 1.6 illustrates the impact of systemic interaction for the financial position g_{short} : Devaluations imply a positive correlation between the SIB's default events and underperforming risky asset prices, causing significant wrong-way risk for short positions. While this is qualitatively apparent, the quantitative size of this effect, in particular relative to the non-SIB counterparty, is remarkable.⁴⁴

In Figure 1.7 we thus decompose the implied SIB spread (i.e., the difference between the value of an otherwise identical financial position with a default-free counterparty and that with an SIB counterparty) into (i) a pure credit risk component (blue), which we identify with the spread between the non-systemic defaultable counterparty and the default-free one, and (ii) a systemic risk component (red), which is present only due to the counterparty's systemic importance. It is apparent that systemic risk is the main driver of the spread, accounting for 80-85% of the total spread across realistic default intensities.

Finally, Figure 1.8 demonstrates that the effect is reversed for long positions, i.e. for g_{long} ; we might refer to this as right-way risk, i.e. negative correlation between the counterparty's default and the value of the financial position. Technically, the systemic risk component turns negative, and the spread becomes significantly smaller than for a non-systemic counterparty. \diamond

⁴²The values of α_m^\pm , $m \in [0:M]$, are computed with MATLAB's pre-implemented routine `polyfit.m` which aims for a best least-squares fit; for details we refer to the online documentation manual <https://de.mathworks.com/help/matlab/ref/polyfit.html> (accessed: 2020-05-31, 1941); note that the obtained values are very close to the polynomial coefficients used in §5 in [Hen12a].

Moreover, the choice of the probability weights p_i , $i \in \mathcal{I}$, is inspired by §4.3 in [Hen12a].

⁴³The stated average running times differ from the ones in Section 1.3 in that simulation runs are *not* repeated; rather we take the average of running times measured for various values of λ and for g_{short} and g_{long} together. Still, as in Footnote 25, all sampled running times are wall-clock times, including the time required by the branching Monte Carlo algorithm itself, its initialization and the estimation of covariances for the optimized control variate.

⁴⁴For further details see also our remarks in Section 1.4.4 below.




Figure 1.6: Valuation of g_{short} as a function of the counterparty's default intensity (with approximate 99%-confidence bands).

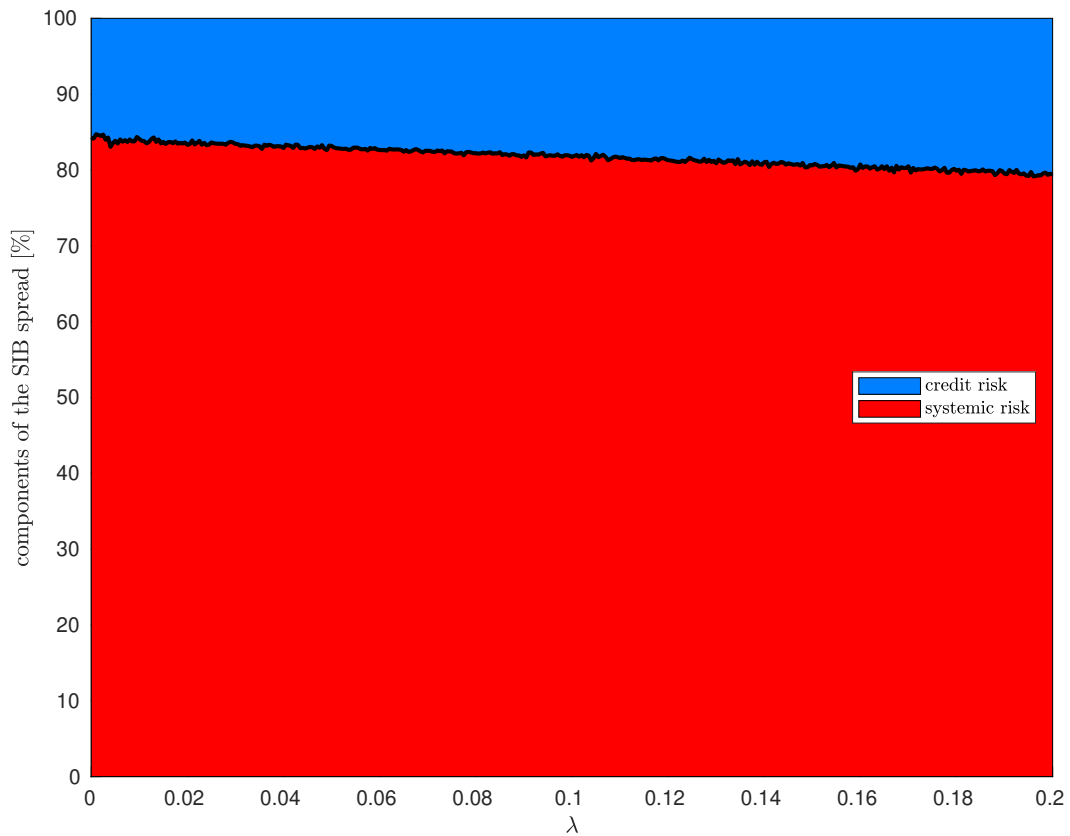


Figure 1.7: Decomposition of the SIB spread into credit risk (blue) and systemic risk (red).




Figure 1.8: Valuation of g_{long} as a function of the counterparty's default intensity (with approximate 99%-confidence bands).

1.4.4 Another High-Dimensional Example

The simulation study in Section 1.3 has already illustrated that the branching Monte Carlo algorithm is able to cope with high-dimensional problems. To further corroborate this result, we adopt the setup of Section 1.4.3 and exemplarily modify the payoff profile g_{short} into a multivariate one; more precisely, we consider a financial position g_{basket} entered into with a defaultable SIB counterparty whose time- T payoff is based on a basket of $d = 50$ underlyings and given by

$$g_{\text{basket}}(X_T) \triangleq \left(K - \frac{1}{d} \cdot \mathbf{1}_d^\top X_T \right)^+ - L = \left(K - \frac{1}{d} \sum_{i=1}^d X_i(T) \right)^+ - L.$$

Mutatis mutandis, the relevant market and simulation parameters are chosen as in Section 1.4.3; the underlyings' price process is driven by a standard (uncorrelated) d -dimensional Brownian motion; and the distribution of jumps at default of the SIB is assumed to be the product measure of the univariate one in the previous section.

We compute the valuation of this financial position at $(t, x) = (0, \mathbf{1}_d)$ for various values of λ and R by means of crude Monte Carlo with parallelization as outlined in Section 1.4.3.⁴¹ Our simulation results are displayed in Figure 1.9; for corresponding estimates of the (relative) standard deviation see Figure 1.10. The average running time of a single simulation run is approximately 103.7 sec.⁴⁵

⁴⁵Similarly as in Footnote 43, we report the average of (wall-clock) running times measured for various values of λ and R .




Figure 1.9: Valuation of g_{basket} for $(\lambda^\dagger, R) \in [0, 0.1] \times [0, 1]$.

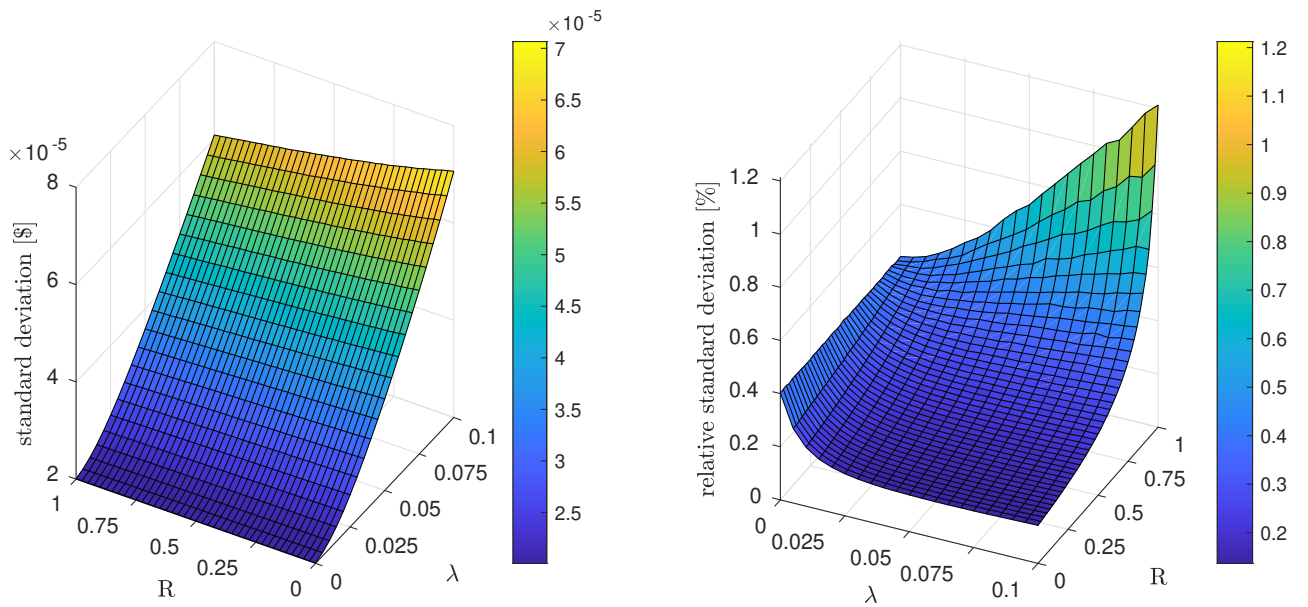


Figure 1.10: Estimated (relative) standard deviation of the branching estimator.⁴⁶

⁴⁶The standard deviation of the Monte Carlo estimator is estimated exactly as in Section 1.3; we estimate its relative counterpart by dividing an estimate of the standard deviation by the absolute value of the corresponding valuation estimate.

The preceding results warrant a brief comment that complements our remarks in Section 1.4.3: Note that, disregarding the fact that the financial position g_{basket} is subject to credit risk, it consists of a long position in a basket put option on the underlying risky assets with strike K and a short position in a zero coupon bond with principal L ; it can thus be regarded as a multivariate analogue of g_{short} , and hence it implements a bearish strategy given the underlying price dynamics (1.27).

Since there is a positive correlation between the counterparty's default events and underperforming risky asset prices, the holder of g_{basket} faces a wrong-way risk the quantitative size of which is, in turn, negatively correlated with the size of the recovery rate R :

For high values of R this risk is less pronounced and valuations either remain almost constant or decline only moderately for increasing values of λ as the holder is rather immune to the impact of a default before maturity.⁴⁷

However, for more realistic low and moderate values of R a rise of λ leads to a considerably tumbling valuation of g_{basket} ; although the event of the SIB's default becomes more likely and, in case of occurrence, simultaneously triggers devaluations in the underlying risky asset prices such that corresponding short positions appreciate in value, the holder of g_{basket} only retrieves a comparably small recovery value if the position's post-default valuation is positive.

As to the standard deviation we observe that it stays below 10^{-4} which is, in terms of scale, in line with what one expects theoretically given an overall number of $N = 8.8 \cdot 10^6$ Monte Carlo samples. Due to the fact that the ratios $|c_i|/p_i$, $i \in \mathcal{I}$, increase as λ increases (see (1.33) and Figure 1.11), for all fixed values of R the standard deviation does so as well; see also Proposition 1.2.3. \diamond

Of course, the preceding comments apply to the financial position g_{short} in equal measure; note that the only new feature in this section is the (artificially introduced) high-dimensional state space. We are aware that the present multivariate price dynamics are not suitable to model market prices. However, the model can readily be extended to include realistic features, such as, e.g., both common and idiosyncratic jumps of asset prices or correlations, but as such extensions are beyond the scope of this chapter, for further details we refer, e.g., to §10 in [Kou07] and the references therein.

Nevertheless, to conclude this section, observe that the results corroborate the findings of Section 1.3 and showcase that our branching Monte Carlo algorithm is able to solve high-dimensional nonlocal PDEs.

⁴⁷As we assume *fractional recovery of post-default mark-to-market value*, note that for $R = 1$ the counterparty's default amounts to an early termination of the contract.

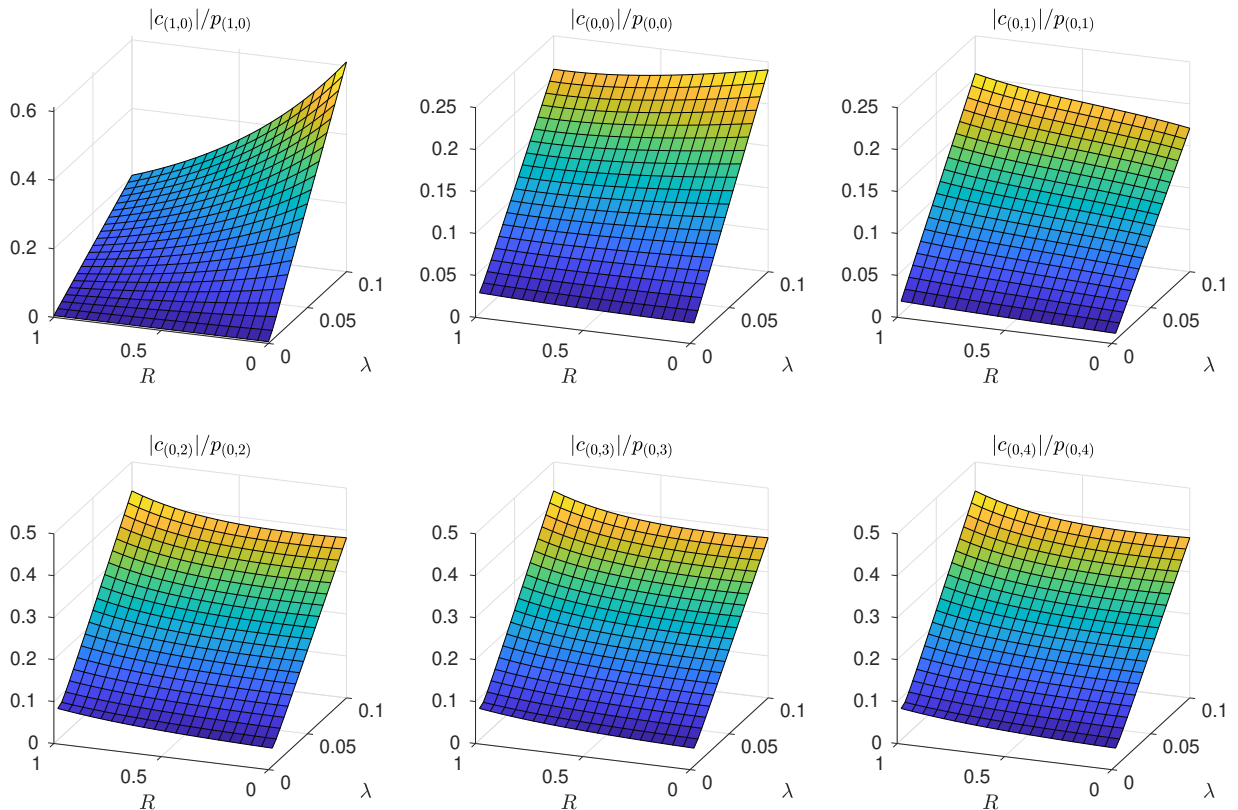


Figure 1.11: Ratios of branching coefficients and offspring probability weights.

Conclusion and Outlook

To conclude, this chapter has developed a branching diffusion with jumps approach to solving parabolic PDEs with nonlocal analytic nonlinearities. We have derived a stochastic representation of a classical solution, and, conversely, we have proved that this representation is a viscosity solution of the considered PDE. Furthermore, we have showcased the performance of the resulting non-nested Monte Carlo methodology in dimensions up to 100, and we have demonstrated how it applies to the valuation of financial positions with systemic counterparties and mark-to-market recovery.

For future research, it would be interesting to investigate a unified approach to *local-nonlocal semilinear*, *fully nonlinear* or, generally speaking, *higher order* PDEs relying on branching processes. For recent developments we refer, e.g., to [HT18] or [War18]. Albeit the branching methodology strongly relies on the polynomial structure of the nonlinearity (see, e.g., Remark 1.1.5 or [HOT⁺19]), we think that it is nevertheless worth attempting to generalize it to other types of nonlinearities.

From a numerical point of view, it would be of interest to optimize branching parameters such as the lifetime or the offspring distribution in the present setup to reduce the variance of the resulting estimator and thus the overall computation complexity. For existing result see, e.g., §4.3 in [Hen12a] for optimized probability weights p_i (offspring distribution), or §5.1 and the appendix in [HOT⁺19] for the choice of the density ρ (lifetime distribution), a discussion of the complexity and an importance sampling scheme.

Moreover, several extensions of the above pricing model are conceivable within the setup of this chapter, including stochastic recovery rates, shocks to state variables (e.g., volatilities), idiosyncratic jumps in asset prices, and margin periods of risk. With regard to modeling jumps we refer, e.g., to §10 in [Kou07]; for margin periods of risk (delays) see, e.g., Remark 5.1 in [Hen12a].

Chapter 2

Continuous-Time Finite-State Mean Field Games with Common Noise

The present chapter and Appendix D/E are a revised and partly extended version of the article [BHS20b] that has already been made publicly available as a preprint and submitted for publication in a scientific journal. ◇

Since the seminal contributions of Lasry and Lions [LL07] and Huang, Malhamé and Caines [HMC06] (see also [Car13] and the references therein), mean field games have become an active field of mathematical research with a wide range of applications, including economics ([CFS15], [CDL17], [KM17], [Nut18], [EIL20], [GS20]), socio-economics ([GVW14]), finance ([LLLL16], [CJ19b]), epidemiology ([LT15], [DGG17], [EHT20]), computer science ([KB16]), crowd dynamics ([LW11]) and cryptocurrencies ([BBLL20]); see also the overview article [GLL11] and the monograph [CD18a].

Mean field games constitute a class of dynamic, multi-player stochastic differential games with identical agents. The key characteristic of the mean field approach is that (i) the payoff and state dynamics of each agent depend on other agents' decisions only through an aggregate statistic (typically, the aggregate distribution of states); and (ii) no individual agent's actions can change the aggregate outcome. Thus, in solving an individual agent's optimization problem, the feedback effect of his own actions on the aggregate outcome can be discarded, breaking the notorious vicious circle ("the optimal strategy depends on the aggregate outcome, which depends on the strategy, which depends ..."). This significantly facilitates the identification of rational expectations equilibria. A standard assumption that further simplifies the analysis is that randomness is idiosyncratic (equivalently, there is no common noise), i.e. that the random variables appearing in one agent's optimization are independent of those in any other's. As a result, all randomness is "averaged out" in the aggregation of individual decisions, and the equilibrium dynamics of the aggregate distribution are deterministic.

In the literature, mean field games are most often studied in settings with a continuous state space and deterministic or diffusive dynamics, i.e. stochastic differential equations (SDEs) driven by Brownian motion. The corresponding dynamic programming equations thus become parabolic

partial differential equations, and the aggregate dynamics are represented by a flow of Borel probability measures; see, e.g., the monographs [BFY13] and [CD18a] and the references therein. Formally, the mean field game is typically formulated in terms of a controlled McKean-Vlasov SDE, where the coefficients depend on the current state and control and the distribution of the solution; intuitively, these McKean-Vlasov dynamics codify the dynamics that pertain to a representative agent. The mathematical link to N -player games is subsequently made through suitable propagation of chaos results in the mean field limit $N \rightarrow \infty$; see, e.g., [Lac15], [Fis17], [CF18], [Lac18] and [DLR19]. In this context, the analysis of McKean-Vlasov SDEs has also seen significant progress recently; see, e.g., [CD13], [BP19], [CP19b] and [MP19a]. In the presence of common noise, i.e. sources of risk that affect all agents and do not average out in the mean field limit, the mathematical analysis becomes even more involved as the dynamics of the aggregate distribution become stochastic, leading to conditional McKean-Vlasov dynamics; see, e.g., [Ahu16], [CDL16], [CW17] and [PW17]. We refer to the monograph [CD18b] for background and further references on continuous-state mean field games with common noise; for discrete-time Markov decision processes with mean field interaction and common noise see also [MP19b].

There is also a strand of literature on mean field games with finite state spaces, including [GMS10], [Gué11], [GMS13], [BHK14], [Gué15], [BC18], [CP19a], [DGG19], [CF20], [Neu20] as well as [CD18a, §7.2]. In a recent article, [CW18b] provide an extension of [GMS13] to mean field interactions that occur not only through the agents' states, but also through their controls. Furthermore, [CW16] and [CW18a] address mean field games including dominating players; see also [Wan19] for a corresponding comprehensive account.

To the best of our knowledge, however, to date there has been no extension of these results to settings that include common noise; in the context of finite-state mean field games, we are only aware of two contributions that include common noise (both via the master equation and with a different focus or setting): [BLL19] formulate the master equation for finite-state mean field games with common noise to include (random) jumps of the agents' states at either deterministic or random jump times, and [BCCD19] add a common Gaussian noise to the aggregate distribution dynamics and study the resulting master equation.

In this chapter, we set up a mathematical framework for finite-state mean field games with common noise.⁴⁸ Our setup extends that of Gomes, Mohr and Souza [GMS13] and Cecchin and Fischer [CF20] by common noise events occurring at fixed points in time. We provide a rigorous formulation of the underlying stochastic dynamics, and we establish a verification theorem for the optimal strategy and an aggregation theorem to determine the resulting aggregate distribution. This leads to a characterization of the mean field equilibrium in terms of a system of (random) forward-backward ordinary differential equations for which we also prove a corresponding existence result. The key insight is that, after conditioning on common noise configurations, we obtain classical piecewise dynamics subject to jump conditions at common noise times.

The remainder of this chapter is organized as follows: In Section 2.1 we set up the mathematical model, provide a probabilistic construction of the state dynamics, and formulate the agent's optimization problem. In Section 2.2 we state the dynamic programming equation and establish

⁴⁸We wish to point out that the focus of this chapter is not on the mean field limit of multi-player games; rather, we directly investigate the mean field equilibrium via the corresponding McKean-Vlasov dynamics (see also Remark 2.3.1 and [CDL13] in that context).

a verification theorem for the agent's optimization, given an *ex ante* aggregate distribution (Theorem 2.2.4). Section 2.3 provides the dynamics of the *ex post* distribution (Theorem 2.3.4) and, on that basis, a system of random forward-backward ODEs for the mean field equilibrium (Definition 2.3.5) as well as a corresponding existence result (Theorem 2.3.6). Finally, in Section 2.4 we showcase our results in three stylized applications: agricultural production, infection control and evacuation.

Three parts of the appendix complement the exposition of this chapter: Appendix A provides some auxiliary technical results and a brief outline of the theory of Carathéodory solutions of ODEs; Appendix D contains the proof of the existence result for a mean field equilibrium (Theorem 2.3.6), relevant assumptions and auxiliary results; and, for the sake of completeness, Appendix E provides the master equation within our framework.

2.1 Mean Field Model

We first provide an *informal* description of the individual agents' state dynamics, optimization problem, and the resulting mean field equilibrium. The agent's state process $X = \{X_t\}$ takes values in the finite set \mathbb{S} . Between common noise events, transitions from state i to state j occur with intensity $Q^{ij}(t, W_t, M_t, \nu_t)$, where W_t represents the common noise events that have occurred up to time t ; M_t is the time- t aggregate distribution of agents; and ν_t represents the agent's control. In addition, upon the realization of a common noise event W_k at time T_k , the state jumps from X_{T_k-} to $X_{T_k} = J^{X_{T_k-}}(T_k, W_{T_k}, M_{T_k-})$. Given these dynamics, the agent aims to maximize

$$\mathbb{E}^\nu \left[\int_0^T \psi^{X_t}(t, W_t, M_t, \nu_t) dt + \Psi^{X_T}(W_T, M_T) \right]$$

where ψ and Ψ are suitable reward functions and the aggregate distribution process $M = \{M_t\}$ is given by

$$M_t \triangleq \mu(t, W_t) \quad \text{for } t \in [0, T].$$

Here the function μ represents the aggregate distribution of states as a function of the common noise factors. We obtain a rational expectations equilibrium by determining μ in such a way that the representative agent's *ex ante* expectations coincide with the *ex post* aggregate distribution resulting from all agents' optimal decisions, i.e.

$$\mathbb{P}^{\hat{\nu}}(X_t \in \cdot \mid W_t) = \hat{\mu}(t, W_t) \quad \text{for all } t \in [0, T],$$

where $\hat{\nu}$ and $\hat{\mu}$ denote the equilibrium strategy and the equilibrium aggregate distribution. In the remainder of this section, we provide a rigorous mathematical formulation of this model.

2.1.1 Probabilistic Setting and Common Noise

Throughout, we fix a time horizon $T > 0$ and a finite set \mathbb{W} and work on a probability space $(\Omega, \mathfrak{A}, \mathbb{P})$ that carries a finite sequence W_1, \dots, W_n of i.i.d. random variables that are uniformly distributed on \mathbb{W} .⁴⁹ We refer to W_1, \dots, W_n as *common noise factors* and to \mathbb{P} as the *reference probability*. The common noise factor W_k is revealed at time T_k , where

$$0 \triangleq T_0 < T_1 < T_2 < \dots < T_n < T_{n+1} \triangleq T.$$

The filtration $\mathfrak{G} = \{\mathfrak{G}_t\}_{t \in [0, T]}$ generated by common noise events is given by

$$\mathfrak{G}_t \triangleq \sigma(W_k : k \in [1:n], T_k \leq t) \vee \mathfrak{N} \quad \text{for } t \in [0, T],$$

where \mathfrak{N} denotes the system of \mathbb{P} -null sets. For later reference, we note that \mathfrak{G} is piecewise constant in the sense that $\mathfrak{G}_t = \mathfrak{G}_{T_k}$ for all $t \in [T_k, T_{k+1})$ and $k \in [0:n]$, where for $0 \leq s \leq t \leq T$ we set

$$[s, t) \triangleq [s, t) \quad \text{if } t < T \quad \text{and} \quad [s, T) \triangleq [s, T).$$




Figure 2.1: Common noise factors and common noise times ($n = 3$).

For each configuration of common noise factors $w \in \mathbb{W}^n$ we write

$$w_t \triangleq (w_1, \dots, w_k) \quad \text{for } t \in [T_k, T_{k+1}), \quad k \in [0:n].$$

With this convention, $W = \{W_t\}_{t \in [0, T]}$ represents a piecewise constant, \mathfrak{G} -adapted process.

Definition 2.1.1 (Non-Anticipative Function). A function $f : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^m$ is said to be *non-anticipative* if for all $t \in [0, T]$

$$f(t, w) = f(t, \bar{w}) \quad \text{whenever } w, \bar{w} \in \mathbb{W}^n \text{ are such that } w_t = \bar{w}_t. \quad \diamond$$

With a slight abuse of notation, if $f : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^m$ is non-anticipative, we write

$$f(t, w_t) \triangleq f(t, w) \quad \text{for } w \in \mathbb{W}^n, \quad t \in [0, T].$$

⁴⁹While the common noise factors are i.i.d. uniformly distributed under \mathbb{P} , the distribution of W_1, \dots, W_n in the agent's optimization problem can be modeled arbitrarily via the functions $\kappa_1, \dots, \kappa_n$ introduced below; see also Lemma 2.1.3.

2.1.2 Optimization Problem

The agent's state and action spaces are given by

$$\mathbb{S} \triangleq [1:d] \quad \text{and} \quad \mathbb{U} \subseteq \mathbb{R}^e, \quad \text{where } d, e \in \mathbb{N} \text{ and } \mathbb{U} \neq \emptyset,$$

and we identify the space of aggregate distributions on \mathbb{S} with the space of probability vectors⁵⁰

$$\mathbb{M} \triangleq \text{Prob}(\mathbb{S}) \triangleq \left\{ m \in [0, \infty)^{1 \times d} : \sum_{i=1}^d m^i = 1 \right\}.$$

We further suppose that $(\Omega, \mathfrak{A}, \mathbb{P})$ supports, for each $i, j \in \mathbb{S}$, $i \neq j$, a standard (i.e., unit intensity) Poisson process $N^{ij} = \{N_t^{ij}\}_{t \in [0, T]}$ and an \mathbb{S} -valued random variable X_0 such that

$$X_0 \quad \text{and} \quad N^{ij}, \quad i, j \in \mathbb{S}, \quad i \neq j, \quad \text{and} \quad W_1, \dots, W_n \quad \text{are independent.}$$

The corresponding full filtration $\mathfrak{F} = \{\mathfrak{F}_t\}_{t \in [0, T]}$ is given by

$$\mathfrak{F}_t \triangleq \sigma(X_0, W_s, N_s^{ij} : s \in [0, t]; i, j \in \mathbb{S}, i \neq j) \vee \mathfrak{N} \quad \text{for } t \in [0, T].$$

Note that $\mathfrak{G}_t \subseteq \mathfrak{F}_t$ for all $t \in [0, T]$, that both \mathfrak{G} and \mathfrak{F} satisfy the usual conditions, and that N^{ij} is a standard $(\mathfrak{F}, \mathbb{P})$ -Poisson process for $i, j \in \mathbb{S}$, $i \neq j$.

The agent's *optimization problem* reads⁵¹

$$\mathbb{E}^\nu \left[\int_0^T \psi^{X_t}(t, W_t, M_t, \nu_t) dt + \Psi^{X_T}(W_T, M_T) \right] \xrightarrow{\nu \in \mathcal{A}} \max! \quad (\text{P}_\mu)$$

where the class of *admissible strategies* for (P_μ) is given by the set of closed-loop controls

$$\begin{aligned} \mathcal{A} \triangleq \{ \nu : [0, T] \times \mathbb{S}^{[0, T]} \times \mathbb{W}^n \rightarrow \mathbb{U} : \nu \text{ is Borel measurable and} \\ \nu(\cdot, x, \cdot) \text{ is non-anticipative for all } x \in \mathbb{S}^{[0, T]} \}. \end{aligned}$$

Note that \mathcal{A} subsumes the class of Markovian feedback controls considered in, e.g., [GMS13] or [Gué11, Gué15], and that each $\nu \in \mathcal{A}$ canonically induces an \mathfrak{F} -adapted \mathbb{U} -valued process via

$$\nu_t \triangleq \nu(t, X_{(\cdot \wedge t)-}, W_t) \quad \text{for } t \in [0, T].$$

The \mathfrak{G} -adapted, \mathbb{M} -valued *ex ante* aggregate distribution $M = \{M_t\}_{t \in [0, T]}$ is given by

$$M_t \triangleq \mu(t, W_t) \quad \text{for } t \in [0, T];$$

$\mathbb{E}^\nu[\cdot]$ denotes the expectation operator with respect to the probability measure \mathbb{P}^ν given by (see Lemma 2.1.3 below)

⁵⁰The set \mathbb{M} is commonly referred to as *probability simplex* in the literature.

⁵¹For notational simplicity, we write X_t instead of X_{t-} , M_t instead of M_{t-} , etc., where it does not make a difference (e.g., in dt -integrals; see Lemma A.1 for a formal justification). In this context, note by Lemma A.4 that X_- and X are modifications of each other on $[0, T] \setminus \{T_1, \dots, T_n\}$, where, here and in the following, we set $X_{0-} \triangleq X_0$.

$$\begin{aligned} \frac{d\mathbb{P}^\nu}{d\mathbb{P}} &= \prod_{\substack{i,j \in \mathbb{S}, \\ i \neq j}} \left(\exp \left\{ \int_0^T (1 - Q^{ij}(t, W_t, M_t, \nu_t)) dt \right\} \cdot \prod_{\substack{t \in (0, T], \\ \Delta N_t^{ij} \neq 0}} Q^{ij}(t, W_t, M_t, \nu_t) \right) \\ &\quad \times |\mathbb{W}|^n \cdot \prod_{k=1}^n \kappa_k(W_k | W_1, \dots, W_{k-1}, M_{T_k-}); \end{aligned} \quad (2.1)$$

and the agent's state process X is given by

$$dX_t = \sum_{\substack{i,j \in \mathbb{S}, \\ i \neq j}} \mathbf{1}_{\{X_{t-}=i\}} (j - i) dN_t^{ij} \quad \text{for } t \in [T_k, T_{k+1}), \quad k \in [0:n], \quad (2.2)$$

subject to the jump conditions

$$X_{T_k} = J^{X_{T_k-}}(T_k, W_{T_k}, M_{T_k-}) \quad \text{for } k \in [1:n]. \quad (2.3)$$

The coefficient functions in the state dynamics and payoff functional are bounded and Borel measurable functions

$$\begin{aligned} Q &: [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U} \rightarrow \mathbb{R}^{d \times d} \\ J &: [0, T] \times \mathbb{W}^n \times \mathbb{M} \rightarrow \mathbb{S}^d \\ \psi &: [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U} \rightarrow \mathbb{R}^d \\ \Psi &: \mathbb{W}^n \times \mathbb{M} \rightarrow \mathbb{R}^d \end{aligned}$$

such that $Q(\cdot, \cdot, m, u)$, $\psi(\cdot, \cdot, m, u)$ and $J(\cdot, \cdot, m)$ are non-anticipative for all fixed $m \in \mathbb{M}$ and $u \in \mathbb{U}$; Q satisfies the intensity matrix conditions $Q^{ij}(t, w, m, u) \geq 0$, $i, j \in \mathbb{S}$, $i \neq j$ and $\sum_{j \in \mathbb{S}} Q^{ij}(t, w, m, u) = 0$, $i \in \mathbb{S}$, for $(t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$; and for each $k \in [1:n]$ the function

$$\kappa_k : \mathbb{W}^k \times \mathbb{M} \rightarrow [0, 1], \quad (w_k, w_1, \dots, w_{k-1}, m) \mapsto \kappa_k(w_k | w_1, \dots, w_{k-1}, m),$$

is Borel measurable and satisfies $\sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, m) = 1$ for all $w_1, \dots, w_{k-1} \in \mathbb{W}$ and $m \in \mathbb{M}$. Finally, the function μ is non-anticipative and regular in the following sense:

Definition 2.1.2 (Regular Function). A function $f : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^m$ is said to be *regular* if $f(\cdot, w)$ is absolutely continuous on $[T_k, T_{k+1})$ for $k \in [0:n]$ and $w \in \mathbb{W}^n$. \diamond

Note that for a regular function f , the one-sided limits $f(T_k-, w) \triangleq \lim_{t \uparrow T_k} f(t, w)$ exist for every $k \in [1:n]$ and all $w \in \mathbb{W}^n$.⁵²

2.1.3 State Dynamics

In the following, we verify that the preceding construction implies the dynamics described informally above:

⁵²This follows since every absolutely continuous function on $[T_k, T_{k+1})$ is uniformly continuous, and thus possesses a continuous extension to $[T_k, T_{k+1}]$ (for the latter see, e.g., [Kö04, p.97]).

Lemma 2.1.3 (\mathbb{P}^ν -Dynamics). For each admissible strategy $\nu \in \mathcal{A}$, \mathbb{P}^ν is a well-defined probability measure on (Ω, \mathfrak{A}) , absolutely continuous with respect to \mathbb{P} , and satisfies

$$\mathbb{P}^\nu = \mathbb{P} \quad \text{on } \sigma(X_0).$$

Moreover, N^{ij} is a counting process with $(\mathfrak{F}, \mathbb{P}^\nu)$ -intensity $\lambda^{ij} = \{\lambda_t^{ij}\}_{t \in [0, T]}$, where

$$\lambda_t^{ij} \triangleq Q^{ij}(t, W_t, M_t, \nu_t) \quad \text{for } t \in [0, T] \text{ and } i, j \in \mathbb{S}, i \neq j.$$

Finally, for all $k \in [1:n]$ we have

$$\mathbb{P}^\nu(W_k = w_k | \mathfrak{G}_{T_k-}) = \kappa_k(w_k | W_1, \dots, W_{k-1}, M_{T_k-}) \quad \text{for all } w_1, \dots, w_k \in \mathbb{W}$$

and, in particular,

$$\mathbb{P}^{\nu_1} = \mathbb{P}^{\nu_2} \quad \text{on } \mathfrak{G}_T \quad \text{for all admissible strategies } \nu_1, \nu_2 \in \mathcal{A}.$$

Proof. We fix $\nu \in \mathcal{A}$ and split the proof into four steps.

Step 1: \mathbb{P}^ν is well-defined by (2.1). Since N^{ij} is a standard Poisson process under \mathbb{P} , the compensated process \bar{N}^{ij} , $\bar{N}_t^{ij} \triangleq N_t^{ij} - t$ for $t \geq 0$, is an $(\mathfrak{F}, \mathbb{P})$ -martingale for all $i, j \in \mathbb{S}$, $i \neq j$. We define the local $(\mathfrak{F}, \mathbb{P})$ -martingale $\theta^\nu = \{\theta_t^\nu\}_{t \in [0, T]}$ via⁵³

$$\theta_t^\nu \triangleq \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \int_0^t (Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) - 1) d\bar{N}_s^{ij}, \quad t \in [0, T],$$

and observe that the Doléans-Dade exponential $\mathcal{E}[\theta^\nu]$ is a local $(\mathfrak{F}, \mathbb{P})$ -martingale⁵⁴ with

$$\mathcal{E}[\theta^\nu]_t = \prod_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \left(\exp \left\{ \int_0^t (1 - Q^{ij}(s, W_s, \mu(s, W_s), \nu_s)) ds \right\} \cdot \prod_{\substack{s \in (0, t], \\ \Delta N_s^{ij} \neq 0}} Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) \right) \quad (2.4)$$

for $t \in [0, T]$. Next, we define $\vartheta = \{\vartheta_t\}_{t \in [0, T]}$ via

$$\vartheta_t \triangleq \sum_{\substack{k \in [1:n], \\ T_k \leq t}} (|\mathbb{W}| \cdot \kappa_k(W_k | W_1, \dots, W_{k-1}, \mu(T_k-, W_{T_k-})) - 1), \quad t \in [0, T],$$

and note that ϑ is an $(\mathfrak{F}, \mathbb{P})$ -martingale. Indeed, for each $k \in [0:n]$ we have $\vartheta_t = \vartheta_{T_k}$ for $t \in [T_k, T_{k+1})$ and, using the fact that W_k is independent of \mathfrak{F}_{T_k-} and uniformly distributed on \mathbb{W} under \mathbb{P} , it follows that

$$\mathbb{E}[\vartheta_{T_k} | \mathfrak{F}_{T_k-}] = \vartheta_{T_k-} + \mathbb{E}[|\mathbb{W}| \cdot \kappa_k(W_k | W_1, \dots, W_{k-1}, \mu(T_k-, W_{T_k-})) - 1 | \mathfrak{F}_{T_k-}]$$

⁵³Note that $\int_0^t Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) - 1) d\bar{N}_s^{ij} = \int_0^t Q^{ij}(s, W_{s-}, \mu(s, W_{s-}), \nu_s) - 1) d\bar{N}_s^{ij}$ for all $t \in [0, T]$ \mathbb{P} -a.s. since $W_{t-} = W_t$ for all $t \in [0, T] \setminus \{T_1, \dots, T_n\}$ and $\mathbb{P}(\Delta N_{T_k}^{ij} \neq 0) = 0$ by Lemma A.4. Hence, since Q is bounded, we obtain, e.g., from Theorem IV.29 in [Pro04] that θ^ν is a local $(\mathfrak{F}, \mathbb{P})$ -martingale.

⁵⁴The local martingale property follows, e.g., from Theorem 6.56.6 in [Med07]. For the subsequent representation see, e.g., Theorem II.37 in [Pro04] or Theorem 6.56.7 in [Med07].

$$\begin{aligned}
&= \vartheta_{T_k^-} - 1 + |\mathbb{W}| \cdot \sum_{w_k \in \mathbb{W}} \mathbb{P}(W_k = w_k | W_1, \dots, W_{k-1}, \mu(T_k^-, W_{T_k^-})) \\
&\quad \times \kappa_k(w_k | W_1, \dots, W_{k-1}, \mu(T_k^-, W_{T_k^-})) \\
&= \vartheta_{T_k^-} - 1 + |\mathbb{W}| \cdot \sum_{w_k \in \mathbb{W}} \frac{1}{|\mathbb{W}|} \cdot \kappa_k(w_k | W_1, \dots, W_{k-1}, \mu(T_k^-, W_{T_k^-})) = \vartheta_{T_k^-}.
\end{aligned}$$

Hence the Doléans-Dade exponential $\mathcal{E}[\vartheta]$ is a local $(\mathfrak{F}, \mathbb{P})$ -martingale⁵⁴, and we have

$$\mathcal{E}[\vartheta]_t = \prod_{s \in (0, t]} (1 + \Delta \vartheta_s) = \prod_{\substack{k \in [1:n], \\ T_k \leq t}} \left(|\mathbb{W}| \cdot \kappa_k(W_k | W_1, \dots, W_{k-1}, \mu(T_k^-, W_{T_k^-})) \right) \quad (2.5)$$

for $t \in [0, T]$. Since $\Delta N_{T_k}^{ij} = 0$ for all $i, j \in \mathbb{S}$, $i \neq j$, and $k \in [1:n]$ \mathbb{P} -a.s. (see Lemma A.4), we have $[\theta^\nu, \vartheta] = 0$; thus, using Yor's formula⁵⁵, the process $Z^\nu \triangleq \mathcal{E}[\theta^\nu + \vartheta] = \mathcal{E}[\theta^\nu] \cdot \mathcal{E}[\vartheta]$, i.e.

$$\begin{aligned}
Z_t^\nu &= \prod_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \left(\exp \left\{ \int_0^t (1 - Q^{ij}(s, W_s, \mu(s, W_s), \nu_s)) ds \right\} \cdot \prod_{\substack{s \in (0, t], \\ \Delta N_s^{ij} \neq 0}} Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) \right) \\
&\quad \times \prod_{\substack{k \in [1:n], \\ T_k \leq t}} \left(|\mathbb{W}| \cdot \kappa_k(W_k | W_1, \dots, W_{k-1}, \mu(T_k^-, W_{T_k^-})) \right) \quad (2.6)
\end{aligned}$$

is a local $(\mathfrak{F}, \mathbb{P})$ -martingale as well.⁵⁴ Since

$$\sup_{t \in [0, T]} |\mathcal{E}[\theta^\nu]_t| \leq e^{d^2 T} \cdot \ell^Y \quad (2.7)$$

where $\ell \triangleq \max_{i, j \in \mathbb{S}, i \neq j} \|Q^{ij}\|_\infty$ and $Y \triangleq \sum_{i, j \in \mathbb{S}, i \neq j} N_T^{ij} \sim_{\mathbb{P}} \text{Poisson}(d(d-1)T)$ and

$$\sup_{t \in [0, T]} |\mathcal{E}[\vartheta]_t| \leq |\mathbb{W}|^n \quad (2.8)$$

it follows that $\sup_{t \in [0, T]} |Z_t^\nu|$ is \mathbb{P} -integrable, so Z^ν is in fact an $(\mathfrak{F}, \mathbb{P})$ -martingale. Since Z^ν is non-negative with $Z_0^\nu = 1$ by construction, we conclude (see also (2.1)) that \mathbb{P}^ν is a well-defined probability measure on \mathfrak{A} , absolutely continuous with respect to \mathbb{P} , with density process

$$\left. \frac{d\mathbb{P}^\nu}{d\mathbb{P}} \right|_{\mathfrak{F}_t} = Z_t^\nu, \quad t \in [0, T].$$

Step 2: \mathbb{P}^ν -intensity of N^{ij} . Let $i, j \in \mathbb{S}$ with $i \neq j$. Since $\mathbb{P}^\nu \ll \mathbb{P}$ it is clear that N^{ij} is a \mathbb{P}^ν -counting process, so it suffices to show that the process ${}^\nu \bar{N}^{ij} = \{{}^\nu \bar{N}_t^{ij}\}_{t \in [0, T]}$,

$${}^\nu \bar{N}_t^{ij} \triangleq N_t^{ij} - \int_0^t Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) ds, \quad t \in [0, T], \quad (2.9)$$

is a local $(\mathfrak{F}, \mathbb{P}^\nu)$ -martingale; see Definition A.2. To show this, by Step 1 it suffices to demonstrate that $Z^\nu \cdot {}^\nu \bar{N}^{ij}$ is a local $(\mathfrak{F}, \mathbb{P})$ -martingale. Noting that

⁵⁵see, e.g, Theorem II.38 in [Pro04] or Proposition 6.55 in [Med07].

$$\begin{aligned}
\triangleright [N^{k\ell}, N^{ij}] &= \sum_{s \in (0, \cdot]} \Delta N_s^{k\ell} \cdot \Delta N_s^{ij} = 0 \text{ whenever } k, \ell \in \mathbb{S} \text{ and } (k, \ell) \neq (i, j),^{56} \\
\triangleright dZ_t^\nu &= Z_{t-}^\nu d\theta_t^\nu + Z_{t-}^\nu d\vartheta_t = \sum_{\substack{k, \ell \in \mathbb{S}, \\ k \neq \ell}} Z_{t-}^\nu (Q^{k\ell}(t, W_t, \mu(t, W_t), \nu_t) - 1) d\bar{N}_t^{k\ell} + Z_{t-}^\nu d\vartheta_t, \\
\triangleright d[Z^\nu, {}^\nu\bar{N}^{ij}]_t &= Z_{t-}^\nu (Q^{ij}(t, W_t, \mu(t, W_t), \nu_t) - 1) dN_t^{ij},
\end{aligned}$$

and using integration by parts we obtain

$$\begin{aligned}
d(Z_t^\nu \cdot {}^\nu\bar{N}_t^{ij}) &= Z_{t-}^\nu d{}^\nu\bar{N}_t^{ij} + {}^\nu\bar{N}_{t-}^{ij} dZ_t^\nu + d[Z^\nu, {}^\nu\bar{N}^{ij}]_t \\
&= Z_{t-}^\nu dN_t^{ij} - Z_{t-}^\nu Q^{ij}(t, W_t, \mu(t, W_t), \nu_t) dt + {}^\nu\bar{N}_{t-}^{ij} dZ_t^\nu \\
&\quad + Z_{t-}^\nu Q^{ij}(t, W_t, \mu(t, W_t), \nu_t) dN_t^{ij} - Z_{t-}^\nu dN_t^{ij} \\
&= {}^\nu\bar{N}_{t-}^{ij} dZ_t^\nu + Z_{t-}^\nu Q^{ij}(t, W_t, \mu(t, W_t), \nu_t) d\bar{N}_t^{ij}.
\end{aligned}$$

Thus ${}^\nu\bar{N}^{ij}$ is a local $(\mathfrak{F}, \mathbb{P}^\nu)$ -martingale, as desired.

Step 3: $\mathbb{P}^\nu = \mathbb{P}$ on $\sigma(X_0)$. Note that for any function $g : \mathbb{S} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}^\nu[g(X_0)] = \mathbb{E}[g(X_0) \cdot Z_T^\nu] = \mathbb{E}[g(X_0) \cdot \mathbb{E}[Z_T^\nu | \mathfrak{F}_0]] = \mathbb{E}[g(X_0) \cdot Z_0^\nu] = \mathbb{E}[g(X_0)]$$

by the tower property of conditional expectation and the $(\mathfrak{F}, \mathbb{P})$ -martingale property of Z^ν . This is equivalent to the desired identity.

Step 4: Distribution of W_k under \mathbb{P}^ν . Let $k \in [1 : n]$ and $w_1, \dots, w_k \in \mathbb{W}$. Since $\mathcal{E}[\theta^\nu]_{T_k} = \mathcal{E}[\theta^\nu]_{T_k-}$ \mathbb{P} -a.s. and W_k is uniformly distributed on \mathbb{W} and independent of \mathfrak{F}_{T_k-} under \mathbb{P} , iterated conditioning yields

$$\begin{aligned}
\mathbb{P}^\nu(W_1 = w_1, \dots, W_k = w_k) &= \mathbb{E}[Z_{T_k}^\nu \cdot \mathbf{1}_{\{W_k = w_k\}} \cdot \mathbf{1}_{\{W_1 = w_1, \dots, W_{k-1} = w_{k-1}\}}] \\
&= \mathbb{E}[Z_{T_k-}^\nu \cdot |\mathbb{W}| \cdot \kappa_k(W_k | W_1, \dots, W_{k-1}, \mu(T_k-, W_{T_k-})) \cdot \mathbf{1}_{\{W_k = w_k\}} \cdot \mathbf{1}_{\{W_1 = w_1, \dots, W_{k-1} = w_{k-1}\}}] \\
&= |\mathbb{W}| \cdot \kappa_k(w_k | w_1, \dots, w_{k-1}, \mu(T_k-, w_{T_k-})) \cdot \mathbb{E}[Z_{T_k-}^\nu \cdot \mathbf{1}_{\{W_1 = w_1, \dots, W_{k-1} = w_{k-1}\}} \cdot \mathbb{P}(W_k = w_k | \mathfrak{F}_{T_k-})] \\
&= \kappa_k(w_k | w_1, \dots, w_{k-1}, \mu(T_k-, w_{T_k-})) \cdot \mathbb{P}^\nu(W_1 = w_1, \dots, W_{k-1} = w_{k-1}).
\end{aligned}$$

Thus we have

$$\mathbb{P}^\nu(W_k = w_k | \mathfrak{G}_{T_k-}) = \kappa_k(w_k | W_1, \dots, W_{k-1}, M_{T_k-})$$

and the proof is complete. \square

Lemma 2.1.3 implies in particular that $\mathbb{P}^\nu(\Delta N_t^{ij} \neq 0) = 0$ for every $t \in [0, T]$ (see Lemma A.4), so as a consequence we have

$$\Delta X_t = 0 \quad \mathbb{P}^\nu\text{-a.s. for all } t \in [0, T] \setminus \{T_1, \dots, T_n\};$$

hence, X_- is a modification of X on $[0, T] \setminus \{T_1, \dots, T_n\}$ under \mathbb{P}^ν .

Moreover, since $\mathbb{P}^{\nu_1} = \mathbb{P}^{\nu_2}$ on \mathfrak{G}_T for all admissible controls $\nu_1, \nu_2 \in \mathcal{A}$ and $M_t = \mu(t, W_t)$

⁵⁶Note that N^{ij} and N^{kl} are independent processes in these cases; hence this follows from Proposition 7.13 in [Med07] in conjunction with [Med07, p.469].

for $t \in [0, T]$, the agent's *ex ante* beliefs concerning the common noise factors are the same, irrespective of his control.

Remark 2.1.4. Our statement of the optimization problem (P_μ) where the control $\nu \in \mathcal{A}$ affects the state process X via its law (rather than changing its paths) is known as a *weak formulation* of the stochastic control problem.⁵⁷ This is natural in diffusive problems where only the drift is controlled, see [CL15]; and in the context of controlled intensities, see [BS05], [BKS13] and [CW18b]. \diamond

2.2 Solution of the Optimization Problem

In the following, we solve the agent's maximization problem (P_μ) using the associated dynamic programming equation (DPE). This is the same methodology as in [GMS13] and [CF20]; we refer to [CW18b] for an alternative approach (to extended mean field games, but without common noise) based on backward stochastic differential equations.

The DPE for the agent's optimization problem reads

$$0 = \sup_{u \in \mathbb{U}} \left\{ \frac{\partial v^i}{\partial t}(t, w) + \psi^i(t, w, \mu(t, w), u) + Q^{i\bullet}(t, w, \mu(t, w), u) \cdot v(t, w) \right\}$$

for $i \in \mathbb{S}$, subject to suitable consistency conditions for $t = T_k$, $k \in [1 : n]$, and the terminal condition

$$v(T, w) = \Psi(w, \mu(T, w)) \quad \text{for all } w \in \mathbb{W}^n.$$

The following crucial assumption will enable us to formulate the latter DPE in a reduced form; see Definition 2.2.2 below.⁵⁸ It is assumed to be valid throughout the remainder of this chapter.

Assumption 2.2.1 (Measurable Maximizer). There exists a measurable function

$$h : [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d \rightarrow \mathbb{U}^d$$

such that for every $i \in \mathbb{S}$ and all $(t, w, m, v) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d$ we have

$$h^i(t, w, m, v) \in \arg \max_{u \in \mathbb{U}} \{ \psi^i(t, w, m, u) + Q^{i\bullet}(t, w, m, u) \cdot v \}. \quad \diamond$$

A sufficient condition for Assumption 2.2.1 to be satisfied is that \mathbb{U} is compact and Q and ψ are continuous with respect to $u \in \mathbb{U}$; see Theorem 18.19 in [AB06].⁵⁹ Note that, since $\psi^i(\cdot, \cdot, m, u)$ and $Q^{i\bullet}(\cdot, \cdot, m, u)$ are non-anticipative for $m \in \mathbb{M}$ and $u \in \mathbb{U}$, we can assume without loss of

⁵⁷For a conceptually different approach where the agents' state process is constructed as a pathwise unique strong solution of a controlled SDE driven by a Poisson random measure we refer, e.g., to [CF20].

⁵⁸The stated assumption is standard given the literature and oftentimes derived from suitable sufficient conditions; we refer, e.g., to §2 in [GMS13], Assumption 3.4(i) in [CW18b] or condition (C) in §2.1 in [CF20].

⁵⁹Note that h is uniquely determined if \mathbb{U} is convex and the *Hamiltonian* function (see also Definition E.1)

$$\mathbb{U} \ni u \mapsto \psi^i(t, w, m, u) + Q^{i\bullet}(t, w, m, u) \cdot v \in \mathbb{R}$$

is strictly concave for each $(i, t, w, m) \in \mathbb{S} \times [0, T] \times \mathbb{W}^n \times \mathbb{M}$ and all relevant $v \in \mathbb{R}^d$.

generality that $h(\cdot, \cdot, m, v)$ is non-anticipative for $m \in \mathbb{M}$ and $v \in \mathbb{R}^d$. In view of Assumption 2.2.1, we can define

$$\begin{aligned}\widehat{Q} &: [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}, & \widehat{Q}^{ij}(t, w, m, v) &\triangleq Q^{ij}(t, w, m, h^i(t, w, m, v)), \\ \widehat{\psi} &: [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d \rightarrow \mathbb{R}^d, & \widehat{\psi}^i(t, w, m, v) &\triangleq \psi^i(t, w, m, h^i(t, w, m, v))\end{aligned}$$

and thus obtain the following reduced-form DPE, which we use in the following:

Definition 2.2.2 (Reduced-Form DPE). Let $\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M}$ be regular and non-anticipative. A function $v : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^d$ is called a *solution* of (DP_μ) subject to (CC_μ) , (TC_μ) if v is non-anticipative and satisfies the ordinary differential equation (ODE)⁶⁰

$$\dot{v}(t, w) = -\widehat{\psi}(t, w, \mu(t, w), v(t, w)) - \widehat{Q}(t, w, \mu(t, w), v(t, w)) \cdot v(t, w) \quad (\text{DP}_\mu)$$

for $t \in [T_k, T_{k+1})$, $k \in [0:n]$, subject to the consistency and terminal conditions

$$v(T_k-, w) = \Psi_k(w, \mu(T_k-, w), v(T_k, \cdot)), \quad (\text{CC}_\mu)$$

$$v(T, w) = \Psi(w, \mu(T, w)) \quad (\text{TC}_\mu)$$

for $k \in [1:n]$ and all $w \in \mathbb{W}^n$. Here for $k \in [1:n]$ the jump operator Ψ_k is defined via

$$\Psi_k^i(w, m, \bar{v}) \triangleq \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, m) \cdot \bar{v}^{J^i(T_k, (w_{-k}, \bar{w}_k), m)}(w_{-k}, \bar{w}_k), \quad i \in \mathbb{S}, \quad (\Psi_k)$$

where $m \in \mathbb{M}$, $\bar{v} : \mathbb{W}^n \rightarrow \mathbb{R}^d$ is a function and $(w_{-k}, \bar{w}_k) \triangleq (w_1, \dots, w_{k-1}, \bar{w}_k, w_{k+1}, \dots, w_n) \in \mathbb{W}^n$ for $\bar{w}_k \in \mathbb{W}$, $w \in \mathbb{W}^n$. \diamond

Observe that (DP_μ) represents a system of (random) ODEs, coupled via $w \in \mathbb{W}^n$. The ODEs run backward in time on each segment $[T_k, T_{k+1}) \times \mathbb{W}^n$, $k \in [0:n]$, and their terminal conditions for $t \uparrow T_{k+1}$ are specified by (TC_μ) for $k = n$ and by (CC_μ) for $k < n$. Note that for $t \in [T_k, T_{k+1})$ the relevant factors W_1, \dots, W_k are known.⁶¹

Remark 2.2.3. While the significance of the DPE (DP_μ) and the terminal condition (TC_μ) are clear, the consistency conditions (CC_μ) warrant a brief comment: For $i \in \mathbb{S}$, $k \in [1:n]$ and $w \in \mathbb{W}^n$ the state process jumps from state i to state $j \triangleq J^i(T_k, (w_{-k}, W_k), \mu(T_k-, w_{T_k-}))$ on the event $\{X_{T_k-} = i\} \cap \{W_{T_k-} = w_{T_k-}\}$ when the common noise factor W_k is revealed at time T_k ; and on this event for every admissible strategy $\nu \in \mathcal{A}$ it holds $\mathbb{P}^\nu(W_k = \bar{w}_k | \mathfrak{G}_{T_k-}) = \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu(T_k-, w_{T_k-}))$, see Lemma 2.1.3. Hence, intuitively, the pre-jump value at time T_k equals the expected post-jump value based on the (conditional) distribution of the common noise factor W_k . \diamond

We next use classical verification arguments⁶² to link the solution of the DPE to the underlying

⁶⁰All ODEs in this chapter are taken in the sense of Carathéodory; see Section A.3 for a brief outline; for further details we refer, e.g., to §I.5 in [Hal80] or Supplement II of III.§10 in [Wal98]. As a consequence, every solution v of (DP_μ) subject to (CC_μ) , (TC_μ) is automatically regular in the sense of Definition 2.1.2.

⁶¹For a corresponding existence result (and relevant assumptions), i.e. for existence of a *partial equilibrium* given the agent's ex ante beliefs μ , we refer to Lemma D.4.

⁶²see also, e.g., Theorem III.8.1 in [FS06] or Theorem 1 in [GMS13].

stochastic control problem.

Theorem 2.2.4 (Verification). Suppose that v is a solution of (DP_μ) subject to (CC_μ) and (TC_μ) , where $\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M}$ is regular and non-anticipative.⁶¹ Then v is the agent's value function for problem (P_μ) , i.e.

$$\sum_{i \in \mathbb{S}} \mathbb{P}(X_0 = i) \cdot v^i(0) = \sup_{\nu \in \mathcal{A}} \mathbb{E}^\nu \left[\int_0^T \psi^{X_t}(t, W_t, M_t, \nu_t) dt + \Psi^{X_T}(W_T, M_T) \right]$$

and the control $\hat{\nu} \in \mathcal{A}$ given by⁶³

$$\hat{\nu}(t, X_{(\cdot \wedge t)^-}, W_t) = h^{X_t^-}(t, W_t, \mu(t, W_t), v(t, W_t)) \quad \text{for } t \in [0, T]$$

is optimal.

Proof. Let $\nu \in \mathcal{A}$ be an admissible strategy. Until further notice we fix $k \in [0:n]$.

Step 1: Dynamics on $[T_k, T_{k+1})$. Using Itô's lemma⁶⁴ and the fact that v is regular, we obtain for $t \in [T_k, T_{k+1})$

$$\begin{aligned} v^{X_t}(t, W_t) &= v^{X_{T_k}}(T_k, W_{T_k}) + \int_{T_k}^t \dot{v}^{X_s}(s, W_s) ds \\ &\quad + \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \left(\int_{(T_k, t]} \mathbb{1}_{\{X_{s-}=i\}} (v^j(s, W_s) - v^i(s, W_s)) d^\nu \bar{N}_s^{ij} \right. \\ &\quad \left. + \int_{T_k}^t \mathbb{1}_{\{X_s=i\}} (v^j(s, W_s) - v^i(s, W_s)) Q^{ij}(s, W_s, \mu(s, W_s), \nu_s) ds \right) \end{aligned}$$

where the processes ${}^\nu \bar{N}^{ij}$, $i, j \in \mathbb{S}$, $i \neq j$, are defined in (2.9). Thus for $t \in [T_k, T_{k+1})$ we have

$$\begin{aligned} v^{X_{T_k}}(T_k, W_{T_k}) &= v^{X_t}(t, W_t) - \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \int_{(T_k, t]} \mathbb{1}_{\{X_{s-}=i\}} (v^j(s, W_s) - v^i(s, W_s)) d^\nu \bar{N}_s^{ij} \\ &\quad - \int_{T_k}^t \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds \end{aligned}$$

and upon sending $t \uparrow T_{k+1}$ we obtain

$$\begin{aligned} v^{X_{T_k}}(T_k, W_{T_k}) &= v^{X_{T_{k+1}}}(T_{k+1}^-, W_{T_{k+1}^-}) \\ &\quad - \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \int_{(T_k, T_{k+1})} \mathbb{1}_{\{X_{s-}=i\}} (v^j(s, W_s) - v^i(s, W_s)) d^\nu \bar{N}_s^{ij} \\ &\quad - \int_{T_k}^{T_{k+1}} \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds. \end{aligned} \tag{2.10}$$

⁶³Formally, this corresponds to the function $\hat{\nu} : [0, T] \times \mathbb{S}^{[0, T]} \times \mathbb{W}^n \rightarrow \mathbb{U}$, $\hat{\nu}(t, x, w) = h^{x_t}(t, w, \mu(t, w), v(t, w))$.

⁶⁴More precisely, we refer to the *fundamental theorem of calculus* (for Lebesgue integrals); see, e.g., Theorem 7.20 in [Rud87] or Theorem VII.4.14(b) in [Els09].

Step 2: Jump dynamics at T_k . We recall that W_k is independent of \mathfrak{F}_{T_k-} and from Lemma 2.1.3 that

$$\begin{aligned} \mathbb{P}^\nu(W_k = \bar{w}_k | X_{T_k-}, W_1, \dots, W_{k-1}) &= \mathbb{P}^\nu(W_k = \bar{w}_k | W_1, \dots, W_{k-1}) \\ &= \mathbb{P}^\nu(W_k = \bar{w}_k | \mathfrak{G}_{T_k-}) \\ &= \kappa_k(\bar{w}_k | W_1, \dots, W_{k-1}, \mu(T_k-, W_{T_k-})). \end{aligned}$$

In view of the jump dynamics (2.3) and the consistency condition (CC $_\mu$), we thus obtain

$$\begin{aligned} &\mathbb{E}^\nu[v^{X_{T_k}}(T_k, W_{T_k}) | \sigma(X_{T_k-}, W_{T_k-})] \\ &= \mathbb{E}^\nu[v^{J^{X_{T_k-}}(T_k, (W_{T_k-}, W_k), \mu(T_k-, W_{T_k-}))}(T_k, (W_{T_k-}, W_k)) | \sigma(X_{T_k-}, W_{T_k-})] \\ &= \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | W_{T_k-}, \mu(T_k-, W_{T_k-})) \\ &\quad \times v^{J^{X_{T_k-}}(T_k, (W_{T_k-}, \bar{w}_k), \mu(T_k-, W_{T_k-}))}(T_k, (W_{T_k-}, \bar{w}_k)) \\ &= \Psi_k^{X_{T_k-}}(W_{T_k-}, \mu(T_k-, W_{T_k-}), v(T_k, \cdot)) = v^{X_{T_k-}}(T_k-, W_{T_k-}). \end{aligned} \tag{2.11}$$

Step 3: Optimality. Recall from Lemma 2.1.3 and Lemma A.4 that $\mathbb{P}^\nu(\Delta N_{T_k}^{ij} \neq 0) = 0$ for every $i, j \in \mathbb{S}$, $i \neq j$ and $k \in [1:n]$. Combining (2.10) and (2.11) for $k = [1:n]$ and using (TC $_\mu$) thus yields

$$\begin{aligned} v^{X_0}(0) &= v^{X_T}(T, W_T) + \sum_{k=1}^n (v^{X_{T_k-}}(T_k-, W_{T_k-}) - v^{X_{T_k}}(T_k, W_{T_k})) \\ &\quad - \sum_{k=0}^n \left[\sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \int_{(T_k, T_{k+1})} \mathbb{1}_{\{X_{s-}=i\}} (v^j(s, W_s) - v^i(s, W_s)) d^\nu \bar{N}_s^{ij} \right. \\ &\quad \left. + \int_{T_k}^{T_{k+1}} \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds \right] \\ &= \Psi^{X_T}(W_T, \mu(T, W_T)) + \sum_{k=1}^n \left(\mathbb{E}^\nu[v^{X_{T_k}}(T_k, W_{T_k}) | \sigma(X_{T_k-}, W_{T_k-})] - v^{X_{T_k}}(T_k, W_{T_k}) \right) \\ &\quad - \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} L_T^{ij} - \int_0^T \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds \end{aligned} \tag{2.12}$$

where for $i, j \in \mathbb{S}$, $i \neq j$ the local $(\mathfrak{F}, \mathbb{P}^\nu)$ -martingale L^{ij} is given by

$$L_t^{ij} \triangleq \int_{(0, t]} \mathbb{1}_{\{X_{s-}=i\}} (v^j(s, W_s) - v^i(s, W_s)) d^\nu \bar{N}_s^{ij} \quad \text{for } t \in [0, T].$$

Since ${}^\nu \bar{N}^{ij}$ is a compensated counting process and v and Q are bounded, it follows from Lemma A.3 that L^{ij} is an $(\mathfrak{F}, \mathbb{P}^\nu)$ -martingale. Hence, by taking \mathbb{P}^ν -expectations in (2.12) and using the tower property of conditional expectation and the fact that \mathbb{P}^ν and \mathbb{P} coincide on $\sigma(X_0)$ by Lemma 2.1.3,

we obtain

$$\begin{aligned}
& \sum_{i \in \mathbb{S}} \mathbb{P}(X_0 = i) v^i(0) = \mathbb{E}[v^{X_0}(0)] = \mathbb{E}^\nu[v^{X_0}(0)] \\
&= \mathbb{E}^\nu \left[\Psi^{X_T}(W_T, \mu(T, W_T)) \right. \\
&\quad \left. - \int_0^T \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds \right] \\
&= \mathbb{E}^\nu \left[\Psi^{X_T}(W_T, \mu(T, W_T)) + \int_0^T \psi^{X_s}(s, W_s, \mu(s, W_s), \nu_s) ds \right. \\
&\quad \left. - \int_0^T \sum_{i=1}^d \mathbb{1}_{\{X_s=i\}} (\dot{v}^i(s, W_s) + \psi^i(s, W_s, \mu(s, W_s), \nu_s) \right. \\
&\quad \left. + Q^{i\bullet}(s, W_s, \mu(s, W_s), \nu_s) \cdot v(s, W_s)) ds \right] \\
&\geq \mathbb{E}^\nu \left[\Psi^{X_T}(W_T, M_T) + \int_0^T \psi^{X_s}(s, W_s, M_s, \nu_s) ds \right]. \tag{2.13}
\end{aligned}$$

If we replace ν with $\hat{\nu}$, the same argument applies with equality in (2.13); we thus conclude that v is the value function of (P_μ) , and that the strategy $\hat{\nu}$ is optimal. \square

The optimal strategy $\hat{\nu}$ is Markovian in the agent's state; this is unsurprising given the literature, see, e.g., Theorem 1 in [GMS13], Proposition 3.9 in [CW18b], or Theorem 4 in [CF20]. Note, however, that the time- t optimal strategy may depend on *all* common noise events that have occurred up to time t , as $W_t = (W_1, \dots, W_k)$ for $t \in [T_k, T_{k+1})$. In the following, we denote by $\hat{\mathbb{P}}$ the probability measure

$$\hat{\mathbb{P}} \triangleq \mathbb{P}^{\hat{\nu}}$$

where $\hat{\nu}$ is the optimal control specified in Theorem 2.2.4. It follows from Lemma 2.1.3 that N^{ij} has $\hat{\mathbb{P}}$ -intensity $\hat{\lambda}_t^{ij} = \{\hat{\lambda}_t^{ij}\}_{t \in [0, T]}$ for $i, j \in \mathbb{S}$, $i \neq j$, where

$$\hat{\lambda}_t^{ij} \triangleq Q^{ij}(t, W_t, \mu(t, W_t), h^{X_{t-}}(t, W_t, \mu(t, W_t), v(t, W_t))) \quad \text{for } t \in [0, T]. \tag{2.14}$$

2.3 Equilibrium

Having solved the agent's optimization problem (P_μ) for a given *ex ante* function μ , we now turn to the resulting mean field equilibrium. We first identify the aggregate distribution resulting from the optimal control.

Remark 2.3.1 (Mean Field Limit). This chapter generally adopts a “representative agent” point of view; a justification of the notion of a *mean field equilibrium* is usually provided either via convergence of Nash equilibria⁶⁵ of symmetric N -player games in the limit $N \rightarrow \infty$, or, reversely, via regarding mean field equilibria as approximate Nash equilibria of corresponding finite-player games with approximation error asymptotically vanishing as $N \rightarrow \infty$; see, among others, [BHK14], [CW16], [DGG16], [Fis17], [BC18], [CF18], [CW18b], [CDLL19, Chapter 6]

⁶⁵This celebrated concept is due to Nash, see [Nas50], and more explicitly spelled out, e.g., in §§1.1.2/2.1.2 in [CD18a].

[CP19a], [CDFP19], [DGG19] and [CF20]. In the setting of this chapter (albeit under additional regularity conditions) a mean field limit justification can be provided along the lines of the proof of Theorem 7 in [GMS13] by conditioning on common noise configurations, similarly as in the proof of Theorem 2.3.4 below. We also refer to [BC18] for a recent alternative approach using the master equation, which is stated in Appendix E for our setup. \diamond

2.3.1 Aggregation

Given an *ex ante* aggregate distribution specified in terms of a regular, non-anticipative function μ and a corresponding solution v of (DP_μ) subject to (CC_μ) , (TC_μ) , Theorem 2.2.4 yields an optimal strategy $\widehat{\nu}$ for the agent's optimization problem (P_μ) . With $\widehat{\mathbb{P}}$ denoting the probability measure associated with $\widehat{\nu}$, the resulting *ex post* aggregate distribution is then given by the \mathbb{M} -valued, \mathfrak{G} -adapted process $\widehat{M} = \{\widehat{M}_t\}_{t \in [0, T]}$ defined via

$$\widehat{M}_t \triangleq \widehat{\mathbb{P}}(X_t \in \cdot \mid \mathfrak{G}_t) \quad \text{for } t \in [0, T].$$

Equilibrium obtains if $\widehat{M}_t = \mu(t, W_t)$ for all $t \in [0, T]$. To proceed, we therefore aim for a more explicit description of \widehat{M} . Thus we define for $k \in [1:n]$ the map

$$\Phi_k : \mathbb{W}^n \times \mathbb{M} \times \mathbb{M} \rightarrow \mathbb{M}, \quad \Phi_k(w, m, \bar{m}) \triangleq m \cdot P_k(w, \bar{m}), \quad (\Phi_k)$$

where $P_k : \mathbb{W}^n \times \mathbb{M} \rightarrow \{0, 1\}^{d \times d}$ is given by

$$P_k^{ij}(w, \bar{m}) \triangleq \mathbf{1}_{\{J^i(T_k, w_1, \dots, w_k, \bar{m}) = j\}} \quad \text{for } i, j \in \mathbb{S},$$

and we set

$$m_0 \triangleq \mathbb{P}(X_0 \in \cdot) = \widehat{\mathbb{P}}(X_0 \in \cdot) \in \mathbb{M}.$$

Lemma 2.3.2. The process \widehat{M} has càdlàg sample paths.

Proof. Since \mathfrak{G} is piecewise constant, this is immediate by dominated convergence and the fact that X is càdlàg. \square

Lemma 2.3.3. Let $\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M}$ and $v : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^d$ be regular and non-anticipative, and suppose that $Y = \{Y_t\}_{t \in [0, T]}$ is an \mathbb{M} -valued stochastic process with dynamics

$$Y_t = Y_{T_k} + \int_{T_k}^t Y_s \cdot \widehat{Q}(s, W_s, \mu(s, W_s), v(s, W_s)) ds \quad \text{for } t \in [T_k, T_{k+1}) \text{ and } k \in [0:n] \quad (2.15)$$

that satisfies the initial condition

$$Y_0 = m_0$$

and the consistency conditions

$$Y_{T_k} = \Phi_k(W_{T_k}, Y_{T_k-}, \mu(T_k-, W_{T_k-})) \quad \text{for } k \in [1:n].$$

Then Y is \mathfrak{G} -adapted.

Proof. We split the proof into two steps.

Step 1: Existence and uniqueness of Carathéodory solutions. For each $k \in [0:n]$ and $w \in \mathbb{W}^n$, since μ and v are regular and Q is bounded, the function

$$f : [T_k, T_{k+1}] \times \mathbb{R}^{1 \times d} \rightarrow \mathbb{R}^{1 \times d}, \quad f(t, y) \triangleq y \cdot \widehat{Q}(t, w, \mu(t, w), v(t, w))$$

is measurable in the first and Lipschitz continuous in the second argument, and bounded on $[T_k, T_{k+1}] \times \overline{B_R(0)}$ for each $R > 0$. Thus, using that μ , v and \widehat{Q} are non-anticipative, classical results, see Theorem I.5.3 in conjunction with Theorem I.5.2 in [Hal80]⁶⁶, imply that there exists a unique Carathéodory solution $\varphi_k^{y, w_{T_k}} : [T_k, T_{k+1}] \rightarrow \mathbb{R}^{1 \times d}$ of

$$\dot{y}(t) = y(t) \cdot \widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v(t, w_{T_k})) \quad \text{for } t \in [T_k, T_{k+1}), \quad y(T_k) = y,$$

for every initial condition $y \in \mathbb{R}^{1 \times d}$.

Step 2: Y is \mathfrak{G} -adapted. First note that $Y_0 = m_0$ is clearly \mathfrak{G}_0 -measurable. Next, suppose that Y_{T_k} is \mathfrak{G}_{T_k} -measurable, and note that for $t \in [T_k, T_{k+1})$ we have $W_t = W_{T_k}$, so

$$Y_t = Y_{T_k} + \int_{T_k}^t Y_s \cdot \widehat{Q}(s, W_{T_k}, \mu(s, W_{T_k}), v(s, W_{T_k})) ds.$$

Thus from uniqueness in part (a) it follows that we have the representation

$$Y_t = \varphi_k^{Y_{T_k}, W_{T_k}}(t) \quad \text{for } t \in [T_k, T_{k+1}).$$

Hence Y_t is \mathfrak{G}_{T_k} -measurable for all $t \in [T_k, T_{k+1})$. Finally, for each $k \in [0:n]$ the consistency condition implies that $Y_{T_{k+1}} = \Phi_{k+1}(W_{T_{k+1}}, Y_{T_{k+1}-}, \mu(T_{k+1}-, W_{T_{k+1}-}))$ is $\mathfrak{G}_{T_{k+1}}$ -measurable, too, so the claim follows by induction on $k \in [0:n]$. \square

We are now in a position to state and prove the promised aggregation result:

Theorem 2.3.4 (Aggregation). Let $\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M}$ be regular and non-anticipative with $\mu(0) = m_0$. Suppose v is a solution of (DP $_{\mu}$) subject to (CC $_{\mu}$), (TC $_{\mu}$), and the representative agent implements his optimal strategy \widehat{v} as defined in Theorem 2.2.4. Then the aggregate distribution \widehat{M} has the $\widehat{\mathbb{P}}$ -dynamics

$$d\widehat{M}_t = \widehat{M}_t \cdot \widehat{Q}(t, W_t, \mu(t, W_t), v(t, W_t)) dt \quad \text{for } t \in [T_k, T_{k+1}), \quad k \in [0:n], \quad (\text{M})$$

and satisfies the initial condition

$$\widehat{M}_0 = m_0 \quad (\text{M}_0)$$

and the jump conditions

$$\widehat{M}_{T_k} = \Phi_k(W_{T_k}, \widehat{M}_{T_k-}, \mu(T_k-, W_{T_k-})) \quad \text{for } k \in [1:n]. \quad (\text{M}_k)$$

Proof. Let $w \in \mathbb{W}^n$ be a common noise configuration. Since X is defined path by path, see (2.2)

⁶⁶see also Theorems A.6 and A.7 as well as Remark A.8(a).

and (2.3), we first note that $X = X^w$ on $\{W_T = w\}$, where X^w satisfies (2.2) and

$$X_{T_k}^w = J^{X_{T_k}^w} (T_k, w_{T_k}, \mu(T_k^-, w_{T_k}^-)) \text{ for } k \in [1:n]. \quad (2.16)$$

We define $\zeta(w) = \{\zeta(w)_t\}_{t \in [0, T]}$ via

$$\begin{aligned} \zeta(w)_t \triangleq & \prod_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \left(\exp \left\{ \int_0^t (1 - Q^{ij}(s, w_s, \mu(s, w_s), h^{X_s^w}(s, w_s, \mu(s, w_s), v(s, w_s)))) ds \right\} \right) \\ & \times \prod_{\substack{s \in (0, t], \\ \Delta N_s^{ij} \neq 0}} Q^{ij}(s, w_s, \mu(s, w_s), h^{X_s^w}(s, w_s, \mu(s, w_s), v(s, w_s))). \end{aligned}$$

Using analogous arguments as in Step 1 of the proof of Lemma 2.1.3 (see in particular (2.4) and (2.7)), it follows that there exists a probability measure $\widehat{\mathbb{P}}^w$ with density process

$$\left. \frac{d\widehat{\mathbb{P}}^w}{d\mathbb{P}} \right|_{\mathfrak{H}_t} \triangleq \zeta(w)_t \quad \text{for } t \in [0, T],$$

where the filtration $\mathfrak{H} = \{\mathfrak{H}_t\}$ is given by⁶⁷

$$\mathfrak{H}_t \triangleq \sigma(X_0, N_s^{ij} : s \in [0, t]; i, j \in \mathbb{S}, i \neq j) \vee \mathfrak{N} \quad \text{for } t \in [0, T].$$

Furthermore, in view of (2.4) and (2.14) we have

$$\zeta(w) = \mathcal{E}[\theta^{\widehat{\nu}}] \quad \text{on } \{W_T = w\}. \quad (2.17)$$

Step 1: Conditional Kolmogorov dynamics. Throughout Step 1, we fix a common noise configuration $w \in \mathbb{W}^n$. It follows exactly as in the proof of Lemma 2.1.3 (with $\widehat{\mathbb{P}}^w$ in place of $\widehat{\mathbb{P}}$) that

$$\widehat{\mathbb{P}}^w \ll \mathbb{P}, \quad \widehat{\mathbb{P}}^w = \mathbb{P} \quad \text{on } \sigma(X_0),$$

and that for $i, j \in \mathbb{S}, i \neq j$, the process N^{ij} is a counting process with $(\mathfrak{H}, \widehat{\mathbb{P}}^w)$ -intensity

$$Q^{ij}(t, w_t, \mu(t, w_t), h^{X_t^w}(t, w_t, \mu(t, w_t), v(t, w_t))) \quad \text{for } t \in [0, T].$$

It follows from Lemma A.3 that for each $z \in \mathbb{R}^d$ the process $L^w[z] = \{L_t^w[z]\}_{t \in [0, T]}$,

$$L_t^w[z] \triangleq \sum_{\substack{i, j \in \mathbb{S}, \\ i \neq j}} \int_{(0, t]} \mathbf{1}_{\{X_{s-}^w = i\}} \cdot (z^j - z^i) d^w \bar{N}_s^{ij} \quad \text{for } t \in [0, T],$$

is an $(\mathfrak{H}, \widehat{\mathbb{P}}^w)$ -martingale, where ${}^w \bar{N}^{ij} = \{{}^w \bar{N}_t^{ij}\}_{t \in [0, T]}$ is given by

$${}^w \bar{N}_t^{ij} \triangleq N_t^{ij} - \int_0^t Q^{ij}(s, w_s, \mu(s, w_s), h^{X_s^w}(s, w_s, \mu(s, w_s), v(s, w_s))) ds, \quad t \in [0, T].$$

⁶⁷Note that $\mathfrak{F}_t = \mathfrak{G}_t \vee \mathfrak{H}_t$ for all $t \in [0, T]$.

Using Itô's lemma and the fact that $\widehat{\lambda}_t^{ij} = \widehat{Q}^{ij}(t, W_t, \mu(t, W_t), v(t, W_t))$ on $\{X_{t-} = i\}$ for $t \in [0, T]$ by (2.14), we have for each $z \in \mathbb{R}^d$, $k \in [0:n]$ and $t \in [T_k, T_{k+1})$

$$\begin{aligned} z^{X_t^w} &= z^{X_{T_k}^w} + \sum_{\substack{i,j \in \mathbb{S}, \\ i \neq j}} \int_{(T_k, t]} \mathbf{1}_{\{X_{s-}^w = i\}} (z^{X_s^w} - z^{X_{s-}^w}) dN_s^{ij} \\ &= z^{X_{T_k}^w} + \sum_{\substack{i,j \in \mathbb{S}, \\ i \neq j}} \int_{(T_k, t]} \mathbf{1}_{\{X_{s-}^w = i\}} (z^j - z^i) d^w \bar{N}_s^{ij} \\ &\quad + \sum_{i \in \mathbb{S}} \sum_{j \in \mathbb{S}, j \neq i} \int_{T_k}^t \mathbf{1}_{\{X_s^w = i\}} \cdot \widehat{Q}^{ij}(s, w_s, \mu(s, w_s), v(s, w_s)) (z^j - z^i) ds \\ &= z^{X_{T_k}^w} + L_t^w[z] - L_{T_k}^w[z] + \sum_{i=1}^d \int_{T_k}^t \mathbf{1}_{\{X_s^w = i\}} \cdot \widehat{Q}^{i\bullet}(s, w_s, \mu(s, w_s), v(s, w_s)) \cdot z \, ds. \end{aligned}$$

Taking expectations with respect to $\widehat{\mathbb{P}}^w$ and using Fubini's theorem yields

$$\widehat{\mathbb{E}}^w[z^{X_t^w}] = \widehat{\mathbb{E}}^w[z^{X_{T_k}^w}] + \sum_{i=1}^d \int_{T_k}^t \widehat{\mathbb{P}}^w(X_s^w = i) \cdot \widehat{Q}^{i\bullet}(s, w_s, \mu(s, w_s), v(s, w_s)) \cdot z \, ds,$$

so with $z = e_i$, $i \in \mathbb{S}$, we get

$$\widehat{\mathbb{P}}^w(X_t^w = i) = \widehat{\mathbb{P}}^w(X_{T_k}^w = i) + \sum_{j=1}^d \int_{T_k}^t \widehat{\mathbb{P}}^w(X_s^w = j) \cdot \widehat{Q}^{ji}(s, w_s, \mu(s, w_s), v(s, w_s)) ds. \quad (2.18)$$

It follows from (2.18) that $\eta(w) = \{\eta(w)_t\}_{t \in [0, T]}$,

$$\eta(w)_t \triangleq \widehat{\mathbb{P}}^w(X_t^w \in \cdot), \quad t \in [0, T] \quad (2.19)$$

satisfies, for all $i \in \mathbb{S}$ and $k \in [0:n]$,

$$\eta(w)_t^i = \eta(w)_{T_k}^i + \int_{T_k}^t \eta(w)_s \cdot \widehat{Q}^{i\bullet}(s, w_s, \mu(s, w_s), v(s, w_s)) ds \quad \text{for } t \in [T_k, T_{k+1}). \quad (2.20)$$

Moreover, since $\widehat{\mathbb{P}}^w = \mathbb{P}$ on $\sigma(X_0)$ and $X_0^w = X_0$, $\eta(w)$ satisfies the initial condition

$$\eta(w)_0 = \widehat{\mathbb{P}}^w(X_0^w \in \cdot) = \mathbb{P}(X_0^w \in \cdot) = \mathbb{P}(X_0 \in \cdot) = m_0. \quad (2.21)$$

Finally, we consider a common noise time $t = T_k$, $k \in [1:n]$, and note that for all $i \in \mathbb{S}$ the jump condition (2.16) implies that

$$\begin{aligned} \eta(w)_{T_k}^i &= \widehat{\mathbb{P}}^w(X_{T_k}^w = i) = \widehat{\mathbb{P}}^w(J^{X_{T_k}^w}(T_k, w_{T_k}, \mu(T_k-, w_{T_k-})) = i) \\ &= \sum_{j=1}^d \widehat{\mathbb{P}}^w(J^j(T_k, w_{T_k}, \mu(T_k-, w_{T_k-})) = i | X_{T_k-}^w = j) \cdot \widehat{\mathbb{P}}^w(X_{T_k-}^w = j) \\ &= \sum_{j=1}^d \mathbf{1}_{\{J^j(T_k, w_{T_k}, \mu(T_k-, w_{T_k-})) = i\}} \cdot \widehat{\mathbb{P}}^w(X_{T_k-}^w = j) \end{aligned}$$

$$= \sum_{j=1}^d P_k^{ji}(w_{T_k}, \mu(T_k^-, w_{T_k^-})) \cdot \eta(w)_{T_k^-}^j = \Phi_k^i(w_{T_k}, \eta(w)_{T_k^-}, \mu(T_k^-, w_{T_k^-})). \quad (2.22)$$

Since $\eta(W_T) = \sum_{w \in \mathbb{W}^n} \mathbf{1}_{\{W_T=w\}} \cdot \eta(w)$, in view of (2.20), (2.21) and (2.22) it follows from Lemma 2.3.3 that the process $\eta(W_T)$ is \mathfrak{G} -adapted.

Step 2: Identification of $\eta(W_T)$. Recall that $\mathfrak{G}_T = \sigma(W_T) \vee \mathfrak{N}$ and let $w \in \mathbb{W}^n$. For $t \in [0, T]$ and $i \in \mathbb{S}$ we have by (2.6) and (2.17)

$$\begin{aligned} \widehat{\mathbb{E}}[\mathbf{1}_{\{W_T=w\}} \cdot \mathbf{1}_{\{X_t=i\}}] &= \mathbb{E}[\mathbf{1}_{\{W_T=w\}} \cdot \mathbf{1}_{\{X_t^w=i\}} \cdot Z_T^w] = \mathbb{E}[\mathbf{1}_{\{W_T=w\}} \cdot \mathbf{1}_{\{X_t^w=i\}} \cdot \zeta(w)_T \cdot \mathcal{E}[\vartheta]_T] \\ &= \prod_{k=1}^n (|\mathbb{W}| \cdot \kappa_k(w_k | w_1, \dots, w_{k-1}, \mu(T_k^-, w_{T_k^-}))) \cdot \mathbb{E}[\mathbf{1}_{\{W_T=w\}} \cdot \mathbf{1}_{\{X_t^w=i\}} \cdot \zeta(w)_T] \\ &= |\mathbb{W}|^n \cdot \widehat{\mathbb{P}}(W_T = w) \cdot \mathbb{P}(W_T = w) \cdot \widehat{\mathbb{P}}^w(X_t^w = i) = \widehat{\mathbb{E}}[\mathbf{1}_{\{W_T=w\}} \cdot \eta(W_T)_t^i], \end{aligned}$$

where in the final line the first identity is due to Lemma 2.1.3 and \mathbb{P} -independence of $(\zeta(w), X^w)$ and \mathfrak{G}_T ; and the second is due to (2.19) and the fact that $\mathbb{P}(W_T = w) = 1/|\mathbb{W}|^n$. Thus

$$\widehat{\mathbb{P}}(X_t \in \cdot | \mathfrak{G}_T) = \eta(W_T)_t \quad \widehat{\mathbb{P}}\text{-a.s. for } t \in [0, T].$$

Step 3: Dynamics of \widehat{M} . By Step 2 and the tower property of conditional expectation, we find that for each $i \in \mathbb{S}$ and $t \in [0, T]$

$$\widehat{M}_t^i = \widehat{\mathbb{P}}(X_t = i | \mathfrak{G}_t) = \widehat{\mathbb{E}}[\widehat{\mathbb{E}}[\mathbf{1}_{\{X_t=i\}} | \mathfrak{G}_T] | \mathfrak{G}_t] = \widehat{\mathbb{E}}[\eta(W_T)_t^i | \mathfrak{G}_t] = \eta(W_T)_t^i \quad \widehat{\mathbb{P}}\text{-a.s.},$$

where the final identity is due to the fact that $\eta(W_T)$ is \mathfrak{G} -adapted by Step 1 and $\widehat{\mathbb{E}}$ denotes $\widehat{\mathbb{P}}$ -expectation. Since both \widehat{M} and $\eta(W_T)$ are càdlàg, it follows that $\widehat{M} = \eta(W_T)$ $\widehat{\mathbb{P}}$ -a.s., and (M), (M₀) and (M_k) follow from (2.20), (2.21) and (2.22). \square

As a by-product, the preceding proof yields the alternative representation

$$\widehat{M}_t = \widehat{\mathbb{P}}(X_t \in \cdot | \mathfrak{G}_T) \quad \text{for } t \in [0, T] \quad \widehat{\mathbb{P}}\text{-a.s.}$$

Our proof relies on stochastic analysis; the traditional method to establish results of this type uses arguments from continuous-time Markov chain theory. We briefly sketch that approach: Since $\widehat{\nu}$ is Markov, it can be argued using Theorem II.2.T6 in [Bré81] that, conditional on the configuration $w \in \mathbb{W}^n$, N^{ij} is an inhomogeneous Poisson process under $\widehat{\mathbb{P}}$ with intensity

$$\widehat{Q}^{ij}(t, w_t, \mu(t, w_t), v(t, w_t)) \quad \text{on the event } \{X_{t-}^w = i\} \text{ for } t \in [0, T].$$

Similar arguments as in the construction of continuous-time homogeneous Markov chains, see, e.g., [Nor97, p.89f.] or §9.1.2 in [Bré99], can be used to conclude that X is a Markov chain under $\widehat{\mathbb{P}}^w$ with the preceding jump intensities. The dynamics (M) are then derived using (time-inhomogeneous versions of) classical results from Markov chain theory; see, e.g., §9.2.2 in [Bré99] or Theorem 12.22 in [Kal02]. \diamond

2.3.2 Mean Field Equilibrium System

As discussed above, equilibrium obtains if the agents' *ex ante* beliefs coincide with the *ex post* outcome. This holds if and only if the *ex post* aggregate distribution process \widehat{M} from (M) satisfies

$$\widehat{\mathbb{P}}(X_t \in \cdot | \mathfrak{G}_t) = \widehat{M}_t \stackrel{!}{=} M_t = \mu(t, W_t) \quad \text{for all } t \in [0, T].$$

Naturally, we arrive at the following⁶⁸

Definition 2.3.5 (Equilibrium System). A pair (μ, v) of regular and non-anticipative functions

$$\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M} \quad \text{and} \quad v : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^d$$

is called a *rational expectations equilibrium*, *mean field equilibrium*, or briefly an *equilibrium*, if for all $w \in \mathbb{W}^n$

$$\dot{\mu}(t, w) = \mu(t, w) \cdot \widehat{Q}(t, w, \mu(t, w), v(t, w)) \quad (\text{E1})$$

$$\dot{v}(t, w) = -\widehat{\psi}(t, w, \mu(t, w), v(t, w)) - \widehat{Q}(t, w, \mu(t, w), v(t, w)) \cdot v(t, w) \quad (\text{E2})$$

for $t \in [T_k, T_{k+1})$, $k \in [0:n]$, subject to the consistency conditions⁶⁹

$$\mu(T_k, w) = \Phi_k(w, \mu(T_k-, w)) \quad (\text{E3})$$

$$v(T_k-, w) = \Psi_k(w, \mu(T_k-, w), v(T_k, \cdot)) \quad (\text{E4})$$

for $k \in [1:n]$, and the initial/terminal conditions

$$\mu(0, w) = m_0 \quad (\text{E5})$$

$$v(T, w) = \Psi(w, \mu(T, w)). \quad (\text{E6})$$

We also refer to (E1)-(E6) as the *equilibrium system*. \diamond

In combination, Theorem 2.2.4 and Theorem 2.3.4 demonstrate that, given a solution (μ, v) of the equilibrium system, v represents the value function of the agent's optimization problem (P_μ) with *ex ante* aggregate distribution μ ; and the *ex post* distribution resulting from the corresponding optimal strategy is given by μ itself. Theorem 2.3.6 (see below) in turn guarantees existence of such a solution (μ, v) , so a mean field equilibrium obtains.

More precisely, under standard Lipschitz conditions, see Assumption D.1 in Appendix D, we have the following existence result for the equilibrium system; the proof is reported in Appendix D and a ramification of that of Theorem 6 in [CF20]; in particular, it is also based on Banach's fixed point theorem.

⁶⁸For the corresponding *initial-terminal value problem* that pertains in the absence of common noise we refer, e.g., to §2.6 in [GMS13].

⁶⁹With a slight abuse of notation, here and in the following we write

$$\Phi_k(w, m) \triangleq \Phi_k(w, m, m) \quad \text{for } k \in [1:n], w \in \mathbb{W}^n \text{ and } m \in \mathbb{M},$$

where the right-hand side is defined in (Φ_k) above.

Theorem 2.3.6 (Existence of Equilibria). Fix $n \in \mathbb{N}_0$ and let Assumption D.1 be satisfied. Then there exists $T^* > 0$ such that for every time horizon $T \leq T^*$ and every choice of common noise times $0 = T_0 < T_1 < \dots < T_n < T_{n+1} = T$ there is a unique solution of the equilibrium system (E1)-(E6).⁷⁰

Proof. See Appendix D. □

In summary, using Theorems 2.2.4 and 2.3.4 we can identify a mean field equilibrium with common noise by producing a solution of the equilibrium system (E1)-(E6); under the conditions of Theorem 2.3.6, a solution exists and can be computed numerically by a fixed-point iteration. We provide some illustrations in Section 2.4; note that, in order to formulate a model within our mean field setting, it suffices to specify:

- the agent's state space \mathbb{S} ,
- the agent's action space \mathbb{U} ,
- the common noise space \mathbb{W} ,
- the common noise times and the terminal time, $0 = T_0 < T_1 < \dots < T_n < T_{n+1} = T$,
- transition intensities $Q(t, w, m, u)$,
- transition kernels $\kappa_k(w_k | w_1, \dots, w_{k-1}, m)$,
- common noise jumps $J(t, w, m)$,
- reward functions $\psi(t, w, m, u)$ and $\Psi(w, m)$.

2.4 Applications

Before we illustrate our results in three showcase examples, we briefly discuss our numerical approach to the equilibrium system (E1)-(E6). (E1)-(E2) is a forward-backward system of $2d$ ODEs with boundary conditions (E3)-(E6), coupled through the parameter $w \in \mathbb{W}^n$ representing common noise configurations. The special case $n = 0$ (no common noise) corresponds to the setting of [GMS13] and [CF20], with the equilibrium system reducing to a single $2d$ -dimensional forward-backward ODE. For $n \geq 1$, the consistency conditions (E3)-(E4) specify initial conditions for μ on $[T_k, T_{k+1})$ and terminal conditions for v on $[T_{k-1}, T_k)$, $k \in [1:n]$; since these conditions are interconnected, there is in general *no* segment $[T_k, T_{k+1}) \times \mathbb{W}^n$ where the equilibrium system yields both an explicit initial condition for μ and an explicit terminal condition for v , so we cannot simply split the problem into subintervals. Rather, the equilibrium system can be regarded as a *multi-point boundary value problem* where for each of the $|\mathbb{W}|^k$ conceivable combinations of common noise factors on $[T_k, T_{k+1})$, $k \in [0:n]$, we have to solve a coupled forward-backward system of ODEs in $2d$ dimensions, resulting in a tree of such systems of size $\sum_{k=0}^n |\mathbb{W}|^k = \frac{|\mathbb{W}|^{n+1} - 1}{|\mathbb{W}| - 1} \in \mathcal{O}(|\mathbb{W}|^n)$; see Figure 2.2.

⁷⁰Note that although Banach's fixed-point theorem also yields uniqueness of the fixed point, the resulting mean field equilibrium need not be unique unless the maximizer h in Assumption 2.2.1 is uniquely determined. For details on uniqueness of mean field equilibria (in setups without a common noise) see also, e.g., §2 in [GMS13] or §4 in [CF20].




Figure 2.2: Illustration of all possible common noise paths ($|\mathbb{W}| = 2$, $n = 3$): The tree of FBODEs.

Our approach to solving (E1)-(E6) is to rely on the probabilistic interpretation as a fixed-point system, based on Theorem 2.3.6. Thus, starting from an initial flow of probability weights $\mu_0(t, w)$, $(t, w) \in [0, T] \times \mathbb{W}^n$, with $\mu_0(0, w) = m_0$ for all $w \in \mathbb{W}^n$, we solve (DP $_{\mu}$) subject to (TC $_{\mu}$) and (CC $_{\mu}$) backward in time for all (non-negligible) common noise configurations $w \in \mathbb{W}^n$ to obtain the value $v_0(t, w)$, $(t, w) \in [0, T] \times \mathbb{W}^n$, of the agents' optimal response to the given belief μ_0 . This, in turn, is used to solve (M) subject to (M $_0$) and (M $_k$) forward in time. As a result, we obtain an ex post aggregate distribution $\mu_1(t, w)$, $(t, w) \in [0, T] \times \mathbb{W}^n$; we then iterate this with μ_1 in place of μ_0 , etc., and thus construct two sequences $\{\mu_j\}_{j \in \mathbb{N}_0}$ and $\{v_j\}_{j \in \mathbb{N}_0}$.⁷¹ Note that Theorem 2.3.6 guarantees the convergence of this methodology for appropriate parameter settings. Our implementation of this fixed-point iteration is in MATLAB;⁷² all relevant ODEs are solved numerically using an explicit Euler scheme with equidistant step size (see below); and the iteration is stopped with iteration index $\iota \in \mathbb{N}$ and an approximate equilibrium $(\mu_{\iota+1}, v_{\iota})$ as soon as the sum of the maximum pointwise difference between the respective current and last iterate of the forward and the backward system is less than or equal to 10^{-4} , i.e. we stop at iteration index $\iota \in \mathbb{N}$ as soon as

$$\sup_{(t,w) \in [0,T] \times \mathbb{W}^n} \|\mu_{\iota}(t, w) - \mu_{\iota-1}(t, w)\|_{\infty} + \sup_{(t,w) \in [0,T] \times \mathbb{W}^n} \|v_{\iota}(t, w) - v_{\iota-1}(t, w)\|_{\infty} \leq 10^{-4}.$$

⁷¹This approach was also chosen in the example treated in §7.2.3 in [CD18a], albeit without common noise.

⁷²We use MATLAB R2019b [Mat19]; no explicit parallelization; OS: Linux Mint 18.1 Serena. Hardware: Intel Core i7-6700 CPU @ 3.40 GHz (4 cores) with 32 GiB RAM capacity.

2.4.1 A Decentralized Agricultural Production Model

As a first (stylized) application we consider a mean field game of agents, each of which owns (an infinitesimal amount of) arable land of identical size and quality within a given area; see Figure 2.3 for a schematic illustration. If it is farmed, each field has a productivity $f(w_k) > 0$ that depends on the common weather condition w_k . For simplicity, we assume that weather is either good, bad or catastrophic, so $w_k \in \mathbb{W} \triangleq \{\uparrow, \downarrow, \downarrow\}$, and changes at given common noise times T_1, \dots, T_n .

Each agent is in exactly one state $i \in \mathbb{S} \triangleq \{0, 1\}$ (hence, $d = 2$) depending on whether he has entered the market and thus grows crops on his field ($i = 1$, the agent is a *farmer*) or not ($i = 0$, the agent is a *non-farmer*).

The selling price p for his harvest depends on the *aggregate* production, and thus in particular on the proportion $m^1 \in [0, 1]$ of farmers; the mean field interaction is transmitted through the market price of the crop. We assume that p is a strictly decreasing function of overall production $f(w_k) \cdot m^1$; see Figure 2.4 for illustration.




Figure 2.3: Stylized area of farmland with $N = 9$ separate fields, $k = 4$ of which are farmed (crossed fields): The overall production within the area is proportional to $k/N \cdot f(w_k) + N-k/N \cdot 0 = k/N \cdot f(w_k)$.




Figure 2.4: Price and instantaneous profit function of the agents (parameters as in Table 2.1).

We assume that $f(\uparrow) \geq f(\downarrow) = f(\downarrow) \geq 0$. Moreover, on the catastrophic event $\{W_k = \downarrow\}$ all agents are reduced to being non-farmers, and thus

$$J^i(t, w, m) \triangleq \begin{cases} 0 & \text{if } t \in \{T_1, \dots, T_n\}, t = T_k, w_k = \downarrow, \\ i & \text{else,} \end{cases}$$

for $(i, t, w, m) \in \mathbb{S} \times [0, T] \times \mathbb{W}^n \times \mathbb{M}$. Each agent can make an effort $u \in \mathbb{U} \triangleq [0, 1]$ to become or stop being a farmer, i.e. u models his propensity to change his current state; the intensity matrix for state transitions is given by

$$Q(t, w, m, u) \triangleq u \cdot \begin{bmatrix} -q_{\text{entry}} & q_{\text{entry}} \\ q_{\text{exit}} & -q_{\text{exit}} \end{bmatrix} \quad \text{for } (t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U},$$

where $q_{\text{entry}}, q_{\text{exit}} \geq 0$ are given maximum rates. The running rewards capture the fact that both efforts to build up farming capacities and production itself are costly, while the sales of the crop generate revenues; thus, for $(t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$,

$$\begin{aligned} \psi^0(t, w, m, u) &\triangleq -\frac{1}{2}c_{\text{entry}} \cdot u^2, \\ \psi^1(t, w, m, u) &\triangleq p(f(w_k) \cdot m^1) \cdot f(w_k) - c_{\text{prod}} \quad \text{if } t \in [T_k, T_{k+1}), k \in [0:n], \end{aligned}$$

where $w_0 \triangleq \uparrow$ and $c_{\text{entry}}, c_{\text{prod}} > 0$; for the profit function see also Figure 2.4. The terminal reward is zero. It follows that the unique maximizer in Assumption 2.2.1 is given by

$$h^0(t, w, m, v) = \left[\frac{q_{\text{entry}}}{c_{\text{entry}}} (v^1 - v^0)^+ \right] \wedge 1 \quad \text{and} \quad h^1(t, w, m, v) = \begin{cases} 0 & \text{if } v^1 \geq v^0, \\ 1 & \text{else,} \end{cases}$$

for $(t, w, m, v) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^2$. ◇

We choose $m_0^1 \triangleq 10\%$ for the initial amount of farmers, and report the relevant coefficients in Table 2.1. Note that weather conditions can change at $n = 4$ common noise times $T_k = k/5$, $k = 1, \dots, 4$.

The fixed-point iteration is initialized by

$$\mu_0(t, w) \triangleq e^{-\lambda(t-T_k)} \cdot \mu_0(T_k, w) + (1 - e^{-\lambda(t-T_k)}) \cdot \bar{m}(w_k) \quad \text{for } (t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n, k \in [0:n],$$

$$\text{where } \mu_0(T_k, w) \triangleq \begin{cases} m_0 & \text{if } k = 0, \\ \Phi_k(w, \mu_0(T_{k-}, w)) & \text{if } k \in [1:n]; \end{cases}$$

it holds $\bar{m}(w_k) = [1 - \bar{m}^1(w_k), \bar{m}^1(w_k)] \in \mathbb{M}$ with $\bar{m}^1(w_k)$ denoting the *break-even* proportion of farmers given by the unique root of $m^1 \mapsto p(f(w_k) \cdot m^1) \cdot f(w_k) - c_{\text{prod}}$ for $w_k \in \mathbb{W}$; ⁷³ and $\lambda \triangleq 10$ determines how fast $\mu_0(\cdot, w)$ approaches $\bar{m}(w_k)$ on each segment $[T_k, T_{k+1}) \times \mathbb{W}^n$, $k \in [0:n]$.

The step size of the Euler scheme is set to 10^{-4} .

Our results for the evolution of the mean field equilibrium are shown in Figures 2.5 through 2.8 for various common noise configurations $w \in \mathbb{W}^n$ and the following two baseline models:

⁷³For parameters as specified in Table 2.1 it holds $\bar{m}^1(\uparrow) = 7/9$ and $\bar{m}^1(\downarrow) = \bar{m}^1(\downarrow) = 4/9$; see Figure 2.4.

Parameter	T	n	T_k	$q_{\text{entry}}, q_{\text{exit}}$	$f(\uparrow)$	$f(\downarrow)$	$p(q)$	c_{prod}	c_{entry}
Value	1	4	$k/5$	0.7	1	0.5	$1/(1+3q)$	0.3	0.1

Table 2.1: Model parameters *Agricultural Production*.

(nC) Catastrophic weather conditions do not occur; we use

$$\kappa_k(\uparrow | w_1, \dots, w_{k-1}, m) \triangleq \kappa_k(\downarrow | w_1, \dots, w_{k-1}, m) \triangleq 0.5$$

for all $w \in \mathbb{W}^n$ and $m \in \mathbb{M}$.

(C) Catastrophic events are likely; we use

$$\kappa_k(\uparrow | w_1, \dots, w_{k-1}, m) \triangleq 0.25,$$

$$\kappa_k(\downarrow | w_1, \dots, w_{k-1}, m) \triangleq 0.25,$$

$$\kappa_k(\downarrow | w_1, \dots, w_{k-1}, m) \triangleq 0.5$$

for all $w \in \mathbb{W}^n$ and $m \in \mathbb{M}$.

Figures 2.5, 2.6 and 2.7 illustrate the resulting equilibrium proportions of farmers, optimal actions, and market prices for some fixed common noise configurations. To illustrate the effect of uncertainty about future weather conditions we also show, for each common noise configuration, the theoretical perfect-foresight equilibria that would pertain if future weather conditions were known; these are plotted using dashed lines in Figures 2.5 through 2.7, and the subscript \circ indicates the relevant *deterministic* common noise path. Finally, Figure 2.8 illustrates the tree of all possible equilibrium evolutions in model (C).

Note that the impact of common noise, i.e. uncertainty about future weather conditions, on the equilibrium proportion of farmers is significant, see Figure 2.5: If catastrophic events cannot occur, see model (nC), in case of consistently good weather conditions fewer agents become a farmer as compared to the corresponding perfect-foresight equilibrium. This is due to the fact that, as long as there are still unrevealed common noise factors, agents cannot foresee whether currently advantageous weather conditions will also prevail in the future, and thus they behave more cautious. Analogously, on the path of consistently bad weather conditions more of them start farming than under perfect foresight since weather conditions might improve.

In comparison, when faced with a high probability of catastrophic events, see model (C), agents are more hesitant to enter the market as their entire farming capacities are wiped out every time a catastrophe occurs and therefore all previous efforts to become a farmer have been in vain in this case. This is also reflected by the non-farmers' optimal actions at times before the final common noise factor is revealed, see Figure 2.6: If catastrophic events are likely, efforts to become a farmer are smaller and decrease more steeply between the common noise times.

Finally, note that equilibrium prices are stochastically modulated by the prevailing weather conditions, see Figure 2.7.

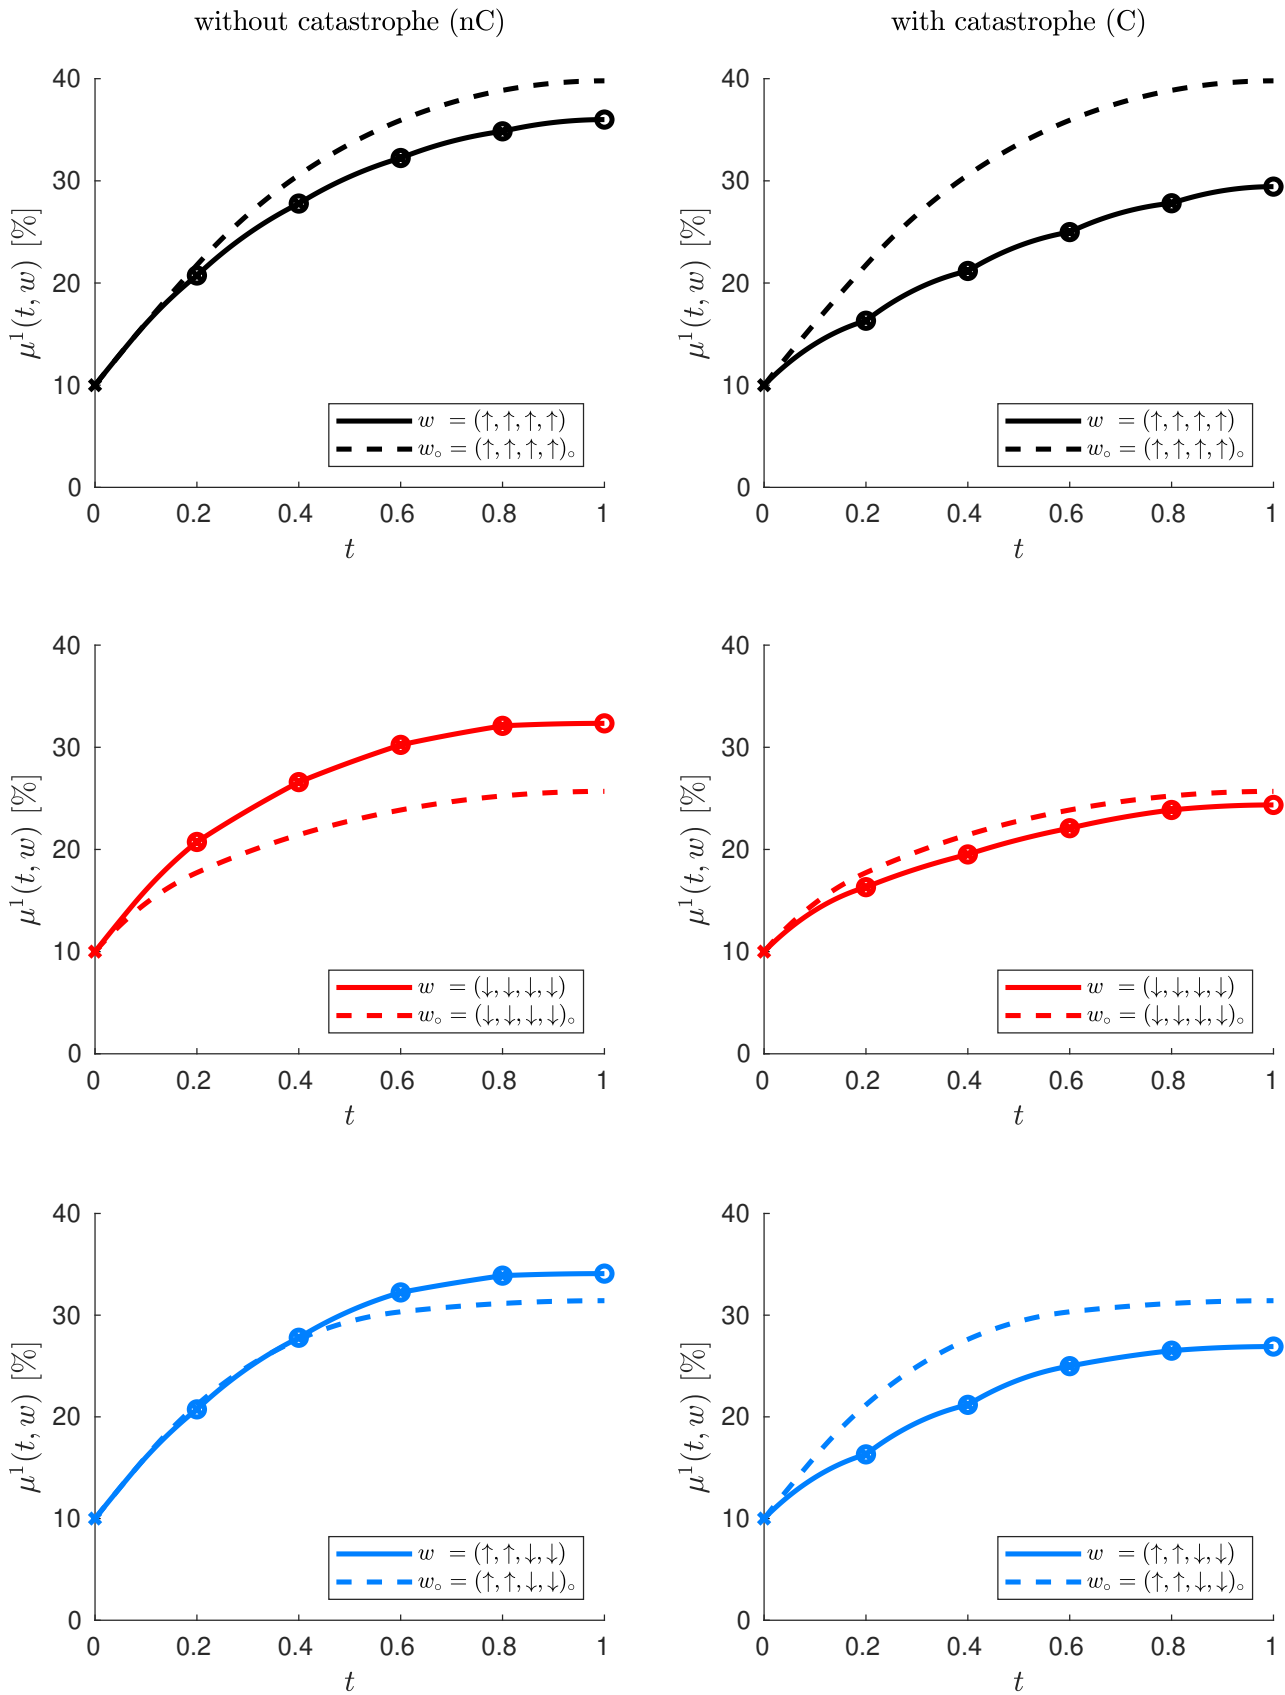


Figure 2.5: Equilibrium proportion of farmers in models (nC) and (C).

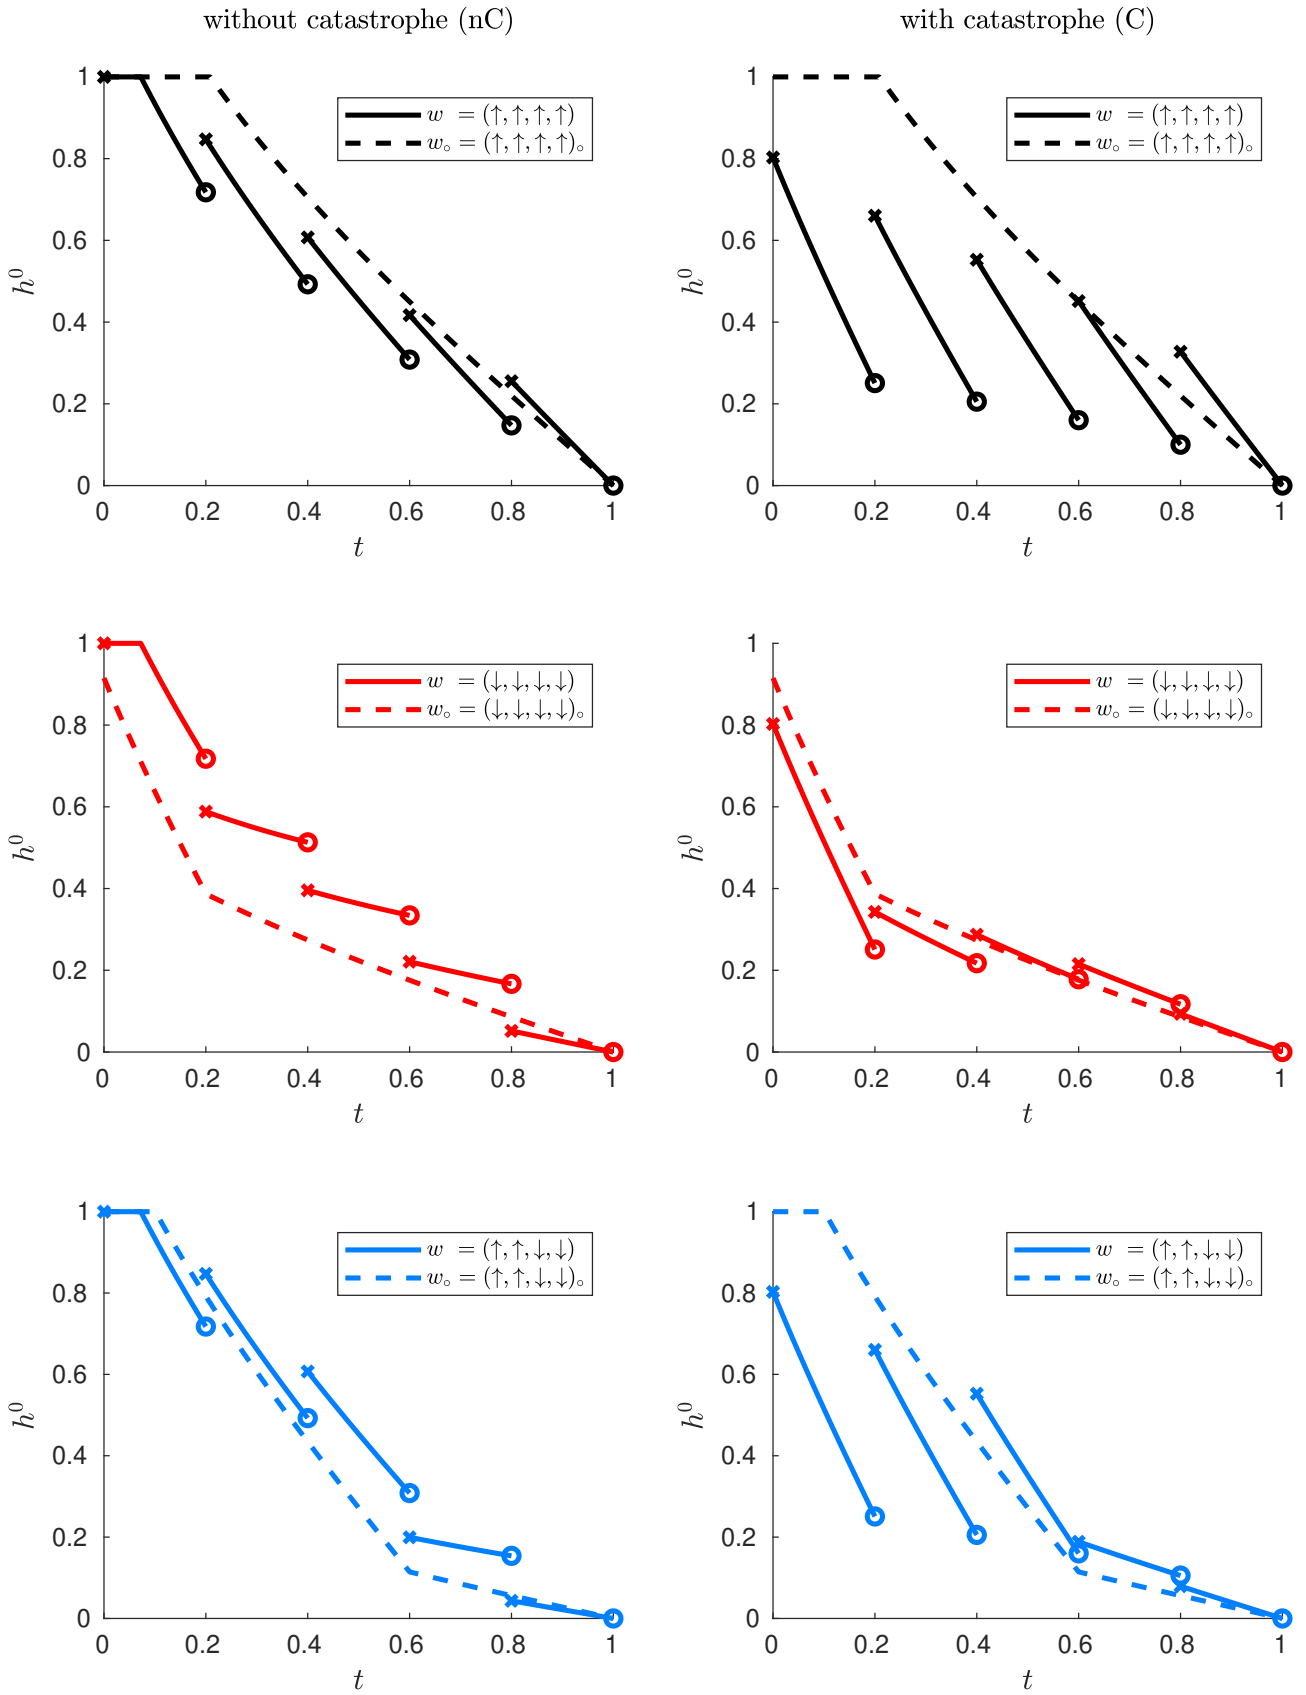


Figure 2.6: Equilibrium optimal action $h^0(t, w, \mu(t, w), v(t, w))$ of non-farmers in models (nC) and (C).⁷⁴

⁷⁴Note that the respective mean field equilibrium is denoted by (μ, v) .

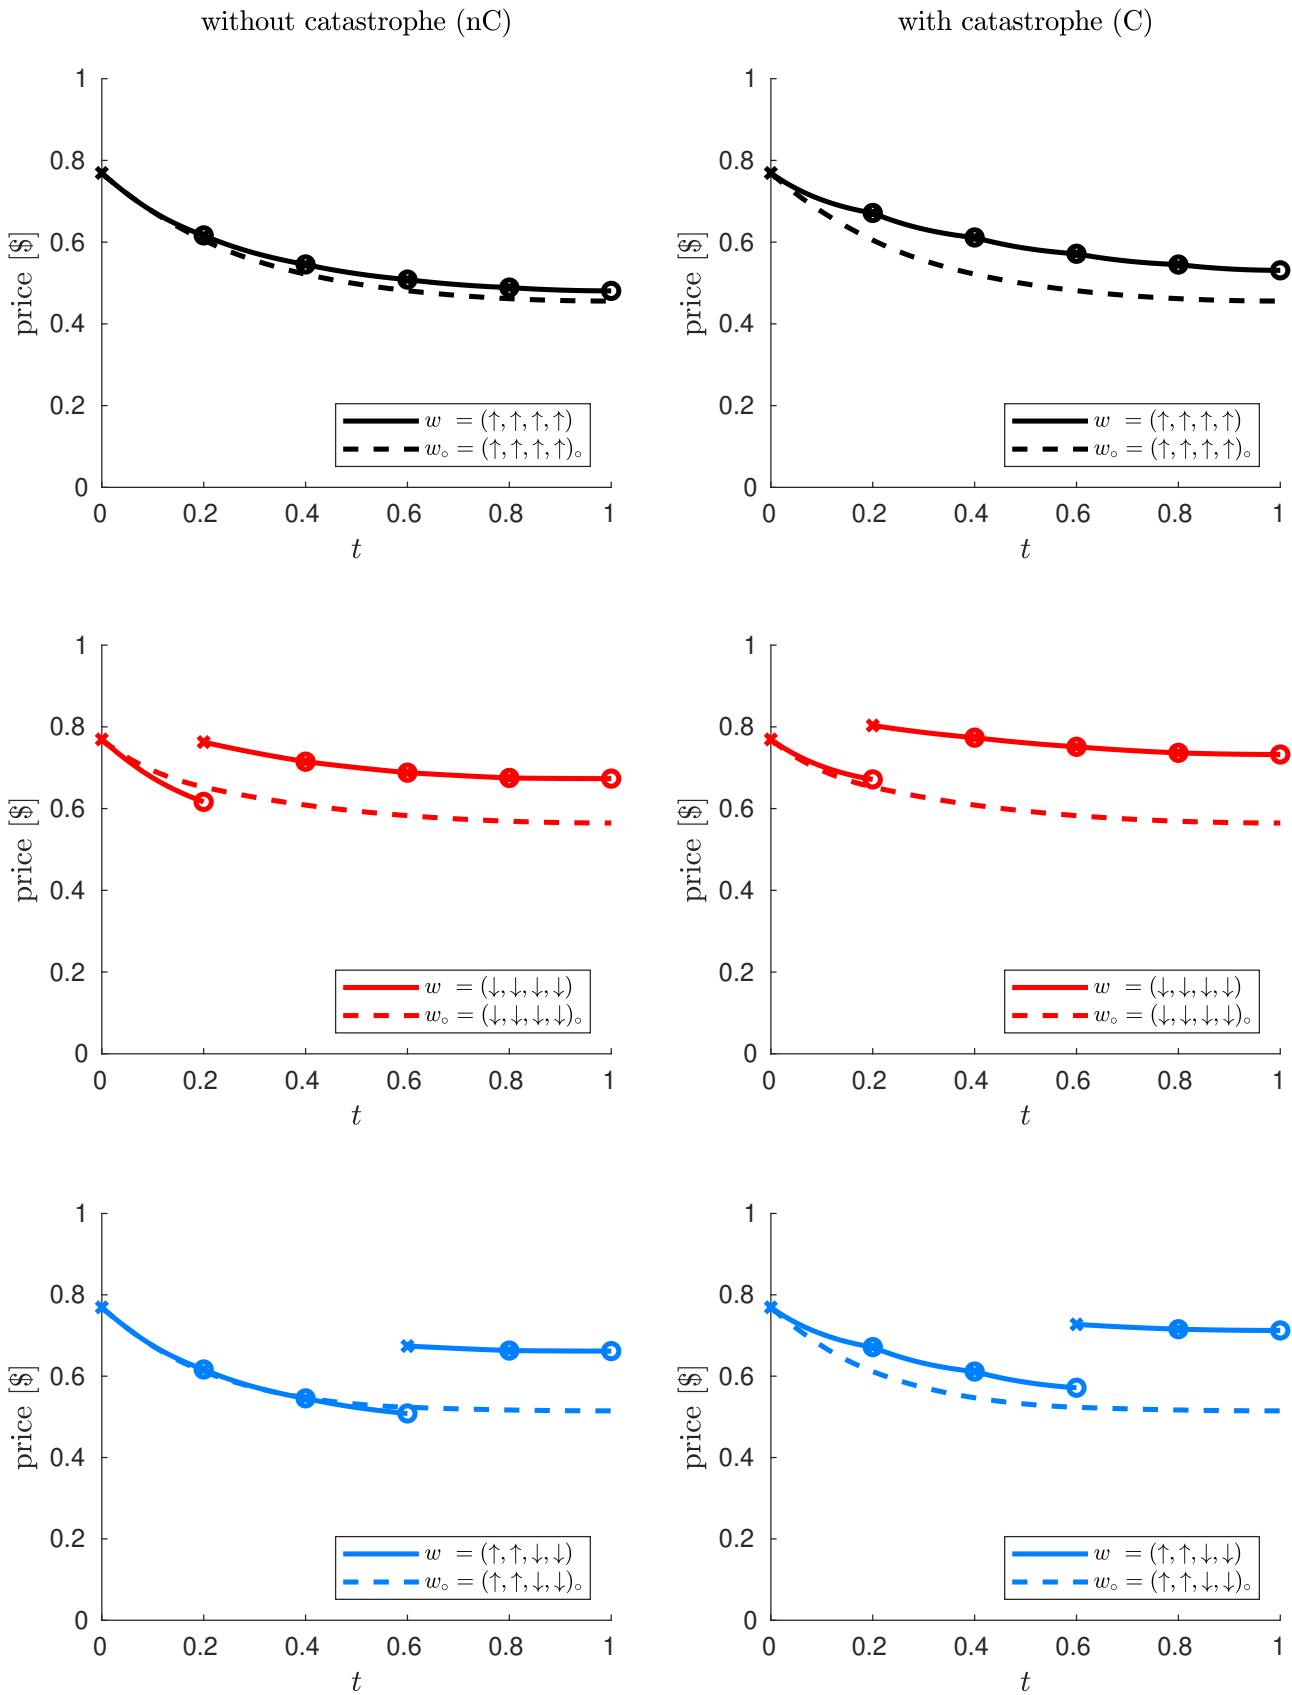


Figure 2.7: Equilibrium market price $p(f(w_k) \cdot \mu^1(t, w))$ for $t \in [T_k, T_{k+1})$, $k \in [1:n]$ in models (nC) and (C).⁷⁴




Figure 2.8: Equilibrium proportion of farmers in model (C) for all possible common noise configurations $w \in \mathbb{W}^n$.

2.4.2 An SIR Model with Random One-Shot Vaccination

Our second application example is a mean field game of agents that are confronted with the spread of an infectious disease where our main focus is to illustrate the qualitative effects of common noise on the equilibrium behavior of the system. We consider a classical SIR model setup⁷⁵ with $\mathbb{S} = \{S, I, R\}$ (hence, $d = 3$): Each agent can be either *susceptible* to infection (S), *infected* and simultaneously *infectious* for other agents (I), or *recovered* and thus immune to (re-)infection (R); see Figure 2.9.




Figure 2.9: State space and transitions in the SIR model.

The infection rate is proportional to the prevalence of the disease, i.e. the percentage of currently infected agents. Susceptible agents can make individual efforts of size $u \in \mathbb{U} \triangleq [0, 1)$ to protect

⁷⁵see Remark 2.4.1 below for further details and related references.

themselves against infection and thus reduce the intensity of infection. The transition intensities are given by

$$Q(t, w, m, u) \triangleq \begin{bmatrix} -q_{\text{inf}}(t, w, m, u) & q_{\text{inf}}(t, w, m, u) & 0 \\ 0 & -q_{\text{IR}} & q_{\text{IR}} \\ 0 & 0 & 0 \end{bmatrix}$$

for $(t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$, where $q_{\text{IR}} \geq 0$ denotes the recovery rate of infected agents and the infection rate is given by

$$q_{\text{inf}}(t, w, m, u) \triangleq q_{\text{SI}} \cdot m^{\text{I}} \cdot (1 - u) \cdot \mathbf{1}_{\{t < \tau^*\}}(w) \quad \text{for } (t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$$

with a given maximum rate $q_{\text{SI}} \geq 0$. The running reward penalizes both protection efforts and time spent in the infected state; correspondingly, with $c_{\text{P}}, \psi_{\text{I}} \geq 0$ we set

$$\psi^{\text{S}}(t, w, m, u) \triangleq -c_{\text{P}} \frac{u}{1 - u}, \quad \psi^{\text{I}}(t, w, m, u) \triangleq -\psi_{\text{I}}, \quad \psi^{\text{R}}(t, w, m, u) \triangleq 0$$

for $(t, w, m, u) \in [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$; see Figure 2.10 for an illustration of ψ^{S} .




Figure 2.10: Running reward of susceptible agents: Penalization for making protection efforts (parameters as in Table 2.2).

In addition, we include the possibility of a one-shot vaccination that becomes available, simultaneously to all agents, at a random point of time $\tau^* \in \{T_1, \dots, T_n\} \subset (0, T)$. We set $\mathbb{W} \triangleq \{0, 1\}$ and identify the k^{th} unit vector $e_k = (\delta_{jk})_{j \in [1:n]} \in \mathbb{W}^n$, $k \in [1:n]$ with the indicator of the event $\{\tau^* = T_k\}$. The event that no vaccine is available until T is represented by $0 \in \mathbb{W}^n$; we set $\tau^* \triangleq +\infty$ in this case.^{76, 77} If and when it is available, all susceptible agents are vaccinated instantaneously, rendering them immune to infection; thus

$$J^{\text{S}}(t, w, m) \triangleq \begin{cases} \text{R} & \text{if } t \in \{T_1, \dots, T_n\}, t = T_k = \tau^*, \\ \text{S} & \text{otherwise} \end{cases} \quad \text{and} \quad J^i(t, w, m) \triangleq i \quad \text{for } i \in \{\text{I}, \text{R}\}$$

⁷⁶Note that, with a slight abuse of notation, we omit the formal dependence on w in the notation τ^* .

⁷⁷The specification of κ_k , $k \in [1:n]$, below implies that $\tau^* = +\infty$ is equivalent to $w = 0 \in \mathbb{W}^n$ \mathbb{P} -a.s., i.e., the configurations in $\mathbb{W}^n \setminus (\{0\} \cup \{e_k : k \in [1:n]\})$ are \mathbb{P} -negligible.

for $(t, w, m) \in [0, T] \times \mathbb{W}^n \times \mathbb{M}$. The probability of vaccination becoming available is proportional to the percentage of agents that have already recovered from the disease. Thus, for $k \in [1:n]$, $w_1, \dots, w_k \in \mathbb{W}$ and $m \in \mathbb{M}$, we set

$$\kappa_k(1 | w_1, \dots, w_{k-1}, m) \triangleq \begin{cases} \alpha \cdot m^{\mathbb{R}} & \text{if } w_1, \dots, w_{k-1} = 0, \\ 0 & \text{otherwise,} \end{cases}$$

and $\kappa_k(0 | w_1, \dots, w_{k-1}, m) \triangleq 1 - \kappa_k(1 | w_1, \dots, w_{k-1}, m)$ with $\alpha \in [0, 1]$. As a consequence, for all $(i, t, w, m, v) \in \mathbb{S} \times [0, T] \times \mathbb{W} \times \mathbb{M} \times \mathbb{R}^3$, a maximizer as required in Assumption 2.2.1 is given by⁷⁸

$$h^{\mathbb{S}}(t, w, m, v) \triangleq \begin{cases} \left[1 - \sqrt{\frac{c_{\mathbb{P}}}{q_{\mathbb{S}\mathbb{I}} \cdot m^{\mathbb{I}} \cdot (v^{\mathbb{S}} - v^{\mathbb{I}})}}\right]^+ & \text{if } v^{\mathbb{S}} > v^{\mathbb{I}} \text{ and } m^{\mathbb{I}} > 0, \\ 0 & \text{otherwise,} \end{cases}$$

and $h^i(t, w, m, v) \triangleq 0$ for $i \in \{\mathbb{I}, \mathbb{R}\}$. ◇

Remark 2.4.1 (SIR Models in the Literature). Note that, given the above specification of the transition matrix Q , the forward dynamics (E1) within the equilibrium system (E1)-(E6) read as follows:

$$\begin{cases} \dot{\mu}^{\mathbb{S}}(t, w) = -q_{\mathbb{S}\mathbb{I}} \cdot \mu^{\mathbb{I}}(t, w) \cdot (1 - h^{\mathbb{S}}(t, w, \mu(t, w), v(t, w))) \cdot \mathbb{1}_{\{t < \tau^*\}}(w) \cdot \mu^{\mathbb{S}}(t, w) \\ \dot{\mu}^{\mathbb{I}}(t, w) = q_{\mathbb{S}\mathbb{I}} \cdot \mu^{\mathbb{I}}(t, w) \cdot (1 - h^{\mathbb{S}}(t, w, \mu(t, w), v(t, w))) \cdot \mathbb{1}_{\{t < \tau^*\}}(w) \cdot \mu^{\mathbb{S}}(t, w) - q_{\mathbb{I}\mathbb{R}} \cdot \mu^{\mathbb{I}}(t, w) \\ \dot{\mu}^{\mathbb{R}}(t, w) = q_{\mathbb{I}\mathbb{R}} \cdot \mu^{\mathbb{I}}(t, w). \end{cases}$$

Disregarding common noise, these constitute a ramification of the classical SIR dynamics, which are a basic building block of numerous compartmental epidemic models in the literature; see, among others, [Het00], [HLM14], [Mil17] or [GMS20] and the references therein. The SIR mean field game with controlled infection rates, albeit without common noise, has recently been studied in the independent article [EHT20];⁷⁹ we also refer, e.g., to [LT15] and [DGG17] for mean field models with controlled (instantaneous) vaccination or controlled vaccination rates, respectively; see also the references therein. Mathematically similar contagion mechanisms also appear in, e.g., [KB16], [KM17], §7.2.3 in [CD18a], §7.1.10 in [CD18b], or §4.4 in [Wan19]. ◇

For our numerical results, the initial distribution of agents is given by $m_0 \triangleq (0.995, 0.005, 0.00)$, and the model coefficients are reported in Table 2.2. Note that there are $n = 1999$ common noise times $T_k = k \cdot 0.01$, $k = 1, \dots, 1999$, at which a vaccine can be administered. The specifications of $q_{\mathbb{S}\mathbb{I}}$ and $q_{\mathbb{I}\mathbb{R}}$ imply a basic reproduction number $R_0 \triangleq q_{\mathbb{S}\mathbb{I}}/q_{\mathbb{I}\mathbb{R}} = 15$ in the absence of vaccination and protection efforts.⁸⁰

The fixed-point iteration⁸¹ is initialized by $\mu_0(t, w) \equiv m_0$ for all $(t, w) \in [0, T] \times \mathbb{W}^n$; the step size of the Euler scheme is set to $5 \cdot 10^{-4}$.

⁷⁸Note that for given $w \in \mathbb{W}^n$ the stated maximizer $h^{\mathbb{S}}$ is unique for times $t < \tau^*$; otherwise its specification is immaterial. The latter applies likewise to $h^{\mathbb{I}}$ and $h^{\mathbb{R}}$.

⁷⁹The authors study in particular the SIR mean field game with a controlled *contact rate* against the backdrop of the COVID-19 pandemic in 2020; this contact rate has an equivalent effect as the above protection efforts since it allows susceptible agents to reduce their social interactions and hence to mitigate their individual infection risk.

⁸⁰see, e.g., §2.2 and §2.3 in [Het00] for further details.

⁸¹Note that in order to cope with the sheer number of $n = 1999$ common noise times our implementation must disregard negligible branches appearing in the tree of FBODEs to be solved (cf. Figure 2.2).

Parameter	T	n	T_k	α	q_{SI}	q_{IR}	c_P	ψ_I
Value	10	1999	$k \cdot 0.01$	0.1	7.5	0.5	25	100

Table 2.2: Coefficients in the SIR model.

Our results for the mean field equilibrium distributions of agents μ and the corresponding optimal protection efforts of susceptible agents h^S are displayed in Figures 2.11 through 2.14 for different common noises configurations, i.e. vaccination times τ^* .⁸² As in Section 2.4.1, we also display the corresponding (theoretical) perfect-foresight equilibria, marked by the subscript \circ .

Note that an agent's running reward is the same in state S with zero protection effort and in state R; agents are penalized relative to these in state I and hence aim to avoid that state. Susceptible agents can reach the state R of immunity by two ways: First, they can become infected and overcome the disease; second, they can be vaccinated and jump instantly from state S to state R. While the first alternative is painful, the second comes at no cost and is hence clearly preferable. However, as the availability of a vaccine cannot be directly controlled by the agents, they can only protect themselves against infection at a certain running cost until the vaccine becomes available. Figures 2.11 to 2.14 demonstrate that the possibility of vaccination as a common noise event can dampen the spread of the disease and lower the peak infection rate. This is due to an increase in agents' protection efforts *during the time period when the proportion of infected agents is high*. By contrast, in the perfect-foresight equilibria where the vaccination date is known, agents do not make substantial protection efforts until the vaccination date is imminent, see Figures 2.12, 2.13 and 2.14; in the scenario without vaccination, see Figure 2.11, protection efforts are only ever made by a very small fraction of the population. With perfect foresight, the agents' main rationale is to avoid being in state I when the vaccine becomes available. This highlights the importance of being able to model the vaccination date as a (random) common noise event. Finally, observe that our numerical results indicate convergence to the stationary distribution $\bar{\mu} = (0, 0, 1) \in \mathbb{M}$, showing that the model is able to capture the entire evolution of an epidemic.

⁸²The displayed protection efforts are given by $h^S(t, w, \mu(t, w), v(t, w))$ where (μ, v) denotes the respective mean field equilibrium.

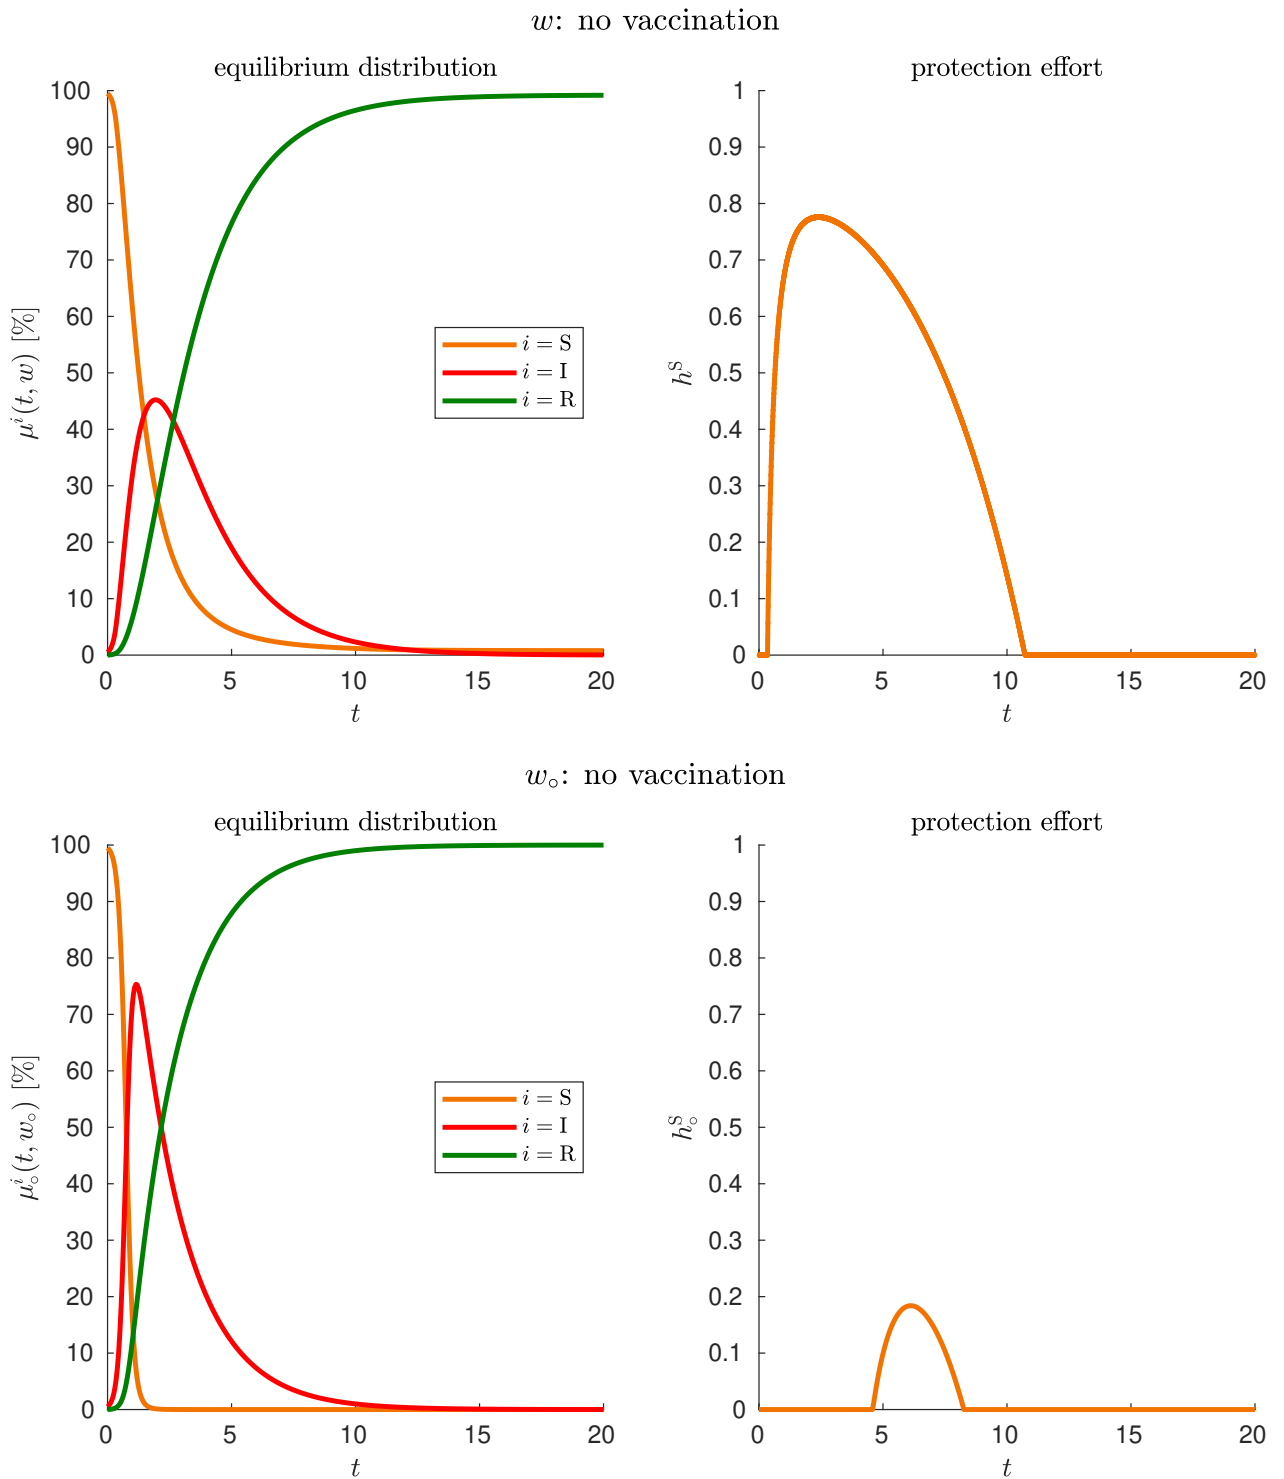


Figure 2.11: Equilibrium distribution and protection effort for $\tau^* = +\infty$: Mean field game with common noise (top) and corresponding perfect-foresight equilibrium (bottom).

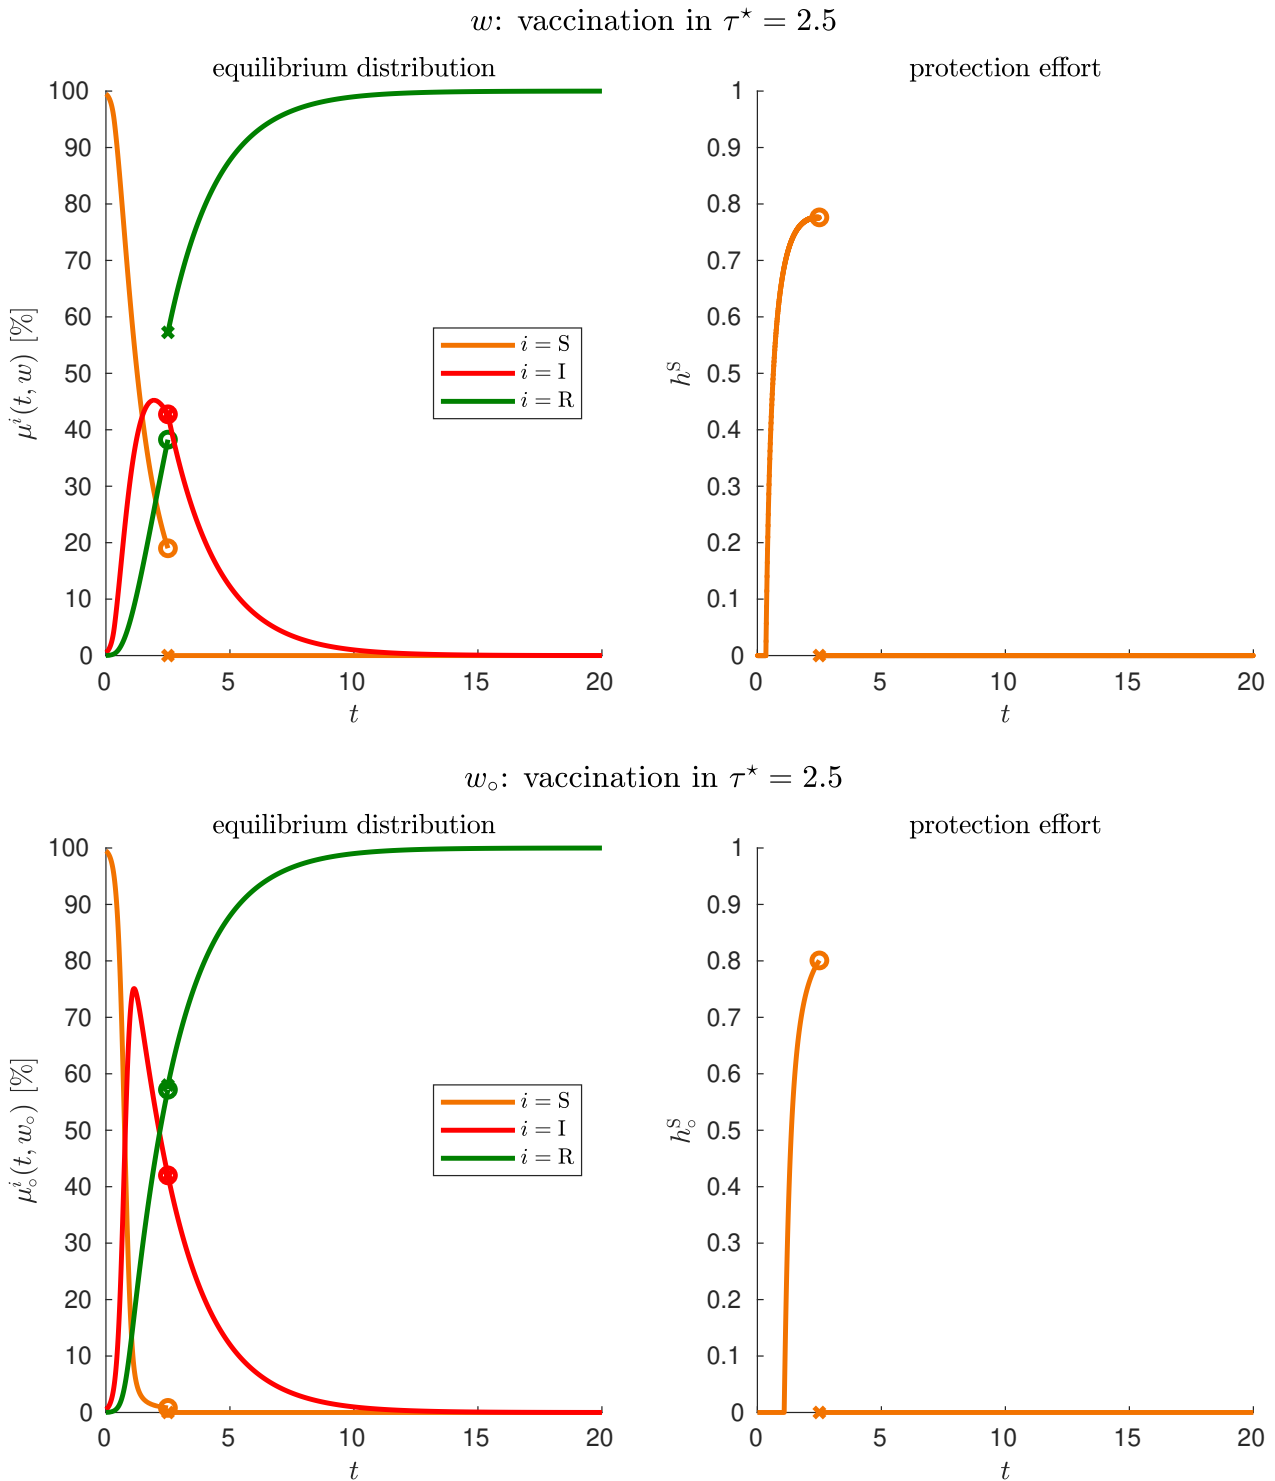


Figure 2.12: Equilibrium distribution and protection effort for $\tau^* = 2.5$: Mean field game with common noise (top) and corresponding perfect-foresight equilibrium (bottom).

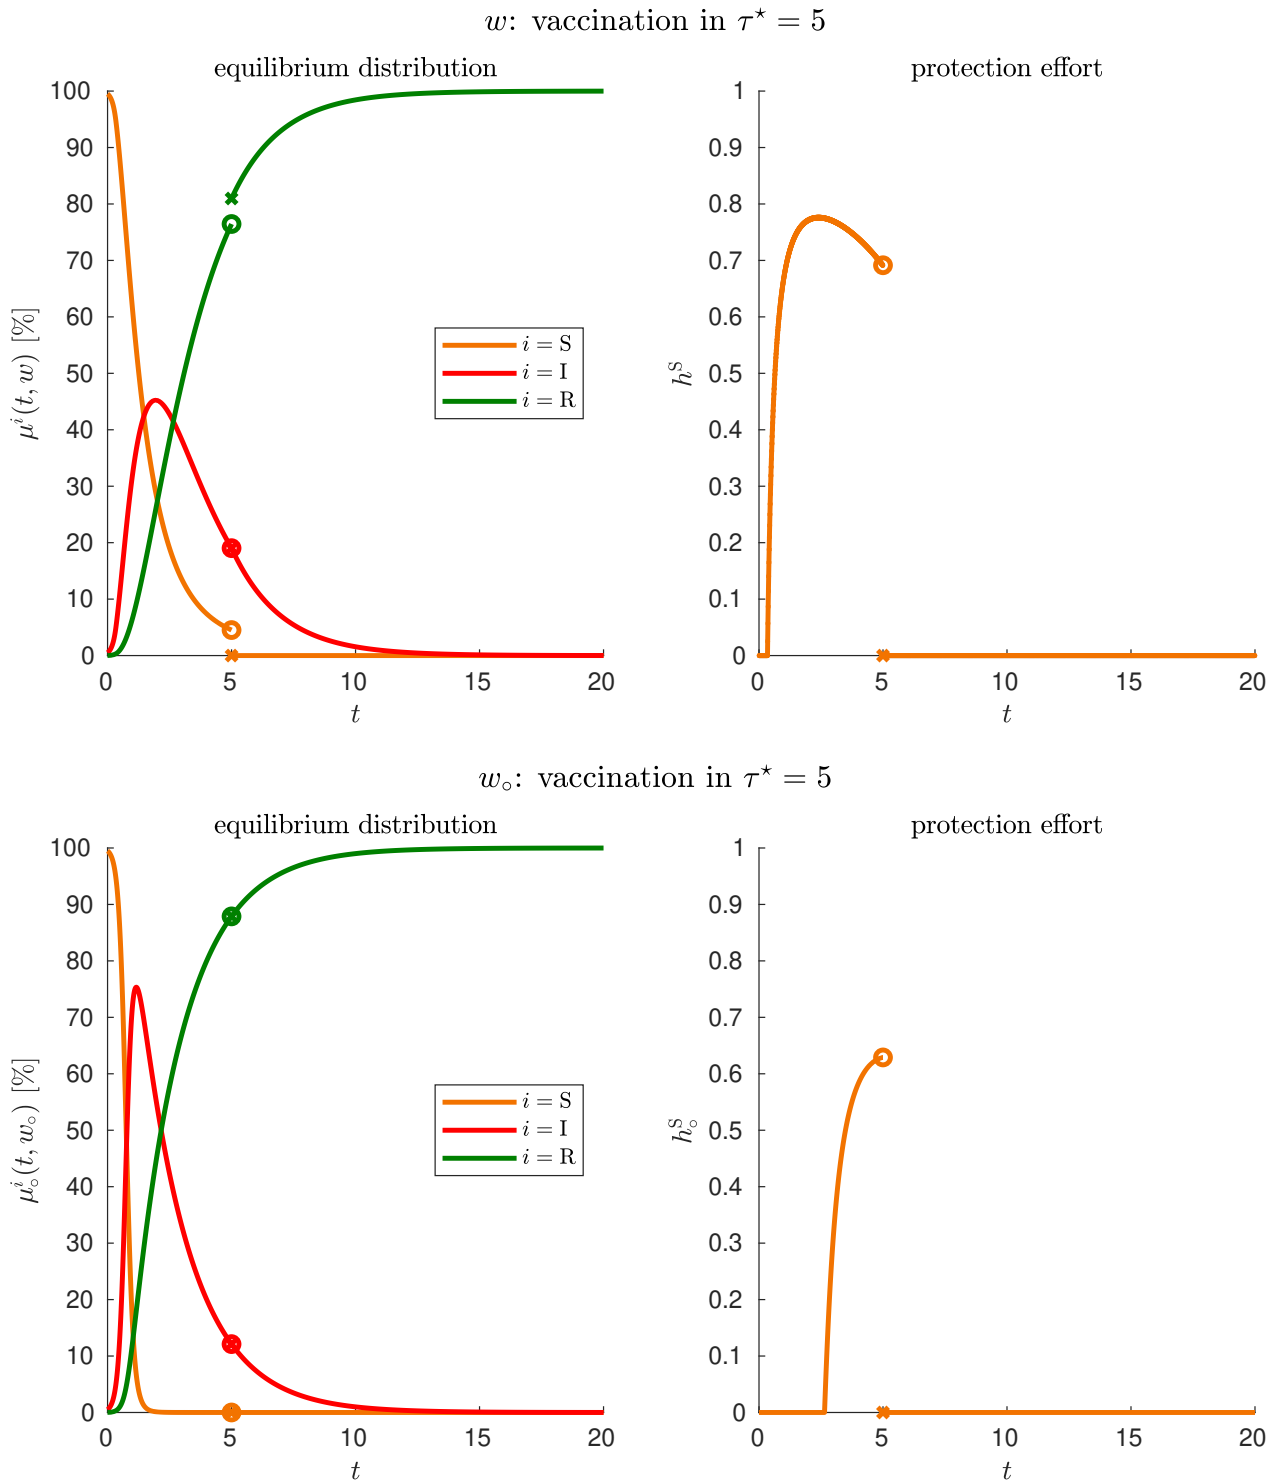


Figure 2.13: Equilibrium distribution and protection effort for $\tau^* = 5$: Mean field game with common noise (top) and corresponding perfect-foresight equilibrium (bottom).

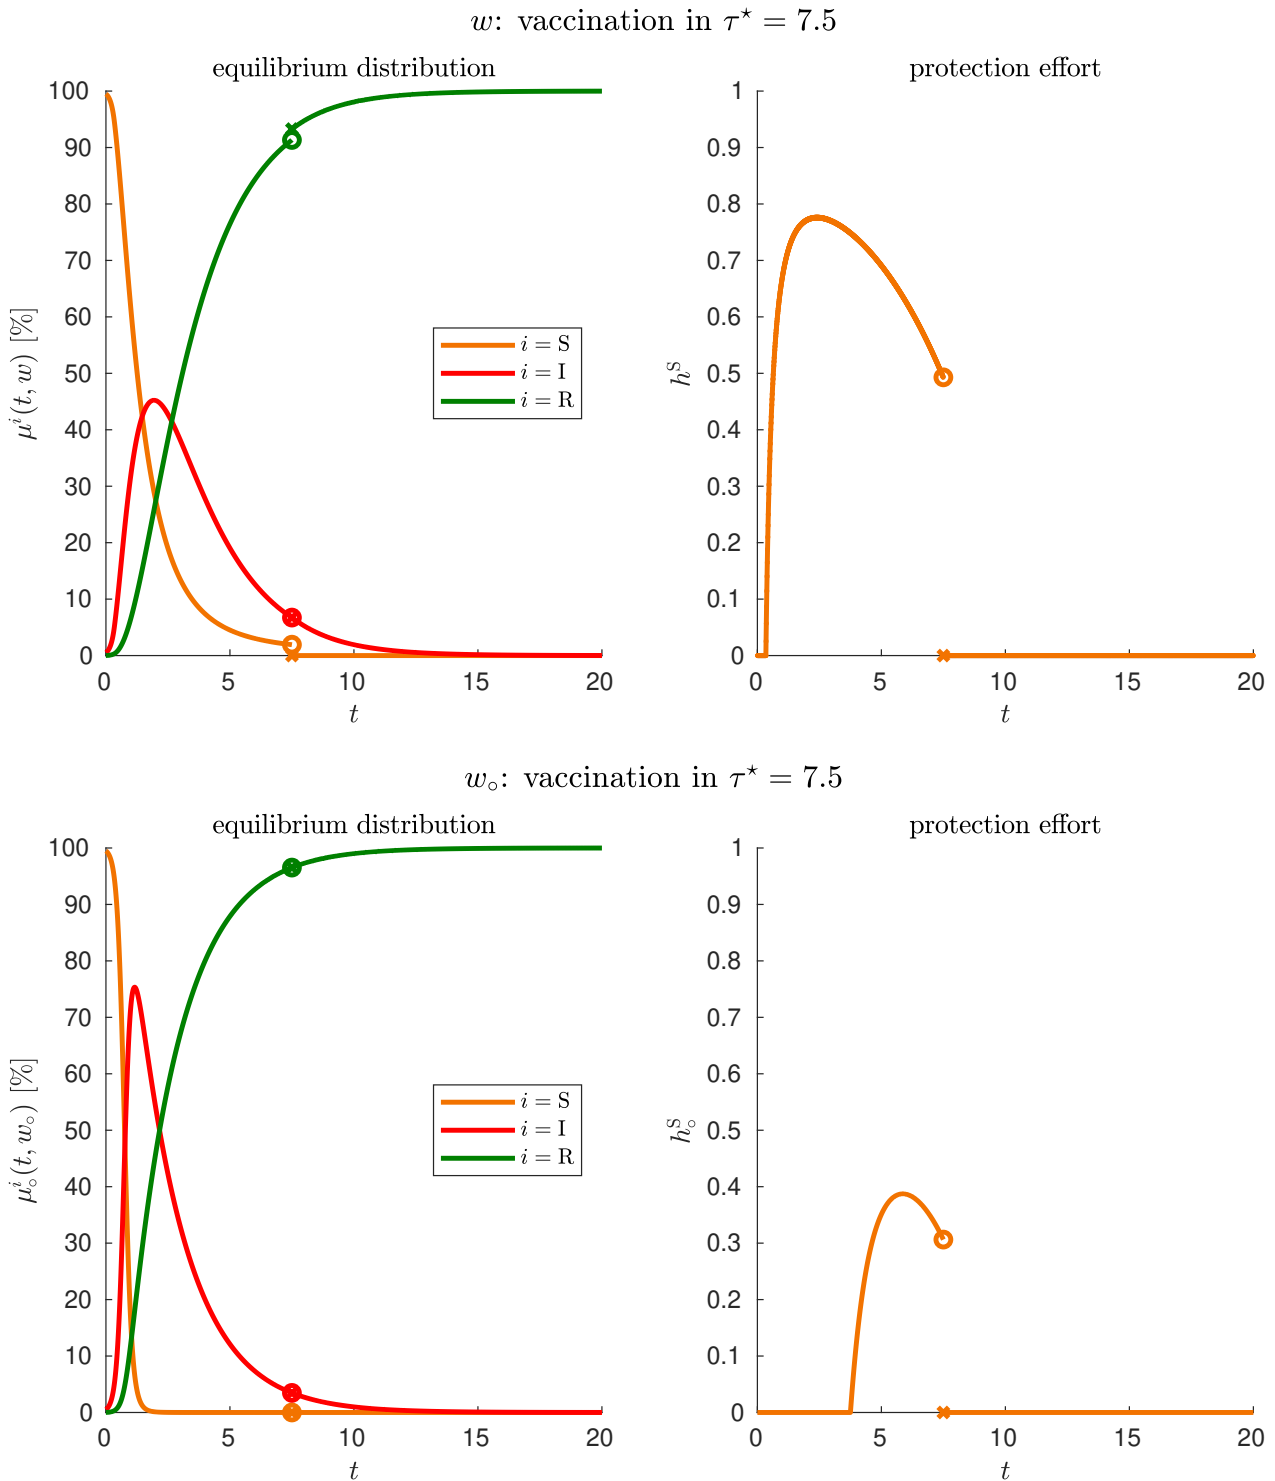


Figure 2.14: Equilibrium distribution and protection effort for $\tau^* = 7.5$: Mean field game with common noise (top) and corresponding perfect-foresight equilibrium (bottom).

2.4.3 Evacuation of a Room with Randomly Opening Doors

As our final application example, we consider a population of agents located inside a room that has two doors. Formally, the agents' positions are represented by the state space

$$\mathbb{S} \triangleq \mathbb{S}_{\square} \cup \{x_{\uparrow}^{\text{out}}, x_{\downarrow}^{\text{out}}\}$$

where $\mathbb{S}_{\square} \triangleq [1:r]^2$ is the set of positions inside the room and $x_{\uparrow}^{\text{out}}, x_{\downarrow}^{\text{out}} \notin \mathbb{S}_{\square}$ are outside positions;⁸³ two distinguished inside positions $x_{\uparrow}^{\text{in}}, x_{\downarrow}^{\text{in}} \in \mathbb{S}_{\square}$ represent the locations of the doors; see Figure 2.15 for an illustration of the setup.⁸⁴




Figure 2.15: State space $\mathbb{S} = \mathbb{S}_{\square} \cup \{x_{\uparrow}^{\text{out}}, x_{\downarrow}^{\text{out}}\}$ where $\mathbb{S}_{\square} = [1:5]^2$ with $x_{\uparrow}^{\text{in}} = (3, 5)$ and $x_{\downarrow}^{\text{in}} = (3, 1)$, $x_{\uparrow}^{\text{out}} = (3, 6)$ and $x_{\downarrow}^{\text{out}} = (3, 0)$.

The two doors of the room are assumed to be stochastically modulated by a common noise. More precisely, whether a door is open or not depends on the outcome of a coin toss: Initially, on $[0, T_1)$ both doors are closed; later on, at each time T_k , a coin is tossed to decide which one of the two is open on $[T_k, T_{k+1})$, $k \in [1:n]$. The set of possible common noise factors is thus set to $\mathbb{W} \triangleq \{\uparrow, \downarrow\}$, such that on the event $\{W_k = \uparrow\}$ the upper door x_{\uparrow}^{in} is open during $[T_k, T_{k+1})$, while on $\{W_k = \downarrow\}$ the lower door $x_{\downarrow}^{\text{in}}$ is open on $[T_k, T_{k+1})$. The coin tosses are assumed to be fair and independent, i.e. i.i.d. Bernoulli(0.5)-distributed, and hence we set

$$\kappa_k(w_k | w_1, \dots, w_{k-1}, m) \triangleq 0.5, \quad \text{for } w_1, \dots, w_k \in \mathbb{W}, m \in \mathbb{M}, k \in [1:n].$$

For two inside positions $x, y \in \mathbb{S}_{\square}$ we write $x \rightsquigarrow y$ if y can be reached from x with one (vertical, horizontal or diagonal) step, see Figure 2.16. For the two outside positions $x_{\uparrow}^{\text{out}}, x_{\downarrow}^{\text{out}}$ we write $x \rightsquigarrow x_{\uparrow}^{\text{out}}$ if and only if $x = x_{\uparrow}^{\text{in}}$ and $x \rightsquigarrow x_{\downarrow}^{\text{out}}$ if and only if $x = x_{\downarrow}^{\text{in}}$.

We let $\mathbb{U} \triangleq [0, 1]^8$ represent the set of feasible actions, where u_{ℓ} represents the effort to move in

⁸³Note that the specification of \mathbb{S} involves a slight abuse of notation and corresponds to a total of $d \triangleq r^2 + 2$ states.

⁸⁴Similar setups have already been considered in the literature on *continuous-state* mean field games for modeling crowd dynamics; see, e.g., [LW11].

direction $\ell \in [1:8]$. The transition intensities are given by

$$Q^{xy}(t, w, m, u) \triangleq \lambda_{\max} \cdot u_\ell \cdot (1 - (m^y)^\alpha \cdot \mathbf{1}_{y \in \mathbb{S}_\square}) \quad \text{if } x \rightsquigarrow y \text{ by action } \ell \in [1:8]$$

and $Q^{xy}(t, w, m, u) \triangleq 0$ otherwise, where $y \in \mathbb{S}$, $y \neq x$, $(x, t, w, m, u) \in \mathbb{S} \times [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$; with $\alpha \in (0, 1]$ the factor $(1 - (m^y)^\alpha \cdot \mathbf{1}_{y \in \mathbb{S}_\square})$ captures a pushback effect for moving to crowded inside positions, and $\lambda_{\max} > 0$ denotes the maximum intensity for a single state transition.

Moreover, at common noise times there are no jumps of the state process, i.e. it holds $J^x(t, w, m) \triangleq x$ for $(x, t, w, m) \in \mathbb{S} \times [0, T] \times \mathbb{W}^n \times \mathbb{M}$.




Figure 2.16: Possible directions of movement.

The agents' preferences are such that they prefer being outside the room, and such that large efforts to change their positions are penalized. Correspondingly, with $C_1^\psi, C_2^\psi, C^\Psi > 0$ we set

$$\psi^x(t, w, m, u) \triangleq -\frac{C_1^\psi}{2} \cdot \sum_{\ell=1}^8 u_\ell^2 + C_2^\psi \cdot \mathbf{1}_{x \notin \mathbb{S}_\square} \quad \text{and} \quad \Psi^x(w, m) \triangleq C^\Psi \cdot \mathbf{1}_{x \notin \mathbb{S}_\square}$$

for $(x, t, w, m, u) \in \mathbb{S} \times [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{U}$. As a consequence, for all $(x, t, w, m, v) \in \mathbb{S} \times [0, T] \times \mathbb{W} \times \mathbb{M} \times \mathbb{R}^d$ the unique maximizer in Assumption 2.2.1 is given by

$$h^x(t, w, m, v)_\ell = \left[\frac{\lambda_{\max}}{C_1^\psi} \cdot (1 - (m^y)^\alpha \cdot \mathbf{1}_{y \in \mathbb{S}_\square}) \cdot (v^y - v^x)^+ \right] \wedge 1 \quad \text{if } x \rightsquigarrow y \text{ by action } \ell$$

and $h^x(t, w, m, v)_\ell \triangleq 0$ otherwise, for $\ell \in [1:8]$. ◇

The relevant model parameters are given in Table 2.3. Note that there are $n = 4$ common noise events revealed at times $T_k \triangleq k/5$, $k = 1, \dots, 4$.

We initialize the fixed-point iteration for (E1)-(E6) by $\mu_0(t, w) \equiv m_0$ for all $(t, w) \in [0, T] \times \mathbb{W}^n$, and we set the step size of the Euler scheme to $5 \cdot 10^{-3}$.

Our results for the mean field equilibrium distributions of agents are displayed in Figures 2.17 through 2.19 and 2.21 through 2.23 for various common noise configurations $w \in \mathbb{W}^n$ and initial distributions $m_0 \in \mathbb{M}$. Exemplarily, for each of the resulting distributions in Figures 2.19 and 2.23 we also provide an illustration of the equilibrium distribution in the corresponding game without incorporation of the pushback effect mentioned above; see Figures 2.20 and 2.24, respectively.⁸⁵

⁸⁵We wish to point out that without the pushback effect there is *no* mean field interaction in the game anymore.

Parameter	T	n	T_k	r	mid	x_{\uparrow}^{in}	$x_{\downarrow}^{\text{in}}$	λ_{max}	α	$C^{r\psi}$	C^{Ψ}
Value	1	4	$k/5$	21	$(r+1)/2$	(mid, r)	$(\text{mid}, 1)$	5	0.3	(2.5, 10)	100

Table 2.3: Model parameters *Evacuation*.

Note that, throughout the entire game, agents have the desire to leave the room since only staying outside is continuously and terminally rewarded. At the same time, due to common noise, they are confronted with doors opening at random, i.e. they cannot foresee through which door they can leave the room after the next coin toss. Therefore and for reasons of symmetry, the initial field of agents either splits, see Figures 2.17 through 2.19 and 2.21, or merges, see Figures 2.22 and 2.23, into two groups each of which aims to leave the room through one specific door.

Furthermore, observe that the specification of the transition intensities includes a pushback effect that renders efforts to move to crowded inside positions less effective than those aimed at less occupied target positions. As a consequence, compared to the corresponding game⁸⁵ without that effect, the majority of agents in each of the two groups moves at a slower pace and the percentage of agents that ultimately reach the outside positions is considerably lower; compare Figures 2.19 and 2.20, or 2.23 and 2.24, respectively.

equilibrium distribution: $w = (\downarrow, \downarrow, \uparrow, \uparrow)$

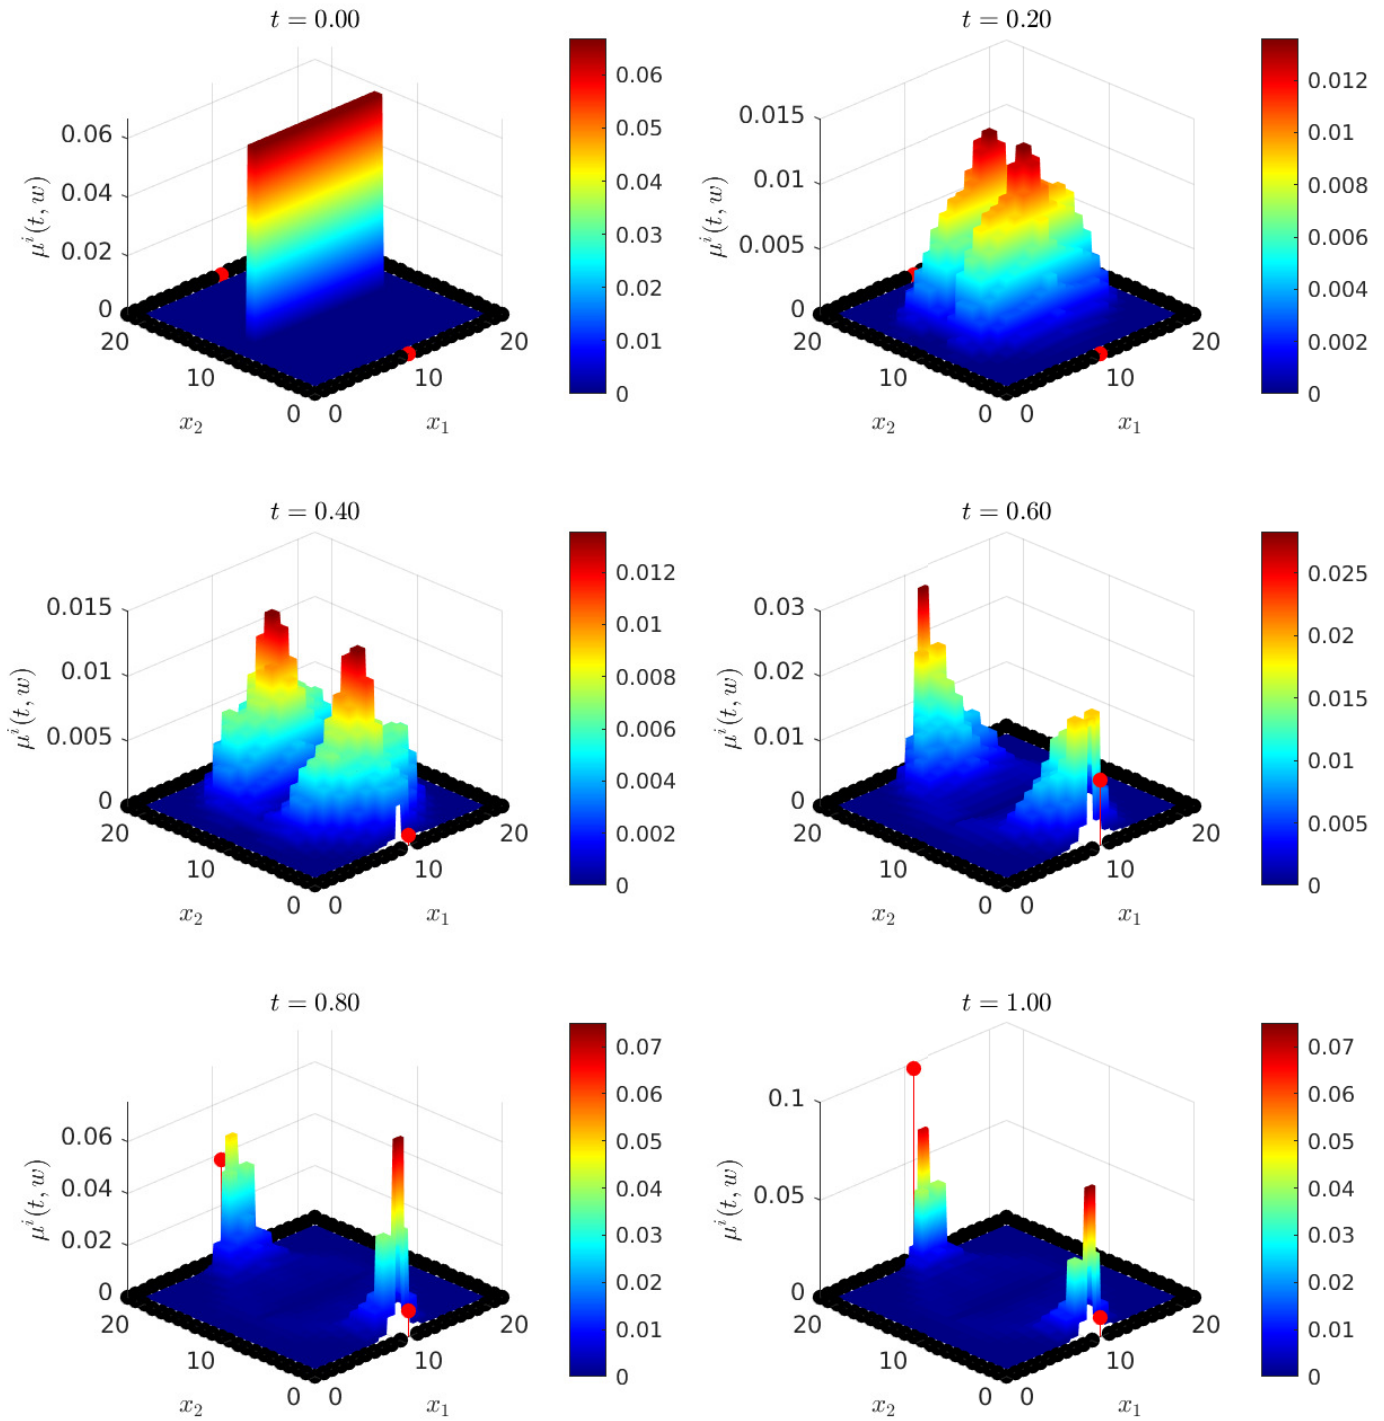


Figure 2.17: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\downarrow, \downarrow, \uparrow, \uparrow)$; m_0 is the (conditional) uniform distribution on $[4:r-3] \times \{\text{mid}\}$.

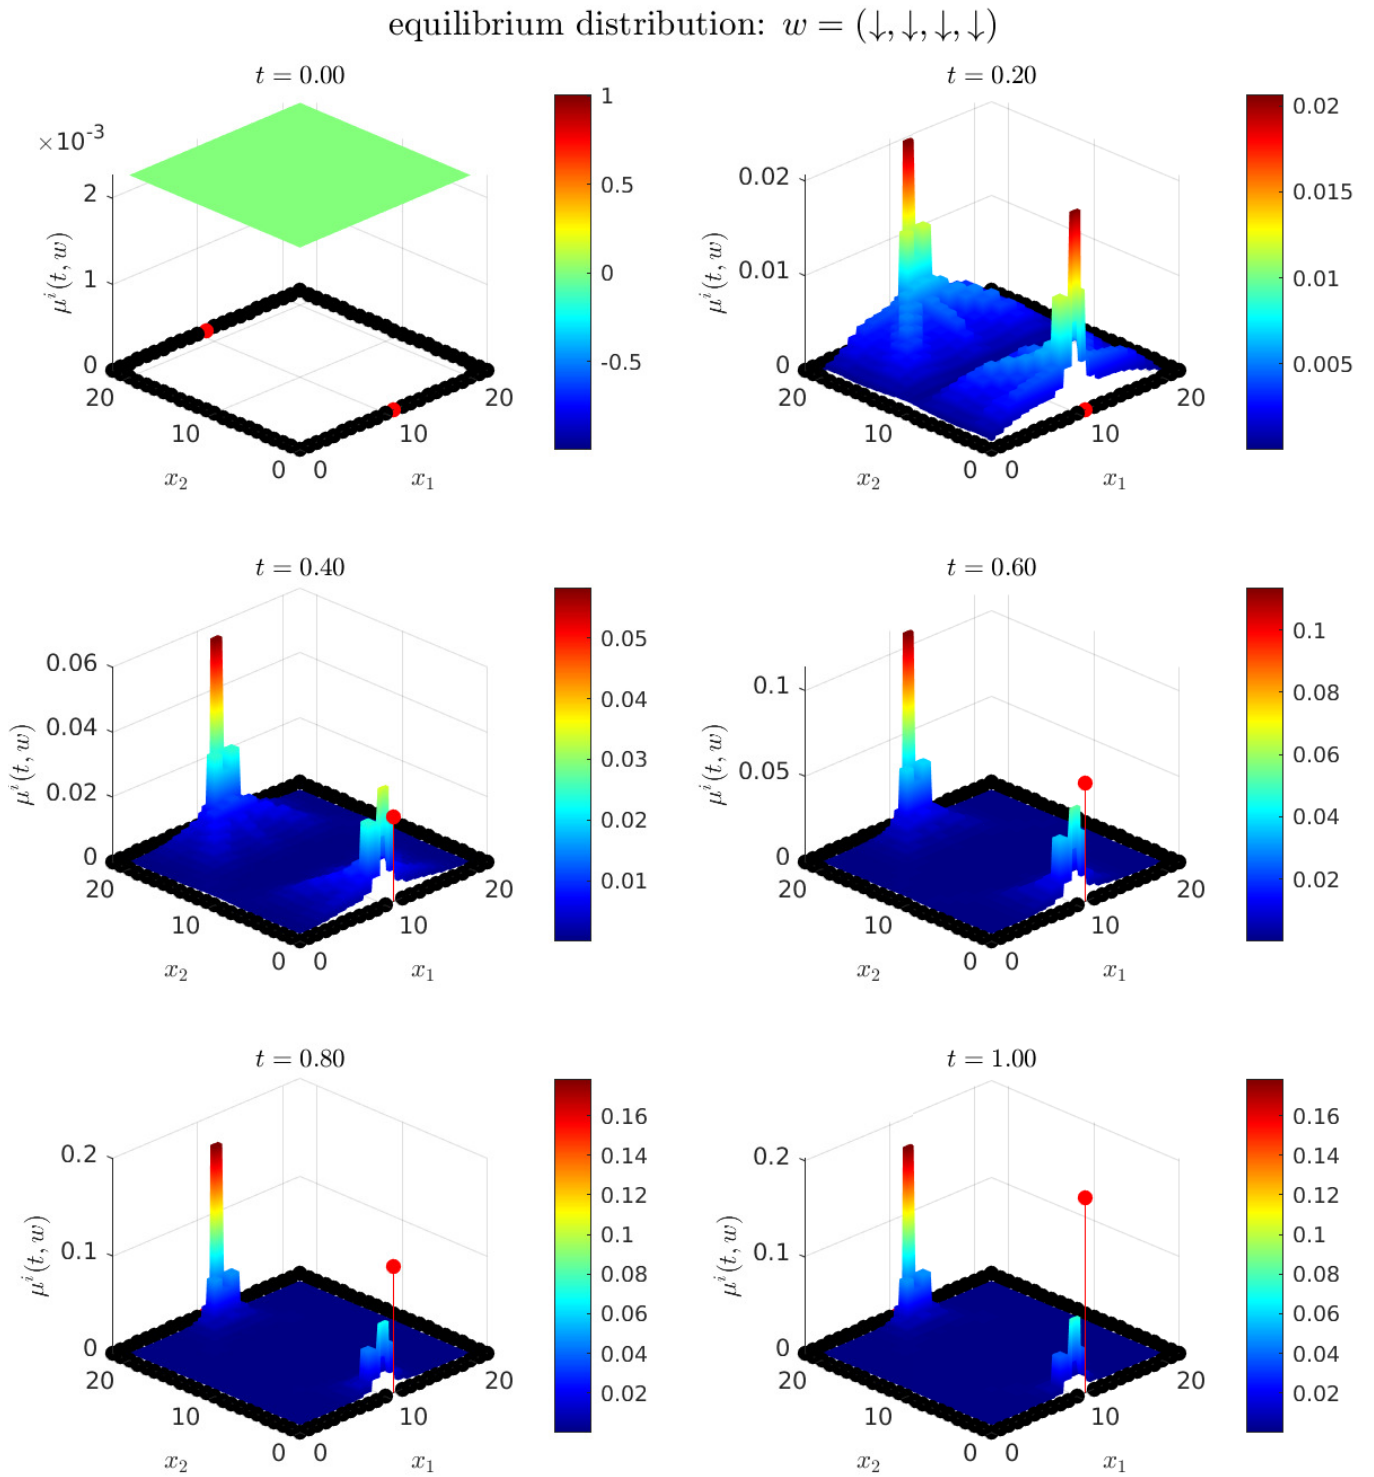


Figure 2.18: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\downarrow, \downarrow, \downarrow, \downarrow)$ and $m_0 = \mathcal{U}(\mathbb{S}_\square)$.

equilibrium distribution: $w = (\uparrow, \downarrow, \uparrow, \downarrow)$




Figure 2.19: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\uparrow, \downarrow, \uparrow, \downarrow)$ and $m_0 = \delta_{(\text{mid}, \text{mid})}$.

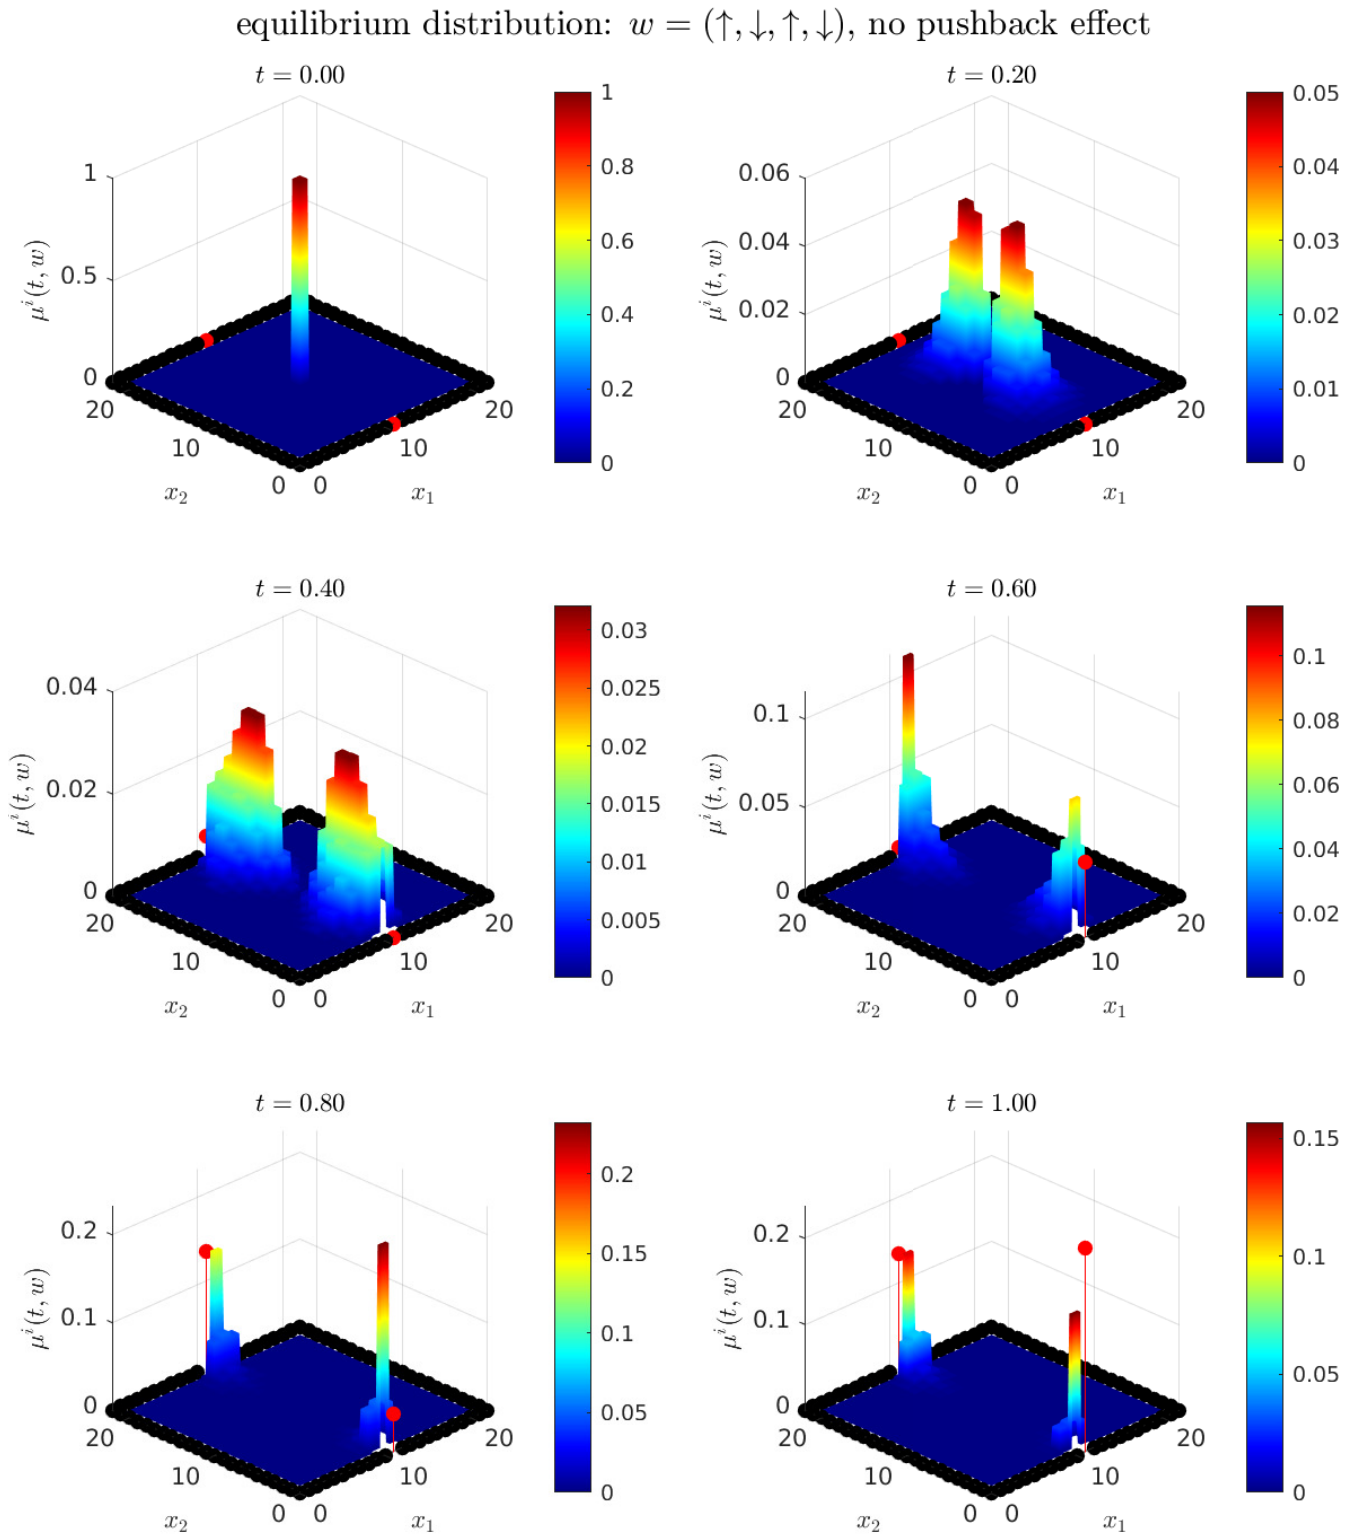


Figure 2.20: Equilibrium aggregate distribution of agents in the (non-mean field) game⁸⁵ without a pushback effect corresponding to Figure 2.19.

equilibrium distribution: $w = (\uparrow, \uparrow, \uparrow, \uparrow)$




Figure 2.21: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\uparrow, \uparrow, \uparrow, \uparrow)$ and $m_0 = \delta_{(\text{mid}, \text{mid})}$.




Figure 2.22: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\uparrow, \uparrow, \uparrow, \downarrow)$; m_0 puts 50% mass each on $(\text{mid} - 3, \text{mid})$ and $(\text{mid} + 3, \text{mid})$.




Figure 2.23: Equilibrium aggregate distribution of agents at $t \in \{T_k : k \in [0:5]\}$ for $w = (\downarrow, \uparrow, \uparrow, \downarrow)$; m_0 puts 25% mass on each of the four corners of the room.

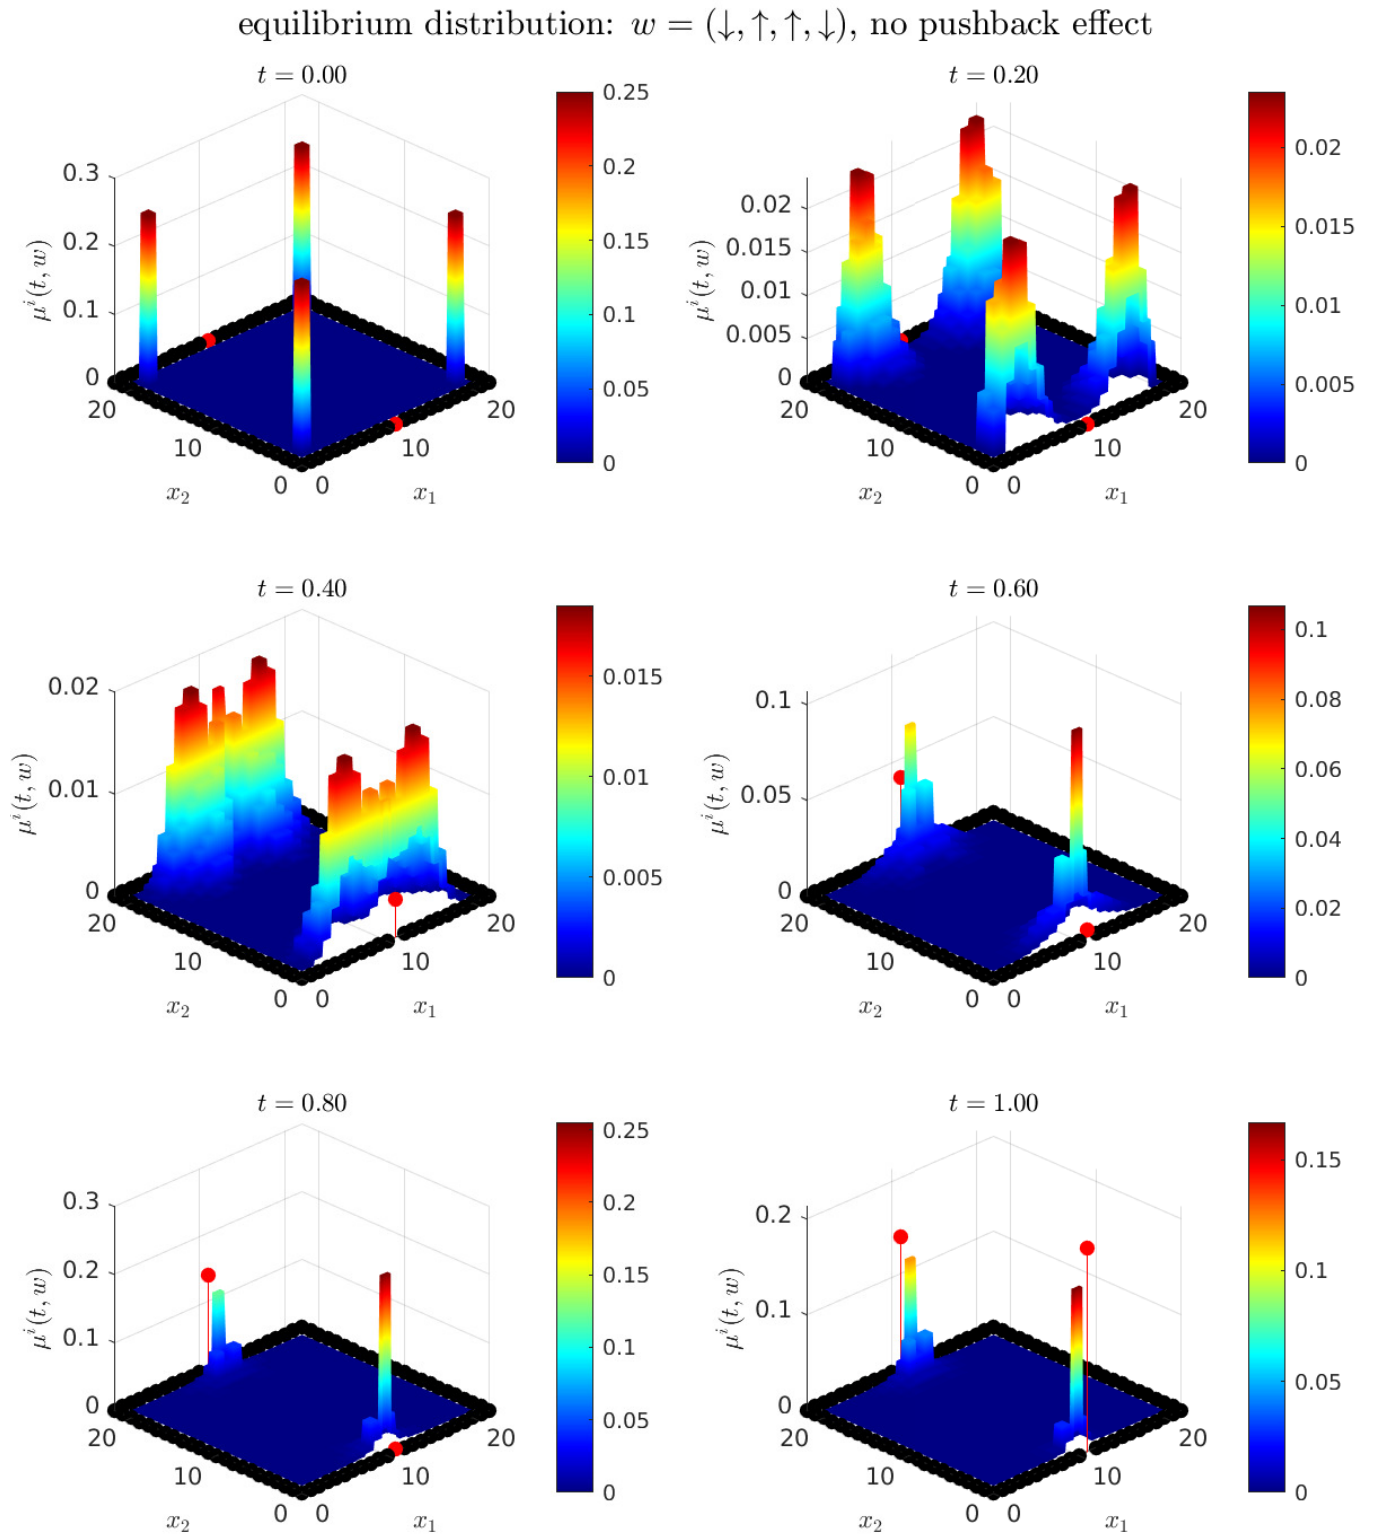


Figure 2.24: Equilibrium aggregate distribution of agents in the (non-mean field) game⁸⁵ without a pushback effect corresponding to Figure 2.23.

Conclusion and Outlook

In summary, this chapter has developed a rigorous framework for continuous-time mean field games with finitely many states and common noise. We have provided a rigorous probabilistic construction of the agents' state dynamics, and we have proved a verification theorem for the representative agents' optimization problems given their *ex ante* beliefs, an aggregation theorem for the resulting dynamics of the *ex post* aggregate distribution of agents and an existence result for the resulting mean field equilibrium system. Moreover, we have highlighted the impact of common noise effects in three benchmark application examples.

For future research we think that combining our model with that of [CW18b] is a natural next step leading to a unified probabilistic approach to *extended finite-state mean field games with common noise*; this would allow agents to interact not only via the aggregate distribution of their states, but rather also via the aggregate distribution of their controls, and conceivably, via a source of common noise, pave the way for finite-state price impact models with mean field interactions.⁸⁶

Moreover, it would be of interest to carry out a probabilistic analysis of finite-state mean field games that include a source of common noise described by an independent continuous-time Markov chain. For related recent developments using the master equation we refer, e.g., to [BLL19].

⁸⁶For further details on *price impact models* we refer, e.g., to §1.3.2 in [CD18a] or [CJ19b] (see also [CJ19a]) and the references therein.

Chapter 3

A Heuristic Approach to Discrete-Time Impulse Control

Impulse control problems constitute a major branch of continuous-time stochastic optimal control theory. As opposed to continuous or singular control, they provide a means to model discrete interventions in a continuous-time stochastic process and naturally include the possibility of charging fixed costs for such *impulses*. Unsurprisingly, this class of stochastic control problems provides *the* appropriate mathematical concept of formalization in many application contexts; examples include portfolio optimization with transaction costs ([EH88], [Kor98], [BP00], [ØS02], [IS06], [PS07], [BC19]), optimal consumption and investment problems ([Liu04], [ARS17]), inventory control and cash management ([CR78], [HST83], [HSZ17]), optimal control of exchange rates ([Kor97], [MØ98], [CZ00]), index tracking ([BK98]), stochastic forest growth and timber harvesting ([Wil98], [Alv04]), optimal control of dividends ([AV06]), and other fields of research in applied probability theory, economics and finance; one may also consult [Kor99] or Part II in [Sto09] for an overview.

From a theoretical point of view, continuous-time impulse control problems are tackled in different ways: As their corresponding dynamic programming equation is given by a quasi-variational inequality (QVI), a nonlocal and highly nonlinear partial differential equation (see, e.g., [BL84], [Dav93, §54] or [ØS05, Chapter 6]), a vast body of literature studies either classical or viscosity solutions thereof; see, e.g., [TY93], [GW09], [Sey09], [DGW10], [BEM13], [CG13]. For numerical methods for related Hamilton-Jacobi-Bellman quasi-variational inequalities (HJBQVIs) we refer to the recent papers [AF16] and [ABL18] as well as the references therein. On the other hand, impulse control is also related to iterated optimal stopping problems (see, e.g., [ØS05, Chapter 7], [Men80], [Men87], [BCS17]), and there is a characterization of the corresponding value function in terms of superharmonic functions available in the literature (see, e.g., [Chr14], [BCS17], [BC19]).

In this chapter, however, we consider an impulse control problem in a *discrete-time* setup with finite time horizon. Such problems have already been addressed in the literature in case of an infinite time horizon; see, e.g., [Ste86] and [Ben08]. Based on the work of Bensoussan [Ben08] who has introduced the concept of a quasi-variational inequality in discrete time, we first reformulate the classical Bellman equation of stochastic optimal control for the considered problem in terms of such a discrete-time QVI. The derivation of this alternative formulation of the dynamic

programming principle strongly relies on the impulsive nature of interventions; more precisely, every non-trivial action the controller takes to intervene in the state dynamics has strictly positive fixed costs that are uniformly bounded away from zero across time, states and (non-trivial) actions. Hereafter, we prove a corresponding verification theorem, including an explicit specification of an optimal impulse control in terms of the standard concepts of a continuation and an intervention region. This strategy, in turn, serves as our starting point to heuristically devise a self-learning algorithm based on policy iteration that aims to approximate the value function of our problem. Regarding its implementation we are most prominently confronted with the following two issues: On the one hand, especially in case of an infinite state space, a both suitable and computationally tractable spatial discretization has to be found; on the other hand, emerging (conditional) expectations have to be estimated.

Both of these problems have already been identified and studied in the literature, but are still the subject of ongoing research: The first one has already been addressed for optimal stopping problems, usually in the context of pricing American or Bermudan option, leading to concepts such as weighted stochastic meshes; see, e.g., [BGH00], [Gla03, §8.5], [BG04a], [LH09], [AJ13] and [BKS19]; moreover, quantization-based techniques have been developed for optimal stopping and control problems; see, e.g., [BP03], [PPP04a] and [PPP04b] as well as the references therein.

The second issue is often approached by regression-based Monte Carlo methods; see, e.g., [LS01], [GY04], [BP17a], [BP17b] and [BLMP19].

Furthermore, note that in recent years stochastic control problems have also attracted interest in the area of machine learning; cf., among others, [HE16] and [HPBL18, BHL19].

Partially inspired by those ideas, the implementation of our theoretical algorithm sets up a hierarchy of levels and combines policy iteration and non-nested Monte Carlo simulation with interpolation based on stochastic nodes adapted to the underlying problem; thus, by exploiting information acquired from previous levels, it explores the state space of the problem by itself and successively learns the behavior of the optimally controlled state process. The developed heuristics of combining policy iteration, non-nested Monte Carlo simulation and an exploration strategy that adapts to the structure of the underlying problem constitutes the core contribution of this chapter.

The remainder of this chapter is organized as follows: In Section 3.1 we introduce the theoretical setup and the considered impulse control problem. Section 3.2 establishes a reformulation of the classical Bellman equation in terms of a discrete-time QVI and provides a corresponding verification theorem. In Section 3.3 we state and discuss our algorithm and sketch its implementation. Finally, Section 3.4 considers the classical problem of *impulse control of Brownian motion* and provides some numerical examples that showcase the effectiveness of our algorithm; underlining the heuristic core contribution of this chapter, we conclude with an outlook on issues to be addressed in future research.

3.1 Theoretical Setup

To start with, we provide an *informal* description of the addressed stochastic control problem. We work on the discrete time interval $[0:T]$ and assume that the controller's state process $\{X_t^\alpha\}$ takes values in a measurable state space \mathbb{X} ; it is driven by a process $\{Z_t\}$ of exogenous stochastic shocks consisting of i.i.d. random variables; and his control $\{\alpha_t\}$ is interpreted to be of *impulsive nature* in that at each time instant $t \in [0:T)$ he can decide to intervene and in turn to change the location of his state process while incurring at least positive fixed *impulse costs*. More precisely, at time $t \in [0:T)$ he can decide to move his state from X_t^α to $\xi(X_t^\alpha, \alpha_t)$ at a cost of $c(t, X_t^\alpha, \alpha_t)$, before the next random shock Z_{t+1} propagates this post-jump state to $X_{t+1}^\alpha = \phi(t, \xi(X_t^\alpha, \alpha_t), Z_{t+1})$ at time $t + 1$. In this way, the controller aims to maximize his *performance functional*

$$\mathbb{E} \left[\sum_{t=0}^{T-1} (\psi(t, \xi(X_t^\alpha, \alpha_t)) - c(t, X_t^\alpha, \alpha_t)) + \Psi(X_T^\alpha) \right]$$

over feasible choices of strategies $\{\alpha_t\}$ where ψ and Ψ are suitable *running* and *terminal reward* functions, respectively.

In the remainder of this section, we provide a rigorous mathematical formulation of this problem.⁸⁷

3.1.1 State Space and Dynamics

Throughout, we fix some *time horizon* $T \in \mathbb{N}$ and consider the discrete time interval $[0:T]$.⁸⁸ We work on a probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ that carries a finite sequence $\{Z_t\}_{t \in [1:T]}$ of i.i.d. random variables that attain their values in a measurable *noise space* $(\mathbb{Z}, \mathfrak{Z})$; each of them is regarded as an *exogenous source of noise* or a *stochastic shock* and their distribution is denoted by $\mu \triangleq \mathbb{P}(Z \in \cdot)$ where we set $Z \triangleq Z_1$ for later reference. The corresponding *flow of information* is thus modeled by the filtration $\{\mathfrak{F}_t\}_{t \in [0:T]}$ defined by

$$\mathfrak{F}_t \triangleq \sigma(Z_s : s \leq t) \quad \text{for } t \in [0:T].$$

The *state space* is given by some non-empty measurable space $(\mathbb{X}, \mathfrak{X})$; the controller's *feasible actions* are chosen to be a subset of a measurable *action space* $(\mathbb{A}, \mathfrak{A})$; see below.

Given some initial configuration $(t, x) \in [0:T] \times \mathbb{X}$, the *state process* $X^{t,x;\alpha} = \{X_s^{t,x;\alpha}\}_{s \in [t:T]}$ starting in state $x \in \mathbb{X}$ at time $t \in [0:T]$ is defined to obey the dynamics⁸⁹

$$X_t^{t,x;\alpha} = x \quad \text{and} \quad X_{s+1}^{t,x;\alpha} = \phi(s, \xi(X_s^{t,x;\alpha}, \alpha_s), Z_{s+1}) \quad \text{for } s \in [t:T), \quad (3.1)$$

⁸⁷Large parts of the model setup to be described below are standard in the literature; see, e.g., Chapter 2 in [Ber76], Chapter 8 in [BS78], Chapter 2 in [BR11] or §1.2 in [Sei11]. With regard to notation we mainly rely on the latter reference.

⁸⁸Of course, upon simple reformulation of the stochastic control problem specified below, this includes the case of a time grid $\Pi = [0 = t_0, t_1, \dots, t_N = T] \subset [0, T]$ for some real-valued time horizon $T > 0$; see also Section 3.4 below.

⁸⁹Whenever it is clear from the context, we drop the part of the upper index concerning the initial configuration and simply write X^α instead of $X^{t,x;\alpha}$.

which is subject to a control process $\alpha \in \mathcal{A}'(t, x)$ that is chosen by the controller. Here,

$$\mathcal{A}'(t, x) \triangleq \left\{ \{\alpha_s\}_{s \in [t:T]} : \alpha \text{ is } \{\mathfrak{F}_s\}_{s \in [t:T]} \text{-adapted, } \alpha_s \in \Gamma(s, X_s^\alpha) \forall s \in [t:T] \right\}$$

denotes the set of *pre-admissible controls* or *strategies*;⁹⁰ the set-valued map

$$\Gamma : [0:T) \times \mathbb{X} \rightarrow 2^{\mathbb{A}}$$

therein specifies the set of *feasible actions* given a time-space configuration, i.e. at $(t, x) \in [0:T) \times \mathbb{X}$ the actions within the set $\Gamma(t, x) \subseteq \mathbb{A}$ are considered feasible. The actions of the controller affect the state process at (t, x) via the measurable jump-to-state map

$$\xi : \mathbb{X} \times \mathbb{A} \rightarrow \mathbb{X},$$

in that, given a feasible action $a \in \Gamma(t, x)$, the state $x \in \mathbb{X}$ is *immediately*, i.e. at time t , transformed into $y = \xi(x, a) \in \mathbb{X}$; correspondingly, an action a is also referred to as an *impulse* to the system. We assume that there is some action $\boxtimes \in \mathbb{A}$ that is always feasible, i.e. $\boxtimes \in \Gamma(t, x)$ for all $(t, x) \in [0:T) \times \mathbb{X}$, and models the controller's *inaction*, i.e.

$$\xi(x, \boxtimes) \equiv x \text{ for all } x \in \mathbb{X}.$$

For later reference we set $\Gamma^\circ(t, x) \triangleq \Gamma(t, x) \setminus \{\boxtimes\}$ and note that $\mathcal{A}'(t, x) \neq \emptyset$ for all $(t, x) \in [0:T) \times \mathbb{X}$. Finally, the function

$$\phi : [0:T) \times \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{X}$$

is assumed to be measurable and describes the *intertemporal transition of states*, i.e. given some instant of time $t \in [0:T)$, some state $y \in \mathbb{X}$ and a random shock $z \in \mathbb{Z}$, the subsequent state at time $(t + 1)$ is computed to be $\phi(t, y, z) \in \mathbb{X}$. Note that the map ϕ is beyond the control of the controller.

The dynamics of the state process is schematically illustrated in Figure 3.1.

3.1.2 Stochastic Control Problem and Value Function

Given an initial configuration $(t, x) \in [0:T) \times \mathbb{X}$, the controller chooses an admissible strategy within the set $\mathcal{A}(t, x)$ to be described below; thereby, he specifies an impulse to intervene in the state process via the map ξ . More precisely, given a feasible action $a \in \Gamma(t, x)$ he moves the state process from x to $\xi(x, a)$ at time t ; see above. In doing so, he incurs *intervention costs* of size $c(t, x, a)$ specified the map

$$c : [0:T) \times \mathbb{X} \times \mathbb{A} \rightarrow [0, \infty),$$

which satisfies

$$c(t, x, \boxtimes) \equiv 0 \quad \text{and} \quad c(t, x, a) > 0 \quad \text{whenever } a \in \Gamma^\circ(t, x)$$

⁹⁰The terms *control* and *strategy* are used synonymously in the following.




Figure 3.1: Dynamics of the state process.

for $(t, x) \in [0: T] \times \mathbb{X}$.⁹¹ As the state process is driven by exogenous random shocks described by the process $\{Z_s\}_{s \in [0: T]}$ via the state transition map ϕ and since both are beyond the control of the controller, his actions thus only affect the current spatial position of the state process at time $t \in [0: T]$ with strictly positive costs in case of intervention. These are precisely the features that render our setup an analogue of continuous-time *impulse control* problems in discrete time.

Given $(t, x) \in [0: T] \times \mathbb{X}$, his goal is to *maximize* the following *performance functional* over the set of admissible controls:

$$\mathcal{J}(t, x; \alpha) \triangleq \mathbb{E} \left[\sum_{s=t}^{T-1} (\psi(s, \xi(X_s^\alpha, \alpha_s)) - c(s, X_s^\alpha, \alpha_s)) + \Psi(X_T^\alpha) \right] \xrightarrow{\alpha \in \mathcal{A}(t, x)} \max! \quad (\text{C})$$

The functions

$$\psi : [0: T] \times \mathbb{X} \rightarrow \mathbb{R} \quad \text{and} \quad \Psi : \mathbb{X} \rightarrow \mathbb{R}$$

model the controller's *running* and *terminal reward*, respectively; both are assumed to be measurable; and the set of *admissible controls* or *strategies*⁹⁰ for a given initial configuration $(t, x) \in [0: T] \times \mathbb{X}$ is defined by

$$\mathcal{A}(t, x) \triangleq \left\{ \alpha \in \mathcal{A}'(t, x) : \mathbb{E} \left[\sum_{s=t}^{T-1} (\psi(s, \xi(X_s^\alpha, \alpha_s)) - c(s, X_s^\alpha, \alpha_s)) + \Psi(X_T^\alpha) \right] < \infty \right\}.$$

Note that the latter refinement of $\mathcal{A}'(t, x)$ guarantees that the emerging expectations in (C) are well-defined and take values in $(-\infty, \infty]$ – provided an admissible strategy α exists at all.

The controller's *value function* is thus given by⁹²

$$v^* : [0: T] \times \mathbb{X} \rightarrow [-\infty, \infty], \quad v^*(t, x) \triangleq \sup_{\alpha \in \mathcal{A}(t, x)} \mathcal{J}(t, x; \alpha). \quad (3.2)$$

⁹¹For a more concrete set of requirements see Assumption 3.2.3 below.

⁹²We stick to the common convention $\sup \emptyset \triangleq -\infty$.

Observe from (3.2) that, due to the sup over a potentially uncountably infinite set, the value function v^* may fail to be measurable. \diamond

As regards the stated integrability conditions in the definition of $\mathcal{A}(t, x)$, we simplify our considerations and require the following assumption to hold throughout this chapter:

Assumption 3.1.1. For all $(t, x) \in [0:T] \times \mathbb{X}$ it holds $\mathcal{A}(t, x) = \mathcal{A}'(t, x)$. \diamond

Note that the latter assumption in particular implies $\mathcal{A}(t, x) \neq \emptyset$ for all $(t, x) \in [0:T] \times \mathbb{X}$ since the trivial strategy $\alpha = \{\boxtimes\}_{s \in [t:T]}$ is clearly admissible, i.e. it is always admissible to let the state dynamics run uncontrolled.⁹³

Sufficient conditions for Assumption 3.1.1 to hold include suitable (growth) conditions on the model coefficients; more precisely, existence of an *upper bounding function* is sufficient; see Definition 2.4.1 in [BR11]; the relevant conditions read as follows in our setup:

Remark 3.1.2 (Upper Bounding Function). Suppose that there is a measurable map $\varpi : \mathbb{X} \rightarrow [0, \infty)$ such that for all $(t, x) \in [0:T] \times \mathbb{X}$ and $a \in \Gamma(t, x)$ it holds

$$\Psi(x)^- \vee \psi(t, \xi(x, a))^- \vee c(t, x, a) \leq \kappa_1 \cdot \varpi(x) \quad \text{and} \quad (3.3)$$

$$\mathbb{E}[\varpi(\phi(t, \xi(x, a), Z))] \leq \kappa_2 \cdot \varpi(x), \quad (3.4)$$

where $\kappa_1, \kappa_2 \in [0, \infty)$ are nonnegative constants. Then one readily shows along the lines of the proof of Proposition 2.4.2 in [BR11] that Assumption 3.1.1 is satisfied. Note in particular that (3.3) and (3.4) are trivially fulfilled if the cost function c is bounded and the reward functions ψ and Ψ are lower bounded. \diamond

As *feedback strategies* will be of frequent use below, we further note in this regard:

Remark 3.1.3 (Feedback Strategies). Let that Assumption 3.1.1 be satisfied and $(t, x) \in [0:T] \times \mathbb{X}$ be an initial configuration. Suppose that $a : [t:T] \times \mathbb{X} \rightarrow \mathbb{A}$ is a measurable map such that $a(s, y) \in \Gamma(s, y)$ for all $(s, y) \in [t:T] \times \mathbb{X}$. We recursively define

$$\begin{aligned} X_t &\triangleq x \quad \text{and} \quad \alpha_t \triangleq a(t, X_t), \quad \text{as well as} \quad \alpha_s \triangleq a(s, X_s) \quad \text{for } s \in (t:T), \\ \text{where } X_s &\triangleq \phi(s-1, \xi(X_{s-1}, \alpha_{s-1}), Z_s) \quad \text{for } s \in (t:T]. \end{aligned} \quad (3.5)$$

Then α is an admissible strategy, i.e. $\alpha = \{\alpha_s\}_{s \in [t:T]} \in \mathcal{A}(t, x)$, and it holds $X = X^{t,x;\alpha}$; see (3.1). We refer to α as the *feedback strategy corresponding to a* .

Whenever we are given a measurable function a as above we just write $\alpha_s = a(s, X_s^\alpha)$ for $s \in [t:T]$ and implicitly assume that the strategy $\alpha \in \mathcal{A}(t, x)$ is constructed recursively as outlined in (3.5). \diamond

To conclude this section, observe that every feedback strategy renders the state process Markov: For $(t, x) \in [0:T] \times \mathbb{X}$ and a feedback strategy given by $\alpha_s = a(s, X_s^\alpha)$, $s \in [t:T]$, for a measurable map a as in Remark 3.1.3, by (3.1) and due to independence of Z_s and $\mathfrak{F}_{s-1} \supseteq \mathfrak{F}_{s-1}^{X_{s-1}^\alpha}$ for $s \in (t:T]$ we obtain for any bounded measurable function $f : \mathbb{X} \rightarrow \mathbb{R}$ the identity

$$\mathbb{E}[f(X_s^\alpha) | X_t^\alpha, \dots, X_{s-1}^\alpha] = \mathbb{E}[f(\phi(s-1, \xi(X_{s-1}^\alpha, a(s-1, X_{s-1}^\alpha)), Z_s)) | X_t^\alpha, \dots, X_{s-1}^\alpha]$$

⁹³One may also consult Proposition 1.9 in [Sei11] for a similar consideration.

$$\begin{aligned}
&= \mathbb{E}[f(\phi(s-1, \xi(X_{s-1}^\alpha, a(s-1, X_{s-1}^\alpha)), Z_s)) | X_{s-1}^\alpha] \\
&= \mathbb{E}[f(X_s^\alpha) | X_{s-1}^\alpha].
\end{aligned}$$

Hence, the state process X^α is a Markov process; correspondingly, in view of problem (C), it is also referred to as a *Markov decision process*.⁹⁴

3.2 Reformulation of the Dynamic Programming Equation

In this section we state the classical Bellman equation and reformulate it to obtain a discrete-time quasi-variational inequality (QVI) in the spirit of Bensoussan [Ben08]; within a suitable space of candidate functions and under standard assumptions, each of them has at most one solution that, in case of existence, coincides with the value function of problem (C), see (3.2).

3.2.1 The Bellman Principle and a Classical Verification Theorem

The classical *Bellman* or *dynamic programming equation* for the given stochastic optimal control problem (C) reads as follows:

$$\begin{aligned}
v(t, x) &= \sup_{a \in \Gamma(t, x)} \left\{ \psi(t, \xi(x, a)) - c(t, x, a) + \int_{\mathbb{Z}} v(t+1, \phi(t, \xi(x, a), z)) \mu(dz) \right\} \\
&\qquad\qquad\qquad \text{for } (t, x) \in [0:T] \times \mathbb{X}, \tag{DP} \\
v(T, x) &= \Psi(x) \quad \text{for } x \in \mathbb{X}.
\end{aligned}$$

For integrals occurring in (DP) to be well-defined, suitable measurability and integrability conditions must be imposed; hence, in the remainder of this chapter, we consider (DP) for functions within the *candidate space*

$$\begin{aligned}
\mathbb{V} \triangleq \left\{ v : [0:T] \times \mathbb{X} \rightarrow \mathbb{R} \mid v \text{ is measurable and} \right. \\
\left. \int_{\mathbb{Z}} v(t+1, \phi(t, \xi(x, a), z))^- \mu(dz) < \infty \right. \\
\left. \text{for all } (t, x) \in [0:T] \times \mathbb{X} \text{ and } a \in \Gamma(t, x) \right\}.
\end{aligned}$$

Thus, the classical *Bellman principle* states that a solution of (DP) can be identified with the value function of (C), provided that the supremum in (DP) is attained; more precisely, the following standard verification theorem holds:

Theorem 3.2.1 (Verification via (DP)). Let Assumption 3.1.1 be satisfied and $v \in \mathbb{V}$ be a function that solves the Bellman equation (DP). Additionally, assume that there exists a measurable maximizer, i.e. a measurable function $a^* : [0:T] \times \mathbb{X} \rightarrow \mathbb{A}$ such that for all $(t, x) \in [0:T] \times \mathbb{X}$ it holds $a^*(t, x) \in \Gamma(t, x)$ and

$$v(t, x) = \psi(t, \xi(x, a^*(t, x))) - c(t, x, a^*(t, x)) + \int_{\mathbb{Z}} v(t+1, \phi(t, \xi(x, a^*(t, x)), z)) \mu(dz). \tag{3.6}$$

⁹⁴For a more detailed account on *Markov decision processes* we refer, e.g., to Chapter 4 in [Put94] or Chapter 2 in [BR11].

Then $v = v^*$ is the value function of problem (C) and a corresponding optimal strategy $\alpha^* \in \mathcal{A}(t, x)$ is given by

$$\alpha_s^* \triangleq a^*(s, X_s^{\alpha^*}) = \begin{cases} a^*(t, x) & \text{for } s = t, \\ a^*(s, \phi(s-1, \xi(X_{s-1}^{\alpha^*}, \alpha_{s-1}^*), Z_s)) & \text{for } s \in (t:T). \end{cases} \quad (3.7)$$

Proof. The stated result is standard given the literature on discrete-time stochastic optimal control; it is readily shown by backward induction. We refer, e.g., to the discussion in §2.3, in particular to Theorem 2.3.7, in [BR11]. \square

Put differently, equation (DP) has at most one solution within \mathbf{V} that admits a measurable maximizer in the sense of (3.6); in case of existence, its unique solution is the (measurable) value function of problem (C). Sufficient conditions for the required measurable maximizer to exist can be obtained in similar fashion as the ones tailored to the second verification equation (QVI) below; see Section 3.2.3.

3.2.2 A Verification QVI in Discrete Time

A reformulation of (DP) for discrete-time impulse control problems with infinite time horizon in terms of a QVI is due to Bensoussan [Ben08]. In order to adapt this result to our setup and for later reference we define the following standard operators in analogy to §4.2 in [Ben08]:

Definition 3.2.2 (Maximum and Continuation Operator). Let $v : [0:T] \times \mathbb{X} \rightarrow [-\infty, \infty]$ be a function and $(t, x) \in [0:T] \times \mathbb{X}$. We define the *maximum operator* \mathcal{M} by⁹⁵

$$\mathcal{M}[v](t, x) \triangleq \begin{cases} \sup_{a \in \Gamma^\circ(t, x)} \{v(t, \xi(x, a)) - c(t, x, a)\} & \text{if } (t, x) \in [0:T] \times \mathbb{X}, \\ \Psi(x) & \text{else;} \end{cases}$$

moreover for $v \in \mathbf{V}$ we define by

$$\mathcal{C}[v](t, x) \triangleq \begin{cases} \psi(t, x) + \int_{\mathbb{Z}} v(t+1, \phi(t, x, z)) \mu(dz) & \text{if } (t, x) \in [0:T] \times \mathbb{X}, \\ \Psi(x) & \text{else} \end{cases}$$

the *continuation operator* \mathcal{C} . \diamond

Observe that \mathcal{M} is a nonlocal operator and that $\mathcal{M}[v]$ may attain the values $\pm\infty$, even if v itself is \mathbb{R} -valued; furthermore, note that the condition $v \in \mathbf{V}$ guarantees that $\mathcal{C}[v] : [0:T] \times \mathbb{X} \rightarrow (-\infty, \infty]$ is well-defined. Clearly, the integral in the definition of \mathcal{C} can be interpreted as an expectation: Since $\mu = \mathbb{P}(Z \in \cdot)$ it holds $\mathcal{C}[v](t, x) = \psi(t, x) + \mathbb{E}[v(t+1, \phi(t, x, Z))]$ for $(t, x) \in [0:T] \times \mathbb{X}$ and $v \in \mathbf{V}$.

⁹⁵Note that the concept of a *maximum/intervention operator* is standard in the literature on impulse control; see also, e.g., (30) in [Ste86], Definition 6.1 in [ØS05] or (4) in [BCS17]. One should bear in mind that the maximum operator \mathcal{M} as defined here excludes \boxtimes and only considers “nontrivial” actions, i.e. those specified by the correspondence Γ° ; see also §4.2 in [Ben08].

Note that with the help of these operators, for a function $v \in \mathbb{V}$ the Bellman equation (DP) can be equivalently rewritten as

$$v(t, x) = \mathcal{M}[\mathcal{C}[v]](t, x) \vee \mathcal{C}[v](t, x) \quad \text{for all } (t, x) \in [0:T] \times \mathbb{X}. \quad (\text{DP})$$

The following set of assumptions, similar to those stated in (4) in [Ben08],⁹⁶ will prove to be sufficient to reformulate (DP) in terms of a discrete-time QVI in the spirit of [Ben08]; see Theorem 3.2.5 below.

Assumption 3.2.3. For every $(t, x) \in [0:T] \times \mathbb{X}$ we assume the following requirements to be satisfied:

- (a) Each action $a \in \Gamma^\circ(t, x)$ has positive, uniformly lower bounded fixed costs, i.e. there is some $c_0 > 0$ and a measurable function $c_1 : [0:T] \times \mathbb{X} \times \mathbb{A} \rightarrow [0, \infty)$ satisfying $c_1(t, x, \boxtimes) = 0$ such that

$$c(t, x, a) = c_0 \cdot \mathbf{1}_{\Gamma^\circ(t, x)}(a) + c_1(t, x, a) \quad \text{for all } a \in \Gamma(t, x).$$

- (b) For every $a \in \Gamma^\circ(t, x)$ and every $a' \in \Gamma^\circ(t, \xi(x, a))$ there is some $a \oplus a' \in \Gamma^\circ(t, x)$ such that it holds

- direct accessibility: $\xi(x, a \oplus a') = \xi(\xi(x, a), a')$ and
- subadditivity: $c_1(t, x, a \oplus a') \leq c_1(t, x, a) + c_1(t, \xi(x, a), a')$.




Figure 3.2: Concatenation of actions.

Condition (b) is illustrated in Figure 3.2. Note that the diagram does not commute when taking costs into account: Due to incurred fixed costs, a single direct action is always more favorable than a concatenation of two separate ones. \diamond

As a direct consequence we obtain the following auxiliary result:

Lemma 3.2.4 (No Multiple Interventions). Let Assumption 3.2.3 be satisfied, $v : [0:T] \times \mathbb{X} \rightarrow \mathbb{R}$ be a function and $(t, x) \in [0:T] \times \mathbb{X}$. Then for all $a \in \Gamma^\circ(t, x)$ and $a' \in \Gamma^\circ(t, \xi(x, a))$ it holds

$$v(t, \xi(\xi(x, a), a')) - c(t, \xi(x, a), a') \leq \mathcal{M}[v](t, x) + c(t, x, a) - c_0.$$

⁹⁶The conditions stated below are widespread in the literature on impulse control; see also, e.g., §1 and (C4) in [Ste86] or Definition 1.7 in [Die09].

In particular, we have

$$\mathcal{M}[v](t, \xi(x, a)) \leq \mathcal{M}[v](t, x) + c(t, x, a) - c_0 \quad \text{for every } a \in \Gamma^\circ(t, x).$$

Proof. Let $a \in \Gamma^\circ(t, x)$ and $a' \in \Gamma^\circ(t, \xi(x, a))$. Using Assumption 3.2.3(b) we fix some action $a \oplus a' \in \Gamma^\circ(t, x)$ and observe that

$$\begin{aligned} & v(t, \xi(\xi(x, a), a')) - c(t, \xi(x, a), a') \\ &= v(t, \xi(x, a \oplus a')) - (c(t, \xi(x, a), a') + c(t, x, a)) + c(t, x, a) \\ &= v(t, \xi(x, a \oplus a')) - (2c_0 + c_1(t, \xi(x, a), a') + c_1(t, x, a)) + c(t, x, a) \\ &\leq v(t, \xi(x, a \oplus a')) - c(t, x, a \oplus a') - c_0 + c(t, x, a) \leq \mathcal{M}[v](t, x) + c(t, x, a) - c_0, \end{aligned} \quad (3.8)$$

as desired. If $\Gamma^\circ(t, \xi(x, a)) = \emptyset$, the second part of the assertion is trivially satisfied since we have $\mathcal{M}[v](t, \xi(x, a)) = -\infty$ in this case; otherwise it follows by taking sup with respect to $a' \in \Gamma^\circ(t, \xi(x, a))$ in the preceding inequality (3.8). \square

Now we are in a position to establish the announced reformulation of the classical dynamic programming equation (DP) in terms of a discrete-time QVI. The proof of the following result is a ramification of the arguments provided in §4.3 in [Ben08].

Theorem 3.2.5 (Quasi-Variational Inequality). Suppose that Assumption 3.2.3 is satisfied and let $v \in \mathcal{V}$. Then v solves (DP) if and only if it solves the equation

$$v(t, x) = \mathcal{M}[v](t, x) \vee \mathcal{C}[v](t, x) \quad \text{for all } (t, x) \in [0:T] \times \mathbb{X}. \quad (\text{QVI})$$

Note that (QVI) can be rewritten as

$$\min\{v(t, x) - \mathcal{C}[v](t, x), v(t, x) - \mathcal{M}[v](t, x)\} = 0 \quad \text{for } (t, x) \in [0:T] \times \mathbb{X}, \quad (\text{QVI})$$

which strongly resembles the usual formulation considered in continuous-time impulse control; see, e.g., (3) in [BCS17]. \diamond

Proof of Theorem 3.2.5. Observe that solutions of (DP) and (QVI) coincide at time $t = T$ as well as at all configurations $(t, x) \in [0:T] \times \mathbb{X}$ for which $\Gamma^\circ(t, x) = \emptyset$ is satisfied; the latter is due to the fact that in this case we have $\mathcal{M}[v](t, x) = \mathcal{M}[\mathcal{C}[v]](t, x) = -\infty$.

Furthermore note that in any case and for every $t \in [0:T]$ we have $\mathcal{C}[v](t, \cdot) \leq v(t, \cdot)$ whence for all $(t, x) \in [0:T] \times \mathbb{X}$ it holds

$$\begin{aligned} \mathcal{M}[\mathcal{C}[v]](t, x) &= \sup_{a \in \Gamma^\circ(t, x)} \{\mathcal{C}[v](t, \xi(x, a)) - c(t, x, a)\} \\ &\leq \sup_{a \in \Gamma^\circ(t, x)} \{v(t, \xi(x, a)) - c(t, x, a)\} = \mathcal{M}[v](t, x). \end{aligned} \quad (3.9)$$

Throughout the remainder of this proof we fix some $(t, x) \in [0:T] \times \mathbb{X}$ such that $\Gamma^\circ(t, x) \neq \emptyset$, and we split the proof into two steps:

Step 1: Sufficiency. Suppose that v solves (QVI). We distinguish two cases:

Case 1: It holds $v(t, x) = \mathcal{M}[v](t, x)$. Let $\varepsilon \in (0, c_0/2)$ and approximate the sup within the definition of \mathcal{M} from below, i.e. let $a_\varepsilon \in \Gamma^\circ(t, x)$ be some action such that

$$v(t, x) < v(t, \xi(x, a_\varepsilon)) - c(t, x, a_\varepsilon) + \varepsilon. \quad (3.10)$$

We claim that it holds

$$v(t, \xi(x, a_\varepsilon)) = \mathcal{C}[v](t, \xi(x, a_\varepsilon)). \quad (3.11)$$

To prove (3.11), suppose to the contrary that it is not valid. Note that we then have $v(t, \xi(x, a_\varepsilon)) > \mathcal{C}[v](t, \xi(x, a_\varepsilon))$ and thus $v(t, \xi(x, a_\varepsilon)) = \mathcal{M}[v](t, \xi(x, a_\varepsilon)) \in \mathbb{R}$ since $v \in \mathbf{V}$ solves (QVI); in particular it follows $\Gamma^\circ(t, \xi(x, a_\varepsilon)) \neq \emptyset$. Hence, by the same token as for (3.10) there is some action $a'_\varepsilon \in \Gamma^\circ(t, \xi(x, a_\varepsilon))$ such that

$$v(t, \xi(x, a_\varepsilon)) < v(t, \xi(\xi(x, a_\varepsilon), a'_\varepsilon)) - c(t, \xi(x, a_\varepsilon), a'_\varepsilon) + \varepsilon. \quad (3.12)$$

As a consequence, we obtain from (3.10), (3.12) and Lemma 3.2.4 that

$$\begin{aligned} \mathcal{M}[v](t, x) &= v(t, x) < v(t, \xi(\xi(x, a_\varepsilon), a'_\varepsilon)) - c(t, \xi(x, a_\varepsilon), a'_\varepsilon) - c(t, x, a_\varepsilon) + 2\varepsilon \\ &\leq \mathcal{M}[v](t, x) - c_0 + 2\varepsilon < \mathcal{M}[v](t, x), \end{aligned}$$

where the final inequality follows since $\varepsilon \in (0, c_0/2)$. This is a contradiction. Thus, equation (3.11) must be valid and together with (3.10) it implies

$$v(t, x) < \mathcal{C}[v](t, \xi(x, a_\varepsilon)) - c(t, x, a_\varepsilon) + \varepsilon \leq \mathcal{M}[\mathcal{C}[v]](t, x) + \varepsilon.$$

Since $\varepsilon \in (0, c_0/2)$ was arbitrary, we can send $\varepsilon \downarrow 0$ and obtain $v(t, x) \leq \mathcal{M}[\mathcal{C}[v]](t, x)$. On the other hand, by (3.9) we have $\mathcal{M}[\mathcal{C}[v]](t, x) \leq \mathcal{M}[v](t, x) = v(t, x)$, and thus $v(t, x) = \mathcal{M}[\mathcal{C}[v]](t, x) = \mathcal{M}[\mathcal{C}[v]](t, x) \vee \mathcal{C}[v](t, x)$, where the latter identity is due to (QVI). Hence (DP) is satisfied at (t, x) .

Case 2: It holds $v(t, x) = \mathcal{C}[v](t, x)$. From (3.9) and (QVI) we conclude that $\mathcal{M}[\mathcal{C}[v]](t, x) \leq \mathcal{M}[v](t, x) \leq v(t, x)$. Thus, we get $v(t, x) = \mathcal{C}[v](t, x) = \mathcal{M}[\mathcal{C}[v]](t, x) \vee \mathcal{C}[v](t, x)$, whence also in this case (DP) is satisfied at (t, x) .

Step 2: Necessity. Suppose that v solves (DP). Moreover assume that $\mathcal{M}[v](t, x) > v(t, x)$. By definition of \mathcal{M} and since $\Gamma^\circ(t, x) \neq \emptyset$ there is some action $a_0 \in \Gamma^\circ(t, x)$ such that

$$\mathcal{M}[v](t, x) \geq v(t, \xi(x, a_0)) - c(t, x, a_0) > v(t, x). \quad (3.13)$$

Suppose that $v(t, \xi(x, a_0)) = \mathcal{C}[v](t, \xi(x, a_0))$. Then it follows from (3.13) and (DP) that

$$v(t, x) < \mathcal{C}[v](t, \xi(x, a_0)) - c(t, x, a_0) \leq \mathcal{M}[\mathcal{C}[v]](t, x) \leq v(t, x),$$

a contradiction. Hence (DP) yields $v(t, \xi(x, a_0)) = \mathcal{M}[\mathcal{C}[v]](t, \xi(x, a_0))$, and Lemma 3.2.4 thus

implies that

$$v(t, \xi(x, a_0)) = \mathcal{M}[\mathcal{C}[v]](t, \xi(x, a_0)) \leq \mathcal{M}[\mathcal{C}[v]](t, x) + c(t, x, a_0) - c_0. \quad (3.14)$$

With the help of (3.13), (3.14) and (DP) we conclude, in turn, that

$$v(t, x) < v(t, \xi(x, a_0)) - c(t, x, a_0) \leq \mathcal{M}[\mathcal{C}[v]](t, x) - c_0 \leq v(t, x) - c_0 < v(t, x),$$

a contradiction. It follows that $\mathcal{M}[v](t, x) \leq v(t, x)$ and we distinguish two cases again:

Case 1: It holds $v(t, x) = \mathcal{M}[\mathcal{C}[v]](t, x)$. Then, using (3.9), we obtain $v(t, x) \leq \mathcal{M}[v](t, x) \leq v(t, x)$ and thus from (DP) further that $v(t, x) = \mathcal{M}[v](t, x) = \mathcal{M}[v](t, x) \vee \mathcal{C}[v](t, x)$ which establishes (QVI) in this case.

Case 2: It holds $v(t, x) = \mathcal{C}[v](t, x)$. In this case, (QVI) follows immediately as $\mathcal{M}[v](t, x) \leq v(t, x) = \mathcal{C}[v](t, x)$.

Hence, as desired, (DP) and (QVI) are equivalent. \square

Our goal is to use Theorem 3.2.5 to identify (QVI) as another verification equation for the considered impulse control problem (C). To this end, we make use of the standard concepts of a *continuation* and an *intervention region*:⁹⁷

Definition 3.2.6 (Continuation and Intervention Region). Given the value function v^* of problem (C), see (3.2), the set

$$\mathfrak{C} \triangleq \{(t, x) \in [0:T] \times \mathbb{X} : v^*(t, x) > \mathcal{M}[v^*](t, x)\}$$

is called its *continuation region*, whereas

$$\mathfrak{J} \triangleq \{(t, x) \in [0:T] \times \mathbb{X} : v^*(t, x) = \mathcal{M}[v^*](t, x)\}$$

is said to be its *intervention region*. \diamond

In view of (QVI), we thus obtain the following verification theorem as an analogue of Theorem 3.2.1:

Corollary 3.2.7 (Verification via (QVI)). Let Assumptions 3.1.1 and 3.2.3 be satisfied and assume that $v \in \mathbb{V}$ is a solution of (QVI). Furthermore suppose that \mathcal{M} admits a measurable maximizer in the sense that there exists a measurable function $g^* : [0:T] \times \mathbb{X} \rightarrow \mathbb{A}$ such that for all $(t, x) \in [0:T] \times \mathbb{X}$ it holds $g^*(t, x) \in \Gamma^\circ(t, x)$ and

$$\mathcal{M}[v](t, x) = v(t, \xi(x, g^*(t, x))) - c(t, x, g^*(t, x)). \quad (3.15)$$

Then $v = v^*$ is the value function of (C) and an optimal impulse control is given by

$$\alpha_s^* \triangleq a^*(s, X_s^{\alpha^*}) = \mathbf{1}_{\mathfrak{C}}(s, X_s^{\alpha^*}) \cdot \boxtimes + \mathbf{1}_{\mathfrak{J}}(s, X_s^{\alpha^*}) \cdot g^*(s, X_s^{\alpha^*}) \quad \text{for } s \in [t:T]. \quad (3.16)$$

⁹⁷These concepts are standard in the literature on continuous-time impulse control; see, e.g., (17) and (18) in [BCS17].

Sufficient conditions for existence of g^* in (3.15) will be stated below; see Assumption 3.2.9 and Proposition 3.2.10. \diamond

Proof of Corollary 3.2.7. Since $v \in \mathbf{V}$ solves (QVI), first note that Theorem 3.2.5 implies that it is also a solution of (DP). Furthermore, by (3.15) the map $\mathcal{M}[v] : [0:T] \times \mathbb{X} \rightarrow \mathbb{R}$ is measurable and thus the following *candidate sets* are measurable as well:

$$\mathfrak{C}_v \triangleq \{(t, x) \in [0:T] \times \mathbb{X} : v(t, x) > \mathcal{M}[v](t, x)\} \quad (3.17)$$

is said to be the *candidate continuation set* based on v , whereas

$$\mathfrak{J}_v \triangleq ([0:T] \times \mathbb{X}) \setminus \mathfrak{C}_v = \{(t, x) \in [0:T] \times \mathbb{X} : v(t, x) = \mathcal{M}[v](t, x)\} \quad (3.18)$$

denotes the corresponding *candidate intervention set*; note that the second identity in (3.18) is valid due to (QVI); by the same token, we have $v(t, x) = \mathcal{C}[v](t, x)$ for $(t, x) \in \mathfrak{C}_v$.

Based on the above candidate sets, we define the measurable action map

$$a^* : [0:T] \times \mathbb{X} \rightarrow \mathbb{A}, \quad a^*(t, x) \triangleq \mathbb{1}_{\mathfrak{C}_v}(t, x) \cdot \boxtimes + \mathbb{1}_{\mathfrak{J}_v}(t, x) \cdot g^*(t, x) \quad (3.19)$$

and aim to prove that it provides a measurable maximizer for (DP) as required by (3.6). To this end, we fix some $(t, x) \in [0:T] \times \mathbb{X}$, note that $a^*(t, x) \in \Gamma(t, x)$ and distinguish the following two cases:

- (a) For $(t, x) \in \mathfrak{C}_v$ we have $a^*(t, x) = \boxtimes$ and $v(t, x) = \mathcal{C}[v](t, x)$. Thanks to (DP) it follows that

$$\begin{aligned} \sup_{a \in \Gamma^\circ(t, x)} \{\mathcal{C}[v](t, \xi(x, a)) - c(t, x, a)\} &= \mathcal{M}[\mathcal{C}[v]](t, x) \leq v(t, x) \\ &= \mathcal{C}[v](t, x) = \mathcal{C}[v](t, \xi(x, a^*(t, x))) - c(t, x, a^*(t, x)). \end{aligned} \quad (3.20)$$

- (b) For $(t, x) \in \mathfrak{J}_v$ we have $a^*(t, x) = g^*(t, x) \in \Gamma^\circ(t, x)$ and $v(t, x) = \mathcal{M}[v](t, x)$. From Step 1, Case 1 in the proof of Theorem 3.2.5 it follows that

$$\sup_{a \in \Gamma^\circ(t, x)} \{\mathcal{C}[v](t, \xi(x, a)) - c(t, x, a)\} = \mathcal{M}[\mathcal{C}[v]](t, x) = v(t, x) = \mathcal{M}[v](t, x).$$

On the other hand, we obtain with Lemma 3.2.4 and (3.15) that

$$\begin{aligned} \mathcal{M}[v](t, \xi(x, a^*(t, x))) &\leq \mathcal{M}[v](t, x) + c(t, x, a^*(t, x)) - c_0 \\ &= v(t, \xi(x, a^*(t, x))) - c_0 < v(t, \xi(x, a^*(t, x))), \end{aligned}$$

whence $(t, \xi(x, a^*(t, x))) \in \mathfrak{C}_v$. As a result and thanks to (QVI) and (3.15) it follows that

$$\begin{aligned} \mathcal{C}[v](t, x) \leq v(t, x) &= \sup_{a \in \Gamma^\circ(t, x)} \{\mathcal{C}[v](t, \xi(x, a)) - c(t, x, a)\} = \mathcal{M}[\mathcal{C}[v]](t, x) \\ &= \mathcal{M}[v](t, x) = v(t, \xi(x, a^*(t, x))) - c(t, x, a^*(t, x)) \\ &= \mathcal{C}[v](t, \xi(x, a^*(t, x))) - c(t, x, a^*(t, x)). \end{aligned} \quad (3.21)$$

Hence, in both cases, see (3.20) and (3.21), we have

$$a^*(t, x) \in \arg \max_{a \in \Gamma(t, x)} \{ \mathcal{C}[v](t, \xi(x, a)) - c(t, x, a) \},$$

and thus, as v solves (DP), it holds

$$\begin{aligned} v(t, x) &= \mathcal{M}[\mathcal{C}[v]](t, x) \vee \mathcal{C}[v](t, x) \\ &= \sup_{a \in \Gamma(t, x)} \{ \mathcal{C}[v](t, \xi(x, a)) - c(t, x, a) \} = \mathcal{C}[v](t, \xi(x, a^*(t, x))) - c(t, x, a^*(t, x)), \end{aligned}$$

as required by (3.6). Thus, the claim follows from Theorem 3.2.1, and we obtain $v = v^*$. As a consequence, it holds $\mathfrak{C}_v = \mathfrak{C}$ as well as $\mathfrak{I}_v = \mathfrak{I}$, and using (3.19) in (3.7) yields that (3.16) defines an optimal impulse control. \square

Note that the specification of an optimal impulse control in (3.16) above literally explains the terms *continuation* and *intervention region*, respectively. However, note also that the following alternative definitions of the candidate continuation and intervention region (replacing (3.17) and (3.18), respectively) serve just as well to establish Corollary 3.2.7:

Remark 3.2.8. Let $v \in \mathbb{V}$ be a solution of (QVI) and suppose that the assumptions of Corollary 3.2.7 are satisfied. Instead of \mathfrak{C}_v and \mathfrak{I}_v , consider the sets

$$\mathfrak{C}'_v \triangleq \{ (t, x) \in [0: T] \times \mathbb{X} : v(t, x) = \mathcal{C}[v](t, x) \} \quad (3.17')$$

and

$$\mathfrak{I}'_v \triangleq ([0: T] \times \mathbb{X}) \setminus \mathfrak{C}'_v = \{ (t, x) \in [0: T] \times \mathbb{X} : v(t, x) > \mathcal{C}[v](t, x) \}. \quad (3.18')$$

Since, thanks to (QVI), for $(t, x) \in [0: T] \times \mathbb{X}$ it holds both

$$(t, x) \in \mathfrak{C}'_v \cup \mathfrak{C}_v \Rightarrow v(t, x) = \mathcal{C}[v](t, x) \quad \text{and} \quad (t, x) \in \mathfrak{I}'_v \cup \mathfrak{I}_v \Rightarrow v(t, x) = \mathcal{M}[v](t, x),$$

the arguments used in the proof of Corollary 3.2.7 remain valid and the corresponding adaption of a^* in (3.19) and α^* in (3.16), respectively, yields an optimal impulse control as well.

Note that $\mathfrak{C}'_v \supseteq \mathfrak{C}_v$ and hence $\mathfrak{I}_v \supseteq \mathfrak{I}'_v$; furthermore we have $(\mathfrak{C}_v \Delta \mathfrak{C}'_v) \cup (\mathfrak{I}_v \Delta \mathfrak{I}'_v) = (\mathfrak{C}'_v \setminus \mathfrak{C}_v) \cup (\mathfrak{I}_v \setminus \mathfrak{I}'_v) \subseteq \{ (t, x) \in [0: T] \times \mathbb{X} : \mathcal{C}[v](t, x) = \mathcal{M}[v](t, x) \}$.

Thus, loosely speaking, the only difference between the two definitions is that the strategy α^* in (3.16) intervenes less often if \mathfrak{C}'_v and \mathfrak{I}'_v are used instead of \mathfrak{C}_v and \mathfrak{I}_v : While for the latter definition interventions are already performed if $(t, x) \in \mathfrak{C}'_v \cap \mathfrak{I}_v$, i.e. if $\mathcal{C}[v](t, x) = v(t, x) = \mathcal{M}[v](t, x)$, for the former one the strict inequality $v(t, x) > \mathcal{C}[v](t, x)$ needs to be satisfied.

However, note that an analogous replacement of \mathfrak{C} and \mathfrak{I} in Definition 3.2.6 is in general not possible as v^* may fail to be measurable and hence $\mathcal{C}[v^*]$ need not be well-defined. \diamond

3.2.3 Excursion: Existence of a Measurable Maximizer

As already promised above and in a setup general enough for our purposes, we provide a set of sufficient conditions to establish existence of a measurable maximizer g^* as required by (3.15) in

Corollary 3.2.7:

Assumption 3.2.9. We suppose that the following conditions are met:

- (a) $\mathbb{X} \in \mathfrak{B}(\mathbb{R}^d)$ is a measurable subset of \mathbb{R}^d , where $d \in \mathbb{N}$, and $\mathfrak{X} = \mathfrak{B}(\mathbb{X})$.
- (b) $\mathbb{A} \in \mathfrak{B}(\mathbb{R}^e)$ is a closed subset of \mathbb{R}^e , where $e \in \mathbb{N}$, and $\mathfrak{A} = \mathfrak{B}(\mathbb{A})$.

Furthermore we assume that for every $t \in [0:T)$ it holds:

- (c) The set-valued function $\Gamma^\circ(t, \cdot) : \mathbb{X} \rightarrow 2^{\mathbb{A}}$ is both nonempty- and compact-valued; furthermore it is continuous in the sense that it is both upper and lower hemicontinuous at every $x \in \mathbb{X}$, i.e.⁹⁸
 - for all sequences $\{x_n\}_{n \in \mathbb{N}} \in \mathbb{X}^{\mathbb{N}}$, $\{a_n\}_{n \in \mathbb{N}} \in \mathbb{A}^{\mathbb{N}}$ satisfying $\lim_{n \rightarrow \infty} x_n = x$ and $a_n \in \Gamma^\circ(t, x_n)$ for all $n \in \mathbb{N}$ there exists a subsequence $\{n_j\}_{j \in \mathbb{N}}$ and some $a \in \Gamma^\circ(t, x)$ such that $\lim_{j \rightarrow \infty} a_{n_j} = a$ (*upper hemicontinuity*), and
 - for every sequence $\{x_n\}_{n \in \mathbb{N}} \in \mathbb{X}^{\mathbb{N}}$ satisfying $\lim_{n \rightarrow \infty} x_n = x$ and every $a \in \Gamma^\circ(t, x)$ there exists some subsequence $\{n_j\}_{j \in \mathbb{N}}$ and elements $a_{n_j} \in \Gamma^\circ(t, x_{n_j})$ for each $j \in \mathbb{N}$ such that $\lim_{j \rightarrow \infty} a_{n_j} = a$ (*lower hemicontinuity*). \diamond

Now standard results imply existence of a measurable maximizer g^* :

Proposition 3.2.10. Let $v : [0 : T] \times \mathbb{X} \rightarrow \mathbb{R}$ be a measurable function and suppose that Assumption 3.2.9 is satisfied. Moreover assume for every $t \in [0:T)$ that the function

$$\mathbf{Graph}(\Gamma^\circ(t, \cdot)) \ni (x, a) \mapsto v(t, \xi(x, a)) - c(t, x, a) \in \mathbb{R}$$

is continuous, where $\mathbf{Graph}(\Gamma^\circ(t, \cdot)) \triangleq \{(x, a) \in \mathbb{X} \times \mathbb{A} : a \in \Gamma^\circ(t, x)\}$. Then there exists a measurable function $g^* : [0 : T) \times \mathbb{X} \rightarrow \mathbb{A}$ such that for all $(t, x) \in [0 : T) \times \mathbb{X}$ it holds $g^*(t, x) \in \Gamma^\circ(t, x)$ and

$$\mathcal{M}[v](t, x) = v(t, \xi(x, g^*(t, x))) - c(t, x, g^*(t, x)).$$

Proof. Until further notice, we fix some $t \in [0 : T)$. Under the stated assumptions, *Berge's Maximum Theorem* (see, e.g., Theorem 17.31 in [AB06]) yields that the function $\mathcal{M}[v](t, \cdot)$ is \mathbb{R} -valued and continuous and satisfies

$$\mathcal{M}[v](t, x) = \max_{a \in \Gamma^\circ(t, x)} \{v(t, \xi(x, a)) - c(t, x, a)\} \quad \text{for } x \in \mathbb{X};$$

moreover, by the same result, the set-valued map of maximizers

$$G^*(t, \cdot) : \mathbb{X} \rightarrow 2^{\mathbb{A}}, \quad G^*(t, x) \triangleq \{a \in \Gamma^\circ(t, x) : v(t, \xi(x, a)) - c(t, x, a) = \mathcal{M}[v](t, x)\}$$

is nonempty- and compact-valued and upper hemicontinuous. This, in turn, allows us to apply the *Kuratowski-Ryll-Nardzewski Selection Theorem* (see, e.g., Theorem 18.13 in [AB06])⁹⁹ to

⁹⁸This version of the definition of hemicontinuity is derived from Definition 17.2 and Theorems 17.20/21 in [AB06].
⁹⁹For its requirement of $G^*(t, \cdot)$ to be *weakly measurable* in the sense of [AB06, Definition 18.1] we resort to both [AB06, Lemma 17.4] and [AB06, Lemma 18.2.1].

conclude that the set-valued map $G^*(t, \cdot)$ admits a measurable selector $g^*(t, \cdot) : \mathbb{X} \rightarrow \mathbb{A}$. As a consequence there is a measurable map $g^* : [0:T] \times \mathbb{X} \rightarrow \mathbb{A}$ satisfying $g^*(t, x) \in G^*(t, x)$ for every $(t, x) \in [0:T] \times \mathbb{X}$, i.e.

$$v(t, \xi(x, g^*(t, x))) - c(t, x, g^*(t, x)) = \mathcal{M}[v](t, x) \quad \text{for all } (t, x) \in [0:T] \times \mathbb{X},$$

as desired. \square

As an immediate consequence of the latter result, we obtain existence of a measurable maximizer g^* in the situation of Corollary 3.2.7 if Assumption 3.2.9 is satisfied and the considered solution $v \in \mathcal{V}$ of (QVI) as well as the cost function c are such that the required continuity condition is fulfilled. \diamond

For a general and comprehensive account of the theory of (measurable) correspondences we refer, e.g., to Chapters 17 and 18 in [AB06]; a *Theorem of the Maximum* is also discussed in §3.3 in [SL89]; and for *measurable selection* one may also consult [Sch74].

3.3 A Self-Learning Monte Carlo Algorithm

In this section we introduce and investigate the announced algorithm to tackle problem (C). We discuss both its theoretical background as well as an approach to a numerical implementation. While the former is essentially based on the concept of policy iteration, see Section 3.3.1, the latter additionally resorts to Monte Carlo simulation and builds up a hierarchy of stochastic interpolation nodes in order to successively learn the behavior of the optimally controlled state process, see Section 3.3.2 for details. Note that the developed heuristics of combining policy iteration, non-nested Monte Carlo simulation and an exploration strategy that adapts to the structure of the underlying problem constitutes the core contribution of this chapter.¹⁰⁰

Throughout this section we suppose that Assumptions 3.1.1 and 3.2.3 are satisfied.

3.3.1 Theoretical Aspects

From a theoretical point of view we propose the following policy iteration-based algorithm (3.22)-(3.24) in order to attack problem (C), i.e. to approximate its value function v^* . It consists of levels $\ell = 0, 1, \dots$, and is inspired by the optimal impulse control (3.16) based on a solution of the verification equation (QVI) (see Corollary 3.2.7): Pretending that $v_{\ell-1}$ is a solution, for every $\ell \in \mathbb{N}$ the iteration steps below construct optimal actions similarly as in (3.19) and evaluate the resulting strategy in order to assign it a value v_ℓ . The individual steps are as follows:

The algorithm is initialized with the function $v_0 : [0:T] \times \mathbb{X} \rightarrow \mathbb{R}$ defined by

$$v_0(t, x) \triangleq \mathcal{C}[v_0](t, x) \quad \text{for } t = T, T-1, \dots, 0, x \in \mathbb{X}, \quad (3.22)$$

and subsequently constructs the functions v_ℓ for $\ell \geq 1$ inductively in two steps:

¹⁰⁰For details on policy iteration we refer, e.g., to Chapter 4 in [How60], §6.4 in [Put94] or §1.2 in [Sei11].

(a) Policy Improvement: For every $(t, x) \in [0:T) \times \mathbb{X}$ compute $g_\ell^*(t, x) \in \Gamma^\circ(t, x)$ such that¹⁰¹

$$g_\ell^*(t, x) \in \arg \max_{a \in \Gamma^\circ(t, x)} \{v_{\ell-1}(t, \xi(x, a)) - c(t, x, a)\}. \quad (3.23)$$

(b) Policy Evaluation: If $t = T$ define $v_\ell(T, x) \triangleq \Psi(x)$ for $x \in \mathbb{X}$; otherwise proceed backward in time, i.e. $t = T - 1, \dots, 0$, and set for $x \in \mathbb{X}$

$$v_\ell(t, x) \triangleq \begin{cases} \mathcal{C}[v_\ell](t, x) & \text{if } v_{\ell-1}(t, x) > \mathcal{M}[v_{\ell-1}](t, x), \\ \mathcal{C}[v_\ell](t, x) \vee (\mathcal{C}[v_\ell](t, \xi(x, g_\ell^*(t, x))) - c(t, x, g_\ell^*(t, x))) & \text{otherwise.} \end{cases} \quad (3.24)$$

Recall from Definition 3.2.2 that $\mathcal{C}[v_\ell](t, x) = \psi(t, x) + \mathbb{E}[v_\ell(t+1, \phi(t, x, Z))]$ for $(t, x) \in [0:T) \times \mathbb{X}$ and $\ell \in \mathbb{N}_0$; thus the backward recursions in (3.22) and (3.24) are both explicit.¹⁰² \diamond

Let us briefly comment on how this (theoretical) algorithm works: It is initialized with the value v_0 of the trivial strategy $\alpha^0 = \{\boxtimes\}$, i.e. $\alpha_s^0 = \boxtimes$ for all relevant s , that leads to the uncontrolled state process X^{α^0} . Then, for $\ell \in \mathbb{N}$, with $v_{\ell-1}$ being the available approximation of the value function of problem (C), the algorithm computes an optimal intervention action $g_\ell^*(t, x) \in \Gamma^\circ(t, x)$ for each $(t, x) \in [0:T) \times \mathbb{X}$. Afterwards, using $v_{\ell-1}$ as a basis of decision-making, it decides either to not carry out this intervention in case that it is considered to result in a lower value, or it implements the computed action in the opposite case. Finally, the resulting policy is consistently evaluated in a manner inspired by classical backward induction or reward iteration,¹⁰³ with an additional comparison of values corresponding to inaction and intervention, respectively, if the policy suggests the latter of these two options. \diamond

From a more abstract point of view the iterative construction of levels $v_{\ell-1} \mapsto v_\ell$, $\ell \in \mathbb{N}$, given by (3.23) and (3.24) can be regarded as the application of an operator $\mathcal{T} : \mathbf{V} \rightarrow \mathbf{V}$,¹⁰⁴ i.e. it holds $v_\ell = \mathcal{T}[v_{\ell-1}]$ for $\ell \in \mathbb{N}$, and hence the above algorithm can be regarded as a fixed-point iteration corresponding to \mathcal{T} . Thus, to justify its correctness, it suffices to show that a fixed point of this operator coincides with the value function of problem (C).

To do so, we make use of the following auxiliary result that is interesting on its own as it rules out multiple interventions on each level of the above algorithm:

Lemma 3.3.1 (No Multiple Interventions). Let Assumption 3.2.3 be fulfilled and suppose that the operator $\mathcal{T} : \mathbf{V} \rightarrow \mathbf{V}$ is well-defined.¹⁰⁴ Then the sequence $\{v_\ell\}_{\ell \in \mathbb{N}_0} \in \mathbf{V}^{\mathbb{N}_0}$ satisfying $v_\ell = \mathcal{T}[v_{\ell-1}]$ for $\ell \in \mathbb{N}$ as specified by (3.22)-(3.24) satisfies

¹⁰¹We assume tacitly that the maximizer maps g_ℓ^* , $\ell \in \mathbb{N}$, exist and that they are measurable; see also Footnote 104 below for further details. Note that then it holds $v_{\ell-1}(t, \xi(x, g_\ell^*(t, x))) - c(t, x, g_\ell^*(t, x)) = \mathcal{M}[v_{\ell-1}](t, x)$.

¹⁰²With regard to potential integrability issues, we refer to Footnote 104 below.

¹⁰³see, e.g., §4.2 in [Put94] or Theorem 2.3.4 in [BR11].

¹⁰⁴To fix ideas and avoid technical subtleties, throughout this section we implicitly *assume* that the operator \mathcal{T} is well-defined; this includes appropriate integrability assumptions as well as the existence of suitable measurable maximizers as required by (3.23). To this end and for instance, note that conditions similar to the ones stated in Remark 3.1.2 (e.g., replace $(\cdot)^-$ with $|\cdot|$ in (3.3)) and Assumption 3.2.9 along with suitable continuity conditions (as, e.g., stated in Proposition 3.2.10) need to be imposed. For a related discussion of such conditions for the classical Bellman equation (DP) we refer, e.g., to §2.4 in [BR11].

$$v_\ell(t, \xi(x, g_\ell^*(t, x))) = \mathcal{C}[v_\ell](t, \xi(x, g_\ell^*(t, x))) \quad \text{whenever } v_{\ell-1}(t, x) \leq \mathcal{M}[v_{\ell-1}](t, x)$$

for $\ell \in \mathbb{N}$ and $(t, x) \in [0:T] \times \mathbb{X}$.

Proof. Let $\ell \in \mathbb{N}$ and $(t, x) \in [0:T] \times \mathbb{X}$ such that $v_{\ell-1}(t, x) \leq \mathcal{M}[v_{\ell-1}](t, x)$. Note from (3.23) that $a^* \triangleq g_\ell^*(t, x) \in \Gamma^\circ(t, x)$ satisfies $v_{\ell-1}(t, \xi(x, a^*)) - c(t, x, a^*) = \mathcal{M}[v_{\ell-1}](t, x)$. Hence, Lemma 3.2.4 yields

$$\begin{aligned} \mathcal{M}[v_{\ell-1}](t, \xi(x, a^*)) &\leq \mathcal{M}[v_{\ell-1}](t, x) + c(t, x, a^*) - c_0 \\ &= v_{\ell-1}(t, \xi(x, a^*)) - c_0 < v_{\ell-1}(t, \xi(x, a^*)), \end{aligned}$$

and thus by (3.24) it holds $v_\ell(t, \xi(x, a^*)) = \mathcal{C}[v_\ell](t, \xi(x, a^*))$, as desired. \square

We are now in a position to prove correctness of our (theoretical) algorithm:

Proposition 3.3.2 (Fixed-Point Property). Let Assumptions 3.1.1 and 3.2.3 be fulfilled and suppose that the operator $\mathcal{T} : \mathbf{V} \rightarrow \mathbf{V}$ is well-defined.¹⁰⁴ If $v \in \mathbf{V}$ is a fixed point of \mathcal{T} then it holds $v = v^*$, i.e. v coincides with the value function of problem (C).

Proof. Using (3.24) first note that $v(T, \cdot) \equiv \Psi$; for $(t, x) \in [0:T] \times \mathbb{X}$ we distinguish the following two cases:

1. Suppose we have $v(t, x) > \mathcal{M}[v](t, x)$. Then it follows from (3.24) that $v(t, x) = \mathcal{C}[v](t, x) = \mathcal{C}[v](t, x) \vee \mathcal{M}[v](t, x)$.
2. Suppose we have $v(t, x) \leq \mathcal{M}[v](t, x)$.¹⁰⁵ Let $a^* \triangleq g^*(t, x) \in \Gamma^\circ(t, x)$ and observe from Lemma 3.3.1 that $v(t, \xi(x, a^*)) = \mathcal{C}[v](t, \xi(x, a^*))$.¹⁰⁶ Thus it follows with (3.23) and (3.24) that

$$\begin{aligned} v(t, x) &= \mathcal{C}[v](t, x) \vee (\mathcal{C}[v](t, \xi(x, a^*)) - c(t, x, a^*)) \\ &= \mathcal{C}[v](t, x) \vee (v(t, \xi(x, a^*)) - c(t, x, a^*)) = \mathcal{C}[v](t, x) \vee \mathcal{M}[v](t, x). \end{aligned}$$

Hence, in both cases the function $v \in \mathbf{V}$ satisfies (QVI) and the claim immediately follows from Corollary 3.2.7. \square

Note that Proposition 3.3.2 only shows that, in case of existence, the unique fixed point of the operator \mathcal{T} , with respect to which the algorithm (3.22)-(3.24) can be regarded as a fixed-point iteration, coincides with the value function of (C). However, although our illustrations in Section 3.4 corroborate the algorithm's correctness, it is by no means clear whether and under what assumptions the stated iteration indeed converges. We leave this open question for further research; see also our remarks on page 118 below.

¹⁰⁵In the following we consider $v_\ell \triangleq v_{\ell-1} \triangleq v$ such that $v_\ell = v = \mathcal{T}[v] = \mathcal{T}[v_{\ell-1}]$ since v is a fixed point of \mathcal{T} .

¹⁰⁶Note that g^* is defined as g_ℓ^* in (3.23) with v in place of $v_{\ell-1}$.

3.3.2 Implementation: Smart Voronoi Tessellations

In the remainder of this chapter we assume that the state space \mathbb{X} is a Borel measurable subset of \mathbb{R}^d , i.e. we consider $\mathbb{X} \in \mathfrak{B}(\mathbb{R}^d)$; moreover we use the Euclidean norm $\|\cdot\| \triangleq \|\cdot\|_2$ to measure distances.

Note that an implementation of the theoretical algorithm as outlined in (3.22)-(3.24) is confronted with the following two major issues:

- (a) The (conditional) expectations emerging in (3.22) and (3.24) are usually not explicitly computable and thus need to be approximated.
- (b) The (finite) number of points $x \in \mathbb{X}$ at which the value function of the problem can be approximated is limited by both time and memory constraints.

A natural stochastic approach to attack issue (a) is to resort to Monte Carlo simulation; we employ (non-nested, plain vanilla) crude Monte Carlo below.¹⁰⁷

As the well-known *curse of dimensionality*-phenomenon usually renders grid-based discretization techniques computationally infeasible already for moderate values of the dimension d , an exploration strategy based on a sparse ensemble of interpolation nodes is required and issue (b) is thus a bit more involved: It is by no means clear a priori which points in space should be selected as interpolation nodes. We address this problem by constructing a hierarchy of stochastic nodes that gradually adapt to the behavior of the optimally controlled state process.

For this, we make use of the concept of a *Voronoi tessellation* of points in \mathbb{R}^d which is standard in the literature on *quantization* and related methods.¹⁰⁸

Definition 3.3.3 (Voronoi Tessellations). Let $\mathbf{X} \triangleq \{x_i\}_{i \in [1:m]} \subset \mathbb{R}^d$ be a finite family of pairwise different points in \mathbb{R}^d .

- (i) The *Voronoi diagram* of \mathbf{X} is defined to be the family $\{\text{Vor}_i^{\mathbf{X}}\}_{i \in [1:m]}$ consisting of the *Voronoi regions*

$$\text{Vor}_i^{\mathbf{X}} \triangleq \{x \in \mathbb{R}^d : \|x - x_i\| \leq \|x - x_j\| \text{ for all } j \in [1:m] \setminus \{i\}\} \quad \text{for } i \in [1:m].$$

- (ii) Whenever $\{C_i^{\mathbf{X}}\}_{i \in [1:m]} \subset \mathfrak{B}(\mathbb{R}^d)$ is a partition of \mathbb{R}^d consisting of Borel measurable sets such that $C_i^{\mathbf{X}} \subseteq \text{Vor}_i^{\mathbf{X}}$ for all $i \in [1:m]$, it is called a *Voronoi partition* of \mathbb{R}^d with respect to \mathbf{X} .

In any case, for $i \in [1:m]$ the point x_i is called (*Voronoi*) *center* of the (*Voronoi*) *region* $\text{Vor}_i^{\mathbf{X}}$ or $C_i^{\mathbf{X}}$, respectively. \diamond

An illustration of the preceding definition can be found in Figure 3.3 for a random sample of $m = 25$ points in \mathbb{R}^2 .¹⁰⁹

¹⁰⁷see, e.g., Algorithm 3.1 in [KKK10] for the *crude* Monte Carlo algorithm.

¹⁰⁸see, e.g., [GL00] and [BP03, PPP04a, PPP04b].

The following definition is based on the notions discussed in §I.1 in [GL00] to which we also refer for further details.

¹⁰⁹The points have been sampled from a $\mathcal{N}_2(0_2, I_{2 \times 2})$ -distribution; the corresponding Voronoi tessellation has been computed using the SciPy-package `spatial.Voronoi` and plotted with `scipy.spatial.voronoi_plot_2d` and `matplotlib` in PYTHON; for details on the language and package versions used, we refer to Footnote 122 below.




Figure 3.3: Example of a Voronoi tessellation in \mathbb{R}^2 with centers given by the blue dots.

Intuitively it is clear that, up to intersecting boundaries, every Voronoi diagram partitions \mathbb{R}^d . For every point $x \in \mathbb{R}^d \setminus \bigcup_{i \neq j} (\partial \text{Vor}_i^{\mathbf{X}} \cap \partial \text{Vor}_j^{\mathbf{X}})$ there is a unique index $\text{idx}(x) = \text{idx}^{\mathbf{X}}(x) \in [1:m]$ such that $x \in \text{Vor}_{\text{idx}(x)}^{\mathbf{X}}$; only for points on the intersecting boundaries there is ambiguity; see also equation (1.6) and Proposition I.1.3 in [GL00].¹¹⁰ Clearly, whenever we are given a Voronoi partition $\{C_i^{\mathbf{X}}\}_{i \in [1:m]}$ of \mathbb{R}^d , idx can be unambiguously defined on the entire space \mathbb{R}^d with $C_i^{\mathbf{X}}$ in place of $\text{Vor}_i^{\mathbf{X}}$, $i \in [1:m]$, above. \diamond

For interpolating discretely sampled functions we employ the following two operators:

Definition 3.3.4 ((Shepard) Interpolation). Let $\mathbf{X} \triangleq \{x_i\}_{i \in [1:m]} \subset \mathbb{R}^d$ be a finite family of pairwise different points in \mathbb{R}^d and $\mathbf{V} \triangleq \{v_i\}_{i \in [1:m]} \subset [-\infty, \infty)$ a corresponding finite sequence of numbers.

- (i) For a given Voronoi partition $\{C_i^{\mathbf{X}}\}_{i \in [1:m]}$ of \mathbb{R}^d with respect to \mathbf{X} we define the *nearest-neighbor interpolation* operator $\mathcal{I}_{\mathbf{X}, \mathbf{V}}^{\text{NN}}$ supported on \mathbf{X} with values specified by \mathbf{V} to be

$$\mathcal{I}_{\mathbf{X}, \mathbf{V}}^{\text{NN}}(y) \triangleq \sum_{i=1}^m v_i \cdot \mathbb{1}_{C_i^{\mathbf{X}}}(y) = v_{\text{idx}(y)} \quad \text{for } y \in \mathbb{R}^d.$$

- (ii) For each $y \in \mathbb{R}^d$ let an *order statistic* with respect to distances to y be given by $\{x_{(i)y}\}_{i \in [1:m]}$,

¹¹⁰However, note that the union of these finitely many region boundaries has Lebesgue measure zero; see, e.g., Theorem I.1.5 in [GL00]. Thus, with $\mathfrak{X} = \mathfrak{B}(\mathbb{X})$, under suitable conditions on ϕ and μ , such that, e.g., it holds $\mathbb{P}(\phi(t, x, Z) \in \cdot) \ll \text{Leb}|_{\mathfrak{X}}$ for all $(t, x) \in [0:T] \times \mathbb{X}$, these intersections are negligible and hence of minor importance in practice.

i.e. $\|x_{(1)y} - y\| \leq \|x_{(2)y} - y\| \leq \dots \leq \|x_{(m)y} - y\|$.¹¹¹ For a number of queried neighbors $\kappa \in [1:m]$ and an exponent $p \geq 1$ we define weights for $i \in [1:\kappa]$ by

$$w_i^{\mathbf{X}}(y) \triangleq \begin{cases} \mathbb{1}_{\{j\}}(i) & \text{if there is some } j \in [1:\kappa] \text{ such that } x_{(j)y} = y, \\ \frac{1}{\|x_{(i)y} - y\|^p} & \text{otherwise,} \end{cases}$$

and set the corresponding *Shepard interpolation operator* $\mathcal{I}_{\mathbf{X},\mathbf{V}}^{\kappa,p}$ supported on \mathbf{X} with node values specified by \mathbf{V} to be

$$\mathcal{I}_{\mathbf{X},\mathbf{V}}^{\kappa,p}(y) \triangleq \frac{\sum_{i=1}^{\kappa} w_i^{\mathbf{X}}(y) \cdot v_{(i)y}}{\sum_{i=1}^{\kappa} w_i^{\mathbf{X}}(y)} \quad \text{for } y \in \mathbb{R}^d. \quad \diamond$$

While the first interpolation operator is standard given the above references on quantization-based methods, the latter interpolation method is a variant of the one commonly referred to as *Shepard's interpolation method* or *inverse distance weighting*; we refer, e.g., to [She68] for further details. Note that it assigns higher weights to those values v_i that correspond to points x_i which are closer to the point of interest y ; for large values of p only the closest points have a noticeable contribution to the overall interpolation result. Further observe that $\mathcal{I}_{\mathbf{X},\mathbf{V}}^{1,p}$, $p \geq 1$, can be identified with $\mathcal{I}_{\mathbf{X},\mathbf{V}}^{\text{NN}}$.¹¹² \diamond

With these tools at hand, we can now turn to a description of a ready-to-implement algorithm that mimics the theoretical outline in (3.22)-(3.24). Note that it successively builds up a hierarchy of levels each of which essentially consists of a finite number of (stochastic) interpolation nodes, referred to as *Voronoi centers* in the following, and corresponding value function estimates.

Throughout the remainder of this section we fix a probability measure γ on $(\mathbb{X}, \mathfrak{B}(\mathbb{X}))$, also referred to as the algorithm's *exploration distribution*; we let $q \in [0, 1]$ be a probability parameter; and we require the following additional assumption to hold:

Assumption 3.3.5. The considered probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ supports

- (a) a family $\{C_t^{\ell,m}\}_{t \in [0:T], \ell, m \in \mathbb{N}_0}$ of i.i.d. Bernoulli(q)-distributed random variables,
- (b) a family $\{Y_t^{\ell,m}\}_{t \in [0:T], \ell, m \in \mathbb{N}_0}$ of i.i.d. random variables with distribution γ , and
- (c) a family $\{Z_t^{\ell,m,n}\}_{t \in [1:T], \ell, m, n \in \mathbb{N}_0}$ of i.i.d. random variables with distribution μ .

These three sequences are assumed to be mutually independent as well as independent of all random variables that have occurred thus far. \diamond

Moreover, we fix

- a number of levels $L \in \mathbb{N}$,

¹¹¹In case that points within \mathbf{X} exhibit equal distance to some point $y \in \mathbb{R}^d$ the problem of ambiguity arises once again; however, as explained above, it does not pose a problem for our implementation below, and thus we take an arbitrary order statistic to be given.

¹¹²Here, the considered order statistics implicitly determines to which cell $C_i^{\mathbf{X}}$ of the underlying Voronoi partition a point on a region boundary belongs.

- a non-decreasing finite sequence $\{M_\ell\}_{\ell \in [0:L]} \subset \mathbb{N}$; M_ℓ denotes the number of considered Voronoi centers at each instant of time within $[0:T]$ on level ℓ ;
- a finite sequence $\{N_\ell\}_{\ell \in [0:L]} \subset \mathbb{N}$; N_ℓ denotes the sample size used for Monte Carlo estimations on level ℓ ;
- and interpolation parameters $\kappa \in [1:M_0]$ and $p \geq 1$ as explained in Definition 3.3.4(ii).

Our practical algorithm thus consists of the following steps:

▷ **(I) Initialization:**

(I-1) For $t = 0, \dots, T$ and $m \in [1:M_0]$ define the initial Voronoi centers by^{113, 114}

$$X_t^{0,m} \triangleq \begin{cases} Y_{0,m} & \text{if } t = 0, \\ \phi(t-1, X_{t-1}^{0,m}, Z_t^{0,m,0}) & \text{if } t \in [1:T], \end{cases}$$

and collect them in $\mathbf{X}_t^0 \triangleq \{X_t^{0,m}\}_{m \in [1:M_0]}$.

(I-2) Define $\hat{v}_0(T, x) \triangleq \Psi(x)$ for $x \in \mathbb{X}$ and set recursively for $t = T-1, \dots, 0$ and $m \in [1:M_0]$ the node values to

$$\hat{v}_t^{0,m} \triangleq \psi(t, X_t^{0,m}) + \frac{1}{N_0} \sum_{n=1}^{N_0} \hat{v}_0(t+1, \phi(t, X_t^{0,m}, Z_{t+1}^{0,m,n})),$$

and collect them in $\mathbf{V}_t^0 \triangleq \{\hat{v}_t^{0,m}\}_{m \in [1:M_0]}$.

(I-3) Set $\hat{v}_0(t, x) \triangleq \mathcal{I}_{\mathbf{X}_t^0, \mathbf{V}_t^0}^{\kappa, p}(x)$ for $(t, x) \in [0:T] \times \mathbb{X}$.

▷ **(II) Policy Iteration:** For $\ell = 1, \dots, L+1$ repeat the following steps:

↪ Policy Improvement:

(II-1) For $t = 0, \dots, T-1$ and $m \in [1:M_{\ell-1}]$ choose

$$\begin{aligned} (\hat{m}_t^{\ell-1,m}, \hat{g}_t^{\ell-1,m}) \in \arg \max \left\{ \hat{v}_t^{\ell-1,m'} - c(t, X_t^{\ell-1,m}, a) : \right. \\ \left. (m', a) \in ([1:M_{\ell-1}] \setminus \{m\}) \times \Gamma^\circ(t, X_t^{\ell-1,m}), \right. \\ \left. \xi(X_t^{\ell-1,m}, a) = X_t^{\ell-1,m'} \right\}, \end{aligned}$$

if it exists; otherwise set $\hat{m}_t^{\ell-1,m} \triangleq 0$ and $\hat{g}_t^{\ell-1,m} \triangleq \boxtimes$. Correspondingly, set

$$\hat{M}_t^{\ell-1,m} \triangleq \begin{cases} -\infty & \text{if } \hat{m}_t^{\ell-1,m} = 0, \\ \hat{v}_t^{\ell-1,m'} - c(t, X_t^{\ell-1,m}, \hat{g}_t^{\ell-1,m}) & \text{else, where } m' = \hat{m}_t^{\ell-1,m}, \end{cases}$$

and collect these values in $\mathbf{M}_t^{\ell-1} \triangleq \{\hat{M}_t^{\ell-1,m}\}_{m \in [1:M_{\ell-1}]}$.

¹¹³Note that we write $t = 0, \dots, T-1$ (or T) to indicate that the corresponding loop runs forward in time, whereas $t = T-1, \dots, 0$ corresponds to a loop running backward in time.

¹¹⁴Without loss of generality, we can assume that all Voronoi centers on a specific level and at a specific time are pairwise different: If in the course of the algorithm two centers with different indexes happen to coincide, we implicitly assume that the algorithm replaces the redundant one in a suitable way.

(II-2) Set $\hat{\mathcal{M}}_t^{\ell-1}[\hat{v}_{\ell-1}](t, x) \triangleq \mathcal{I}_{\mathbf{X}_t^{\ell-1}, \mathbf{M}_t^{\ell-1}}^{\kappa, p}(x)$ for $(t, x) \in [0:T] \times \mathbb{X}$.

\rightsquigarrow If $\ell = L + 1$ stop; otherwise continue with the following steps:

\rightsquigarrow Adjustment of Voronoi Centers:

(II-3) For $m \in [1:M_{\ell-1}]$ set the preliminary initial Voronoi centers to $\tilde{X}_0^{\ell, m} \triangleq X_T^{\ell-1, m}$.

(II-4) For $t = 0, \dots, T-1$ and $m \in [1:M_{\ell-1}]$ let $m_1 \triangleq \text{id}_{\mathbf{X}_t^{\ell-1}}(\tilde{X}_t^{\ell, m})$ be the index of the Voronoi center closest to $\tilde{X}_t^{\ell, m}$ on level $\ell-1$ and $m_2 \triangleq \hat{m}_t^{\ell-1, m_1}$ be the index of the corresponding best target Voronoi center starting from center m_1 . Define $J_t^{\ell, m} \triangleq 0$ and set

$$\tilde{X}_{t+1}^{\ell, m} \triangleq \begin{cases} \phi(t, \tilde{X}_t^{\ell, m}, Z_{t+1}^{\ell, m, 0}) & \text{if } \hat{v}_{\ell-1}(t, \tilde{X}_t^{\ell, m}) > \hat{\mathcal{M}}_t^{\ell-1}[\hat{v}_{\ell-1}](t, \tilde{X}_t^{\ell, m}) \text{ or } m_2 = 0, \\ \phi(t, \tilde{X}_t^{\ell-1, m_2}, Z_{t+1}^{\ell, m, 0}) & \text{else. In this case redefine } J_t^{\ell, m} \triangleq 1. \end{cases}$$

(II-5) For $t = 0, \dots, T$ and $m \in [1:M_{\ell-1}]$ define the new Voronoi centers as follows: Set

$$X_t^{\ell, m} \triangleq \tilde{X}_t^{\ell, m} \quad \text{if } J_t^{\ell, m} = 1 \text{ or } C_t^{\ell, m} = 1 \text{ or } t = T;$$

otherwise identify indexes $m_1, \dots, m_\kappa \in [1:M_{\ell-1}]$ corresponding to κ nearest neighbors of $\tilde{X}_t^{\ell, m}$ among preliminary Voronoi centers on level ℓ , i.e. $0 = \|\tilde{X}_t^{\ell, m_1} - \tilde{X}_t^{\ell, m}\| \leq \dots \leq \|\tilde{X}_t^{\ell, m_\kappa} - \tilde{X}_t^{\ell, m}\|$, and set

$$X_t^{\ell, m} \triangleq \begin{cases} Y_t^{\ell, m} & \text{if } \sum_{i=1}^{\kappa} J_t^{\ell, m_i} = 0, \\ \tilde{X}_t^{\ell, m} & \text{else.} \end{cases}$$

(II-6) If $M_{\ell-1} < M_\ell$, for $t = 0, \dots, T$ and $m \in (M_{\ell-1}:M_\ell]$ set the additional Voronoi centers, i.e. define

$$X_t^{\ell, m} \triangleq Y_t^{\ell, m}.$$

(II-7) For each $t \in [0:T]$ collect all corresponding Voronoi centers in $\mathbf{X}_t^\ell \triangleq \{X_t^{\ell, m}\}_{m \in [1:M_\ell]}$.

\rightsquigarrow Policy Evaluation:

(II-8) Define $\hat{v}_\ell(T, x) \triangleq \Psi(x)$ for $x \in \mathbb{X}$ and set, recursively for $t = T-1, \dots, 0$ and $m \in [1:M_\ell]$, the node values as follows: Let $m_1 \triangleq \text{id}_{\mathbf{X}_t^{\ell-1}}(X_t^{\ell, m})$ be the index of the Voronoi center closest to $X_t^{\ell, m}$ on level $\ell-1$ and $m_2 \triangleq \hat{m}_t^{\ell-1, m_1}$ be the index of the corresponding best target Voronoi center starting from center m_1 . Set

$$\hat{v}_t^{\ell, m} \triangleq \psi(t, X_t^{\ell, m}) + \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} \hat{v}_\ell(t+1, \phi(t, X_t^{\ell, m}, Z_{t+1}^{\ell, m, n}))$$

if $\hat{v}_{\ell-1}(t, X_t^{\ell, m}) > \hat{\mathcal{M}}_t^{\ell-1}[\hat{v}_{\ell-1}](t, X_t^{\ell, m})$, or if $m_2 = 0$, or if there is no $a \in \Gamma^\circ(t, X_t^{\ell, m})$ such that $\xi(X_t^{\ell, m}, a) = X_t^{\ell-1, m_2}$; otherwise determine¹¹⁵ $\bar{a} \in \Gamma^\circ(t, X_t^{\ell, m})$ such that

¹¹⁵In case that there is more than one feasible action one should identify and consider the ‘‘cheapest’’ one.

$\xi(X_t^{\ell,m}, \bar{a}) = X_t^{\ell-1,m_2}$ and define

$$\hat{v}_t^{\ell,m} \triangleq \left[\psi(t, X_t^{\ell-1,m_2}) - c(t, X_t^{\ell,m}, \bar{a}) + \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} \hat{v}_\ell(t+1, \phi(t, X_t^{\ell-1,m_2}, Z_{t+1}^{\ell,m,n})) \right] \\ \vee \left[\psi(t, X_t^{\ell,m}) + \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} \hat{v}_\ell(t+1, \phi(t, X_t^{\ell,m}, Z_{t+1}^{\ell,m,n})) \right];$$

and collect the resulting values in $\mathbf{V}_t^\ell \triangleq \{\hat{v}_t^{\ell,m}\}_{m \in [1:M_\ell]}$.

(II-9) Set $\hat{v}_\ell(t, x) \triangleq \mathcal{I}_{\mathbf{X}_t^\ell, \mathbf{V}_t^\ell}^{\kappa, p}(x)$ for $(t, x) \in [0:T] \times \mathbb{X}$. ◇

Albeit the functioning of the preceding algorithm is in principle in line with the theoretical outline given in (3.22)-(3.24), some of the steps warrant a brief comment: In contrast to its theoretical counterpart and as necessarily required to be run on a computer, the algorithm estimates the relevant functions v_ℓ and $\mathcal{M}[v_\ell]$ only at *finitely* many (stochastic) interpolation nodes which are referred to as *Voronoi centers*; for function interpolation the operators introduced in Definition 3.3.4 are employed; see Steps (I-3),(II-2) and (II-9).

Initially, these Voronoi centers are determined using independent copies of the uncontrolled state process, the initial conditions of which are distributed according to γ ; see Step (I-1).¹¹⁶ For each subsequent level $\ell \in [1:L]$, the algorithm similarly derives preliminary centers based on pre-intervention states of (conditionally) independent copies of the state process controlled by the respective level's prevailing intervention strategy; see Step (II-4).¹¹⁷ Depending on ϕ and μ , this approach tends to move interpolation nodes into the continuation region; see also the discussion on page 119 below. However, to avoid their accumulation therein, Step (II-5) randomly discards preliminary centers that are located in the currently available approximation of the continuation region if all of their relevant (preliminary) neighbors are also located therein; those centers are then replaced with independent samples from the exploration distribution γ . The parameter $q \in [0, 1]$ thus controls the algorithm's propensity to stick to preliminary Voronoi centers.

Furthermore, note that for policy improvement only interventions that start from and lead to Voronoi centers belonging to the respective preceding level are taken into account; see Step (II-1). Accordingly, for a given Voronoi center $X_t^{\ell,m}$ on level ℓ , in case of a suggested intervention, i.e. if $\hat{v}_{\ell-1}(t, X_t^{\ell,m}) \leq \hat{\mathcal{M}}_t^{\ell-1}[\hat{v}_{\ell-1}](t, X_t^{\ell,m})$, policy evaluation first determines the nearest Voronoi center among those on the previous level ($\ell - 1$) and then assesses whether and how the subsequently determined corresponding target point can be reached by a feasible jump; see Step (II-8).

Finally, observe that expectations appearing in the evaluation of the continuation operator $\mathcal{C}[v_\ell]$, cf. (3.22) and (3.24), are estimated by means of non-nested Monte Carlo simulation; see Steps (I-2) and (II-8), where the latter makes use of common random numbers for both arguments of the computed maximum. ◇

¹¹⁶Note that the construction of interpolation nodes is conceptually similar to the notion of a *stochastic mesh*, as considered, e.g., in [BGH00] or [BG04a].

¹¹⁷With regard to the choice of initial preliminary Voronoi centers in Step (II-3) we also refer to our remarks on page 119 below.

3.4 Numerical Results

In this section we showcase the effectiveness of the proposed algorithm. As an example we consider some variants of the problem of *impulse control of Brownian motion* in discrete time and for a finite time horizon. This problem is classical given the literature, but usually addressed in continuous time and for an infinite time horizon; see, e.g., [CR78], [HST83] or §1.2/3 in [Die09].

3.4.1 Impulse Control of Brownian Motion towards Zero

With a slight abuse of notation, we consider the stochastic impulse control problem¹¹⁸

$$\mathbb{E} \left[\sum_{k=0}^{n-1} (c_\psi \|W_{t_k}^\alpha + \alpha_{t_k}\|_2^2 + (c_0 \cdot \mathbf{1}_{\alpha_{t_k} \neq 0} + \|C_1 \cdot \alpha_{t_k}\|_2)) + c_\Psi \|W_T^\alpha\|_2^2 \right] \xrightarrow{\alpha \in \mathcal{A}(0,x)} \min! \quad (\text{BM})$$

on an equidistant grid $\Pi \triangleq [0 = t_0, t_1, \dots, t_n = T]$ with step size $h \triangleq T/n$, i.e. $t_k \triangleq k \cdot T/n$ for $k \in [0:n]$, and $T > 0$; $c_\psi, c_\Psi \geq 0$ and $c_0 > 0$ are constants; $C_1 \in \mathbb{R}^{d \times d}$ is a matrix; $\mathcal{A}(0, x)$ is defined as above – with feasible actions chosen to be unrestricted; and W^α denotes a d -dimensional *controlled Brownian motion in discrete time* (see below).

Using the notation of Section 3.1, the stated problem (BM) corresponds to the following specification of the model ingredients:

$$\begin{aligned} (\mathbb{X}, \mathfrak{X}) &\triangleq (\mathbb{Z}, \mathfrak{Z}) \triangleq (\mathbb{A}, \mathfrak{A}) \triangleq (\mathbb{R}^d, \mathfrak{B}(\mathbb{R}^d)), & \phi(t, x, z) &\triangleq x + \sqrt{h}z, & \mu &\triangleq \mathcal{N}_d(0, \mathbb{I}_{d \times d}), \\ \psi(t, x) &\triangleq -c_\psi \|x\|_2^2, & c(t, x, a) &\triangleq c_0 \cdot \mathbf{1}_{\{a \neq 0\}} + \|C_1 \cdot a\|_2, & \Psi(x) &\triangleq -c_\Psi \|x\|_2^2, \\ \Gamma(t, x) &\triangleq \mathbb{R}^d, & \xi(x, a) &\triangleq x + a, & \boxtimes &\triangleq 0. \end{aligned} \quad (3.25)$$

Unlike before, the considered impulse control problem (BM) is now formulated as a *minimization* problem; hence we speak of (running and terminal) *costs* rather than of (negative) *rewards*.¹¹⁹

The state process W^α starting in $W_0^\alpha = x \in \mathbb{R}^d$ at time $t_0 = 0$ and given an admissible strategy $\alpha = \{\alpha_{t_k}\}_{k \in [0:n]} \in \mathcal{A}(0, x)$ is thus recursively defined by

$$W_{t_k}^\alpha = W_{t_{k-1}}^\alpha + \alpha_{t_{k-1}} + \sqrt{h} \cdot Z_k \quad \text{for } k \in [1:n], \quad (3.26)$$

where $Z_1, \dots, Z_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_d(0, \mathbb{I}_{d \times d})$.¹²⁰

In view of the relevant model assumptions note that Assumption 3.2.3 is automatically satisfied; a sufficient condition for Assumption 3.1.1 in the sense of Remark 3.1.2 could be obtained, e.g., by restricting the feasible actions to the set $\Gamma(t, x) = \overline{B_{\|x\|_2}(-x)}$.¹²¹ Nevertheless, our numerical results below are based on the specification of Γ as stated in (3.25).

All numerical computations are based on the implementation of our algorithm as outlined in

¹¹⁸As already pointed out in Footnote 88, the stated problem is considered on a time grid Π in place of $[0:T]$, but can easily be reformulated in terms of the setup of Section 3.1.

¹¹⁹Nevertheless, to be consistent with the preceding sections of this chapter, the signs of these costs are reversed in the specification of the relevant parameters; see (3.25).

¹²⁰Note that the discrete dynamics (3.26) can be regarded as the Euler-Maruyama discretization of a (controlled) Brownian motion.

¹²¹Note that Assumption 3.2.3 is also valid for this restrictive choice of feasible actions.

Section 3.3.2; our implementation is in PYTHON.¹²²

Our results are displayed in Sections 3.4.1.1 through 3.4.1.3 and 3.4.2; each section contains a table reporting all relevant parameters.

The plots displayed below show heat maps; for each problem instance we compute an approximation of its continuation (green) and intervention region (red). More precisely, in dimension 1 and 2 the plots consist of a rectangular grid: First, for each level $\ell \in [0:L]$ a cell with midpoint (t, x) is colored red if and only if $\hat{v}_\ell(t, x) \leq \hat{\mathcal{M}}_\ell^\ell[\hat{v}_\ell](t, x)$; otherwise the cell is colored green. Afterwards, we compute the equally weighted average of the resulting $(L + 1)$ RGB values per cell to obtain a heat map as an approximation of the regions. In higher dimensions, all but two coordinates within $(t, x) \in [0:T] \times \mathbb{R}^d$ are fixed to a priori specified values and the corresponding figures show cuts along the non-fixed coordinates; see Section 3.4.1.3 below.

For the sake of comparison, we also plotted the boundaries between the theoretical continuation and intervention region in dimensions 1 and 2; see the *white lines* in 1D or the **bold white lines** in 2D, respectively; these boundaries are obtained by classical Bellman backward induction and a numerical integration scheme, i.e. via solving (DP) numerically.¹²² The *white dots* (1D) or *thin white lines* (2D), respectively, in the interior of each connected component of the continuation region correspond to the target points to which actions taken at points in the intervention region lead; *black crosses* (\times) indicate target Voronoi centers of *our* algorithm based on the centers of the final level $\ell = L$.

Finally, we attached the running time of a single run of our algorithm to the respective figure.¹²³ Note that each problem instance is solved on the entire time grid Π .

3.4.1.1 A One-Dimensional Example without Proportional Costs

We consider the problem (BM) without proportional costs in dimension 1. Our first instance does not have terminal costs, the second one does. Their relevant common parameters are reported in Table 3.1; for individual parameters we refer to the caption of the respective figure; see Figures 3.4 and 3.5.

Parameter	d	T	n	c_ψ	c_0	C_1	L	M_ℓ	N_ℓ	κ	p	q	γ
Value	1	1	10	1	1	0	10	50	100	5	2	0.5	$\mathcal{U}([-5, 5])$

Table 3.1: Model and numerical parameters for Figures 3.4 and 3.5.

¹²² We use PYTHON 3.7.3 ([Pyt19, Oli07, MA11], ANACONDA package manager conda (v4.7.12, [Ana19])) with NumPy (v1.17.3, [Oli06, vCV11]) and IPython (v7.9.0, [PG07]); IDE: SPYDER (v3.3.6, [Spy19]). Plots are generated with matplotlib (v3.1.1, [Hun07]).

For some tasks we use packages from SciPy (v1.3.1, [VGO⁺20]): Nearest-neighbor search queries are performed using k -d trees via `spatial.cKDTree(.query)` (for details see <https://docs.scipy.org/doc/scipy-1.3.1/reference/generated/scipy.spatial.cKDTree.html> (accessed: 2020-05-31, 2251) and the references therein), and numerical integration for the classical Bellman equation (DP) in 1D and 2D makes use of the convolution procedures `signal.convolve` and `signal.convolve2d`, respectively.

Our implementation does not involve explicit parallelization; OS: Linux Mint 18.1 Serena.

Hardware: Intel Core i7-6700 CPU @ 3.40 GHz (4 cores) with 32 GiB RAM capacity.

¹²³ All stated running times are wall-clock times measured as a `timedelta` using `datetime` in PYTHON; they only include the time span our algorithm required to solve the respective problem instance; problem initialization, computation and plotting of heat maps as well as the numerical solution of (DP) (in 1D and 2D) are excluded.

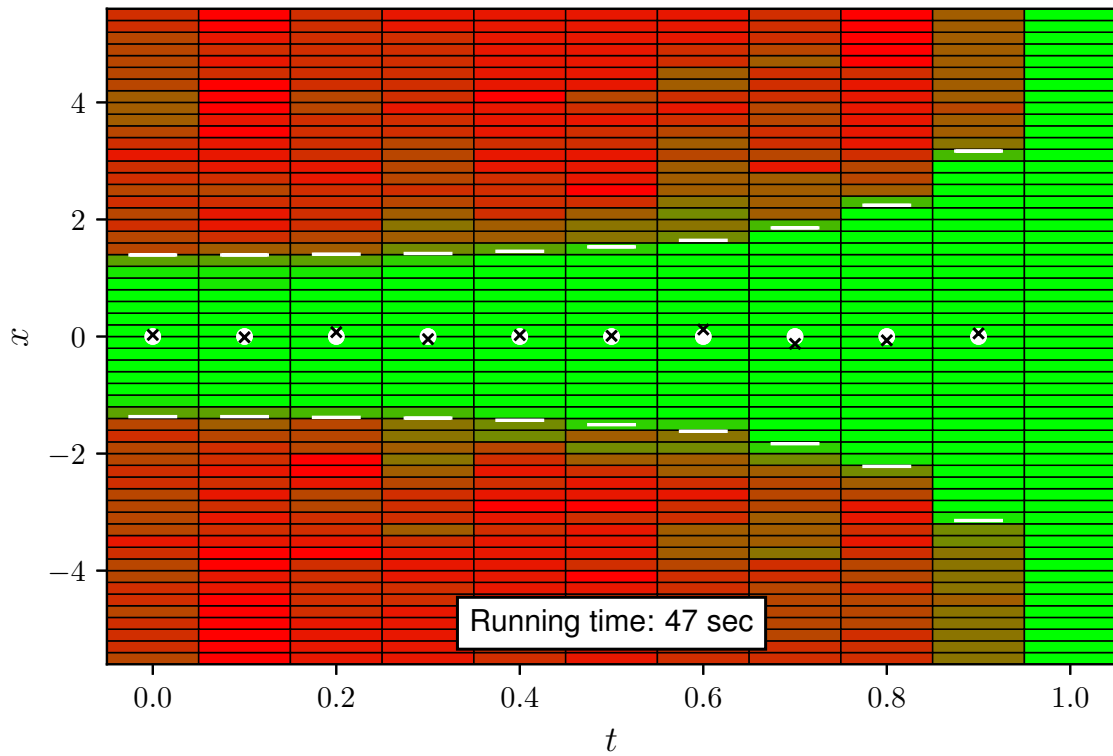


Figure 3.4: Controlled Brownian motion (BM) in 1D without proportional and without terminal costs; $C_1 = 0$, $c_\Psi = 0$.




Figure 3.5: Controlled Brownian motion (BM) in 1D without proportional, but with terminal costs; $C_1 = 0$, $c_\Psi = 0.2$.

3.4.1.2 A One-Dimensional Example with Fixed and Proportional Costs

Once again, we consider (BM) in dimension 1, but with both non-vanishing fixed and proportional costs; see Table 3.2 for the relevant common parameters, and Figures 3.6 and 3.7 for our results.

Parameter	d	T	n	c_ψ	c_Ψ	c_0	C_1	L	M_ℓ	N_ℓ	p	q	γ
Value	1	1	10	1	0	0.1	0.2	10	50	100	2	0.5	$\mathcal{U}([-5, 5])$

Table 3.2: Model and numerical parameters for Figures 3.6 and 3.7.

Note that in Figure 3.7 we also plotted all Voronoi centers on level $L = 10$ in grey.




Figure 3.6: Controlled Brownian motion (BM) in 1D with fixed and proportional costs; $\kappa = 5$.




Figure 3.7: Controlled Brownian motion (BM) in 1D with fixed and proportional costs; $\kappa = 1$.

3.4.1.3 Some Examples in Higher Dimensions

Problem (BM) can also be considered in higher dimensions and with an asymmetric specification of proportional costs; more precisely, we let the corresponding parameter matrix $C_1 \in \mathbb{R}^{d \times d}$ be given by the diagonal matrix $C_1 = \text{diag}(C_{11} \cdot \mathbf{I}_{\frac{d}{2} \times \frac{d}{2}}, C_{22} \cdot \mathbf{I}_{\frac{d}{2} \times \frac{d}{2}})$ where $C_{11}, C_{22} \geq 0$.

The relevant parameters can be found in Tables 3.3, 3.4 and 3.5; our results are displayed in Figures 3.8, 3.9 and 3.10.

A Two-Dimensional Example




Figure 3.8: Controlled Brownian motion (BM) in 2D with asymmetric costs.

Parameter	d	T	n	c_ψ	c_Ψ	c_0	C_{11}	C_{22}	L	M_ℓ	N_ℓ	κ	p	q	γ
Value	2	1	10	1	0	0.5	1	0.5	10	100	100	20	5	0.5	$\mathcal{U}([-10, 10]^2)$

Table 3.3: Model and numerical parameters for Figure 3.8.

A Four-Dimensional Example




Figure 3.9: Controlled Brownian motion (BM) in 4D with asymmetric costs (cuts along x_1 and x_4 , remaining spatial coordinates fixed to 0).

Parameter	d	T	n	c_ψ	c_Ψ	c_0	C_{11}	C_{22}	L	M_ℓ	N_ℓ	κ	p	q	γ
Value	4	1	10	1	0	0.5	1	0.5	10	200	100	20	5	0.2	$\mathcal{U}([-10, 10]^4)$

Table 3.4: Model and numerical parameters for Figure 3.9.

An Eight-Dimensional Example




Figure 3.10: Controlled Brownian motion (BM) in 8D with asymmetric costs (cuts along x_3 and x_6 , remaining spatial coordinates fixed to 0).

Parameter	d	T	n	c_ψ	c_Ψ	c_0	C_{11}	C_{22}	L	M_ℓ	N_ℓ	κ	p	q	γ
Value	8	1	10	1	0	0.5	1	0.5	10	600	100	20	5	0.2	$\mathcal{U}([-12.5, 12.5]^8)$

Table 3.5: Model and numerical parameters for Figure 3.10.

3.4.2 A Problem with Disconnected Continuation Region

We alter problem (BM) in dimension 1 to obtain an example whose continuation region exhibits two connected components for all times before T ; more specifically, we consider the following control problem:

$$\mathbb{E} \left[\sum_{k=0}^{n-1} (c_\psi \cdot \mathbf{1}_{[b_*, b^*]^c}(|W_{t_k}^\alpha|) + (c_0 \cdot \mathbf{1}_{\alpha_{t_k} \neq 0} + c_1 \cdot |\alpha_{t_k}|)) + c_\Psi \cdot \mathbf{1}_{[b_*, b^*]^c}(|W_T^\alpha|) \right]_{\alpha \in \mathcal{A}(0, x)} \longrightarrow \min! \quad (\text{BM}^b)$$

Mutatis mutandis, the explanations in Section 3.4.1 remain valid; note in particular that c_ψ , b_* and b^* are nonnegative constants and feasible actions are still taken to be unconstrained. The relevant parameters are reported in Table 3.6; our result is displayed in Figure 3.11.

Parameter	d	T	n	c_ψ	c_Ψ	c_0	c_1	b_*	b^*	L	M_ℓ	N_ℓ	κ	p	q	γ
Value	1	1	10	10	10	1	1	2.5	5	10	100	100	5	2	0.2	$\mathcal{U}([-10, 10])$

Table 3.6: Model and numerical parameters for Figure 3.11.




Figure 3.11: Controlled Brownian motion (BM^b) in 1D: Penalization when being outside of a certain region.

3.4.3 A Short Remark on the Numerical Results

First recall that our algorithm has no a priori information about the solution at all; only the relevant model ingredients as described in Section 3.1 are provided as input; therefore, it solely relies on knowledge that it acquires by itself during its iterations.

For all low-dimensional problem instances considered above, the algorithm is able to estimate the continuation and intervention region within a reasonable period of time. Note that in dimensions 1 and 2 the accuracy of the results is remarkably high compared to the benchmark solution of the classical Bellman equation (DP) obtained by numerical integration.

In higher dimensions, results are less accurate, albeit they point in the right direction: There is obviously room for improvement and future research; some related remarks and ideas will be sketched in the next and final section of this chapter.

Conclusion and Outlook

This chapter has derived a reformulation of the classical Bellman equation for discrete-time impulse control problems with finite time horizon in terms of a QVI in the spirit of Bensoussan [Ben08], including a corresponding verification theorem. Inspired by the resulting optimal impulse control, we have devised a non-nested Monte Carlo algorithm based on policy iteration which gradually learns the behavior of the optimally controlled state process and constructs a spatial discretization to explore the state space of its input problem instance by itself. Finally, we have showcased the effectiveness of our algorithm in several variants of a classical example problem.

There are many open questions for future research to address among which the following appear most prominent to us:

- It is of utmost interest to provide sufficient conditions for the convergence of the stated algorithm – both for the theoretical outline (3.22)-(3.24) in Section 3.3.1 and its implementation in Section 3.3.2. Furthermore, it would be interesting to investigate whether the stabilization via averaging of levels that we used for our numerical results is really necessary to obtain convergence, whether the sequence $\{v_\ell\}_{\ell \in \mathbb{N}_0}$ (or a corresponding sequence of weighted averages) exhibits monotonicity in some sense, and at what speed the relevant sequence converges. Additionally, for more precise numerical results, an explicit relation between the numerical parameters L , M_ℓ and N_ℓ and the obtained output accuracy would be desirable.
- Our algorithm is open to a combination with variance reduction techniques.¹²⁴ Aside from the obvious potential to use parallelization to distribute independent simulations across multiple cores, concepts such as control variates or importance sampling seem intuitively worth considering in order to induce faster convergence and hence to reduce the overall running time.

¹²⁴For details on these (and other) variance reduction principles we refer, e.g., to Chapter 4 in [Gla03] or §3.3/4 in [KKK10].

- As the numerical results indicate that our algorithm suffers from an error propagation that renders estimates at earlier points in time less accurate (see in particular the examples in dimensions beyond 1 in Figures 3.8, 3.9 and 3.10), further investigations should, e.g., attempt to replace the Bellman-style backward recursions in (3.22) and (3.24) by forward simulation. Note that (3.24) leads to a nested simulation that rapidly becomes computationally intractable. However, there might be a reformulation of the latter evaluation equation that allows for non-nested simulation and still yields a high accuracy of results; we think, e.g., of dropping the term “ $\mathcal{C}[v_\ell](t, x) \vee$ ” in the second case of (3.24) which we “only” used to derive the “fixed-point property” in Proposition 3.3.2.
- Our examples (BM) and (BM^b) can be considered as toy problems. Beyond the fact that they are symmetric, in particular they do not encounter any complications resulting from constrained action sets $\Gamma(t, x)$. However, as the latter prevail in realistic models – one may, e.g., think of admissibility and self-financing conditions that limit the set of feasible transactions in case of portfolio optimization with transaction costs – further investigations should try to broaden the applicability of our algorithm to such more involved application problems.
- The validity of the choice of the preliminary initial Voronoi centers in Step (II-3) should be taken with a grain of salt as the selected initial centers tend to be located in the continuation region. We are aware that this choice could be misleading in general – especially for problems with non-symmetric continuation and intervention region; see also the next item of this list. Hence, we think that the aforementioned step of our algorithm merits some further investigations: The choice of centers made in this step should be replaced by a rule that depends on a priori knowledge of the solution (if available) or exploits already gathered information from previous levels. Beyond the generic approach that draws independent samples from the exploration distribution γ at each level, i.e. $\tilde{X}_0^{\ell, m} \triangleq Y_0^{\ell, m}$, it is, e.g., conceivable that one uses the approximation of the intervention region from the respective preceding level to intensify exploration efforts therein.
- Step (II-5) of our algorithm can be regarded as a means of triggering new exploration efforts as soon as preliminary Voronoi centers do not promise any gain of information in subsequent iterations. Without this mechanism, exploration of the intervention region is usually insufficient as centers tend to accumulate within the continuation region – potentially leading to poor results; see, e.g., Figures 3.12 and 3.13. Although this accumulation effect is desirable to some extent since it stabilizes the algorithm by providing accurate interpolation results within the continuation region, exploiting mostly information from centers in the interior of the latter is likely to be misleading in later iterations as, e.g., more complex geometric shapes of its boundary could not be detected. Furthermore, note that in its current implementation our *thinning strategy* merely discards non-promising Voronoi centers and replaces them by samples from the exploration distribution γ . Additionally, it is also conceivable to transfer knowledge from those removed centers to close neighbors leading to some kind of *coalescence of centers*.
- Last but not least, it appears tempting to us to combine our algorithm with well-known

techniques from *machine learning* that successively *update* their respective value function approximations by newly acquired information. For instance, the concept of *Q-learning* (see, e.g., [WD92] or §8.2 in [Pow07]) could help to retain more information from previous levels than just using the preceding one as a basis of decision-making.¹²⁵

Note that the latter three items somewhat reflect a well-known issue in the context of machine/reinforcement learning, namely the problem of finding a balance between *exploration* and *exploitation*.¹²⁶ For our algorithm we leave a thorough analysis of related issues for future research; we expect it to yield substantial improvements of the results, especially for problems in higher dimensions.




Figure 3.12: Controlled Brownian motion (BM) in 1D with fixed and proportional costs; $\kappa = 1$; $q = 1$; the remaining parameters are chosen as in Table 3.2 (Voronoi centers on level $L = 10$ in grey): The effect of insufficient exploration.¹²⁷

¹²⁵For a comprehensive account of *approximate dynamic programming* we refer, e.g., to the monograph [Pow07].

¹²⁶see, e.g., Chapter 10 in [Pow07] for an introduction.

We also refer to the recent paper [CS19] that addresses *exploration vs. exploitation* issues for continuous-time impulse control problems in a diffusive setup with unknown drift.

¹²⁷compare Figures 3.7 and 3.12.

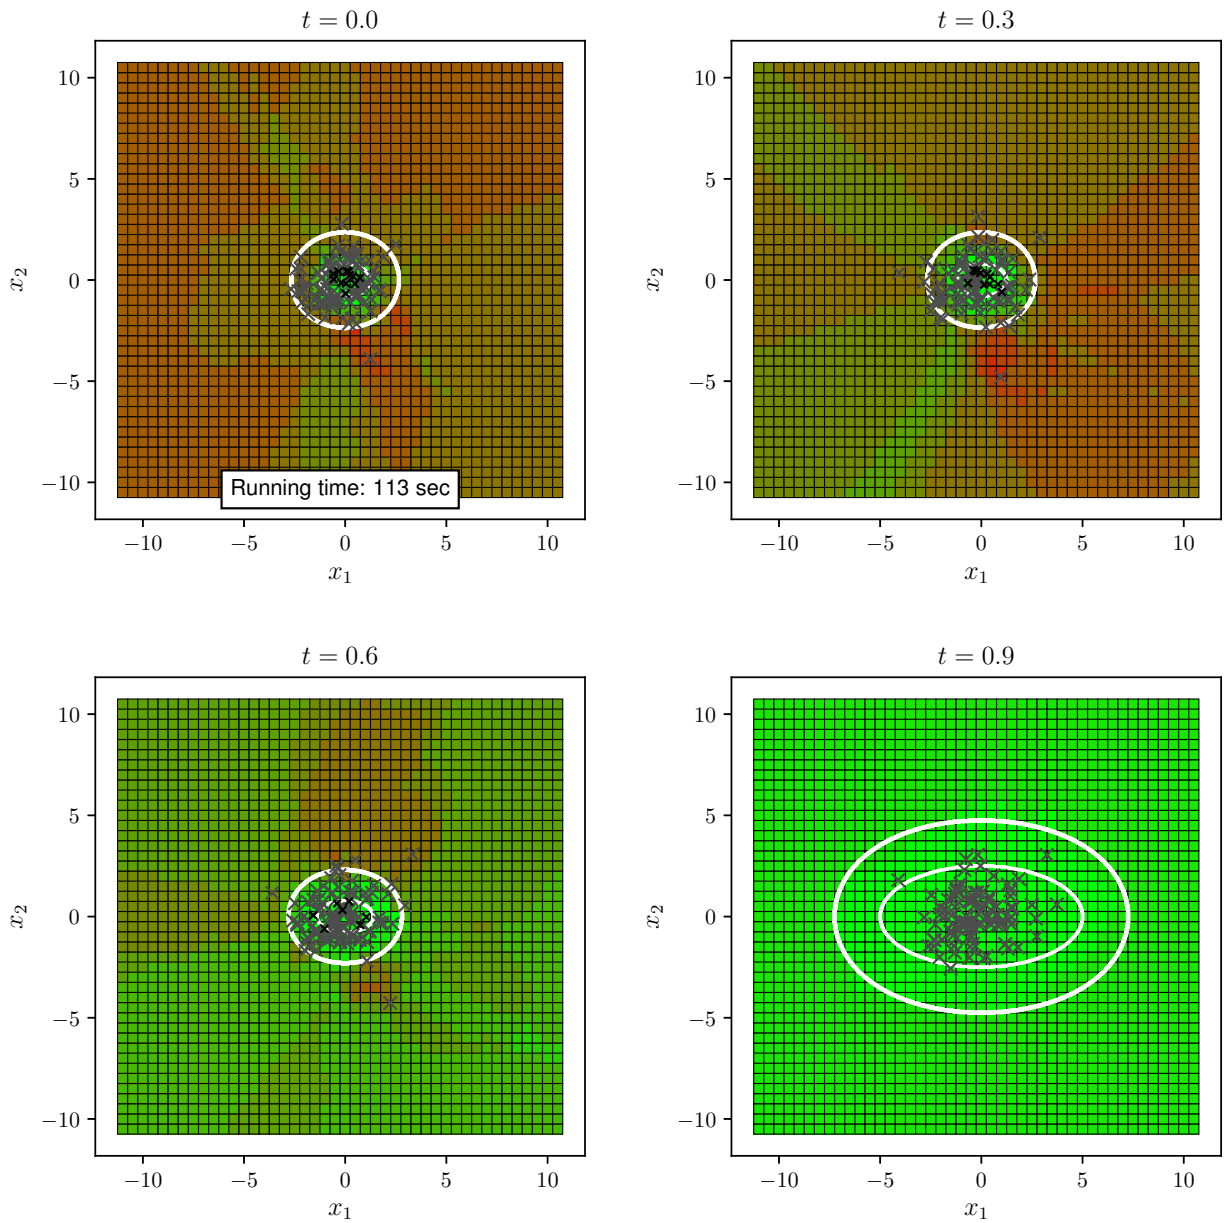


Figure 3.13: Controlled Brownian motion (BM) in 2D with asymmetric costs; $q = 1$; the remaining parameters are chosen as in Table 3.3 (Voronoi centers on level $L = 10$ in grey): The effect of insufficient exploration.¹²⁸

¹²⁸compare Figures 3.8 and 3.13.

Appendix A

Some Auxiliary Results

This part of the appendix collects some auxiliary results referred to in Chapters 1 and 2 (as well as Appendix D). Let $(\Omega, \mathfrak{A}, \{\mathfrak{F}_t\}_{t \in [0, T]}, \mathbb{P})$ be a filtered probability space such that the filtration \mathfrak{F} satisfies the usual conditions of \mathbb{P} -completeness and right-continuity. Furthermore, we fix some $T > 0$ and denote by Leb the one-dimensional Lebesgue measure on $([0, \infty), \mathfrak{B}([0, \infty)))$.

A.1 Modifications and Integrals

In Sections 1.4 and 2.1-2.3 we frequently switch between a càdlàg integrand process and its respective left-limit counterpart if the latter is a modification of the former and whenever it is more convenient or notationally less cumbersome. A rigorous justification for that is given by the following auxiliary result:

Lemma A.1 (Modification of Integrands). Let $\{X_t\}_{t \in [0, T]}$ and $\{Y_t\}_{t \in [0, T]}$ be \mathfrak{F} -progressively measurable stochastic processes taking values in some measurable space (S, \mathfrak{S}) . Assume that X and Y are modifications of each other, i.e.

$$X_t = Y_t \quad \mathbb{P}\text{-a.s. for all } t \in [0, T].$$

Then it holds

$$X_t(\omega) = Y_t(\omega) \quad \text{for } \text{Leb} \otimes \mathbb{P}\text{-a.e. } (t, \omega) \in [0, T] \times \Omega. \quad (\text{A.1})$$

Furthermore, for every bounded $\mathfrak{B}([0, T]) \otimes \mathfrak{S}$ -measurable function $f : [0, T] \times S \rightarrow \mathbb{R}$ we have

$$\int_0^t f(s, X_s) ds = \int_0^t f(s, Y_s) ds \quad \text{for all } t \in [0, T] \text{ } \mathbb{P}\text{-a.s.}$$

and

$$\int_0^t f(s, X_s) dW_s = \int_0^t f(s, Y_s) dW_s \quad \text{for all } t \in [0, T] \text{ } \mathbb{P}\text{-a.s.},$$

where $W = \{W_t\}_{t \in [0, T]}$ is a one-dimensional $(\mathfrak{F}, \mathbb{P})$ -Brownian motion.

Proof. Observe that Tonelli's theorem¹²⁹ yields

$$\begin{aligned}
& (\text{Leb} \otimes \mathbb{P})(\{(t, \omega) \in [0, T] \times \Omega : X_t(\omega) \neq Y_t(\omega)\}) \\
&= \int_{[0, T] \times \Omega} \mathbb{1}_{\{(t, \omega) \in [0, T] \times \Omega : X_t(\omega) \neq Y_t(\omega)\}}(t, \omega) (\text{Leb} \otimes \mathbb{P})(d(t, \omega)) \\
&= \int_0^T \underbrace{\mathbb{P}(X_t \neq Y_t)}_{=0} dt = 0; \tag{A.1}
\end{aligned}$$

this proves the first assertion.

Next, let $f : [0, T] \times \mathbb{S} \rightarrow \mathbb{R}$ be a bounded $\mathfrak{B}([0, T]) \otimes \mathfrak{G}$ -measurable function and note that for all $t \in [0, T]$ and $A \in \mathfrak{A}$ it holds

$$\begin{aligned}
\mathbb{E} \left[\mathbb{1}_A \cdot \int_0^t f(s, X_s) ds \right] &= \int_{[0, T] \times \Omega} \mathbb{1}_A(\omega) \cdot f(s, X_s(\omega)) (\text{Leb} \otimes \mathbb{P})(d(s, \omega)) \\
&= \int_{[0, T] \times \Omega} \mathbb{1}_A(\omega) \cdot f(s, Y_s(\omega)) (\text{Leb} \otimes \mathbb{P})(d(s, \omega)) \\
&= \mathbb{E} \left[\mathbb{1}_A \cdot \int_0^t f(s, Y_s) ds \right].
\end{aligned}$$

Thus, upon considering the cases $A \in \{A_1, A_2\}$ separately, where

$$A_1 = \left\{ \int_0^t f(s, X_s) ds < \int_0^t f(s, Y_s) ds \right\} \quad \text{and} \quad A_2 = \left\{ \int_0^t f(s, X_s) ds > \int_0^t f(s, Y_s) ds \right\},$$

we have that $\int_0^t f(s, X_s) ds = \int_0^t f(s, Y_s) ds$ \mathbb{P} -a.s. for all $t \in [0, T]$. As due to dominated convergence both sides of this identity define continuous processes, their indistinguishability follows. Since $\text{Leb} \otimes \mathbb{P}$ is the Doléans measure of W the final assertion follows immediately from (A.1) and the construction of Brownian stochastic integrals; see, e.g., §3.2 in [KS98] or Exkurs 2 in [KK01]. \square

A.2 Counting Processes

Unless explicitly stated otherwise, whenever *counting processes* are involved in Chapters 1 and 2, we resort to the following definition:

Definition A.2 (Counting Process, Intensity). (a) A càdlàg and \mathfrak{F} -adapted stochastic process $\{N_t\}_{t \in [0, T]}$ is called a *counting process* if it is \mathbb{P} -a.s. \mathbb{N}_0 -valued, nondecreasing and satisfies $N_0 = 0$.

(b) If $\lambda = \{\lambda_t\}_{t \in [0, T]}$ is a nonnegative, \mathfrak{F} -progressively measurable process such that

$$\bar{N}_t \triangleq N_t - \int_0^t \lambda_s ds, \quad t \in [0, T]$$

is a local $(\mathfrak{F}, \mathbb{P})$ -martingale, then λ is called an $(\mathfrak{F}, \mathbb{P})$ -intensity of N . In that regard, \bar{N} is also referred to as a *compensated counting process*. \diamond

¹²⁹In fact, this argument can be regarded as an application of *Cavalieri's principle*; see, e.g., §V.1.3 in [Els09].

The preceding definition is a ramification of the one of a *non-explosive counting process* as stated in Definition II.3.D7 in [Bré81]; see also [Bré81, p.18/19] and Theorems II.3.T8/9 in [Bré81].

It is a natural question to ask whether an integral with respect to a compensated counting process is an honest martingale. The following result turns out to be sufficient for our purposes:

Lemma A.3. Let N be an $(\mathfrak{F}, \mathbb{P})$ -counting process with intensity $\lambda = \{\lambda_t\}_{t \in [0, T]}$. If $H = \{H_t\}_{t \in [0, T]}$ is \mathfrak{F} -predictable and with $\mathbb{E}[\int_0^T |H_s| \lambda_s ds] < +\infty$, then the process $M = \{M_t\}_{t \in [0, T]}$,

$$M_t \triangleq \int_{(0, t]} H_s d\bar{N}_s \quad \text{for } t \in [0, T]$$

is an $(\mathfrak{F}, \mathbb{P})$ -martingale.

Proof. See Theorem II.3.T8(β) in [Bré81].

However, in this thesis we require only the case where λ and H are bounded where we can argue as follows: Since $\bar{N} = N - \int_0^\cdot \lambda_s ds$ is a local $(\mathfrak{F}, \mathbb{P})$ -martingale and H is (locally) bounded and \mathfrak{F} -predictable, Theorem IV.29 in [Pro04] implies that both $M = H \bullet \bar{N}$ and $H^2 \bullet \bar{N}$ are local $(\mathfrak{F}, \mathbb{P})$ -martingales. Since N is a counting process, it is nondecreasing and moreover we have $[\bar{N}] = [N] = N$; in particular, $[\bar{N}] - \int_0^\cdot \lambda_s ds = N - \int_0^\cdot \lambda_s ds$ is a local $(\mathfrak{F}, \mathbb{P})$ -martingale and thus by Proposition 3.64 and Corollaries 3.53 and 3.65 in [Med07] the predictable quadratic variation of N is given by $\langle \bar{N} \rangle = \int_0^\cdot \lambda_s ds$. As a consequence, with the help of Theorem 3.52 in [Med07], we obtain

$$\mathbb{E}[[M]_T] = \mathbb{E}\left[\int_{(0, T]} H_s^2 d[\bar{N}]_s\right] = \mathbb{E}\left[\int_{(0, T]} H_s^2 d\langle \bar{N} \rangle_s\right] = \mathbb{E}\left[\int_0^T H_s^2 \lambda_s ds\right] < \infty.$$

Thus, M is an L^2 -martingale by Corollary 3 in [Pro04, p.73]. \square

Finally, in the setup of Sections 1.4 and 2.1-2.3 we make use of the following result that renders Lemma A.1 applicable:

Lemma A.4. Let $\{N_t\}_{t \in [0, T]}$ be a counting process with progressively measurable intensity process $\{\lambda_t\}_{t \in [0, T]}$ such that $\mathbb{E}[\int_0^T \lambda_s ds] < +\infty$. Then it holds $\Delta N_t = 0$ \mathbb{P} -a.s. for all $t \in [0, T]$.

In particular, the preceding result applies if N has a *bounded* intensity process λ . \diamond

Proof of Lemma A.4. Recall that $N_{0-} = N_0 = 0$ and let $t \in (0, T]$ in the following. With the help of monotone and dominated convergence and Lemma A.3 we obtain:

$$\begin{aligned} \mathbb{P}(\Delta N_t \neq 0) &= \mathbb{E}[\Delta N_t] = \mathbb{E}[N_t - N_{t-}] = \lim_{s \uparrow t} \mathbb{E}[N_t - N_s] \\ &= \lim_{s \uparrow t} \mathbb{E}\left[\int_s^t \lambda_u du\right] = \mathbb{E}\left[\lim_{s \uparrow t} \int_s^t \lambda_u du\right] = 0. \end{aligned}$$

This proves the stated assertion. \square

A.3 Carathéodory Solutions of ODEs

As already pointed out in Footnote 60, all ODEs in Chapter 2 (and Appendix D) are taken in the sense of Carathéodory. Therefore, as a point of reference and for the sake of completeness, we provide a brief outline of the relevant definitions and results. For details we refer, e.g., to §I.5 in [Hal80] or Supplement II of III.§10 in [Wal98]; note that this section is mainly based on and collects the results stated in the first of these two references.

For an open set $D \subseteq \mathbb{R}^{d+1}$, some initial configuration $(t_0, x_0) \in D$ and a Borel measurable function $f : D \rightarrow \mathbb{R}^d$ we consider the initial value problem

$$\dot{x}(t) = f(t, x(t)) \quad \text{for } t \in I, \quad x(t_0) = x_0, \quad (\text{IVP})$$

where $I \subseteq \mathbb{R}$ denotes some real interval with $t_0 \in I$. A *Carathéodory solution* of (IVP) is defined as follows (see [Hal80, p.28]):

Definition A.5 (Carathéodory Conditions, Carathéodory Solution). Let $D \subseteq \mathbb{R}^{d+1}$ be an open set and $f : D \rightarrow \mathbb{R}^d$ be a Borel measurable function.

(a) The function f is said to satisfy the *Carathéodory conditions* on D if the following three conditions are met:

(i) $f(\cdot, x)$ is Borel measurable for each fixed x .

(ii) $f(t, \cdot)$ is continuous for each fixed t .

(iii) For each compact set $U \subset D$ there is an integrable function m_U such that

$$\|f(t, x)\| \leq m_U(t) \quad \text{for all } (t, x) \in U.$$

(b) An absolutely continuous function $x : I \rightarrow \mathbb{R}^d$, defined on some real interval $I \subseteq \mathbb{R}$ with $t_0 \in I$ satisfying $(t, x(t)) \in D$ for every $t \in I$, is said to be a *Carathéodory solution* of (IVP) if

$$x(t) = x_0 + \int_{t_0}^t f(s, x(s)) ds \quad \text{for all } t \in I. \quad (\text{IVP})$$

holds. ◇

Note that the concept of a *Carathéodory solution* considers the differential equation (IVP) solely as an *integral equation*.¹³⁰ Existence and uniqueness of Carathéodory solutions is established by the following result (see Theorem I.5.3 in [Hal80]):

Theorem A.6 (Existence and Uniqueness). Let $D \subseteq \mathbb{R}^{d+1}$ be open set and suppose that $f : D \rightarrow \mathbb{R}^d$ is a Borel measurable function that satisfies the Carathéodory conditions on D . Moreover assume that for every compact set $U \subset D$ there is an integrable function k_U such that the Lipschitz condition

$$\|f(t, x_1) - f(t, x_2)\| \leq k_U(t) \cdot \|x_1 - x_2\| \quad \text{for all } (t, x_1), (t, x_2) \in U \quad (\text{A.2})$$

¹³⁰For differentiability (Leb-a.e.) of absolutely continuous functions and a corresponding version of the *fundamental theorem of calculus* we refer, e.g., to Theorem 7.20 in [Rud87] or Theorem VII.4.14 in [Els09].

is satisfied. Then, for any $(t_0, x_0) \in U$, there is a unique Carathéodory solution of (IVP) passing through (t_0, x_0) .¹³¹ \diamond

Note that if one omits the stated Lipschitz condition (A.2) in the preceding result, one can still derive existence of a Carathéodory solution of (IVP); see Theorem I.5.1 in [Hal80]. However, uniqueness might fail to hold. \diamond

Furthermore, Carathéodory solutions can be extended up to the boundary of the domain D of the right-hand side f (see Theorem I.5.2 in [Hal80]):

Theorem A.7 (Continuation/Maximal Interval of Existence). Let $D \subseteq \mathbb{R}^{d+1}$ be open set and suppose that $f : D \rightarrow \mathbb{R}^d$ is a Borel measurable function that satisfies the Carathéodory conditions on D . Furthermore, let x denote a Carathéodory solution of (IVP) on some interval $I \subseteq \mathbb{R}$. Then there is a *continuation* of x to a *maximal interval of existence* $(\alpha, \beta) \subseteq \mathbb{R}$; this extension tends to the boundaries of D as soon as $t \rightarrow \alpha$ or $t \rightarrow \beta$. \diamond

For our purposes the following remark will be of interest:

Remark A.8 (Global Existence and Uniqueness). Let $\emptyset \neq I \subset \mathbb{R}$ be a compact interval.

- (a) Suppose that $f : I \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a Borel measurable function that satisfies the assumptions of Theorem A.6, i.e. f fulfills the Carathéodory conditions and a Lipschitz condition (A.2). Moreover, assume that f is *linearly bounded*, i.e. for some constants $a, b \geq 0$ it holds

$$\|f(t, x)\| \leq a + b \cdot \|x\| \quad \text{for all } (t, x) \in I \times \mathbb{R}^d. \quad (\text{A.3})$$

Clearly, by Theorem A.6 in conjunction with Theorem A.7 there is a unique Carathéodory solution x of (IVP) on some maximal interval of existence $(\alpha, \beta) \subseteq \overset{\circ}{I}$. Moreover, by (A.3) it holds

$$\|x(t)\| \leq \|x_0\| + \int_{t_0}^t \|f(s, x(s))\| ds \leq \|x_0\| + a(t - t_0) + b \cdot \int_{t_0}^t \|x(s)\| ds \quad \text{for } t \in (\alpha, \beta)$$

and thus Gronwall's inequality¹³² implies

$$\|x(t)\| \leq (\|x_0\| + a(t - t_0)) \cdot e^{b(t-t_0)} \quad \text{for } t \in (\alpha, \beta).$$

In particular, it holds $\sup_{t \in (\alpha, \beta)} \|x(t)\| < +\infty$ and, as a result, we obtain from Theorem A.7 that $(\alpha, \beta) = \overset{\circ}{I}$ which, in turn, implies global existence and uniqueness of x on I .

- (b) Let $E \subseteq \mathbb{R}^d$ be a compact set and define $D \triangleq I \times E$. Moreover, suppose that $f : D \rightarrow \mathbb{R}^d$ is a Borel measurable function which satisfies the assumptions of Theorem A.6 with a global Lipschitz condition (A.2), i.e. the functions k_U reduce to a positive constant (uniform across t and U). Then, Theorem A.6 can be proved in the same way as the celebrated Picard-Lindelöf theorem for classical solutions of ODEs; see, e.g., Theorem II.6.I in [Wal98] for details on the proof of this standard result. In particular, one obtains global existence and uniqueness of a solution of (IVP) in the sense of Carathéodory on I . \diamond

¹³¹Moreover, continuous dependence on the initial data $(t_0, x_0) \in D$ can be shown.

¹³²see, e.g., Lemma I.6.2 and p.37 in [Hal80]

Appendix B

Branching Diffusions and Nonlocal Dirichlet Problems

In Section 1.1 we have derived a branching diffusion representation of classical solutions of nonlocal PDEs of type (PDE) subject to a terminal condition (TC); see Theorem 1.1.6. Inspired by [AC18], it is also possible to tackle a certain class of nonlocal Dirichlet problems with analytic nonlinearities on *bounded* spatial domains with the help of our branching methodology. In this part of the appendix we provide a brief outline of a corresponding representation result.

Let $D \subsetneq \mathbb{R}^d$ be a non-empty, open and bounded domain and $\mu : \bar{D} \rightarrow \mathbb{R}^d$ and $\sigma : \bar{D} \rightarrow \mathbb{R}^{d \times d}$ be Lipschitz continuous functions. Keeping notation, mutatis mutandis, as in Section 1.1, we let \mathcal{A} be the infinitesimal generator of the time-homogeneous Itô diffusion with dynamics

$$d\bar{X}_t^x = \mu(\bar{X}_t^x)dt + \sigma(\bar{X}_t^x)d\bar{W}_t, \quad \bar{X}_0^x = x, \quad \text{for } x \in D,$$

and consider the Dirichlet problem

$$\begin{aligned} \mathcal{A}[u](x) + \int_{\Xi} f(x, \xi, \mathcal{J}[u](x, \xi))\gamma(d\xi) &= 0 & \text{for } x \in D, \\ u(x) &= g(x) & \text{for } x \in \partial D \cup D_1; \end{aligned} \tag{D}$$

the jump maps $\Gamma_\ell : \bar{D} \times \Xi \rightarrow \mathbb{R}^d$, $\ell \in [1:m]$, and the coefficients $c_i : \bar{D} \times \Xi \rightarrow \mathbb{R}$, $i \in \mathcal{I}$, are assumed to be jointly measurable; c_i , $i \in \mathcal{I}$, is furthermore assumed to be bounded and continuous with respect to $x \in \bar{D}$; the *interaction domain* is defined by $D_1 \triangleq \bigcup_{\ell=1}^m (\Gamma_\ell(D \times \Xi) \setminus D)$,¹³³ and the function $g : \partial D \cup D_1 \rightarrow \mathbb{R}$ is assumed to be bounded and continuous.

For a point $x \in D$ and a Brownian motion \bar{W} we define the first exit time of \bar{X}^x from D by

$$\tau_D^{x, \bar{W}} \triangleq \inf\{t \geq 0 : \bar{X}_t^x \notin D\},$$

¹³³Note that D_1 comprises all target points $y = \Gamma_\ell(x, \xi) \notin D$ that can be reached by a jump from a point $x \in D$ (and given a jump mark in $\xi \in \Xi$) via some jump map Γ_ℓ , $\ell \in [1:m]$. In other words, the point $y \notin D$ *interacts* with some point $x \in D$, and thus the set of all such outside points is commonly referred to as *interaction domain* in the literature on nonlocal models and operators; for details we refer, e.g., to [DGLZ12], [DGLZ13] or [SV19] and the references therein.

and assume that $\mathbb{E}[\tau_D^{x, \bar{W}}] < +\infty$.¹³⁴ By continuity of \bar{X}^x we have

$$\bar{X}_t^x \in D \quad \text{for } t \in [0, \tau_D^{x, \bar{W}}) \quad \text{and} \quad \bar{X}_{\tau_D^{x, \bar{W}}}^x \in \partial D.$$

Let $u \in \mathcal{C}^2(D) \cap \mathcal{C}^0(\bar{D})$ be a classical solution of (D) and fix some $x \in D$. From Itô's lemma¹³⁵ we obtain the Feynman-Kač representation

$$\begin{aligned} u(x) &= \mathbb{E} \left[g(\bar{X}_{\tau_D^{x, \bar{W}}}^x) + \int_0^{\tau_D^{x, \bar{W}}} \int_{\Xi} f(\bar{X}_s^x, \xi, \mathcal{J}[u](\bar{X}_s^x, \xi)) \gamma(d\xi) ds \right] \\ &= \mathbb{E} \left[g(\bar{X}_{\tau_D^{x, \bar{W}}}^x) + \int_0^{\tau_D^{x, \bar{W}}} \int_{\Xi} \sum_{i \in \mathcal{I}} c_i(\bar{X}_s^x, \xi) \mathcal{J}[u](\bar{X}_s^x, \xi)^i \gamma(d\xi) ds \right]. \end{aligned} \quad (\text{B.1})$$

Accordingly, under suitable conditions,¹³⁶ with the same ingredients and by proceeding similarly as in Section 1.1, we thus introduce particles $k \in \mathbf{N}$ with branching times $T^{(k)}$ recursively given by $T^\emptyset \triangleq 0$ and

$$T^{(k)} \triangleq \begin{cases} T^{(k-)} + \tau^{(k)} & \text{if } \tau^{(k)} < \tau_D^{(k)}, \\ T^{(k-)} + \tau_D^{(k)} & \text{else,} \end{cases} \quad \text{where } \tau_D^{(k)} \triangleq \tau_D^{X_{T^{(k-)}}^{(k)}, W^{(k)}}.$$

A particle k only comes into existence if its respective initial condition is feasible: Its dynamics start in either $X_0^{(1)} \triangleq x \in D$ if $k = (1)$, or $X_{T^{(k-)}}^{(k)} \triangleq \Gamma_{J^{(k)}}(X_{T^{(k-)}}^{(k-)}, \Delta^{(k-)})$ whenever $k \neq (1)$ and $X_{T^{(k-)}}^{(k-)}, X_{T^{(k-)}}^{(k)} \in D$, respectively, and – in these cases – obey

$$dX_t^{(k)} = \mu(X_t^{(k)})dt + \sigma(X_t^{(k)})dW_t^{(k)} \quad \text{for } t \in [T^{(k-)}, T^{(k)}];$$

a particle k branches into $|I^{(k)}|$ shadow offspring particles with jump marks according to $I^{(k)}$ (see Section 1.1.2) whenever $X_{T^{(k)}}^{(k)} \in D$; otherwise, it ceases to exist without branching; shadow offspring particles with infeasible initial conditions are immediately discarded (see above), and only the remaining ones turn into particles.¹³⁷

We define the set of all particles to be $\bar{\mathcal{K}}_x$ and introduce the set $\mathcal{K}_x^\partial \triangleq \{k \in \bar{\mathcal{K}}_x : X_{T^{(k)}}^{(k)} \in \partial D\}$ of particles hitting the boundary of D before branching; for each particle k in the complement, i.e. if $k \in \bar{\mathcal{K}}_x \setminus \mathcal{K}_x^\partial$, we define the set of infeasible jump marks by $\mathcal{L}_x^{\text{out}}(k) \triangleq \{\ell \in [1 : m] : \Gamma_\ell(X_{T^{(k)}}^{(k)}, \Delta^{(k)}) \in D_1\}$.

Then, under analogous (integrability) assumptions,¹³⁸ we can mimic the proof of Theorem 1.1.6 above, see in particular (1.14), and rewrite (B.1) to obtain the representation

¹³⁴Note that $\tau_D^{x, \bar{W}}$ is an $\mathfrak{F}^{\bar{X}^x}$ -stopping time since \bar{X}^x is continuous and $\mathbb{R}^d \setminus D$ is closed.

¹³⁵For technical details we refer, e.g., to Theorem II.2.1 in [Fre85].

¹³⁶see also, e.g., Assumptions 2.1-2.3 in [AC18].

¹³⁷Note that, unlike in Section 1.1, the branching mechanism itself is now affected by the branching diffusion dynamics. We refer to §2.2 in [AC18] for a similar construction of a *marked branching diffusion with absorption* (albeit without jumps); see also, e.g., §3.1.1 in [BCL⁺15].

¹³⁸For further details see also the related discussions in [AC18].

$$\begin{aligned}
u(x) = \mathbb{E} & \left[\prod_{k \in \mathcal{K}_x^\partial} \frac{g(X_{T^{(k)}}^{(k)})}{F(T^{(k)} - T^{(k-)})} \right. \\
& \times \left. \prod_{k \in \bar{\mathcal{K}}_x \setminus \mathcal{K}_x^\partial} \frac{c_{I^{(k)}}(X_{T^{(k)}}^{(k)}, \Delta^{(k)})}{\rho(T^{(k)} - T^{(k-)}) p_{I^{(k)}}} \prod_{\ell \in \mathcal{L}_x^{\text{out}}(k)} g(\Gamma_\ell(X_{T^{(k)}}^{(k)}, \Delta^{(k)}))^{I_\ell^{(k)}} \right]. \quad (\text{B.2})
\end{aligned}$$

Note that this representation is in line with and partially extends existing results on local Dirichlet problems to the nonlocal case; see, e.g., Propositions 3.1 and 4.1 in [AC18]. \diamond

To conclude this part of the appendix, let us observe that the numerical efficiency of Monte Carlo simulation of (B.2) heavily relies on the ability to decide *computationally fast* whether a given point $x \in \mathbb{R}^d$ satisfies $x \in D$ or not; in other words, one needs to be able to compute the indicator $\mathbb{1}_D$ of the underlying domain D efficiently. This is, e.g., the case for special domains such as hyperrectangles or balls.

Note also that such a simulation has to be formally justified by deriving a viscosity solution property of (B.2) similarly as in Section 1.2; we leave this issue for future research and refer, e.g., to Theorems 3.1 and 4.1 in [AC18] for corresponding results in the case of local Dirichlet problems.

Appendix C

Counterparties and Pricing PDEs

In this part of the appendix, we consider the setup of Section 1.4. For the sake of completeness and comparison we state the relevant pricing PDEs for the considered benchmark scenarios in Section 1.4.3 in synopsis: Let $\lambda^\dagger \triangleq \lambda$ denote the default intensity function of the SIB counterparty. Supplemented with the terminal condition $u(T, x) = g_\circ(x)$, $x \in [0, \infty)^d$, where $\circ \in \{\text{short}, \text{long}\}$, a fair price of the derivative $g_\circ(X_T)$ is obtained by solving one of the following equations on $[0, T] \times [0, \infty)^d$:

(i) in case of a defaultable SIB counterparty:

$$\begin{aligned} \partial_t u(t, x) + \mathcal{A}^*[u](t, x) - r(t, x)u(t, x) \\ + \lambda^\dagger(t, x) \int_E \left(h \circ u(t, (\mathbf{I}_{d \times d} + \text{diag}(z))x) - u(t, x) \right) \gamma(dz) = 0. \end{aligned} \quad (\text{C.1})$$

(ii) in case of a defaultable non-SIB counterparty with distinct¹³⁹ default intensity function $\lambda^\# : [0, T] \times [0, \infty)^d \rightarrow [0, \infty)$:

$$\partial_t u(t, x) + \mathcal{A}^\dagger[u](t, x) - r(t, x)u(t, x) + \lambda^\#(t, x)(h \circ u(t, x) - u(t, x)) = 0. \quad (\text{C.1})$$

(iii) in case of a default-free counterparty:

$$\partial_t u(t, x) + \mathcal{A}^\dagger[u](t, x) - r(t, x)u(t, x) = 0. \quad (\text{C.2})$$

Here, \mathcal{A}^\dagger denotes the infinitesimal generator of the underlying risky asset prices X (see (1.27)), i.e.

$$\begin{aligned} \mathcal{A}^\dagger[u](t, x) &\triangleq \mu(t, x)^\top \text{diag}(x) \nabla_x u(t, x) + \frac{1}{2} \text{tr} \left[\text{diag}(x) \sigma(t, x) \sigma(t, x)^\top \text{diag}(x) \nabla_x^2 u(t, x) \right] \\ &\quad + \lambda^\dagger(t, x) \int_E \left(u(t, (\mathbf{I}_{d \times d} + \text{diag}(z))x) - u(t, x) \right) \gamma(dz) \\ &= \mathcal{A}^*[u](t, x) + \lambda^\dagger(t, x) \int_E \left(u(t, (\mathbf{I}_{d \times d} + \text{diag}(z))x) - u(t, x) \right) \gamma(dz). \end{aligned}$$

¹³⁹Note that, for reasons of comparability, we choose $\lambda^\dagger \equiv \lambda^\# \equiv \lambda$ in all our considerations; see Section 1.4.3.

The link between each of these equations and a suitable risk-neutral pricing formula is established via a Feynman-Kač representation of a corresponding classical solution.¹⁴⁰ The proof of Theorem 1.4.1 considers case (i) of a defaultable SIB counterparty and links PDE (1.28) to formula (1.25); its arguments are readily adapted to cover the other two cases.

Note that while in the case of a default-free counterparty pricing boils down to solving a (standard, albeit nonlocal) linear PDE¹⁴¹ (see (C.2)), the corresponding equation becomes nonlinear as soon as the counterparty is defaultable and credit risk needs to be taken into account (see (C.1)); if the contract's counterparty is furthermore modeled as systemically important the additional complication of a *nonlocal* nonlinearity emerges (see (1.28)).

With regard to branching and our simulation results in Section 1.4.3, note that (1.28) is solved as outlined in Section 1.4.2; analogously, the nonlinearity of (C.1) is approximated by a polynomial as discussed in §5 in [Hen12a]; and the one of the linear PDE (C.2) is of polynomial type anyway. Note that the branching estimators of the latter two equations lead to branching *jump*-diffusions – without any conceptual difference with regard to the branching methodology.

¹⁴⁰This approach is standard given the literature; see, e.g., §4.8 in [Eth02], Chapter 6 in [Shr04] or [BJR05].

¹⁴¹In fact, in the specific setup of Section 1.4.3 this PDE can be solved explicitly; see, e.g., [Kou02].

Appendix D

Existence of Mean Field Equilibria

In this part of the appendix we provide a proof of Theorem 2.3.6 together with relevant assumptions and auxiliary results. Throughout, we consider the setup of Chapter 2 and fix some time horizon $T > 0$ and a number of common noise events $n \in \mathbb{N}_0$. We set $\Pi \triangleq \{\pi = [T_0, T_1, \dots, T_n, T_{n+1}] : 0 = T_0 < T_1 < \dots < T_n < T_{n+1} = T\}$.

Assumption D.1 (Lipschitz Conditions).

- (i) The terminal reward function Ψ is Lipschitz continuous with respect to m , i.e.

$$\|\Psi(w, m_1) - \Psi(w, m_2)\| \leq L_\Psi \cdot \|m_1 - m_2\| \quad (\text{L}_\Psi)$$

for all $w \in \mathbb{W}^n$ and $m_1, m_2 \in \mathbb{M}$.

- (ii) The reduced-form running reward function $\hat{\psi}$ is jointly Lipschitz continuous with respect to m and v , i.e.

$$\|\hat{\psi}(t, w, m_1, v_1) - \hat{\psi}(t, w, m_2, v_2)\| \leq L_{\hat{\psi}} \cdot (\|m_1 - m_2\| + \|v_1 - v_2\|) \quad (\text{L}_{\hat{\psi}})$$

for all $t \in [0, T]$, $w \in \mathbb{W}^n$, $m_1, m_2 \in \mathbb{M}$ and $v_1, v_2 \in \mathbb{R}^d$.

- (iii) The reduced-form intensity matrix function \hat{Q} is jointly Lipschitz continuous with respect to m and v , i.e.

$$\|\hat{Q}(t, w, m_1, v_1) - \hat{Q}(t, w, m_2, v_2)\| \leq L_{\hat{Q}} \cdot (\|m_1 - m_2\| + \|v_1 - v_2\|) \quad (\text{L}_{\hat{Q}})$$

for all $t \in [0, T]$, $w \in \mathbb{W}^n$, $m_1, m_2 \in \mathbb{M}$ and $v_1, v_2 \in \mathbb{R}^d$.

- (iv) For each $\pi \in \Pi$ and $k \in [1:n]$ the transition kernels κ_k satisfy the Lipschitz condition

$$\left| \sum_{\bar{w}_k \in \mathbb{W}} \left(\kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, m_1) \cdot v^{J^i(T_k, (w_{-k}, \bar{w}_k), m_1)} - \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, m_2) \cdot v^{J^i(T_k, (w_{-k}, \bar{w}_k), m_2)} \right) \right| \leq L_\kappa \cdot \|m_1 - m_2\| \quad (\text{L}_\kappa)$$

for all $i \in \mathbb{S}$, $w \in \mathbb{W}^n$, $m_1, m_2 \in \mathbb{M}$ and $v \in \mathbb{R}^d$ with $\|v\| \leq v_{\max}$, where

$$v_{\max} \triangleq (\Psi_{\max} + T \cdot \psi_{\max}) \cdot e^{Q_{\max} \cdot T}. \quad (v_{\max})$$

(v) For each $\pi \in \Pi$, $k \in [1:n]$ and $w \in \mathbb{W}^n$ the map⁶⁹ $\Phi_k(w, \cdot)$ is Lipschitz continuous, i.e.

$$\|\Phi_k(w, m_1) - \Phi_k(w, m_2)\| \leq L_{\Phi} \cdot \|m_1 - m_2\| \quad (L_{\Phi})$$

for all $w \in \mathbb{W}^n$ and $m_1, m_2 \in \mathbb{M}$. \diamond

Since all norms on \mathbb{R}^d are equivalent, the concrete specification is immaterial for Assumption D.1. However, the following results partially depend on the sizes of the relevant constants; to be concrete, in the following we use the maximum norm on \mathbb{R}^d and a compatible¹⁴² matrix norm on $\mathbb{R}^{d \times d}$, moreover, we suppose that $(L_{\widehat{Q}})$ holds for both \widehat{Q} and \widehat{Q}^{\top} .¹⁴³ The constants in (v_{\max}) are given by¹⁴⁴

$$\begin{aligned} Q_{\max} &\triangleq \sup_{\substack{t \in [0, T], w \in \mathbb{W}^n, \\ m \in \mathbb{M}, v \in \mathbb{R}^d}} \|\widehat{Q}(t, w, m, v)\| \vee \|\widehat{Q}(t, w, m, v)^{\top}\|, \\ \psi_{\max} &\triangleq \sup_{\substack{t \in [0, T], w \in \mathbb{W}^n, \\ m \in \mathbb{M}, v \in \mathbb{R}^d}} \|\widehat{\psi}(t, w, m, v)\| \quad \text{and} \quad \Psi_{\max} \triangleq \sup_{\substack{m \in \mathbb{M}, \\ w \in \mathbb{W}^n}} \|\Psi(w, m)\|. \end{aligned} \quad (D.1)$$

Remark D.2. Sufficient conditions for Assumption D.1(i)-(iii) in terms of the model's primitives can be found in, e.g., [GMS13] or [CF20]. Furthermore, in the special case where the jump map J is independent of $m \in \mathbb{M}$, Assumption D.1(v) is trivially satisfied, and a sufficient condition for Assumption D.1(iv) is given by the condition that the map

$$\mathbb{M} \ni m \mapsto \sum_{\bar{w}_k \in \cdot} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, m) \in \text{Prob}(\mathbb{W})$$

is Lipschitz continuous in total variation norm. \diamond

Let $p \in \mathbb{N}$ and $E \subseteq \mathbb{R}^p$. We consider the Skorokhod space

$$D(E) \triangleq \{f : [0, T] \times \mathbb{W}^n \rightarrow E : f \text{ is càdlàg}\}$$

together with the norm $\|f\|_{\text{sup}} \triangleq \sup_{t \in [0, T], w \in \mathbb{W}^n} \|f(t, w)\|$ for $f \in D(E)$, and for $\pi \in \Pi$ we define

⁶⁹With a slight abuse of notation, here and in the following we write

$$\Phi_k(w, m) \triangleq \Phi_k(w, m, m) \quad \text{for } k \in [1:n], w \in \mathbb{W}^n \text{ and } m \in \mathbb{M},$$

where the right-hand side is defined in (Φ_k) above.

¹⁴²A matrix norm is said to be *compatible* if $\|Ax\| \leq \|A\| \cdot \|x\|$ for every $A \in \mathbb{R}^{d \times d}$ and $x \in \mathbb{R}^d$; for example, one could consider the operator norm corresponding to $\|\cdot\|_{\infty}$, i.e. the maximum absolute row sum.

¹⁴³For the concrete specification of the norm, i.e. for $\|\cdot\| = \|\cdot\|_{\infty}$, and if one considers the corresponding operator norm on $\mathbb{R}^{d \times d}$, i.e. the maximum absolute row sum, one can easily derive the Lipschitz condition for \widehat{Q}^{\top} from the one for \widehat{Q} by using equivalence of $\|\cdot\|_{\infty}$ and $\|\cdot\|_1$ on \mathbb{R}^d ; recall that the operator norm corresponding to $\|\cdot\|_1$ is given by the maximum absolute column sum.

¹⁴⁴We implicitly assume $\psi_{\max} > 0$; otherwise this constant must be replaced by an arbitrary positive number.

the linear subspace

$$D_\pi(E) \triangleq \{f : [0, T] \times \mathbb{W}^n \rightarrow E : f \text{ is càdlàg and non-anticipative}\} \subseteq D(E).$$

By elementary results, it is clear that $D(E)$ is a Banach space provided $E \subseteq \mathbb{R}^p$ is closed; correspondingly, as the property of being non-anticipative with respect to a fixed grid $\pi \in \Pi$ is preserved by pointwise limits, $D_\pi(E)$ is also a Banach space in this case; the linear subspace of regular non-anticipative functions is denoted by $\text{Reg}_\pi(E) \subseteq D_\pi(E)$.¹⁴⁵

Lemma D.3 (Forward/Backward Gronwall Estimates). Let $\pi \in \Pi$, $f \in D_\pi([0, \infty))$ and $\delta \geq 0$.

(a) Let $\vec{\alpha}, \vec{\beta} \geq 0$ and $\vec{\gamma} \geq 1$, and suppose that $f(0) = 0$,

$$f(t, w) \leq f(T_k, w) + \vec{\alpha}(t - T_k) \cdot \delta + \int_{T_k}^t \vec{\beta} \cdot f(s, w) ds, \quad \text{for } t \in [T_k, T_{k+1}), w \in \mathbb{W}^n, \quad (\text{D.2})$$

for $k \in [0:n]$, and

$$f(T_k, w) \leq \vec{\gamma} \cdot f(T_{k-}, w) \quad \text{for } w \in \mathbb{W}^n, \quad (\text{D.3})$$

for $k \in [1:n]$. Then we have

$$f(t, w) \leq \vec{C} \cdot \delta \quad \text{for all } (t, w) \in [0, T] \times \mathbb{W}^n,$$

where

$$\vec{C} \triangleq \vec{\alpha} \cdot (\vec{\gamma})^n \cdot e^{\vec{\beta}T} \cdot T.$$

(b) Let $\overleftarrow{\rho}, \overleftarrow{\alpha}, \overleftarrow{\beta}, \overleftarrow{\varepsilon} \geq 0$ and suppose that $f(T-, w) = f(T, w) \leq \overleftarrow{\rho} \cdot \delta$ for $w \in \mathbb{W}^n$,

$$f(t, w) \leq f(T_{k+1-}, w) + \overleftarrow{\alpha}(T_{k+1} - t) \cdot \delta + \int_t^{T_{k+1}} \overleftarrow{\beta} \cdot f(s, w) ds, \quad \text{for } t \in [T_k, T_{k+1}), w \in \mathbb{W}^n, \quad (\text{D.4})$$

for $k \in [0:n]$, and

$$f(T_{k-}, w) \leq \sum_{\bar{w}_k \in \mathbb{W}} \overleftarrow{\gamma}_k(w_1, \dots, w_{k-1}, \bar{w}_k) \cdot f(T_k, (w_{-k}, \bar{w}_k)) + \overleftarrow{\varepsilon} \cdot \delta, \quad \text{for } w \in \mathbb{W}^n, \quad (\text{D.5})$$

for $k \in [1:n]$, where for all $w_1, \dots, w_{k-1} \in \mathbb{W}$ the family $\{\overleftarrow{\gamma}_k(w_1, \dots, w_{k-1}, \bar{w}_k)\}_{\bar{w}_k \in \mathbb{W}}$ consists of probability weights on \mathbb{W} . Then we have

$$f(t, w) \leq \overleftarrow{C} \cdot \delta \quad \text{for all } (t, w) \in [0, T] \times \mathbb{W}^n,$$

where

$$\overleftarrow{C} \triangleq (\overleftarrow{\rho} + \overleftarrow{\alpha} \cdot T + n \cdot \overleftarrow{\varepsilon}) \cdot e^{\overleftarrow{\beta}T}.$$

Importantly, the constants \vec{C} and \overleftarrow{C} do not depend on $\pi \in \Pi$. In particular, we have

$$C \triangleq \vec{C} \cdot \overleftarrow{C} < 1 \quad \text{whenever } T > 0 \text{ is sufficiently small.} \quad (\text{D.6})$$

¹⁴⁵cf. Definitions 2.1.1 and 2.1.2.

Proof. First we prove part (a): We recursively define $\vec{C}_0 \triangleq 0$ and

$$\vec{C}_{k+1} \triangleq \begin{cases} \vec{\gamma} \cdot (\vec{C}_k + \vec{\alpha}(T_{k+1} - T_k)) \cdot e^{\vec{\beta}(T_{k+1} - T_k)} & \text{for } k \in [0:n), \\ (\vec{C}_n + \vec{\alpha}(T - T_n)) \cdot e^{\vec{\beta}(T - T_n)} & \text{for } k = n. \end{cases}$$

Note that the sequence $\{\vec{C}_k\}_{k \in [0:(n+1)]}$ is non-decreasing and, by iterating the above recursion, it follows that $\vec{C}_{n+1} = \vec{\alpha} \cdot \sum_{k=0}^n (\vec{\gamma})^{n-k} \cdot e^{\vec{\beta}(T - T_k)} \cdot (T_{k+1} - T_k)$. Hence it holds $\vec{C}_{n+1} \leq \vec{C}$. As a consequence it suffices to show that

$$f(t, w) \leq \vec{C}_{k+1} \cdot \delta \quad \text{for all } (t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n, k \in [0:n]. \quad (\text{D.7})$$

It is clear that $f(T_0) = f(0) = 0$. Next let $k \in [0:n]$ and

$$f(T_k, w) \leq \vec{C}_k \cdot \delta \quad \text{for all } w \in \mathbb{W}^n.$$

It follows from (D.2) and Gronwall's inequality¹⁴⁶ that for $(t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n$

$$\begin{aligned} f(t, w) &\leq (f(T_k, w) + \vec{\alpha}(t - T_k) \cdot \delta) \cdot e^{\vec{\beta}(t - T_k)} \\ &\leq (\vec{C}_k + \vec{\alpha}(T_{k+1} - T_k)) \cdot e^{\vec{\beta}(T_{k+1} - T_k)} \cdot \delta \\ &\leq \vec{C}_{k+1} \cdot \delta. \end{aligned} \quad (\text{D.8})$$

In particular, if $k \in [0:n)$, (D.3) and (D.8) yield

$$\begin{aligned} f(T_{k+1}, w) &\leq \vec{\gamma} \cdot f(T_{k+1}^-, w) \\ &\leq \vec{\gamma} \cdot (\vec{C}_k + \vec{\alpha}(T_{k+1} - T_k)) \cdot e^{\vec{\beta}(T_{k+1} - T_k)} \cdot \delta = \vec{C}_{k+1} \cdot \delta, \end{aligned}$$

and consequently it holds

$$f(t, w) \leq \vec{C}_{k+1} \cdot \delta \quad \text{for all } (t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n.$$

Thus (D.7) follows by induction on $k \in [0:n]$ and the proof of part (a) is finished.

Part (b) follows by analogous arguments upon time reversal; for the sake of completeness we nevertheless state its proof in full detail: We recursively define $\overleftarrow{C}_{n+1} \triangleq \overleftarrow{\rho}$ and

$$\overleftarrow{C}_k \triangleq \begin{cases} (\overleftarrow{C}_{k+1} + \overleftarrow{\alpha}(T_{k+1} - T_k) + \overleftarrow{\varepsilon}) \cdot e^{\overleftarrow{\beta}(T_{k+1} - T_k)} & \text{for } k \in (0:n], \\ (\overleftarrow{C}_1 + \overleftarrow{\alpha}(T_1 - T_0)) \cdot e^{\overleftarrow{\beta}(T_1 - T_0)} & \text{for } k = 0. \end{cases}$$

Note that the sequence $\{\overleftarrow{C}_k\}_{k \in [0:n]}$ is non-increasing and, by iterating the above recursion it follows that $\overleftarrow{C}_0 = \overleftarrow{\rho} \cdot e^{\overleftarrow{\beta}T} + \overleftarrow{\alpha} \cdot \sum_{k=0}^n (T_{k+1} - T_k) \cdot e^{\overleftarrow{\beta}T_{k+1}} + \overleftarrow{\varepsilon} \cdot \sum_{k=1}^n e^{\overleftarrow{\beta}T_{k+1}}$. Hence it holds $\overleftarrow{C}_0 \leq \overleftarrow{C}$.

¹⁴⁶see, e.g., Lemma I.6.2 and p.37 in [Hal80].

As a consequence it suffices to show that

$$f(t, w) \leq \overleftarrow{C}_k \cdot \delta \quad \text{for all } (t, w) \in [T_k, T_{k+1}] \times \mathbb{W}^n, k \in [0:n]. \quad (\text{D.9})$$

It is clear that $f(T_{n+1}^-, w) = f(T, w) \leq \overleftarrow{\rho} \cdot \delta$ for all $w \in \mathbb{W}^n$. Next let $k \in [0:n]$ and

$$f(T_{k+1}^-, w) \leq \overleftarrow{C}_{k+1} \cdot \delta \quad \text{for all } w \in \mathbb{W}^n.$$

It follows from (D.4) and Gronwall's inequality¹⁴⁶ (applied backward in time) that for all $(t, w) \in [T_k, T_{k+1}] \times \mathbb{W}^n$

$$\begin{aligned} f(t, w) &\leq (f(T_{k+1}^-, w) + \overleftarrow{\alpha}(T_{k+1} - t) \cdot \delta) \cdot e^{\overleftarrow{\beta}(T_{k+1} - t)} \\ &\leq (\overleftarrow{C}_{k+1} + \overleftarrow{\alpha}(T_{k+1} - T_k)) \cdot e^{\overleftarrow{\beta}(T_{k+1} - T_k)} \cdot \delta \\ &\leq \overleftarrow{C}_k \cdot \delta. \end{aligned} \quad (\text{D.10})$$

In particular, if $k \in (0:n]$, (D.5) and (D.10) yield

$$\begin{aligned} f(T_k^-, w) &\leq \sum_{\bar{w}_k \in \mathbb{W}} \overleftarrow{\gamma}_k(w_1, \dots, w_{k-1}, \bar{w}_k) \cdot f(T_k, (w_{-k}, \bar{w}_k)) + \overleftarrow{\varepsilon} \cdot \delta \\ &\leq (\overleftarrow{C}_{k+1} + \overleftarrow{\alpha}(T_{k+1} - T_k) + \overleftarrow{\varepsilon}) \cdot e^{\overleftarrow{\beta}(T_{k+1} - T_k)} \cdot \delta \\ &= \overleftarrow{C}_k \cdot \delta \end{aligned}$$

and hence (D.9) follows by backward induction on $k = n, n-1, \dots, 0$; this finishes the proof of part (b). \square

In the following, we first consider the backward system (E2), (E4), (E6) and subsequently the forward system (E1), (E3), (E5).

Lemma D.4 (Backward System: Existence and Uniqueness). Suppose that Assumptions D.1(ii) and (iii) hold and let $\pi \in \Pi$ and $\mu \in \mathbf{D}_\pi(\mathbb{M})$. Then there exists a unique solution \bar{v} of (E2) subject to (E4) and (E6). Moreover, we have $\bar{v} \in \text{Reg}_\pi(\mathbb{R}^d)$ and $\|\bar{v}(t, w)\| \leq v_{\max}$ for all $(t, w) \in [0, T] \times \mathbb{W}^n$ where v_{\max} is given by (v_{max}).

In view of Definition 2.2.2 and Theorem 2.2.4, note that Lemma D.4 in particular establishes existence of a *partial equilibrium* $\bar{v} \in \text{Reg}_\pi(\mathbb{R}^d)$ given the representative agent's ex ante beliefs $\mu \in \mathbf{D}_\pi(\mathbb{M})$. \diamond

Proof of Lemma D.4. We divide the proof into two steps.

Step 1: Construction of \bar{v} . We construct \bar{v} by backward induction on $k \in [0:n]$ on each segment $[T_k, T_{k+1}] \times \mathbb{W}^n$. First, we set $\bar{v}(T, w) \triangleq \Psi(w, \mu(T, w))$ for $w \in \mathbb{W}^n$.

Suppose that $k \in [0:n]$, fix $w \in \mathbb{W}^n$, and let $\tilde{v}(T_{k+1}, w_{T_k}) \in \mathbb{R}^d$ be given and independent of w_{k+1}, \dots, w_n . Using (L $_{\hat{\psi}}$), (L $_{\hat{Q}}$) and (D.1) we obtain for $t \in [T_k, T_{k+1}]$ and $v_1, v_2 \in \overline{B_R(0)}$ (for an arbitrary radius $R > 0$)

$$\begin{aligned} \|\widehat{\psi}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) + \widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) \cdot v_1\| &\leq \psi_{\max} + Q_{\max} \cdot \|v_1\| \\ &\leq \psi_{\max} + Q_{\max} \cdot R \end{aligned} \quad (\text{D.11})$$

and

$$\begin{aligned} &\|\widehat{\psi}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) + \widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) \cdot v_1 \\ &\quad - \widehat{\psi}(t, w_{T_k}, \mu(t, w_{T_k}), v_2) - \widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_2) \cdot v_2\| \\ &\leq \|\widehat{\psi}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) - \widehat{\psi}(t, w_{T_k}, \mu(t, w_{T_k}), v_2)\| \\ &\quad + \|\widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_1)\| \cdot \|v_1 - v_2\| \\ &\quad + \|\widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_1) - \widehat{Q}(t, w_{T_k}, \mu(t, w_{T_k}), v_2)\| \cdot \|v_2\| \\ &\leq (L_{\widehat{\psi}} + Q_{\max} + L_{\widehat{Q}} \cdot R) \cdot \|v_1 - v_2\|. \end{aligned}$$

Thus it follows that the Carathéodory conditions¹⁴⁷ are satisfied, so Theorem I.5.3 in conjunction with Theorem I.5.2 in [Hal80]¹⁴⁸ yields the unique Carathéodory solution $\tilde{v}(\cdot, w_{T_k}) : [T_k, T_{k+1}] \rightarrow \mathbb{R}^d$ of

$$\begin{aligned} \tilde{v}(t, w_{T_k}) &= \tilde{v}(T_{k+1}, w_{T_k}) + \int_t^{T_{k+1}} \left(\widehat{\psi}(s, w_{T_k}, \mu(s, w_{T_k}), \tilde{v}(s, w_{T_k})) \right. \\ &\quad \left. + \widehat{Q}(s, w_{T_k}, \mu(s, w_{T_k}), \tilde{v}(s, w_{T_k})) \cdot \tilde{v}(s, w_{T_k}) \right) ds \\ &= \tilde{v}(T_{k+1}, w_{T_k}) + \int_t^{T_{k+1}} \left(\widehat{\psi}(s, w, \mu(s, w), \tilde{v}(s, w_{T_k})) \right. \\ &\quad \left. + \widehat{Q}(s, w, \mu(s, w), \tilde{v}(s, w_{T_k})) \cdot \tilde{v}(s, w_{T_k}) \right) ds, \quad t \in [T_k, T_{k+1}], \end{aligned}$$

where the final identity is due to the fact that μ and for all $(\bar{m}, \bar{v}) \in \mathbb{M} \times \mathbb{R}^d$ both $\widehat{\psi}(\cdot, \cdot, \bar{m}, \bar{v})$ and $\widehat{Q}(\cdot, \cdot, \bar{m}, \bar{v})$ are non-anticipative. We define

$$\bar{v}(t, w) \triangleq \tilde{v}(t, w_{T_k}) \quad \text{for } t \in [T_k, T_{k+1}].$$

By construction, $\bar{v}(\cdot, w)$ solves (E2) on $[T_k, T_{k+1}]$ and does not depend on w_{k+1}, \dots, w_n . Having constructed \bar{v} on $[T_k, T_{k+1}] \times \mathbb{W}^n$, in case of $k > 0$, we use (E4) and define

$$\tilde{v}(T_k, w_{T_{k-1}}) \triangleq \Psi_k(w, \mu(T_k-, w), \bar{v}(T_k, \cdot)) \quad \text{for } w \in \mathbb{W}^n.$$

By construction, (Ψ_k) and the fact that μ and J are non-anticipative, it follows that this definition does not depend on w_k, \dots, w_n . Consequently, the above construction can be iterated, and hence we obtain \bar{v} as the unique solution of (E2) subject to (E4) and (E6). By definition, \bar{v} is non-anticipative and regular, i.e. $\bar{v} \in \text{Reg}_\pi(\mathbb{R}^d)$.

¹⁴⁷see [Hal80, p.28] and Definition A.5(a).

¹⁴⁸see also Theorem A.6 and A.7 as well as Remark A.8(a). Note that the right-hand side of the ODE is linearly bounded by (D.11).

Step 2: A priori bound. For $k \in [0:n]$ and $(t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n$ we have

$$\begin{aligned} \|\bar{v}(t, w)\| &\leq \|\bar{v}(T_{k+1}^-, w)\| \\ &\quad + \int_t^{T_{k+1}} (\|\widehat{\psi}(s, w, \mu(s, w), \bar{v}(s, w))\| + \|\widehat{Q}(s, w, \mu(s, w), \bar{v}(s, w))\| \cdot \|\bar{v}(s, w)\|) ds \\ &\leq \|\bar{v}(T_{k+1}^-, w)\| + \psi_{\max} \cdot (T_{k+1} - t) + Q_{\max} \cdot \int_t^{T_{k+1}} \|\bar{v}(s, w)\| ds. \end{aligned} \quad (\text{D.12})$$

On the other hand, for $k \in [1:n]$, $w \in \mathbb{W}^n$ and $i \in \mathbb{S}$ we observe from (Ψ_k) that

$$\|\bar{v}(T_k^-, w)\| \leq \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu(T_k^-, w_{T_k^-})) \cdot \|\bar{v}(T_k, (w_{-k}, \bar{w}_k))\|. \quad (\text{D.13})$$

Since $\|\bar{v}(T^-, w)\| = \|\bar{v}(T, w)\| = \|\Psi(w, \mu(T, w))\| \leq \Psi_{\max}$ it follows from (D.12), (D.13) and Lemma D.3(b) (see (D.4) and (D.5)) with $\bar{\rho} \triangleq \Psi_{\max}/\psi_{\max}$, $\bar{\alpha} = 1$, $\bar{\beta} = Q_{\max}$, $\bar{\varepsilon} = 0$ and $\delta \triangleq \psi_{\max}$ that

$$\|\bar{v}(t, w)\| \leq \bar{C} \cdot \delta = (\Psi_{\max} + T \cdot \psi_{\max}) \cdot e^{Q_{\max} \cdot T} = v_{\max} \quad \text{for all } (t, w) \in [0, T] \times \mathbb{W}^n. \quad \square$$

Lemma D.5 (Forward System: Existence and Uniqueness). Suppose that Assumption D.1(iii) is satisfied and let $\pi \in \Pi$ and $v \in \mathbf{D}_\pi(\mathbb{R}^d)$. Then there is a unique solution $\bar{\mu}$ of (E1) subject to (E3) and (E5), and we have $\bar{\mu} \in \text{Reg}_\pi(\mathbb{M})$.

Proof. The proof is analogous to (but somewhat simpler than) that of Lemma D.4. For the sake of completeness we nevertheless state it in full detail: We construct $\bar{\mu}$ by forward induction on $k \in [0:n]$ on each segment $[T_k, T_{k+1}) \times \mathbb{W}^n$. First, we set $\bar{\mu}(0, w) \triangleq m_0 \in \mathbb{M}$ for $w \in \mathbb{W}^n$. Suppose that $k \in [0:n]$, fix $w \in \mathbb{W}^n$ and let $\tilde{\mu}(T_k, w_{T_k}) \in \mathbb{M}$ be given and independent of w_{k+1}, \dots, w_n . By $(L_{\widehat{Q}})$ and (D.1) we obtain for $t \in [T_k, T_{k+1})$ and $m_1, m_2 \in \mathbb{M}$

$$\|\widehat{Q}(t, w_{T_k}, m_1, v(t, w_{T_k}))^\top \cdot m_1^\top\| \leq \|\widehat{Q}(t, w_{T_k}, m_1, v(t, w_{T_k}))^\top\| \cdot \|m_1\| \leq Q_{\max}$$

and

$$\begin{aligned} &\|\widehat{Q}(t, w_{T_k}, m_1, v(t, w_{T_k}))^\top \cdot m_1^\top - \widehat{Q}(t, w_{T_k}, m_2, v(t, w_{T_k}))^\top \cdot m_2^\top\| \\ &\leq \|\widehat{Q}(t, w_{T_k}, m_1, v(t, w_{T_k}))^\top - \widehat{Q}(t, w_{T_k}, m_2, v(t, w_{T_k}))^\top\| \cdot \|m_1\| \\ &\quad + \|\widehat{Q}(t, w_{T_k}, m_2, v(t, w_{T_k}))^\top\| \cdot \|m_1 - m_2\| \\ &\leq (L_{\widehat{Q}} + Q_{\max}) \cdot \|m_1 - m_2\|, \end{aligned} \quad (\text{D.14})$$

and thus the Carathéodory conditions¹⁴⁷ are satisfied, so Theorem I.5.3 in [Hal80]¹⁴⁹ yields the unique Carathéodory solution $\tilde{\mu}(\cdot, w_{T_k}) : [T_k, T_{k+1}] \rightarrow \mathbb{R}^d$ of

$$\tilde{\mu}(t, w_{T_k}) = \tilde{\mu}(T_k, w_{T_k}) + \int_{T_k}^t \tilde{\mu}(s, w_{T_k}) \cdot \widehat{Q}(s, w_{T_k}, \tilde{\mu}(s, w_{T_k}), v(s, w_{T_k})) ds$$

¹⁴⁹see also Theorem A.6 and Remark A.8(b). Note that the right-hand side of the ODE is globally Lipschitz continuous by (D.14).

$$= \tilde{\mu}(T_k, w_{T_k}) + \int_{T_k}^t \tilde{\mu}(s, w_{T_k}) \cdot \widehat{Q}(s, w, \tilde{\mu}(s, w_{T_k}), v(s, w)) ds \quad \text{for } t \in [T_k, T_{k+1}],$$

where the final identity follows as v and, for all $(\bar{m}, \bar{v}) \in \mathbb{M} \times \mathbb{R}^d$, the map $\widehat{Q}(\cdot, \cdot, \bar{m}, \bar{v})$ are non-anticipative. Observe that both $\tilde{\mu}(T_k, w_{T_k}) \in \mathbb{M}$ and the intensity matrix condition guarantee $\tilde{\mu}(t, w_{T_k}) \in \mathbb{M}$ for $t \in [T_k, T_{k+1}]$; indeed, we have

$$\begin{aligned} & \sum_{i=1}^d \sum_{j=1}^d \tilde{\mu}^j(s, w_{T_k}) \cdot \widehat{Q}^{ji}(s, w_{T_k}, \tilde{\mu}(s, w_{T_k}), \bar{v}(s, w_{T_k})) \\ &= \sum_{j=1}^d \tilde{\mu}^j(s, w_{T_k}) \cdot \sum_{i=1}^d \widehat{Q}^{ji}(s, w_{T_k}, \tilde{\mu}(s, w_{T_k}), \bar{v}(s, w_{T_k})) = 0 \quad \text{for all } s \in [T_k, T_{k+1}]. \end{aligned}$$

We define

$$\bar{\mu}(t, w) \triangleq \tilde{\mu}(t, w_{T_k}) \in \mathbb{M} \quad \text{for } t \in [T_k, T_{k+1}].$$

By construction, $\bar{\mu}(\cdot, w)$ solves (E1) on $[T_k, T_{k+1}]$ and does not depend on w_{k+1}, \dots, w_n . Having constructed $\bar{\mu}$ on $[T_k, T_{k+1}] \times \mathbb{W}^n$, in case of $k < n$, we use (E3) and define

$$\tilde{\mu}(T_{k+1}, w_{T_{k+1}}) \triangleq \Phi_{k+1}(w, \bar{\mu}(T_{k+1}^-, w)) \quad \text{for } w \in \mathbb{W}^n.$$

By construction, $(\Phi_k)^{69}$ and since J is non-anticipative, it follows that this definition does not depend on w_l for $l > k + 1$. Moreover, note that $\tilde{\mu}(T_{k+1}, w_{T_{k+1}}) \in \mathbb{M}$ since, thanks to (Φ_k) , it holds $\sum_{i=1}^d \tilde{\mu}^i(T_{k+1}, w_{T_{k+1}}) = \sum_{j=1}^d \bar{\mu}^j(T_{k+1}, w_{T_k}) \sum_{i=1}^d P^{ji}(w, \bar{\mu}(T_{k+1}, w_{T_k})) = 1$. Hence, the above construction can be iterated and we obtain $\bar{\mu}$ as the unique solution of (E1) subject to (E3) and (E5). By construction, $\bar{\mu}$ is non-anticipative, regular and \mathbb{M} -valued, i.e. $\bar{\mu} \in \text{Reg}_\pi(\mathbb{M})$. \square

We are now in a position to establish existence of a mean field equilibrium, i.e. to prove Theorem 2.3.6. For convenience we repeat its statement:

Theorem 2.3.6 (Existence of Equilibria). Fix $n \in \mathbb{N}_0$ and let Assumption D.1 be satisfied.¹⁵⁰ Then there exists $T^* > 0$ such that for every time horizon $T \leq T^*$ and every choice of common noise times $0 = T_0 < T_1 < \dots < T_n < T_{n+1} = T$ there is a unique solution of the equilibrium system (E1)-(E6).

As already pointed out above, the following proof is a ramification of the one of Theorem 6 in [CF20]; it extends the latter result to the present setup which includes common noise. \diamond

Proof of Theorem 2.3.6. Let $T > 0$ be as in (D.6) below Lemma D.3, where the relevant coefficients are given by

$$\begin{aligned} \vec{\alpha} &= L_{\widehat{Q}}, & \vec{\beta} &= Q_{\max} + L_{\widehat{Q}}, & \vec{\gamma} &= L_{\Phi} + 1, \\ \overleftarrow{\rho} &= L_{\Psi}, & \overleftarrow{\alpha} &= L_{\widehat{\psi}} + L_{\widehat{Q}} \cdot v_{\max}, & \overleftarrow{\beta} &= L_{\widehat{\psi}} + Q_{\max} + L_{\widehat{Q}} \cdot v_{\max}, & \overleftarrow{\varepsilon} &= L_{\kappa}, \end{aligned}$$

and fix some arbitrary $\pi \in \Pi$.

¹⁵⁰Formally, we suppose that Assumption D.1 is satisfied for the given fixed $n \in \mathbb{N}_0$ and some $T > 0$.

Step 1: Solution operators. We define

$$\overleftarrow{\chi} : \mathbf{D}_\pi(\mathbb{M}) \rightarrow \text{Reg}_\pi(\mathbb{R}^d), \quad \overleftarrow{\chi}[\mu] \triangleq \bar{v},$$

where $\bar{v} \in \text{Reg}_\pi(\mathbb{R}^d)$ is the unique solution of (E2) subject to (E4) and (E6) given $\mu \in \mathbf{D}_\pi(\mathbb{M})$; $\overleftarrow{\chi}$ is well-defined by Lemma D.4 and it holds $\|\overleftarrow{\chi}[\mu](t, w)\| \leq v_{\max}$ for all $(t, w) \in [0, T] \times \mathbb{W}^n$ and $\mu \in \mathbf{D}_\pi(\mathbb{M})$. Moreover, let

$$\overrightarrow{\chi} : \mathbf{D}_\pi(\mathbb{R}^d) \rightarrow \text{Reg}_\pi(\mathbb{M}), \quad \overrightarrow{\chi}[v] \triangleq \bar{\mu},$$

where $\bar{\mu} \in \text{Reg}_\pi(\mathbb{M})$ is the unique solution of (E1) subject to (E3) and (E5) given $v \in \mathbf{D}_\pi(\mathbb{R}^d)$; $\overrightarrow{\chi}$ is well-defined by Lemma D.5.

Step 2: Lipschitz continuity of $\overleftarrow{\chi}$. Let $\mu_1, \mu_2 \in \mathbf{D}_\pi(\mathbb{M})$ and set $\bar{v}_1 \triangleq \overleftarrow{\chi}[\mu_1]$ and $\bar{v}_2 \triangleq \overleftarrow{\chi}[\mu_2]$. Using $(L_{\widehat{\psi}})$, $(L_{\widehat{Q}})$ and (D.1) it follows that for all $k \in [0:n]$ and $(t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n$ we have

$$\begin{aligned} \|\bar{v}_1(t, w) - \bar{v}_2(t, w)\| &\leq \|\bar{v}_1(T_{k+1}^-, w) - \bar{v}_2(T_{k+1}^-, w)\| \\ &\quad + \int_t^{T_{k+1}} \|\widehat{\psi}(s, w, \mu_1(s, w), \bar{v}_1(s, w)) - \widehat{\psi}(s, w, \mu_2(s, w), \bar{v}_2(s, w))\| ds \\ &\quad + \int_t^{T_{k+1}} \|\widehat{Q}(s, w, \mu_1(s, w), \bar{v}_1(s, w)) \cdot \bar{v}_1(s, w) - \widehat{Q}(s, w, \mu_2(s, w), \bar{v}_2(s, w)) \cdot \bar{v}_2(s, w)\| ds \\ &\leq \|\bar{v}_1(T_{k+1}^-, w) - \bar{v}_2(T_{k+1}^-, w)\| \\ &\quad + \int_t^{T_{k+1}} L_{\widehat{\psi}} \cdot (\|\mu_1(s, w) - \mu_2(s, w)\| + \|\bar{v}_1(s, w) - \bar{v}_2(s, w)\|) ds \\ &\quad + \int_t^{T_{k+1}} \left(L_{\widehat{Q}} \cdot (\|\mu_1(s, w) - \mu_2(s, w)\| + \|\bar{v}_1(s, w) - \bar{v}_2(s, w)\|) \cdot \|\bar{v}_1(s, w)\| \right. \\ &\quad \quad \quad \left. + \|\widehat{Q}(s, w, \mu_2(s, w), \bar{v}_2(s, w))\| \cdot \|\bar{v}_1(s, w) - \bar{v}_2(s, w)\| \right) ds \\ &\leq \|\bar{v}_1(T_{k+1}^-, w) - \bar{v}_2(T_{k+1}^-, w)\| + (L_{\widehat{\psi}} + L_{\widehat{Q}} \cdot v_{\max}) \cdot (T_{k+1} - t) \cdot \|\mu_1 - \mu_2\|_{\text{sup}} \\ &\quad + (L_{\widehat{\psi}} + L_{\widehat{Q}} \cdot v_{\max} + Q_{\max}) \cdot \int_t^{T_{k+1}} \|\bar{v}_1(s, w) - \bar{v}_2(s, w)\| ds. \end{aligned} \tag{D.15}$$

On the other hand for $k \in [1:n]$ we obtain from (Ψ_k) and (L_κ) that for every $i \in \mathbb{S}$ and $w \in \mathbb{W}^n$

$$\begin{aligned} &|\bar{v}_1^i(T_k^-, w) - \bar{v}_2^i(T_k^-, w)| \\ &= \left| \sum_{\bar{w}_k \in \mathbb{W}} \left(\kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_1(T_k^-, w_{T_k^-})) \cdot \bar{v}_1^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_1(T_k^-, w))}(T_k, (w_{-k}, \bar{w}_k)) \right. \right. \\ &\quad \left. \left. - \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_2(T_k^-, w_{T_k^-})) \cdot \bar{v}_2^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_2(T_k^-, w))}(T_k, (w_{-k}, \bar{w}_k)) \right) \right| \\ &\leq \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_1(T_k^-, w_{T_k^-})) \cdot \left| \bar{v}_1^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_1(T_k^-, w_{T_k^-}))}(T_k, (w_{-k}, \bar{w}_k)) \right. \\ &\quad \left. - \bar{v}_2^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_1(T_k^-, w_{T_k^-}))}(T_k, (w_{-k}, \bar{w}_k)) \right| \\ &\quad + \left| \sum_{\bar{w}_k \in \mathbb{W}} \left(\kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_1(T_k^-, w_{T_k^-})) \cdot \bar{v}_2^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_1(T_k^-, w_{T_k^-}))}(T_k, (w_{-k}, \bar{w}_k)) \right. \right. \end{aligned}$$

$$\begin{aligned}
& - \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_2(T_k-, w_{T_k-})) \cdot \bar{v}_2^{J^i(T_k, (w_{-k}, \bar{w}_k), \mu_2(T_k-, w_{T_k-}))}(T_k, (w_{-k}, \bar{w}_k)) \Big| \\
\leq & \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_1(T_k-, w_{T_k-})) \cdot \|\bar{v}_1(T_k, (w_{-k}, \bar{w}_k)) - \bar{v}_2(T_k, (w_{-k}, \bar{w}_k))\| \\
& + L_\kappa \cdot \|\mu_1(T_k-, w_{T_k-}) - \mu_2(T_k-, w_{T_k-})\|.
\end{aligned}$$

As a consequence it holds for every $w \in \mathbb{W}^n$

$$\begin{aligned}
& \|\bar{v}_1(T_k-, w) - \bar{v}_2(T_k-, w)\| \\
\leq & \sum_{\bar{w}_k \in \mathbb{W}} \kappa_k(\bar{w}_k | w_1, \dots, w_{k-1}, \mu_1(T_k-, w_{T_k-})) \cdot \|\bar{v}_1(T_k, (w_{-k}, \bar{w}_k)) - \bar{v}_2(T_k, (w_{-k}, \bar{w}_k))\| \\
& + L_\kappa \cdot \|\mu_1 - \mu_2\|_{\text{sup}}. \tag{D.16}
\end{aligned}$$

Finally, for $w \in \mathbb{W}^n$ by (L_Ψ) we have

$$\begin{aligned}
\|\bar{v}_1(T-, w) - \bar{v}_2(T-, w)\| &= \|\bar{v}_1(T, w) - \bar{v}_2(T, w)\| \\
&= \|\Psi(w, \mu_1(T, w)) - \Psi(w, \mu_2(T, w))\| \leq L_\Psi \cdot \|\mu_1 - \mu_2\|_{\text{sup}}. \tag{D.17}
\end{aligned}$$

In view of (D.15), (D.16) and (D.17) it follows from Lemma D.3(b) with $\delta \triangleq \|\mu_1 - \mu_2\|_{\text{sup}}$ that

$$\|\bar{v}_1(t, w) - \bar{v}_2(t, w)\| \leq \bar{C} \cdot \|\mu_1 - \mu_2\|_{\text{sup}} \quad \text{for all } (t, w) \in [0, T] \times \mathbb{W}^n,$$

and thus

$$\|\bar{\chi}[\mu_1] - \bar{\chi}[\mu_2]\|_{\text{sup}} = \|\bar{v}_1 - \bar{v}_2\|_{\text{sup}} \leq \bar{C} \cdot \|\mu_1 - \mu_2\|_{\text{sup}}. \tag{D.18}$$

Step 3: Lipschitz continuity of $\bar{\chi}$. Let $v_1, v_2 \in \mathbf{D}_\pi(\mathbb{R}^d)$ and set $\bar{\mu}_1 \triangleq \bar{\chi}[v_1]$ and $\bar{\mu}_2 \triangleq \bar{\chi}[v_2]$. By $(L_{\widehat{Q}})$ and (D.1) we have for $k \in [0:n]$ and $(t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n$

$$\begin{aligned}
\|\bar{\mu}_1(t, w) - \bar{\mu}_2(t, w)\| &\leq \|\bar{\mu}_1(T_k, w) - \bar{\mu}_2(T_k, w)\| \\
&+ \int_{T_k}^t \|\bar{\mu}_1(s, w) \cdot \widehat{Q}(s, w, \bar{\mu}_1(s, w), v_1(s, w)) - \bar{\mu}_2(s, w) \cdot \widehat{Q}(s, w, \bar{\mu}_2(s, w), v_2(s, w))\| ds \\
\leq & \|\bar{\mu}_1(T_k, w) - \bar{\mu}_2(T_k, w)\| \\
&+ \int_{T_k}^t L_{\widehat{Q}} \cdot (\|\bar{\mu}_1(s, w) - \bar{\mu}_2(s, w)\| + \|v_1(s, w) - v_2(s, w)\|) \cdot \|\bar{\mu}_1(s, w)\| ds \\
&+ \int_{T_k}^t \|\widehat{Q}(s, w, \bar{\mu}_2(s, w), v_2(s, w))^\top\| \cdot \|\bar{\mu}_1(s, w) - \bar{\mu}_2(s, w)\| ds \\
\leq & \|\bar{\mu}_1(T_k, w) - \bar{\mu}_2(T_k, w)\| + L_{\widehat{Q}} \cdot (t - T_k) \cdot \|v_1 - v_2\|_{\text{sup}} \\
&+ (L_{\widehat{Q}} + Q_{\text{max}}) \cdot \int_{T_k}^t \|\bar{\mu}_1(s, w) - \bar{\mu}_2(s, w)\| ds. \tag{D.19}
\end{aligned}$$

On the other hand, by (L_Φ) we have for $k \in [1:n]$ and $w \in \mathbb{W}^n$

$$\begin{aligned}
\|\bar{\mu}_1(T_k, w) - \bar{\mu}_2(T_k, w)\| &= \|\Phi_k(w, \bar{\mu}_1(T_k-, w)) - \Phi_k(w, \bar{\mu}_2(T_k-, w))\| \\
&\leq L_\Phi \cdot \|\bar{\mu}_1(T_k-, w) - \bar{\mu}_2(T_k-, w)\|. \tag{D.20}
\end{aligned}$$

Since $\bar{\mu}_1(0) = \bar{\mu}_2(0) = m_0$, it follows from (D.19), (D.20)¹⁵¹ and Lemma D.3(a) with $\delta \triangleq \|v_1 - v_2\|_{\text{sup}}$ that

$$\|\bar{\mu}_1(t, w) - \bar{\mu}_2(t, w)\| \leq \vec{C} \cdot \|v_1 - v_2\|_{\text{sup}} \quad \text{for all } (t, w) \in [0, T] \times \mathbb{W}^n,$$

and consequently

$$\|\vec{\chi}[v_1] - \vec{\chi}[v_2]\|_{\text{sup}} = \|\bar{\mu}_1 - \bar{\mu}_2\|_{\text{sup}} \leq \vec{C} \cdot \|v_1 - v_2\|_{\text{sup}}. \quad (\text{D.21})$$

Step 4: Construction of the fixed point. Let $\chi : \mathbf{D}_\pi(\mathbb{M}) \rightarrow \mathbf{Reg}_\pi(\mathbb{M})$, $\chi \triangleq \vec{\chi} \circ \overleftarrow{\chi}$ and note that by (D.18) and (D.21) we have

$$\|\chi[\mu_1] - \chi[\mu_2]\|_{\text{sup}} \leq C \cdot \|\mu_1 - \mu_2\|_{\text{sup}} \quad \text{for all } \mu_1, \mu_2 \in \mathbf{D}_\pi(\mathbb{M}),$$

where $C = \vec{C} \cdot \overleftarrow{C} < 1$ by (D.6). Thus Banach's fixed point theorem yields a unique fixed point $\mu \in \mathbf{D}_\pi(\mathbb{M})$ of χ . Finally, note that by setting $v \triangleq \overleftarrow{\chi}[\mu] \in \mathbf{Reg}_\pi(\mathbb{R}^d)$ it follows that $\mu = \vec{\chi}[v] \in \mathbf{Reg}_\pi(\mathbb{M})$ and that (μ, v) is a solution of (E1)-(E6). \square

¹⁵¹Note that $\vec{\gamma} = L_\Phi + 1 \geq 1$.

Appendix E

The Master Equation

An alternative characterization of a mean field equilibrium (see Definition 2.3.5) can be obtained via the so-called *master equation*, a first-order partial differential equation. For the sake of completeness, we briefly outline this approach within the framework of Chapter 2 in this part of the appendix.

The master equation is oftentimes considered in a diffusive continuous-state setup; see, e.g., Chapters 4 and 5 in [CD18b] or [CDLL19] for corresponding general discussions. Nevertheless, it has also been addressed in the literature on finite-state mean field games; we refer, e.g., to [Gué11], [GVW14] or §7.2 in [CD18a] for the case without common noise; and for recent developments including common noise see, e.g., [BCCD19], [BLL19] or [BLL20]. Note in particular that the master equation in Definition E.2 below is related to the ones stated in [GVW14, (1.1)] and [BLL19, (11)].

Definition E.1 (Hamiltonian). The function $\mathcal{H} : [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d \times \mathbb{U} \rightarrow \mathbb{R}^d$,

$$\mathcal{H}^i(t, w, m, v, u) \triangleq \psi^i(t, w, m, u) + Q^{i\bullet}(t, w, m, u) \cdot v$$

is called the *Hamiltonian*. The *reduced Hamiltonian* $\widehat{\mathcal{H}} : [0, T] \times \mathbb{W}^n \times \mathbb{M} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined via

$$\widehat{\mathcal{H}}^i(t, w, m, v) \triangleq \sup_{u \in \mathbb{U}} \mathcal{H}^i(t, w, m, v, u) = \widehat{\psi}^i(t, w, m, v) + \widehat{Q}^{i\bullet}(t, w, m, v) \cdot v. \quad \diamond$$

Note that using the Hamiltonian, the reduced DPE (DP $_{\mu}$) for the agent's optimization problem (P $_{\mu}$) can be rewritten as follows:

$$\dot{v}(t, w) + \widehat{\mathcal{H}}(t, w, \mu(t, w), v(t, w)) = 0, \quad (t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n, \quad k \in [0:n].$$

Definition E.2 (Solution of the Master Equation). Let $\mathcal{U} : [0, T] \times \mathbb{W}^n \times \mathbb{M} \rightarrow \mathbb{R}^d$ be such that

- (i) $\mathcal{U}(\cdot, \cdot, m)$ is non-anticipative for every $m \in \mathbb{M}$,
- (ii) $\mathcal{U}(\cdot, w, \cdot)$ is continuous on $[T_k, T_{k+1}) \times \mathbb{M}$, extends continuously to $[T_k, T_{k+1}] \times \mathbb{M}$, and is continuously differentiable on $(T_k, T_{k+1}) \times \mathbb{M}$ for $k \in [1:n]$ and every $w \in \mathbb{W}^n$.¹⁵²

¹⁵²For a discussion of differentiability with respect to probability distributions, i.e. elements $m \in \mathbb{M}$, we refer, e.g., to [CD18a, p.653/654] or §3.2.1 in [BCCD19]. As this part of the appendix is only intended to give a brief outline, for the sake of simplicity, we require \mathcal{U} to be defined in an open neighborhood of $\mathbb{M} \subset \mathbb{R}^{1 \times d}$.

If for every $i \in \mathbb{S}$ and $(t, w, m) \in (T_k, T_{k+1}) \times \mathbb{W}^n \times \mathbb{M}$, $k \in [0:n]$, we have

$$\partial_t \mathcal{U}^i(t, w, m) + \widehat{\mathcal{H}}^i(t, w, m, \mathcal{U}(t, w, m)) + m \cdot \widehat{Q}(t, w, m, \mathcal{U}(t, w, m)) \cdot \nabla_m \mathcal{U}^i(t, w, m) = 0 \quad (\text{M-E})$$

and \mathcal{U} further satisfies the consistency and terminal conditions

$$\mathcal{U}(T_k-, w, m) = \Psi_k(w, m, \bar{v}_k(\cdot, m)) \quad \text{with } \bar{v}_k(w, m) \triangleq \mathcal{U}(T_k, w, \Phi_k(w, m)) \quad (\text{M-CC})$$

$$\mathcal{U}(T, w, m) = \Psi(w, m) \quad (\text{M-TC})$$

for $(w, m) \in \mathbb{W}^n \times \mathbb{M}$ and $k \in [1:n]$, where the functions Ψ_k and Φ_k , $k \in [1:n]$, are defined in (Ψ_k) and (Φ_k) ⁶⁹, respectively, then \mathcal{U} is called a *solution of the master equation*. \diamond

Similarly as in Propostion 7.7 in [CD18a] in a setting without common noise, a solution of the master equation induces a mean field equilibrium via the following result:

Theorem E.3 (Master Equation). Suppose that $\mathcal{U} : [0, T] \times \mathbb{W}^n \times \mathbb{M} \rightarrow \mathbb{R}^d$ is a solution of the master equation. Let $\mu : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{M}$ be a piecewise smooth, non-anticipative solution of the Kolmogorov forward equation

$$\dot{\mu}(t, w) = \mu(t, w) \cdot \widehat{Q}(t, w, \mu(t, w), \mathcal{U}(t, w, \mu(t, w))) \quad (\text{M-K})$$

for $(t, w) \in [T_k, T_{k+1}) \times \mathbb{W}^n$, $k \in [0:n]$, subject to the initial condition

$$\mu(0, w) = m_0, \quad (\text{M-M}_0)$$

and the consistency conditions

$$\mu(T_k, w) = \Phi_k(w, \mu(T_k-, w)) \quad (\text{M-M}_k)$$

for $w \in \mathbb{W}^n$ and $k \in [1:n]$, where the functions Φ_k , $k \in [1:n]$, are given by (M_k) ⁶⁹. Moreover, define $v : [0, T] \times \mathbb{W}^n \rightarrow \mathbb{R}^d$ via

$$v(t, w) \triangleq \mathcal{U}(t, w, \mu(t, w)) \quad \text{for } (t, w) \in [0, T] \times \mathbb{W}^n. \quad (\text{M-V})$$

Then, the pair (μ, v) is a solution of the equilibrium system (E1)-(E6), i.e. it constitutes a rational expectations equilibrium.

Proof. First note that v is non-anticipative and regular (in fact, piecewise smooth) by construction. Using (M-E) and (M-K) we compute for each $i \in \mathbb{S}$ with the chain rule

$$\begin{aligned} \dot{v}^i(t, w) &= \partial_t \mathcal{U}^i(t, w, \mu(t, w)) + \nabla_m \mathcal{U}^i(t, w, \mu(t, w))^\top \cdot \partial_t \mu(t, w)^\top \\ &= \partial_t \mathcal{U}^i(t, w, \mu(t, w)) + \mu(t, w) \cdot \widehat{Q}(t, w, \mu(t, w), \mathcal{U}(t, w, \mu(t, w))) \cdot \nabla_m \mathcal{U}^i(t, w, \mu(t, w)) \\ &= -\widehat{\mathcal{H}}^i(t, w, \mu(t, w), \mathcal{U}(t, w, \mu(t, w))) = -\widehat{\mathcal{H}}^i(t, w, \mu(t, w), v(t, w)) \end{aligned}$$

for $(t, w) \in (T_k, T_{k+1}) \times \mathbb{W}^n$, $k \in [0:n]$, and thus v satisfies (DP_μ) . The consistency conditions (CC_μ) follow immediately from part (ii) of Definition E.2 and (M-CC), as for $w \in \mathbb{W}^n$ and

$k \in [1:n]$ we have

$$\begin{aligned}
 v(T_k-, w) &= \lim_{t \uparrow T_k} v(t, w) = \lim_{t \uparrow T_k} \mathcal{U}(t, w, \mu(t, w)) = \lim_{t \uparrow T_k} \mathcal{U}(t, w, \mu(T_k-, w)) \\
 &= \Psi_k(w, \mu(T_k-, w), \bar{v}_k(\cdot, \mu(T_k-, w))) \\
 &= \Psi_k(w, \mu(T_k-, w), v(T_k, \cdot))
 \end{aligned}$$

where the last identity follows from (M-M_k) via

$$\bar{v}_k(w, \mu(T_k-, w)) = \mathcal{U}(T_k, w, \Phi_k(w, \mu(T_k-, w))) = \mathcal{U}(T_k, w, \mu(T_k, w)) = v(T_k, w).$$

The terminal condition (TC_μ) is readily derived from (M-TC) as it holds

$$v(T, w) = \mathcal{U}(T, w, \mu(T, w)) = \Psi(w, \mu(T, w)).$$

Thus v solves (DP_μ) and is subject to (CC_μ), (TC_μ) with μ defined by (M-K) subject to (M-M₀), (M-M_k); hence the pair (μ, v) is a solution of the equilibrium system (E1)-(E6). \square

List of Figures

1.1	Illustration of the branching mechanism.	5
1.2	Sample path of a one-dimensional branching Brownian motion with jumps.	7
1.3	Simulation results (Simulation Study).	22
1.4	Polynomial approximation of the recovery value function h (SIB).	27
1.5	Payoff profiles g_{short} and g_{long}	28
1.6	Valuation of g_{short}	31
1.7	Decomposition of the SIB spread into credit risk and systemic risk.	31
1.8	Valuation of g_{long}	32
1.9	Valuation of g_{basket}	33
1.10	Estimated (relative) standard deviation for g_{basket}	33
1.11	Ratios of branching coefficients and offspring probability weights (SIB).	35
2.1	Illustration of common noise factors and times	40
2.2	Illustration of common noise paths: The Tree of FBODEs.	58
2.3	Stylized area of farmland (<i>Agricultural Production</i>).	59
2.4	Price and instantaneous profit function (<i>Agricultural Production</i>).	59
2.5	Equilibrium proportion of farmers (<i>Agricultural Production</i>).	62
2.6	Equilibrium optimal action h^0 of non-farmers (<i>Agricultural Production</i>).	63
2.7	Equilibrium market price (<i>Agricultural Production</i>).	64
2.8	Equilibrium proportion of farmers in model (C) for all possible common noise configurations (<i>Agricultural Production</i>).	65
2.9	State space SIR <i>Model</i>	65
2.10	Running reward of susceptible agents: Penalization for making protection efforts (SIR <i>Model</i>).	66
2.11	Equilibrium distribution and protection effort for $\tau^* = +\infty$ (SIR <i>Model</i>).	69
2.12	Equilibrium distribution and protection effort for $\tau^* = 2.5$ (SIR <i>Model</i>).	70
2.13	Equilibrium distribution and protection effort for $\tau^* = 5$ (SIR <i>Model</i>).	71
2.14	Equilibrium distribution and protection effort for $\tau^* = 7.5$ (SIR <i>Model</i>).	72
2.15	State space <i>Evacuation</i>	73

2.16	Possible directions of movement (<i>Evacuation</i>).	74
2.17	Example 1 (<i>Evacuation</i>): Aggregate distribution.	76
2.18	Example 2 (<i>Evacuation</i>): Aggregate distribution.	77
2.19	Example 3 (<i>Evacuation</i>): Aggregate distribution.	78
2.20	Example 3 (<i>Evacuation</i>): Aggregate distribution without a pushback effect.	79
2.21	Example 4 (<i>Evacuation</i>): Aggregate distribution.	80
2.22	Example 5 (<i>Evacuation</i>): Aggregate distribution.	81
2.23	Example 6 (<i>Evacuation</i>): Aggregate distribution.	82
2.24	Example 6 (<i>Evacuation</i>): Aggregate distribution without a pushback effect.	83
3.1	Dynamics of the state process.	89
3.2	Concatenation of actions.	93
3.3	Example of a Voronoi tessellation in \mathbb{R}^2	104
3.4	Controlled Brownian motion (BM) in 1D without proportional and terminal costs.	111
3.5	Controlled Brownian motion (BM) in 1D without proportional, but with terminal costs.	111
3.6	Controlled Brownian motion (BM) in 1D with fixed and proportional costs I.	112
3.7	Controlled Brownian motion (BM) in 1D with fixed and proportional costs II.	113
3.8	Controlled Brownian motion (BM) in 2D with asymmetric costs.	114
3.9	Controlled Brownian motion (BM) in 4D with asymmetric costs.	115
3.10	Controlled Brownian motion (BM) in 8D with asymmetric costs.	116
3.11	Controlled Brownian motion (BM ^b) in 1D.	117
3.12	Controlled Brownian motion (BM) in 1D with fixed and proportional costs: The effect of insufficient exploration I.	120
3.13	Controlled Brownian motion (BM) in 2D with asymmetric costs: The effect of insufficient exploration II.	121

List of Tables

1.1	Parameters of the branching diffusion (Simulation Study).	20
1.2	Simulation results (Simulation Study).	21
1.3	Relative error (Simulation Study).	21
1.4	Running time (Simulation Study).	22
1.5	Market coefficients (SIB)	29
1.6	Simulation parameters (SIB)	30
2.1	Model parameters <i>Agricultural Production</i>	61
2.2	Model parameters <i>SIR Model</i>	68
2.3	Model parameters <i>Evacuation</i>	75
3.1	Model and numerical parameters (BM) – 1D, no proportional costs.	110
3.2	Model and numerical parameters (BM) – 1D, fixed and proportional costs.	112
3.3	Model and numerical parameters (BM) – 2D.	114
3.4	Model and numerical parameters (BM) – 4D.	115
3.5	Model and numerical parameters (BM) – 8D.	116
3.6	Model and numerical parameters (BM ^p) – 1D.	117

References

- [AB06] C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition, 2006.
- [ABL18] P. Azimzadeh, E. Bayraktar, and G. Labahn. Convergence of implicit schemes for Hamilton-Jacobi-Bellman quasi-variational inequalities. *SIAM Journal on Control and Optimization*, 56(6):3994–4016, 2018.
- [AC18] A. Agarwal and J. Claisse. Branching diffusion representation of semi-linear elliptic PDEs and estimation using Monte Carlo method. Preprint, available at <https://arxiv.org/abs/1704.00328v2>, 2018.
- [AF16] P. Azimzadeh and P. A. Forsyth. Weakly chained matrices, policy iteration, and impulse control. *SIAM Journal on Numerical Analysis*, 54(3):1341–1364, 2016.
- [Ahu16] S. Ahuja. Wellposedness of mean field games with common noise under a weak monotonicity condition. *SIAM Journal on Control and Optimization*, 54(1):30–48, 2016.
- [AJ13] A. Agarwal and S. Juneja. Comparing optimal convergence rate of stochastic mesh and least squares method for Bermudan option pricing. In *2013 Winter Simulations Conference (WSC)*. IEEE, 2013.
- [Alv04] L. H. R. Alvarez. Stochastic forest stand value and optimal timber harvesting. *SIAM Journal on Control and Optimization*, 42(6):1972–1993, 2004.
- [AN72] K. B. Athreya and P. E. Ney. *Branching Processes*, volume 196 of *Die Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, 1972.
- [Ana19] Anaconda, Inc. *Anaconda Software Distribution*, 2019. Available at <https://www.anaconda.com/products/individual> (accessed: 2020-05-31, 2126). Documentation of conda package manager available at <https://docs.conda.io/en/latest/> (accessed: 2020-05-31, 2126).
- [ARS17] A. Altarovici, M. Reppen, and H. M. Soner. Optimal consumption and investment with fixed and proportional transaction costs. *SIAM Journal on Control and Optimization*, 55(3):1673–1710, 2017.
- [AV06] L. H. R. Alvarez and J. Virtanen. A class of solvable stochastic dividend optimization problems: On the general impact of flexibility on valuation. *Economic Theory*, 28(2):373–398, 2006.

- [BLL20] C. Bertucci, L. Bertucci, J.-M. Lasry, and P.-L. Lions. Mean field game approach to Bitcoin mining. Preprint, available at <https://arxiv.org/abs/2004.08167v1>, 2020.
- [BC18] E. Bayraktar and A. Cohen. Analysis of a finite state many player game using its master equation. *SIAM Journal on Control and Optimization*, 56(5):3538–3568, 2018.
- [BC19] C. Belak and S. Christensen. Utility maximisation in a factor model with constant and proportional transaction costs. *Finance and Stochastics*, 23(1):29–96, 2019.
- [BCCD19] E. Bayraktar, A. Cecchin, A. Cohen, and F. Delarue. Finite state mean field games with Wright-Fisher common noise. Preprint, available at <https://arxiv.org/abs/1912.06701v1>, 2019.
- [BCL⁺15] M. Bossy, N. Champagnat, H. Leman, S. Maire, L. Violeau, and M. Yvinec. Monte Carlo methods for linear and non-linear Poisson-Boltzmann equation. *ESAIM: Proceedings and Surveys*, 48:420–446, 2015.
- [BCS17] C. Belak, S. Christensen, and F. T. Seifried. A general verification result for stochastic impulse control problems. *SIAM Journal on Control and Optimization*, 55(2):627–649, 2017.
- [BEM13] E. Bayraktar, T. Emmerling, and J.-L. Menaldi. On the impulse control of jump diffusions. *SIAM Journal on Control and Optimization*, 51(3):2612–2637, 2013.
- [Ben08] A. Bensoussan. Impulse control in discrete time. *Georgian Mathematical Journal*, 15(3):439–454, 2008.
- [Ber76] D. P. Bertsekas. *Dynamic Programming and Stochastic Control*, volume 125 of *Mathematics in Science and Engineering*. Academic Press, New York, NY, 1976.
- [BFY13] A. Bensoussan, J. Frehse, and P. Yam. *Mean Field Games and Mean Field Type Control Theory*. SpringerBriefs in Mathematics. Springer, New York, NY, 2013.
- [BG04a] M. Broadie and P. Glasserman. A stochastic mesh method for pricing high-dimensional American options. Working Paper, available at http://www.columbia.edu/~mnb2/broadie/Assets/mesh_working_paper.pdf (accessed: 2020-05-31, 1954); for a published version see [BG04b], 2004.
- [BG04b] M. Broadie and P. Glasserman. A stochastic mesh method for pricing high-dimensional American options. *The Journal of Computational Finance*, 7(4):35–72, 2004.
- [BGH00] M. Broadie, P. Glasserman, and Z. Ha. Pricing American options by simulation using a stochastic mesh with optimized weights. In S. P. Uryasev, editor, *Probabilistic Constrained Optimization: Methodology and Applications*, pages 26–44. Springer US, 2000.
- [BHK14] R. Basna, A. Hilbert, and V. N. Kolokoltsov. An epsilon-nash equilibrium for non-linear Markov games of mean-field-type on finite spaces. *Communications on Stochastic Analysis*, 8(4):449–468, 2014.

- [BHLP19] A. Bachouch, C. Huré, N. Langrené, and H. Pham. Deep neural networks algorithms for stochastic control problems on finite horizon: numerical applications. Preprint, available at <https://arxiv.org/abs/1812.05916v2>, 2019.
- [BHS20a] C. Belak, D. Hoffmann, and F. T. Seifried. Branching diffusions with jumps and valuation with systemic counterparties. Preprint, available at <https://ssrn.com/abstract=3451280>, submitted, 2020.
- [BHS20b] C. Belak, D. Hoffmann, and F. T. Seifried. Continuous-time mean field games with finite state space and common noise. Preprint, available at <https://ssrn.com/abstract=3458336>, submitted, 2020.
- [BJR05] T. R. Bielecki, M. Jeanblanc, and M. Rutkowski. PDE approach to valuation and hedging of credit derivatives. *Quantitative Finance*, 5(3):257–270, 2005.
- [BK98] I. R. C. Buckley and R. Korn. Optimal index tracking under transaction costs and impulse control. *International Journal of Theoretical and Applied Finance*, 1(3):315–330, 1998.
- [BKS13] M. Busch, R. Korn, and F. T. Seifried. Optimal consumption and investment for a large investor: An intensity-based control framework. *Mathematical Finance*, 23(4):687–717, 2013.
- [BKS19] D. Belomestny, M. Kaledin, and J. Schoenmakers. Semi-tractability of optimal stopping problems via a weighted stochastic mesh algorithm. Preprint, available at <https://arxiv.org/abs/1906.09431v1>, 2019.
- [BL84] A. Bensoussan and J.-L. Lions. *Impulse Control and Quasi-Variational Inequalities*. Gauthier-Villars, Paris, 1984.
- [BLL19] C. Bertucci, J.-M. Lasry, and P.-L. Lions. Some remarks on mean field games. *Communications in Partial Differential Equations*, 44(3):205–227, 2019.
- [BLL20] C. Bertucci, J.-M. Lasry, and P.-L. Lions. Master equation for the finite state space planning problem. Preprint, available at <https://arxiv.org/abs/2002.09330v1>, 2020.
- [BLMP19] A. Balata, M. Ludkovski, A. Maheshwari, and J. Palczewski. Statistical learning for probability-constrained stochastic optimal control. Preprint, available at <https://arxiv.org/abs/1905.00107v1>, 2019.
- [BMP13] D. Brigo, M. Morini, and A. Pallavicini. *Counterparty Credit Risk, Collateral and Funding: With Pricing Cases For All Asset Classes*. Wiley, Chichester, UK, 2013.
- [BP00] T. R. Bielecki and S. R. Pliska. Risk sensitive asset management with transaction costs. *Finance and Stochastics*, 4(1):1–33, 2000.
- [BP03] V. Bally and G. Pagès. A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli*, 9(6):1003–1049, 2003.

- [BP17a] A. Balata and J. Palczewski. Regress-later Monte Carlo for optimal control of Markov processes. Preprint, available at <https://arxiv.org/abs/1712.09705v1>, 2017.
- [BP17b] A. Balata and J. Palczewski. Regress-later Monte Carlo for optimal inventory control with applications in energy. Preprint, available at <https://arxiv.org/abs/1703.06461v2>, 2017.
- [BP18] D. Brigo and N. Pede. Examples of wrong-way risk in CVA induced by devaluations on default. In K. Glau, D. Linders, A. Min, M. Scherer, L. Schneider, and R. Zagst, editors, *Innovations in Insurance, Risk- and Asset-Management*, chapter 4, pages 95–115. World Scientific, 2018.
- [BP19] M. Basei and H. Pham. A weak martingale approach to linear-quadratic McKean-Vlasov stochastic control problems. *Journal of Optimization Theory and Applications*, 181(2):347–382, 2019.
- [BPP18] D. Brigo, N. Pede, and A. Petrelli. Multi currency credit default swaps: Quanto effects and FX devaluation jumps. Preprint, available at <https://arxiv.org/abs/1512.07256v2>, 2018.
- [BR11] N. Bäuerle and U. Rieder. *Markov Decision Processes with Applications to Finance*. Universitext. Springer, Berlin, 2011.
- [Bré81] P. Brémaud. *Point Processes and Queues: Martingale Dynamics*. Springer Series in Statistics. Springer, New York, NY, 1981.
- [Bré99] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*, volume 31 of *Texts in Applied Mathematics*. Springer, New York, NY, 1999.
- [BS73] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [BS78] D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*, volume 139 of *Mathematics in Science and Engineering*. Academic Press, New York, NY, 1978.
- [BS05] D. Becherer and M. Schweizer. Classical solutions to reaction-diffusion systems for hedging problems with interacting Itô and point processes. *The Annals of Applied Probability*, 15(2):1111–1144, 2005.
- [BTW19] B. Bouchard, X. Tan, and X. Warin. Numerical approximation of general Lipschitz BSDEs with branching processes. *ESAIM: Proceedings and Surveys*, 65:309–329, 2019.
- [BTWZ17] B. Bouchard, X. Tan, X. Warin, and Y. Zou. Numerical approximation of BSDEs using local polynomial drivers and branching processes. *Monte Carlo Methods and Applications*, 23(4):241–263, 2017.
- [Car13] P. Cardaliaguet. Notes on mean field games. Lecture Notes, available at <https://www.ceremade.dauphine.fr/~cardaliaguet/MFG20130420.pdf> (accessed: 2020-05-31, 2000), 2013.

- [CD13] R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM Journal on Control and Optimization*, 51(4):2705–2734, 2013.
- [CD18a] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, Cham, CH, 2018.
- [CD18b] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*, volume 84 of *Probability Theory and Stochastic Modelling*. Springer, Cham, CH, 2018.
- [CDFP19] A. Cecchin, P. Dai Pra, M. Fischer, and G. Pelino. On the convergence problem in mean field games: A two state model without uniqueness. *SIAM Journal on Control and Optimization*, 57(4):2443–2466, 2019.
- [CDL13] R. Carmona, F. Delarue, and A. Lachapelle. Control of McKean-Vlasov dynamics versus mean field games. *Mathematics and Financial Economics*, 7(2):131–166, 2013.
- [CDL16] R. Carmona, F. Delarue, and D. Lacker. Mean field games with common noise. *The Annals of Probability*, 44(6):3740–3803, 2016.
- [CDL17] R. Carmona, F. Delarue, and D. Lacker. Mean field games of timing and models for bank runs. *Applied Mathematics & Optimization*, 76(1):217–260, 2017.
- [CDLL19] P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions. *The Master Equation and the Convergence Problem in Mean Field Games*, volume 201 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2019.
- [CF18] L. Campi and M. Fischer. N -player games and mean-field games with absorption. *The Annals of Applied Probability*, 28(4):2188–2242, 2018.
- [CF20] A. Cecchin and M. Fischer. Probabilistic approach to finite state mean field games. *Applied Mathematics & Optimization*, 81(2):253–300, 2020. Published online: 2018-03-12.
- [CFS15] R. Carmona, J.-P. Fouque, and L.-H. Sun. Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4):911–933, 2015.
- [CG13] Y.-S. A. Chen and X. Guo. Impulse control of multidimensional jump diffusions in finite time horizon. *SIAM Journal on Control and Optimization*, 51(3):2638–2663, 2013.
- [CG15] P. Carr and S. Ghamami. Derivatives pricing under bilateral counterparty risk. 2015. Preprint, available at <https://ssrn.com/abstract=2602879>; for a published version see [CG17].
- [CG17] P. Carr and S. Ghamami. Derivatives pricing under bilateral counterparty risk. *Journal of Risk*, 20(1):77–107, 2017.
- [Chr14] S. Christensen. On the solution of general impulse control problems using superharmonic functions. *Stochastic Processes and their Applications*, 124(1):709–729, 2014.

- [CJ19a] P. Casgrain and S. Jaimungal. Algorithmic trading in competitive markets with mean field games. *SIAM News*, 52(2), March 2019.
Available at <https://sinews.siam.org/Details-Page/algorithmic-trading-in-competitive-markets-with-mean-field-games> (accessed: 2020-05-31, 1958).
- [CJ19b] P. Casgrain and S. Jaimungal. Mean field games with partial information for algorithmic trading. Preprint, available at <https://arxiv.org/abs/1803.04094v2>, 2019.
- [CL15] R. Carmona and D. Lacker. A probabilistic weak formulation of mean field games and applications. *The Annals of Applied Probability*, 25(3):1189–1231, 2015.
- [CP19a] A. Cecchin and G. Pelino. Convergence, fluctuations and large deviations for finite state mean field games via the master equation. *Stochastic Processes and their Applications*, 129(11):4510–4555, 2019.
- [CP19b] A. Cosso and H. Pham. Zero-sum stochastic differential games of generalized McKean–Vlasov type. *Journal de Mathématiques Pures et Appliquées*, 129:180–212, 2019.
- [CR78] G. M. Constantinides and S. F. Richard. Existence of optimal simple policies for discounted-cost inventory and cash management in continuous time. *Operations Research*, 26(4):620–636, 1978.
- [CS19] S. Christensen and C. Strauch. Nonparametric learning for impulse control problems, 2019. Preprint, available at <https://arxiv.org/abs/1909.09528v1>.
- [CW16] R. Carmona and P. Wang. Finite state mean field games with major and minor players. Preprint, available at <https://arxiv.org/abs/1610.05408v1>, 2016.
- [CW17] R. Carmona and P. Wang. An alternative approach to mean field game with major and minor players, and applications to herders impacts. *Applied Mathematics & Optimization*, 76(1):5–27, 2017.
- [CW18a] R. Carmona and P. Wang. Finite-state contract theory with a principal and a field of agents. Preprint, available at <https://arxiv.org/abs/1808.07942v1>, 2018.
- [CW18b] R. Carmona and P. Wang. A probabilistic approach to extended finite state mean field games. Preprint, available at <https://arxiv.org/abs/1808.07635v1>, 2018.
- [CZ00] A. Cadenillas and F. Zapatero. Classical and impulse stochastic control of the exchange rate using interest rates and reserves. *Mathematical Finance*, 10(2):141–156, 2000.
- [Dav93] M.H.A. Davis. *Markov Models & Optimization*, volume 49 of *Monographs on Statistics and Applied Probability*. Chapman and Hall/CRC, Boca Raton, FL, 1993.
- [DGG16] J. Doncel, N. Gast, and B. Gaujal. Are mean-field games the limits of finite stochastic games? *ACM SIGMETRICS Performance Evaluation Review*, 44(2):18–20, 2016.
- [DGG17] J. Doncel, N. Gast, and B. Gaujal. A mean-field game analysis of SIR dynamics with vaccination. Preprint, available at <https://hal.inria.fr/hal-01496885> (hal-01496885, version 1, accessed: 2020-05-31, 2002), 2017.

- [DGG19] J. Doncel, N. Gast, and B. Gaujal. Discrete mean field games: Existence of equilibria and convergence. *Journal of Dynamics & Games*, 6(3):221–239, 2019.
- [DGLZ12] Q. Du, M. Gunzburger, R. B. Lehoucq, and K. Zhou. Analysis and approximation of nonlocal diffusion problems with volume constraints. *SIAM Review*, 54(4):667–696, 2012.
- [DGLZ13] Q. Du, M. Gunzburger, R. B. Lehoucq, and K. Zhou. A nonlocal vector calculus, nonlocal volume-constrained problems, and nonlocal balance laws. *Mathematical Models and Methods in Applied Sciences*, 23(03):493–540, 2013.
- [DGW10] M. H. A. Davis, X. Guo, and G. Wu. Impulse control of multidimensional jump diffusions. *SIAM Journal on Control and Optimization*, 48(8):5276–5293, 2010.
- [DH96] D. Duffie and M. Huang. Swap rates and credit quality. *The Journal of Finance*, 51(3):921–949, 1996.
- [Die09] P. M. Diesinger. *Stochastic Impulse Control and Asset Allocation with Liquidity Breakdowns*. Dissertation, Technische Universität Kaiserslautern, 2009. Available at <http://nbn-resolving.de/urn:nbn:de:hbz:386-kluedo-23468>.
- [DLR19] F. Delarue, D. Lacker, and K. Ramanan. From the master equation to mean field game limit theory: A central limit theorem. *Electronic Journal of Probability*, 24(51):54 pp., 2019.
- [DS99] D. Duffie and K. Singleton. Modeling term structure of defaultable bonds. *Review of Financial Studies*, 12:687–720, 1999.
- [EH88] J. F. Eastham and K. J. Hastings. Optimal impulse control of portfolios. *Mathematics of Operations Research*, 13(4):588–605, 1988.
- [EHT20] R. Elie, E. Hubert, and G. Turinici. Contact rate epidemic control of COVID-19: an equilibrium view. Preprint, available at <https://arxiv.org/abs/2004.08221v1>, 2020.
- [EIL20] R. Elie, T. Ichiba, and M. Laurière. Large banking systems with default and recovery: A mean field game model. Preprint, available at <https://arxiv.org/abs/2001.10206v1>, 2020.
- [Els09] J. Elstrodt. *Maß- und Integrationstheorie*. Springer, Berlin, 6th, corrected edition, 2009.
- [Eth02] A. Etheridge. *A Course in Financial Calculus*. Cambridge University Press, Cambridge, UK, 2002. 5th (corrected) printing 2008.
- [Fis17] M. Fischer. On the connection between symmetric n -player games and mean field games. *The Annals of Applied Probability*, 27(2):757–810, 2017.
- [Fre85] M. Freidlin. *Functional Integration and Partial Differential Equations*, volume 109 of *The Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1985.

- [FS06] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer, New York, NY, 2nd edition, 2006.
- [GL00] S. Graf and H. Luschgy. *Foundations of Quantization for Probability Distributions*, volume 1730 of *Lecture Notes in Mathematics*. Springer, Berlin, 2000.
- [Gla03] P. Glasserman. *Monte Carlo Methods in Financial Engineering*, volume 53 of *Stochastic Modelling and Applied Probability*. Springer, New York, NY, 2003.
- [GLL11] O. Guéant, J.-M. Lasry, and P.-L. Lions. Mean field games and applications. In R. Carmona, E. Çinlar, I. Ekeland, E. Jouini, J.A. Scheinkman, and N. Touzi, editors, *Paris-Princeton Lectures on Mathematical Finance 2010*, Lecture Notes in Mathematics, pages 205–266. Springer, 2011.
- [GMS10] D. A. Gomes, J. Mohr, and R. R. Souza. Discrete time, finite state space mean field games. *Journal de Mathématiques Pures et Appliquées*, 93(3):308–328, 2010.
- [GMS13] D. A. Gomes, J. Mohr, and R. R. Souza. Continuous time finite state mean field games. *Applied Mathematics & Optimization*, 68(1):99–143, 2013.
- [GMS20] V. Grimm, F. Mengel, and M. Schmidt. Extensions of the SEIR model for the analysis of tailored social distancing and tracing approaches to cope with COVID-19. *medRxiv*, 2020. Preprint, available at <https://www.medrxiv.org/content/10.1101/2020.04.24.20078113v1> (accessed: 2020-05-31, 2003).
- [GS20] D. A. Gomes and J. Saúde. A mean-field game approach to price formation. *Dynamic Games and Applications*, to appear, 2020.
- [Gué11] O. Guéant. From infinity to one: The reduction of some mean field games to a global control problem. Preprint, available at <https://arxiv.org/abs/1110.3441v2>, 2011.
- [Gué15] O. Guéant. Existence and uniqueness result for mean field games with congestion effect on graphs. *Applied Mathematics & Optimization*, 72(2):291–303, 2015.
- [GVW14] D. A. Gomes, R. M. Velho, and M.-T. Wolfram. Socio-economic applications of finite state mean field games. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028), 2014.
- [GW09] X. Guo and G. Wu. Smooth fit principle for impulse control of multidimensional diffusion processes. *SIAM Journal on Control and Optimization*, 48(2):594–617, 2009.
- [GY04] P. Glasserman and B. Yu. Simulation for American options: Regression now or regression later? In H. Niederreiter, editor, *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pages 213–226. Springer, Berlin, 2004.
- [Hal80] J. K. Hale. *Ordinary Differential Equations*. Robert E. Krieger Publishing Company, Inc., Malabar, FL, 2nd edition, 1980.
- [Har63] T. E. Harris. *The Theory of Branching Processes*. Springer, Berlin, 1963.

- [HE16] J. Han and W. E. Deep learning approximation for stochastic control problems. Preprint, available at <https://arxiv.org/abs/1611.07422v1>, 2016.
- [Hen12a] P. Henry-Labordère. Counterparty risk valuation: A marked branching diffusion approach. Preprint, available at <https://arxiv.org/abs/1203.2369v1>; for a published version see [Hen12b], 2012.
- [Hen12b] P. Henry-Labordère. Cutting CVA’s complexity. *Risk Magazine*, 25(7):67–73, July 2012. For an extended version see [Hen12a].
- [Het00] H. W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.
- [HL99] B. Hugu and D. Lando. Swap pricing with two-sided default risk in a rating-based model. *Review of Finance*, 3(3):239–268, 1999.
- [HLM14] T. Harko, F. S. N. Lobo, and M. K. Mak. Exact analytical solutions of the susceptible-infected-recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Applied Mathematics and Computation*, 236:184–194, 2014.
- [HMC06] M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221–252, 2006.
- [HOT⁺19] P. Henry-Labordère, N. Oudjane, X. Tan, N. Touzi, and X. Warin. Branching diffusion representation of semilinear PDEs and Monte Carlo approximation. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 55(1):184–210, 2019.
- [How60] R. A. Howard. *Dynamic Programming and Markov Processes*. Technology Press of the Massachusetts Institute of Technology and John Wiley & Sons, Inc., Cambridge, MA, and New York, NY, 1960.
- [HPBL18] C. Huré, H. Pham, A. Bachouch, and N. Langrené. Deep neural networks algorithms for stochastic control problems on finite horizon, part I: convergence analysis. Preprint, available at <https://arxiv.org/abs/1812.04300v1>, 2018.
- [HST83] J. M. Harrison, T. M. Sellke, and A. J. Taylor. Impulse control of Brownian motion. *Mathematics of Operations Research*, 8(3):454–466, 1983.
- [HSZ17] K. L. Helmes, R. H. Stockbridge, and C. Zhu. Continuous inventory models of diffusion type: Long-term average cost criterion. *The Annals of Applied Probability*, 27(3):1831–1885, 2017.
- [HT18] P. Henry-Labordère and N. Touzi. Branching diffusion representation for nonlinear Cauchy problems and Monte Carlo approximation. Preprint, available at <https://arxiv.org/abs/1801.08794v1>, 2018.
- [HTT14] P. Henry-Labordère, X. Tan, and N. Touzi. A numerical algorithm for a class of BSDEs via the branching process. *Stochastic Processes and Their Applications*, 124(2):1112–1140, 2014.

- [Hun07] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. Documentation of `matplotlib` available at <https://matplotlib.org/3.1.1/contents.html> (accessed: 2020-05-31, 2304).
- [IS06] A. Irle and J. Sass. Optimal portfolio policies under fixed and proportional transaction costs. *Advances in Applied Probability*, 38(4):916–942, 2006.
- [Kal02] O. Kallenberg. *Foundations of Modern Probability*. Probability and Its Applications. Springer, New York, NY, 2nd edition, 2002.
- [KB16] V. N. Kolokoltsov and A. Bensoussan. Mean-field-game model for botnet defense in cyber-security. *Applied Mathematics & Optimization*, 74(3):669–692, 2016.
- [KK01] R. Korn and E. Korn. *Optionsbewertung und Portfolio-Optimierung - Moderne Methoden der Finanzmathematik*. Vieweg+Teubner Verlag, Wiesbaden, 2nd edition, 2001. Nachdruck 2009.
- [KKK10] R. Korn, E. Korn, and G. Kroisandt. *Monte Carlo Methods and Models in Finance and Insurance*. CRC Press, Boca Raton, FL, 2010.
- [KL16] J. Kim and T. Leung. Pricing derivatives with counterparty risk and collateralization: A fixed point approach. *European Journal of Operational Research*, 249(2):525–539, 2016.
- [KM17] V. N. Kolokoltsov and O. A. Malafeyev. Mean-field-game model of corruption. *Dynamic Games and Applications*, 7(1):34–47, 2017.
- [Kor97] R. Korn. Optimal impulse control when control actions have random consequences. *Mathematics of Operations Research*, 22(3):639–667, 1997.
- [Kor98] R. Korn. Portfolio optimisation with strictly positive transaction costs and impulse control. *Finance and Stochastics*, 2(2):85–114, 1998.
- [Kor99] R. Korn. Some applications of impulse control in mathematical finance. *Mathematical Methods of Operations Research*, 50(3):493–518, 1999.
- [Kou02] S. G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002.
- [Kou07] S. G. Kou. Jump-diffusion models for asset pricing in financial engineering. In J. R. Birge and V. Linetsky, editors, *Financial Engineering*, volume 15 of *Handbooks in Operations Research and Management Science*, chapter 2, pages 73 – 116. Elsevier, 2007.
- [KS98] I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer, New York, NY, 2nd edition, 1998.
- [Kus99] S. Kusuoka. A remark on default risk models. In S. Kusuoka and T. Maruyama, editors, *Advances in Mathematical Economics*, volume 1, pages 69–82. Springer, Tokyo, 1999.

- [Kö04] K. Königsberger. *Analysis 1*. Springer, Berlin, 6th edition, 2004.
- [Lac15] D. Lacker. A general characterization of the mean field limit for stochastic differential games. *Probability Theory and Related Fields*, 165(3-4):581–648, 2015.
- [Lac18] D. Lacker. On a strong form of propagation of chaos for McKean-Vlasov equations. *Electronic Communications in Probability*, 23(45):11 pp., 2018.
- [LH09] G. Liu and L. J. Hong. Revisit of stochastic mesh method for pricing American options. *Operations Research Letters*, 37(6):411 – 414, 2009.
- [Liu04] H. Liu. Optimal consumption and investment with transaction costs and multiple risky assets. *The Journal of Finance*, 59(1):289–338, 2004.
- [LL07] J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- [LLLL16] A. Lachapelle, J.-M. Lasry, C.-A. Lehalle, and P.-L. Lions. Efficiency of the price formation process in presence of high frequency participants: a mean field game analysis. *Mathematics and Financial Economics*, 10(3):223–262, 2016.
- [LS01] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001.
- [LT15] L. Laguzet and G. Turinici. Individual vaccination as Nash equilibrium in a SIR model with application to the 2009-2010 Influenza A (H1N1) epidemic in France. *Bulletin of Mathematical Biology*, 77(10):1955–1984, 2015.
- [LW11] A. Lachapelle and M.-T. Wolfram. On a mean field game approach modeling congestion and aversion in pedestrian crowds. *Transportation Research Part B: Methodological*, 45(10):1572–1589, 2011.
- [MA11] K. J. Millman and M. Aivazis. Python for scientists and engineers. *Computing in Science & Engineering*, 13(2):9–12, 2011.
- [Mat19] The Mathworks, Inc., Natick, MA. MATLAB R2019b (v9.7.0.1190202), 2019. Documentation available at <https://de.mathworks.com/help/matlab/> (accessed: 2020-05-31, 2046). Our MATLAB programs make use of linked lists imported from JAVA (JDK 8u202, v1.8.0_202-b08) via “import java.util.LinkedList” (documentation available at <https://docs.oracle.com/javase/8/docs/api/java/util/LinkedList.html> (accessed: 2020-05-31, 2046)).
- [McK75] H. P. McKean. Application of Brownian motion to the equation of Kolmogorov-Petrovskii-Piskunov. *Communications on Pure and Applied Mathematics*, 28(3):323–331, 1975.
- [Med07] P. Medvegyev. *Stochastic Integration Theory*, volume 14 of *Oxford Graduate Texts in Mathematics*. Oxford University Press Inc., New York, NY, 2007. Reprint 2009.

- [Men80] J. L. Menaldi. On the optimal impulse control problem for degenerate diffusions. *SIAM Journal on Control and Optimization*, 18(6):722–739, 1980.
- [Men87] J.-L. Menaldi. Optimal impulse control problems for degenerate diffusions with jumps. *Acta Applicandae Mathematicae*, 8(2):165–198, 1987.
- [Mer73] R. C. Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183, 1973.
- [Mil17] J. C. Miller. Mathematical models of SIR disease spread with combined non-sexual and sexual transmission routes. *Infectious Disease Modelling*, 2(1):35–55, 2017.
- [ML15a] F. Mercurio and M. Li. Jumping with default: Wrong-way risk modeling for credit valuation adjustment. 2015. Preprint, available at <https://ssrn.com/abstract=2605648>; for a published version see [ML15b].
- [ML15b] F. Mercurio and M. Li. Jumping with default: Wrong-way risk modeling for CVA. *Risk*, pages 58–63, December 2015.
- [MØ98] G. Mundaca and B. Øksendal. Optimal stochastic intervention control with application to the exchange rate. *Journal of Mathematical Economics*, 29(2):225–243, 1998.
- [MP19a] E. Miller and H. Pham. Linear-quadratic McKean-Vlasov stochastic differential games. In G. Yin and Q. Zhang, editors, *Modeling, Stochastic Control, Optimization, and Applications (The IMA Volumes in Mathematics and its Applications)*, volume 164, pages 451–481. Springer International Publishing, Cham, CH, 2019.
- [MP19b] M. Motte and H. Pham. Mean-field Markov decision processes with common noise and open-loop controls. Preprint, available at <https://arxiv.org/abs/1912.07883v1>, 2019.
- [Nas50] J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [Neu20] B. A. Neumann. Stationary equilibria of mean field games with finite state and action space. *Dynamic Games and Applications*, to appear, 2020.
- [Nor97] J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK, 1997.
- [Nut18] M. Nutz. A mean field game of optimal stopping. *SIAM Journal on Control and Optimization*, 56(2):1206–1221, 2018.
- [Oli06] T. E. Oliphant. *Guide to NumPy*. 2006. E-book, available at <https://web.mit.edu/dvp/Public/numpybook.pdf> (accessed: 2020-05-31, 1949). Documentation of NumPy available at <https://numpy.org/doc/1.17/> (accessed: 2020-05-31, 1949).
- [Oli07] Travis E. Oliphant. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007.

- [ØS02] B. Øksendal and A. Sulem. Optimal consumption and portfolio with both fixed and proportional transaction costs. *SIAM Journal on Control and Optimization*, 40(6):1765–1790, 2002.
- [ØS05] B. Øksendal and A. Sulem. *Applied Stochastic Control of Jump Diffusions*. Universitext. Springer, Berlin, 2005.
- [PG07] F. Perez and B. E. Granger. IPython: A system for interactive scientific computing. *Computing in Science & Engineering*, 9(3):21–29, 2007. Documentation of IPython available at <https://ipython.readthedocs.io/en/7.x/> (accessed: 2020-05-31, 2021).
- [Pow07] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics. Wiley-Interscience, Hoboken, NJ, 2007.
- [PPP04a] G. Pagès, H. Pham, and J. Printems. Optimal quantization methods and applications to numerical problems in finance. In S. T. Rachev, editor, *Handbook of Computational and Numerical Methods in Finance*, pages 253–297. Birkhäuser Boston, Boston, MA, 2004.
- [PPP04b] G. Pagès, H. Pham, and J. Printems. An optimal Markovian quantization algorithm for multi-dimensional stochastic control problems. *Stochastics and Dynamics*, 04(04):501–545, 2004.
- [PR14] E. Pardoux and A. Răşcanu. *Stochastic Differential Equations, Backward SDEs, Partial Differential Equations*, volume 69 of *Stochastic Modelling and Applied Probability*. Springer, Cham, CH, 2014.
- [Pro04] P. E. Protter. *Stochastic Integration and Differential Equations*, volume 21 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, 2nd edition, 2004.
- [PS07] J. Palczewski and L. Stettner. Impulsive control of portfolios. *Applied Mathematics and Optimization*, 56(1):67–103, 2007.
- [PS13] M. Pykhtin and A. Sokol. Exposure under systemic risk impact. *Risk*, pages 88–93, September 2013.
- [Put94] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability & Mathematical Statistics: Applied Probability and Statistics Section. John Wiley & Sons, Inc., Hoboken, NJ, 1994.
- [PW17] H. Pham and X. Wei. Dynamic programming for optimal control of stochastic McKean-Vlasov dynamics. *SIAM Journal on Control and Optimization*, 55(2):1069–1101, 2017.
- [Pyt19] Python Software Foundation. *The Python Language Reference*, 2019. Available at <https://docs.python.org/3.7/reference/> (accessed: 2020-05-31, 2011). Documentation of PYTHON 3.7 available at <https://docs.python.org/3.7/> (accessed: 2020-05-31, 2011).

- [RRM10] A. Rasulov, G. Raimova, and M. Mascagni. Monte Carlo solution of Cauchy problem for a nonlinear parabolic equation. *Mathematics and Computers in Simulation*, 80(6):1118–1123, 2010.
- [Rud87] W. Rudin. *Real and Complex Analysis*. McGraw-Hill Book Co., New York, NY, 3rd edition, 1987. 2nd printing.
- [Sch74] M. Schäl. A selection theorem for optimization problems. *Archiv der Mathematik*, 25(1):219–224, 1974.
- [Sei11] F. T. Seifried. Stochastic control and financial applications. Lecture Notes – TU Kaiserslautern, 2011.
- [Sey09] R. C. Seydel. Existence and uniqueness of viscosity solutions for QVI associated with impulse control of jump-diffusions. *Stochastic Processes and their Applications*, 119(10):3719–3748, 2009.
- [She68] D. Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM National Conference*, ACM '68, pages 517–524, New York, NY, 1968. ACM.
- [Shr04] S. Shreve. *Stochastic Calculus for Finance II - Continuous-Time Models*. Springer Finance. Springer, New York, NY, 2004.
- [Sko64] A. V. Skorokhod. Branching diffusion processes. *Theory of Probability & Its Applications*, 9(3):445–449, 1964. (translated from Russian into English by B. Seckler).
- [SL89] N. L. Stokey and R. E. Lucas. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA, 1989. 7th printing (2004).
- [Spy19] Spyder project contributors. *Spyder: The Scientific Python Development Environment*, 2019. Documentation available at <https://docs.spyder-ide.org/> (accessed: 2020-05-31, 1956).
- [Ste86] L. Stettner. Discrete time adaptive impulsive control theory. *Stochastic Processes and their Applications*, 23(2):177–197, 1986.
- [Sto09] N. L. Stokey. *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press, Princeton, NJ, 2009.
- [SV19] V. Schulz and C. Vollmann. Exploiting multilevel Toeplitz structures in high dimensional nonlocal diffusion. *Computing and Visualization in Science*, 20(1-2):29–46, 2019.
- [Tou13] N. Touzi. *Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE*, volume 29 of *Fields Institute Monographs*. Springer, New York, NY, 2013.
- [TY93] S. Tang and J. Yong. Finite horizon stochastic optimal switching and impulse controls with a viscosity solution approach. *Stochastics and Stochastic Reports*, 45(3-4):145–176, 1993.

- [vCV11] S. van der Walt, S. C. Colbert, and G. Varoquaux. The NumPy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.
- [VGO⁺20] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020. Documentation of SciPy available at <https://docs.scipy.org/doc/scipy-1.3.1/reference/> (accessed: 2020-05-31, 2253).
- [Wal98] W. Walter. *Ordinary Differential Equations (Graduate Texts in Mathematics)*, volume 182 of *Graduate Texts in Mathematics*. Springer, New York, NY, 1998.
- [Wan19] P. Wang. *Finite State Mean Field Games*. Dissertation, Princeton University, Princeton, NJ, 2019. Available at <http://arks.princeton.edu/ark:/88435/dsp01zw12z808g>.
- [War18] X. Warin. Variations on branching methods for non linear PDEs. Preprint, available at <https://arxiv.org/abs/1701.07660v1>, 2018.
- [Wat65] S. Watanabe. On the branching process for Brownian particles with an absorbing boundary. *Journal of Mathematics of Kyoto University*, 4(2):385–398, 1965.
- [WD92] C. J. C. H. Watkins and P. Dayan. Q-Learning. *Machine Learning*, 8(3/4):279–292, 1992.
- [Wil98] Y. Willassen. The stochastic rotation problem: A generalization of Faustmann’s formula to stochastic forest growth. *Journal of Economic Dynamics and Control*, 22(4):573–596, 1998.
- [Zha17] J. Zhang. *Backward Stochastic Differential Equations: From Linear to Fully Nonlinear Theory*, volume 86 of *Probability Theory and Stochastic Modelling*. Springer, New York, NY, 2017.