

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Optimal control under uncertain model and horizon in financial engineering

Permalink

<https://escholarship.org/uc/item/2rd574z6>

ISBN

9798189272158

Author

Gennaro, Alberto

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Optimal control under uncertain model and horizon in financial engineering

by

Alberto Gennaro

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thibaut Mastrolia, Chair

Professor Steven N. Evans

Professor Xin Guo

Spring 2026

Optimal control under uncertain model and horizon in financial engineering

Copyright 2026
by
Alberto Gennaro

Abstract

Optimal control under uncertain model and horizon in financial engineering

by

Alberto Gennaro

Doctor of Philosophy in Engineering – Industrial Engineering and Operations Research

University of California, Berkeley

Professor Thibaut Mastrolia, Chair

This dissertation studies stochastic optimal control problems in which time, information, and model specification are uncertain. This manuscript can be broadly divided into two different parts. The first part develops second-order backward stochastic differential equations with an exogenous random horizon and applies them to control problems under volatility ambiguity and erratic terminal times. As a real world application, we use the theory developed in a principal agent setting for portfolio delegation when the investment horizon is random, linking continuous-time principal-agent methods with default-time modeling and numerical approximation. The second part is related to model uncertainty: the machinery of exact optimal control heavily depends on model assumptions, especially in non-markovian setting where the (2)bsde theory is deployed. A data driven alternative is provided by the signature transform. We therefore study another well known financial engineering problem, the market making, and we develop signature-based methods for approximating controls in limit-order-book environments. Together, these chapters show how modern stochastic control combines martingale methods, weak formulations, backward equations, and approximation tools to treat financial problems where both the dynamics and the horizon of optimization are uncertain.

“Per la mia famiglia, siete la mia colonna.”

— Alberto Gennaro

Acknowledgments

This dissertation closes a chapter of my life, and it would not exist without the people who shaped my last five years in California. I owe them more than I can write here, but I want to record some of it.

First of all, I am deeply grateful to my advisor, Professor Thibaut Mastrolia, for taking me on as a doctoral student and for the years of patient, generous mentorship that followed. When we met, I was facing a challenging time, and I will never forget the belief, encouragement and constant support you have given me. You gave me the mathematical tools and the questions that have driven this thesis, the freedom to pursue them in my own way, and the precision and intellectual honesty without which any of the results that follow would have remained vague intuitions. Your supervision style is everything I have ever envisioned in a great mentor: freedom to explore, paired with constant availability (and coffees) and the discipline of making our work 1% better every day.

I would like to thank the members of my dissertation committee, Professor Steven N. Evans and Professor Xin Guo, for the time you have invested in reading this work and for the conversations that improved it. Your comments on earlier presentations of these results were both incisive and kind, and I greatly enjoyed your classes and the seminars you organized.

I have been very fortunate to do this work at the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. I am grateful to the IEOR faculty whose courses and seminars shaped the early years of this PhD, to the administrative staff who quietly make the department function, and to my fellow doctoral students for staying in the trenches with me in Etcheverry 4176, pretending the picture of the Golden Gate on the wall was an actual view of the Bay. In particular, I want to thank my office mate Denise for endless conversations on the most random topics; Grace, for being a great co-TA, co-Ambassador, co-Intern and co-complainer (!); and my whole research group — William, Hao, Isabel and Julius. As the oldest member of the team, I hope I left a positive impact on you while being your TA; I know I learned a great deal from you when we took classes together and in our conversations.

I want to thank my friends: the ones I left (together with a piece of my heart) when I flew away from Turin; the ones I met at International House, a place I learned to call home for the better part of this journey and which gave me friendships that will survive any geographical barrier; and the ones I met on the tennis courts (thanks Sergio, for being an incredible flatmate, a great mathematician, and a tennis player slightly worse than I am).

An exhaustive list would make this thesis a thousand pages long; you know who you are.

In particular, I want to thank the brothers-for-life I made along the way: thank you William, thank you Mehdi, thank you Jorge, and thank you Ughets. You are the reason I am graduating: you kept me sane, you made me insane sometimes, but you were always there for me, for adventures, crazy travels, chill nights, listening to my complaints, hosting me. Life would not have been the same without you in these five years.

To Sophia: I tried to beat you, but you were faster! You make me want to be a better person; your determination, your work ethic and your passion kept pushing me, and your love and support cross the ocean at the speed of light. You mean the world to me and I cannot wait to see what life has in store for us next.

To my family. Leaving you was the toughest decision I have ever made, and I do not think I know the words to explain how hard it was. To my brother Carlo: I am so proud of what you are achieving and I hope to have pushed you to be better at least as much as you pushed me. To my parents Luisa and Vittorio: thank you for your unwavering support and love, for all the time you worried about me, for all the phone calls. And to my grandparents Paola, Carlo and Marika: I know I spent a lot of important years far away from you, and I hope to have made you proud.

Finally, I want to thank God: for the opportunity to grow as a human being, for giving me faith in the tough moments, for giving me health and the privilege of a great education, for the chance to travel, to discover new countries and cultures, and to meet new people, for everything and everyone He has put on my path, because that is the reason I am who I am.

Alberto Gennaro
Berkeley, Spring 2026

Contents

Abstract	1
Acknowledgments	ii
Unified notation	ix
1 A pedagogical introduction to continuous-time stochastic optimal control	1
1.1 Overview and organizing theme	1
1.2 Strong and weak formulations: Pontryagin versus martingale optimality	3
1.2.1 The strong formulation and Pontryagin's maximum principle	3
1.2.2 The weak formulation and the change of measure	4
1.3 BSDEs as value processes	6
1.3.1 Main results: existence, uniqueness and comparison	6
1.4 From BSDEs to 2BSDEs: volatility ambiguity and volatility control	9
1.4.1 The aggregation principle behind 2BSDEs	10
1.5 Random horizons and default times	11
1.5.1 Enlargement of filtrations and the density hypothesis	11
1.5.2 The before/after-default decomposition for BSDEs	12
1.6 Portfolio delegation and principal-agent control	13
1.7 Market making and signature-based approximation	14
1.8 How the chapters fit together	15
2 2BSDEs with uncertain horizon and stochastic control in erratic environments	16
2.1 Introduction	16
2.1.1 Backward Stochastic Differential Equations and stochastic control	16
2.1.2 Second order BSDE and volatility-controlled or ambiguous optimization	17
2.1.3 BSDE with erratic horizon	18
2.1.4 Contribution: 2BSDE with erratic horizon and erratic stochastic control	19
2.2 Mathematical framework and Definitions	21
2.2.1 Probabilistic framework	21
2.2.2 Enlargement of filtration	22

2.2.3	Spaces	26
2.2.4	Formulation and Definition	28
2.3	Existence, uniqueness and comparison	30
2.3.1	Existence via auxiliary 2BSDE	30
2.3.2	Uniqueness of solution to Lipschitz 2BSDE with erratic horizon	34
2.3.3	Comparison theorem for Lipschitz 2BSDE with erratic horizon	38
2.3.4	Markovian 2BSDE with erratic horizon and connection with fully non linear PDE	39
2.4	Application to stochastic control in erratic environments	41
2.4.1	A controlled model to unify volatility controlled and ambiguity	41
2.4.2	Supporting 2BSDEs and driver definitions	45
2.4.3	The volatility control case: from full control to 2BSDE	47
2.4.4	The volatility ambiguity case: from adversarial nature to 2BSDE	47
2.5	Discussion and outlook	50
3	Portfolio delegation under random horizon	51
3.1	Introduction	51
3.2	The model and the optimization problem	56
3.2.1	Risky assets and portfolio dynamics	56
3.2.2	Default time and enlargement of filtration	57
3.2.3	Admissible strategy and contracts	58
3.2.4	Delegated portfolio management and bi-level stochastic programming	62
3.3	Optimal contract and investment strategy	63
3.3.1	Optimal investment with random horizon: solving the agent problem	63
3.3.2	The optimal contract and verification results	66
3.4	Numerical solutions	77
3.5	From portfolio delegation to market making under default	88
4	Bridge: market making under a random terminal horizon	90
4.1	The market-making problem	90
4.1.1	Stylised facts and the dealer's setup	90
4.1.2	The Avellaneda–Stoikov benchmark	91
4.1.3	Bounded inventory: the GLFT linearisation	92
4.2	Random terminal horizon	93
4.2.1	Method 1 —running-density formulation with the GLFT ansatz	94
4.2.2	Method 2 —direct AS HJB and an equivalent change of variable	96
4.2.3	Scalar warmup: zero inventory, decoupled from the system	96

4.2.4	Numerical verification	97
4.3	Open points and connection with the next chapter	98
5	Signature methods in optimal market making	100
5.1	Introduction	100
5.2	Related literature	102
5.3	Preliminaries	103
5.3.1	Notation	103
5.3.2	Signature of a càdlàg semimartingale	104
5.4	Signature market-making optimization	105
5.4.1	Financial modeling and optimal quotes	105
5.4.2	Signature methods	106
5.5	Mean–variance objective and controlled dynamics	110
5.5.1	Signature-linear feedback and shuffle reduction	111
5.5.2	Semi-linear structure under the controlled law	112
5.5.3	Likelihood-ratio gradient and the role of the intensity model	112
5.5.4	Pseudo-algorithm: likelihood-ratio optimization with tilde/shuffle payoff	115
5.6	Numerical experiments	116
5.6.1	Setup and benchmarks	116
5.6.2	Policy optimization versus expert approximation	116
5.6.3	Choice of the signature truncation level	118
5.6.4	Poisson order flow	119
5.6.5	Hawkes order flow	121
5.6.6	Discussion	126
5.7	Conclusion	127
A	Extra results from the portfolio-delegation paper	130
A.1	First-best optimal delegated contract	130
A.2	Portfolio manager’s problem and second-order BSDE	132
A.2.1	Probabilistic framework and weak formulation	133
A.2.2	The portfolio manager’s optimization: 2BSDE with erratic termination	134
A.3	Extension to Merton’s portfolio problem with consumption	135
A.4	The agent optimal strategy with power utility	140
	Bibliography	144

List of Figures

3.1	Random default with uniform distribution on $[0, T]$ and exponential distribution. Top left: portfolio value for bounded and unbounded distributions. Top right: average Z^X for bounded and unbounded distributions. Bottom left: average U for bounded and unbounded distributions. Bottom right: optimal investment strategy for bounded and unbounded distributions.	84
3.2	Random default with uniform distribution on $[0, T]$. Top left: average portfolio value for different risk-aversion ratios. Top right: average Z^X for different risk-aversion ratios. Bottom left: average U for different risk-aversion ratios. Bottom right: average optimal investment strategy for different risk-aversion ratios. The blue curve corresponds to $\eta = 0.5 \varsigma$ (agent less risk averse than the principal), the red curve corresponds to $\eta = 2 \varsigma$ (agent more risk averse than the principal), the orange curve corresponds to $\eta = \varsigma$ with a low level of risk aversion, and the green curve corresponds to $\eta = \varsigma$ with a high level of risk aversion.	86
3.3	Random default with 3 different beta distributions on $[0, T]$: in green, we have a left-skewed Beta(4, 2), in blue a right-skewed Beta(2, 4), and in orange a uniform distribution.	87
3.4	Linear incentive on portfolio growth under bounded uniform default on $[0, T]$ (so that the effective horizon is $T \wedge \tau$). Top left: portfolio value for different percentages of the incentive. Top right: average Z^X for different percentages of the incentive. Bottom left: average U for different percentages of the incentive. Bottom right: optimal investment strategy for different percentages of the incentive.	88
4.1	Recovery of GLFT as $\mu \rightarrow 0$	98
5.1	R^2 of the ridge fit to the expert policy and out-of-sample reward of the fitted linear-in-signature rule, as a function of the signature truncation level K . Top panel: Poisson order flow. Bottom panel: Hawkes order flow.	119

5.2	Training history in the Poisson experiment. The top panel reports the average reward, while the lower panels decompose the objective into the execution-revenue component and the inventory-penalty component.	120
5.3	Matched-path comparison in the Poisson experiment. The left panels compare the total spread $\delta^a + \delta^b$ under the learned policy and the expert. The right panels compare the corresponding inventory paths.	121
5.4	Poisson experiment: distribution of the mean–variance reward, left, and terminal execution component, right, for the learned policy and the expert. .	121
5.5	Poisson experiment: inventory-risk statistics for the learned policy and the expert. The learned policy is slightly more conservative in this run.	122
5.6	Training history in the Hawkes experiment. The learned policy stabilizes quickly despite the self-exciting order-flow dynamics.	123
5.7	Hawkes experiment: sample paths of ask and bid half-spreads together with inventory paths. Solid curves are quotes and dotted curves are inventories. .	124
5.8	Hawkes experiment: distribution of the mean–variance reward, left, and terminal execution component, right, for the learned policy and the expert. .	125
5.9	Hawkes experiment: inventory-risk statistics for the learned policy and the expert. In this path-dependent setting, the learned policy achieves a better risk profile.	125
5.10	Hawkes experiment, main configuration $(T, k) = (300, 0.2)$: pathwise comparison of the learned signature policy and the expert benchmark on a representative out-of-sample trajectory.	126
5.11	Hawkes experiment, shorter horizon and stronger self-excitation: the learned policy departs more visibly from a linear-in-inventory rule, reflecting the additional value of path-dependent features in this regime.	127

Unified notation

This page consolidates the most-used symbols across the four research chapters. Where the same symbol denotes different objects in different chapters, the dominant chapter convention is given first and the local exception is flagged. Symbols used only inside a single chapter (e.g. tensor-algebra notation in Chapter 5) are defined locally and not repeated here.

Symbol	Meaning	Chapter / context
<i>Probability and filtration</i>		
$(\Omega, \mathcal{F}, \mathbb{P})$	Reference probability space	all
W	d -dimensional Brownian motion	all
$\mathbb{F} = (\mathcal{F}_t)$	Augmented natural filtration of W	all
τ	Exogenous random termination time	Ch. 2–4
$H_t = \mathbf{1}_{\tau \leq t}$	Default indicator process	Ch. 2–4
$\mathbb{G} = (\mathcal{G}_t)$	Progressive enlargement of $\mathbb{F}$ with τ	Ch. 2–4
$\gamma(t, u)$	Conditional density of τ given $\mathcal{F}_t$	Ch. 1–4
λ^τ, λ	$\mathbb{G}$ -intensity of τ on $\{t < \tau\}$	Ch. 2–4
$\Lambda_t = \int_0^t \lambda_s^\tau ds$	Cumulative compensator of τ	Ch. 2
$\mathcal{P}_0$	Non-dominated semimartingale measure family	Ch. 2
$\hat{a}$	Pathwise quadratic-variation density of canonical X	Ch. 2
<i>BSDE/2BSDE objects</i>		
(Y, Z, U)	BSDE solution: value, Brownian sensitivity, jump mark	all
$K^\mathbb{P}$	Minimal nondecreasing process aggregating across $\mathcal{P}_0$	Ch. 2
F	BSDE driver / generator	all (Ch. 2–3)
ξ	Terminal payoff of the BSDE	Ch. 2–3
ξ^a, ξ^b	Post- τ / pre- τ pieces of ξ in the decomposition	Ch. 2

Symbol	Meaning	Chapter / context
<i>Control</i>		
$\alpha \in A$	Drift control / agent's action	Ch. 1–3
$\beta \in B$	Volatility control / Nature's scenario	Ch. 2
π	Portfolio strategy of the agent	Ch. 3
α (benchmark)	Tracked benchmark strategy in the PA problem	Ch. 3 only
δ^a, δ^b	Dealer's ask/bid half-spreads	Ch. 4, 5
T	Deterministic horizon	all
$T \wedge \tau$	Effective (random) horizon	Ch. 2–4
<i>Financial state</i>		
X^π	Wealth process under strategy π	Ch. 3
Q	Dealer's inventory	Ch. 4, 5
S	Mid-price process	Ch. 4, 5
N^a, N^b	Buy/sell market-order point processes	Ch. 4, 5
η	CARA risk-aversion coefficient	Ch. 3–5
U^A, U^P	Agent / principal utility functions	Ch. 3
R	Agent's reservation utility / participation level	Ch. 3
<i>Numerics</i>		
v, ϕ	Value function (HJB unknown)	Ch. 3–4
$\hat{z}, \hat{z}^x, \hat{\Gamma}, \hat{U}$	Optimizers in the Hamiltonian (PA HJB)	Ch. 3
ρ_m	Learning-rate schedule (signature MM algorithm)	Ch. 5
$\mathbb{Y}_t$	Truncated signature of the augmented path up to t	Ch. 5

Notes on collisions.

- The symbol α denotes (i) the drift control in Chapters 1–2, (ii) the agent's investment strategy in Chapter 3, and (iii) the benchmark strategy in Chapter 3 when an index-

tracking penalty $\frac{\varepsilon}{2}\|\pi - \alpha\|^2$ is in force. Local context disambiguates.

- The CARA risk-aversion coefficient was originally written γ in the published version of the signature-market-making paper that underlies Chapter 5; throughout this thesis we have renamed it to η to match Chapters 3–4. The conditional density of τ keeps the name $\gamma(t, u)$ in all chapters in which it appears.
- The learning rate of the policy-gradient algorithm in Chapter 5 (Algorithm 3) was originally written η_m ; we have renamed it to ρ_m to avoid the clash with the CARA η inside the same chapter.

A pedagogical introduction to continuous-time stochastic optimal control

1.1 Overview and organizing theme

The unifying theme of this dissertation is the question of *which classical assumption is the binding constraint* in a continuous-time stochastic control problem, and how to relax it without losing tractability. The textbook formulation of stochastic optimal control rests on three pillars: (i) the state is a low-dimensional Markov process; (ii) the family of admissible probability measures is dominated by a single reference measure; (iii) the terminal date is deterministic and known in advance. Each pillar can fail in financial applications, and each failure forces a different change in the underlying mathematical machinery. This dissertation studies three such failures, in increasing distance from the classical setting:

- **Deterministic horizon $\rightarrow$ exogenous random termination.** The controller no longer chooses when to stop; an external default time τ ends the problem. The natural object becomes a backward equation on the random interval $[0, T \wedge \tau]$, and well-posedness requires enlargement-of-filtration tools and a density hypothesis on τ . Chapters 2 and 3 treat this failure.
- **Dominated measure family $\rightarrow$ non-dominated family.** As soon as the controller (or Nature) acts on the volatility, the laws of the canonical process become mutually singular and classical BSDE methods fail. The right object is a second-order BSDE that aggregates BSDE values across a non-dominated family of measures. Chapter 2 addresses this failure simultaneously with the random-horizon one.
- **Low-dimensional Markovian state $\rightarrow$ path-valued state.** When the state is the recent history of a high-frequency limit order book, no finite-dimensional summary captures the relevant information. The right object is a path representation rich enough to encode order-flow imbalances and self-excitation; signatures provide one. Chapter 5 pursues this direction as a methodological complement: the analytical machinery of

the first three chapters is exchanged for a data-driven, approximation-based approach. Chapter 4 sits between the two, applying the random-horizon machinery to a market-making problem that motivates the signature approach.

The common object of this dissertation is a decision made continuously in time under uncertainty. A controller observes a stochastic system, chooses an action, and evaluates the resulting terminal payoff or running reward. In financial applications the state may be an investor's wealth, a risky asset, a default indicator, a dealer's inventory, or a path-valued description of a limit order book. The mathematical question is how to characterize an optimal policy when the dynamics are stochastic, when the terminal time may be random, when the volatility may itself be controlled or ambiguous, and when the state variable is too rich to be represented by a low-dimensional Markov process.

A canonical model starts from a controlled state process

$$dX_t = b(t, X_t, \alpha_t)dt + \sigma(t, X_t, \alpha_t)dW_t, \quad X_0 = x, \quad (1.1)$$

where W is a Brownian motion and α is the control. The objective is typically of the form

$$V_0 = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[\int_0^T f(t, X_t, \alpha_t)dt + g(X_T) \right]. \quad (1.2)$$

This deceptively simple formulation hides several different viewpoints. If the probability space, Brownian motion, and filtration are fixed in advance, one speaks of a strong formulation. If the control changes the law of the canonical process, so that optimization is performed over a family of probability measures, one speaks of a weak formulation. If the volatility can be optimized or if the modeler is uncertain about volatility, the relevant probability measures are typically mutually singular and classical BSDE arguments must be replaced by second-order BSDEs. If the terminal date is random, the backward equation must interact with default-time or enlargement-of-filtration methods. Finally, if the state is the history of a high-frequency order book, one needs an approximation language rich enough to encode path dependence; this motivates the use of signatures.

This chapter is meant to provide the conceptual bridge for the four research chapters that follow. Section 1.2 contrasts the Pontryagin maximum principle with the weak formulation and the martingale optimality principle. Section 1.3 introduces BSDEs as adjoint equations and as value processes. Section 1.4 explains why volatility control and volatility ambiguity naturally lead to 2BSDEs. Section 1.5 explains how random horizons fit into this picture, preparing Chapter 2. Section 1.6 introduces the continuous-time principal-agent problem behind portfolio delegation, preparing Chapter 3. Section 1.7 introduces market making and

recalls the Avellaneda–Stoikov and Guéant–Lehalle–Fernandez-Tapia perspective. It then previews the bridge Chapter 4, which embeds a random terminal horizon into a market-making problem, and motivates the signature-based approach of Chapter 5.

Notation and standing assumptions. Throughout this chapter we work on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ carrying a d -dimensional Brownian motion W and, when needed, a random time $\tau : \Omega \rightarrow [0, T] \cup \{+\infty\}$. We denote by $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ the augmentation of the natural filtration of W , and by $\mathbb{G} = (\mathcal{G}_t)_{t \geq 0}$ its progressive enlargement with τ , defined by $\mathcal{G}_t = \bigcap_{s > t} (\mathcal{F}_s \vee \sigma(\tau \wedge s))$. We write $H_t = \mathbf{1}_{\{\tau \leq t\}}$ for the default indicator process. The set $\mathbb{S}_{\mathbb{F}}^p(0, T)$ denotes $\mathbb{F}$ -adapted continuous processes Y with $\mathbb{E}[\sup_{t \leq T} |Y_t|^p] < \infty$, and $\mathbb{H}_{\mathbb{F}}^p(0, T)$ denotes $\mathbb{F}$ -predictable processes Z with $\mathbb{E}[(\int_0^T |Z_s|^2 ds)^{p/2}] < \infty$ (these norm-space symbols are local to this chapter and are not added to the Unified Notation page). Constants depending only on the data are generically denoted C , and may change from line to line.

1.2 Strong and weak formulations: Pontryagin versus martingale optimality

1.2.1 The strong formulation and Pontryagin’s maximum principle

In the strong formulation, the stochastic basis is fixed. Controls are processes adapted to a given filtration, and the Brownian motion remains the same for all admissible controls. The controller changes the coefficients of the stochastic differential equation, but not the underlying probability space. For the controlled diffusion (1.1), one seeks a control α^* maximizing (1.2).

The dynamic programming principle studies this problem through the value function

$$v(t, x) = \sup_{\alpha \in \mathcal{A}} \mathbb{E}_{t,x} \left[\int_t^T f(s, X_s, \alpha_s) ds + g(X_T) \right],$$

which, in a Markovian setting, formally solves the Hamilton–Jacobi–Bellman equation

$$\begin{cases} -\partial_t v(t, x) - \sup_{a \in A} \left\{ b(t, x, a) \cdot Dv(t, x) + \frac{1}{2} \text{Tr} \left[\sigma \sigma^\top(t, x, a) D^2 v(t, x) \right] + f(t, x, a) \right\} = 0 \\ v(T, x) = g(x). \end{cases} \quad (1.3)$$

The HJB equation gives a feedback rule by maximizing the Hamiltonian pointwise in (t, x, Dv, D^2v) .

Pontryagin’s maximum principle gives a complementary local characterization. Instead of

differentiating the value function, one perturbs the control and derives first-order optimality conditions. The adjoint process is a backward equation. In a simple diffusion-control problem, the Hamiltonian has the schematic form

$$H(t, x, a, p, q) = b(t, x, a) \cdot p + \text{Tr}(\sigma^\top(t, x, a) : q) + f(t, x, a),$$

with adjoint variables $p \in \mathbb{R}^n$ and $q \in \mathbb{R}^{n \times d}$, so that the second term is the Frobenius product of σ and q and the optimal control satisfies a maximum condition

$$H(t, X_t^*, \alpha_t^*, p_t, q_t) = \sup_{a \in A} H(t, X_t^*, a, p_t, q_t).$$

The pair (p, q) solves a backward stochastic equation driven by the derivatives of H . Thus the maximum principle naturally leads to forward-backward stochastic differential equations: the state equation runs forward, while the adjoint equation runs backward from a terminal condition involving $Dg(X_T^*)$. This method is particularly natural for open-loop controls and for deriving necessary conditions even when the value function is not smooth; see the classical treatments of stochastic control and maximum principles in [156, 127, 157].

1.2.2 The weak formulation and the change of measure

The weak formulation begins from a different point of view. Rather than fixing a controlled drift in the state equation, one often fixes a canonical process and lets the control select a probability measure under which this canonical process has the desired dynamics. This approach is especially convenient in moral hazard, contract theory, and Stackelberg problems, because the principal may observe only the output process while the agent's action is hidden.

Consider the simple output model

$$dX_t = \sigma_t (\lambda(t, X, \alpha_t) dt + dW_t^\alpha),$$

where W^α is Brownian under a probability measure $\mathbb{P}^\alpha$. If one starts from a reference measure $\mathbb{P}^0$ under which X has zero drift, Girsanov's theorem constructs $\mathbb{P}^\alpha$ by an exponential martingale. The agent's choice of drift is then represented by a choice of measure. Under this measure, the continuation utility can often be represented as the Y -component of a BSDE. The optimal action is obtained by maximizing the generator of this BSDE, and incentive compatibility is expressed by the condition that the agent's chosen action attains this maximizer.

This weak viewpoint is not merely a technical alternative. It changes the objects of the

problem. In the strong formulation, one optimizes controls on a fixed stochastic basis. In the weak formulation, one optimizes over laws. When only the drift changes and Girsanov applies, the corresponding measures are equivalent. When the volatility changes, however, the laws of the canonical process are generally mutually singular. This is the point at which the theory must pass from BSDEs to second-order BSDEs.

The Girsanov change of measure underpinning the weak formulation. Because the weak formulation will be used throughout Chapters 3 and 4, it is worth recalling the precise mechanism. Let $\lambda : [0, T] \times \Omega \times A \rightarrow \mathbb{R}^d$ be a measurable map and let $\alpha = (\alpha_t)_{t \leq T}$ be an $\mathbb{F}$ -predictable A -valued process such that the Doléans-Dade (stochastic) exponential

$$\mathcal{E}_t^\alpha := \exp \left\{ \int_0^t \lambda(s, \omega, \alpha_s) \cdot dW_s - \frac{1}{2} \int_0^t |\lambda(s, \omega, \alpha_s)|^2 ds \right\} = \mathcal{E} \left(\int_0^\cdot \lambda(s, \omega, \alpha_s) \cdot dW_s \right)_t$$

(that is, the unique solution of $d\mathcal{E}_t^\alpha = \mathcal{E}_t^\alpha \lambda(t, \omega, \alpha_t) \cdot dW_t$ with $\mathcal{E}_0^\alpha = 1$) is a true $(\mathbb{F}, \mathbb{P}^0)$ -martingale on $[0, T]$. Define the new measure $d\mathbb{P}^\alpha = \mathcal{E}_T^\alpha d\mathbb{P}^0$. Girsanov's theorem (see, e.g., [81, Theorem III.3.24]) implies that

$$W_t^\alpha := W_t - \int_0^t \lambda(s, \omega, \alpha_s) ds, \quad t \in [0, T],$$

is an $(\mathbb{F}, \mathbb{P}^\alpha)$ -Brownian motion, and the canonical output process satisfies

$$dX_t = \sigma_t \left(\lambda(t, \omega, \alpha_t) dt + dW_t^\alpha \right), \quad \mathbb{P}^\alpha\text{-a.s.}$$

This rewriting allows the controller's action α to be encoded through the choice of measure $\mathbb{P}^\alpha$, while the canonical process X stays the same on the underlying space. As long as σ is independent of α , the family $\{\mathbb{P}^\alpha : \alpha \in \mathcal{A}\}$ is mutually absolutely continuous, and BSDE methods apply. As soon as σ depends on α , a Girsanov change of measure cannot reach all admissible laws: the family becomes non-dominated, and one must work with a “second-order” object aggregating the BSDE solutions across the non-equivalent measures. This is the structural reason why volatility control or volatility ambiguity forces the theory from BSDEs to 2BSDEs (Section 1.4).

For a careful treatment of weak formulations in stochastic control we refer to [156, 157, 154]; the principal-agent literature uses this viewpoint in [54, 115, 53, 45].

1.3 BSDEs as value processes

A backward stochastic differential equation is specified by a terminal condition ξ and a generator F . In its simplest Brownian form, a solution is a pair (Y, Z) satisfying

$$Y_t = \xi + \int_t^T F(s, Y_s, Z_s) ds - \int_t^T Z_s dW_s, \quad 0 \leq t \leq T. \quad (1.4)$$

The process Y_t is interpreted as a time- t value, continuation utility, nonlinear expectation, or risk measure; the process Z_t is the martingale sensitivity. In Markovian settings, if $\xi = g(X_T)$ and $Y_t = u(t, X_t)$, then Z_t is related to the gradient of u , and the BSDE is connected to a semilinear PDE through the nonlinear Feynman–Kac formula. Foundational results include [129, 64, 112, 154].

BSDEs enter stochastic control in two complementary ways. First, in the maximum principle, the adjoint equation is a BSDE. Second, in the weak formulation, the value process itself is often represented by a BSDE. For example, if an agent chooses effort α and receives terminal compensation ξ , the agent’s continuation utility (*not* the principal-agent value, which is obtained only after the principal’s outer optimization) may be represented under $\mathbb{P}^\alpha$ as the Y -component of

$$Y_t = \xi + \int_t^T F(s, Z_s, \alpha_s) ds - \int_t^T Z_s dW_s^\alpha.$$

After maximizing the generator over the action, the same object can characterize the agent’s optimal response. This is the BSDE version of the martingale optimality principle: under every admissible action the candidate value is a supermartingale, and under the optimizer it is a martingale.

For this dissertation, the most important lesson is that BSDEs are a language for dynamic optimization under a fixed or equivalent family of measures. They encode the backward propagation of utility from a terminal condition, the martingale representation of risk, and the first-order or incentive-compatibility conditions of the controller. Chapter 3 uses this idea in a principal-agent setting with a random investment horizon.

1.3.1 Main results: existence, uniqueness and comparison

We collect here, in textbook form, the three classical building blocks of BSDE theory that will be used throughout the dissertation: the Pardoux–Peng a-priori estimate, the Picard-iteration existence and uniqueness theorem, and the comparison theorem. The proofs are sketched at the level needed by the rest of the manuscript; full details can be found in the original article

[129] and in the textbook treatments [64, 157, 154, 36].

Assumption 1.3.1 (Lipschitz generator). The generator $F : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is $\mathbb{F}$ -progressively measurable and there exists $L \geq 0$ such that, $dt \otimes \mathbb{P}$ -a.e.,

$$|F(t, \omega, y, z) - F(t, \omega, y', z')| \leq L(|y - y'| + |z - z'|).$$

The driver also satisfies the integrability condition

$$\mathbb{E} \left[\int_0^T |F(t, \omega, 0, 0)|^2 dt \right] < \infty.$$

The terminal condition $\xi \in L^2(\mathcal{F}_T)$.

Theorem 1.3.2 (A-priori estimate). *Under Assumption 1.3.1, any solution $(Y, Z) \in \mathbb{S}^2 \times \mathbb{H}^2$ of (1.4) satisfies, for some constant $C = C(L, T)$,*

$$\mathbb{E} \left[\sup_{t \leq T} |Y_t|^2 + \int_0^T |Z_s|^2 ds \right] \leq C \mathbb{E} \left[|\xi|^2 + \int_0^T |F(s, 0, 0)|^2 ds \right].$$

Proof sketch. Apply Itô's formula to $e^{\beta t} |Y_t|^2$, where $\beta > 0$ will be fixed below. Using the BSDE (1.4), which is written in the convention $dY_t = -F(t, Y_t, Z_t) dt + Z_t dW_t$,

$$d(e^{\beta t} |Y_t|^2) = e^{\beta t} (\beta |Y_t|^2 + |Z_t|^2) dt - 2e^{\beta t} Y_t F(t, Y_t, Z_t) dt + 2e^{\beta t} Y_t Z_t dW_t.$$

Integrating from t to T , taking conditional expectations, using the Lipschitz bound

$$2yF(s, y, z) \leq 2|y| |F(s, 0, 0)| + 2L|y|(|y| + |z|)$$

and Young's inequality $2|y| |z| \leq \frac{1}{\eta} |y|^2 + \eta |z|^2$ for small η , one absorbs the $|Z|^2$ term on the left provided β is chosen large enough. Burkholder–Davis–Gundy on the local martingale part gives the supremum estimate. See [157, Section 7.2] for the full computation. $\square$

Theorem 1.3.3 (Existence and uniqueness, Pardoux–Peng 1990). *Under Assumption 1.3.1, the BSDE (1.4) admits a unique solution $(Y, Z) \in \mathbb{S}_{\mathbb{F}}^2(0, T) \times \mathbb{H}_{\mathbb{F}}^2(0, T)$.*

Proof sketch. Define the Picard iteration $(Y^{n+1}, Z^{n+1}) = \Phi(Y^n, Z^n)$ on $\mathbb{S}^2 \times \mathbb{H}^2$ by

$$Y_t^{n+1} = \xi + \int_t^T F(s, Y_s^n, Z_s^n) ds - \int_t^T Z_s^{n+1} dW_s,$$

where Z^{n+1} is obtained by applying the martingale representation theorem to the $\mathbb{F}$ -martingale

$$M_t^n := \mathbb{E} \left[\xi + \int_0^T F(s, Y_s^n, Z_s^n) ds \middle| \mathcal{F}_t \right],$$

and Y^{n+1} is then defined by the displayed equation. Using Theorem 1.3.2 on the difference $(Y^{n+1} - Y^n, Z^{n+1} - Z^n)$, one obtains the contraction

$$\|(Y^{n+1} - Y^n, Z^{n+1} - Z^n)\|_\beta \leq \kappa \|(Y^n - Y^{n-1}, Z^n - Z^{n-1})\|_\beta$$

in a β -weighted norm, with $\kappa < 1$ for β large. The Banach fixed-point theorem then gives existence and uniqueness, and the resulting (Y, Z) inherits the integrability of Theorem 1.3.2. $\square$

Theorem 1.3.4 (Comparison theorem). *Let (ξ^i, F^i) , $i = 1, 2$, satisfy Assumption 1.3.1 and let (Y^i, Z^i) be the corresponding solutions. Assume $\xi^1 \leq \xi^2$, $\mathbb{P}$ -a.s., and $F^1(t, Y_t^1, Z_t^1) \leq F^2(t, Y_t^1, Z_t^1)$, $dt \otimes \mathbb{P}$ -a.e. Then $Y_t^1 \leq Y_t^2$, $\mathbb{P}$ -a.s., for every $t \in [0, T]$. If, in addition, either the slack $\text{slack}_s := F^2(s, Y_s^1, Z_s^1) - F^1(s, Y_s^1, Z_s^1)$ is strictly positive on a set of positive $dt \otimes \mathbb{P}$ -measure, or $\xi^1 < \xi^2$ on a set of positive probability and one has the additional structural assumption that ensures strict propagation (e.g. F^i is monotone in z in the sense of Royer [144], or non-trivial slack), then $Y_0^1 < Y_0^2$.*

Proof sketch. Set $(\Delta Y, \Delta Z) = (Y^2 - Y^1, Z^2 - Z^1)$ and $\Delta \xi = \xi^2 - \xi^1 \geq 0$. Subtracting the two BSDEs and using the Lipschitz property,

$$\Delta Y_t = \Delta \xi + \int_t^T (\beta_s \Delta Y_s + \gamma_s \cdot \Delta Z_s + \text{slack}_s) ds - \int_t^T \Delta Z_s dW_s,$$

with β_s, γ_s bounded by L and $\text{slack}_s = F^2(s, Y_s^1, Z_s^1) - F^1(s, Y_s^1, Z_s^1) \geq 0$. A Girsanov change of measure that absorbs γ_s into the Brownian motion, followed by a multiplication by the explicit linear factor $\exp(\int_0^s \beta_r dr)$, gives

$$\Delta Y_t = \tilde{\mathbb{E}} \left[e^{\int_t^T \beta_r dr} \Delta \xi + \int_t^T e^{\int_t^s \beta_r dr} \text{slack}_s ds \middle| \tilde{\mathcal{F}}_t \right] \geq 0.$$

Strict inequality on a set of positive probability propagates by positivity of the conditional expectation. See [64, Theorem 2.2] or [157, Theorem 7.4.1]. $\square$

Remark 1.3.5 (Markovian Feynman–Kac connection). When $\xi = g(X_T)$ and $F(t, \omega, y, z) = f(t, X_t(\omega), y, z)$ for a forward diffusion X solving $dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t$, the unique BSDE solution has the form $(Y_t, Z_t) = (u(t, X_t), v(t, X_t)\sigma(t, X_t))$, where u is the (viscosity)

solution of the semilinear PDE

$$\partial_t u + \mathcal{L}u + f(t, x, u, \sigma^\top(t, x) Du) = 0, \quad u(T, x) = g(x),$$

with $\mathcal{L}u = \frac{1}{2}\text{Tr}(\sigma\sigma^\top D^2u) + b \cdot Du$. This is the nonlinear Feynman–Kac formula [129, 112, 64]. It is the bridge linking the BSDE and HJB pictures of stochastic control.

1.4 From BSDEs to 2BSDEs: volatility ambiguity and volatility control

When the volatility is controlled, the quadratic-variation density of the canonical process depends on the control. In continuous time this typically changes the probability measure in a mutually singular way: two Brownian laws with different volatility processes do not generally admit a common equivalent reference measure. Classical Girsanov arguments therefore cannot reduce volatility control to a drift change.

Second-order BSDEs provide the corresponding backward language. A schematic 2BSDE is

$$Y_t = \xi + \int_t^T \widehat{F}_s(Y_s, Z_s, \widehat{a}_s) ds - \int_t^T Z_s dX_s + K_T^\mathbb{P} - K_t^\mathbb{P}, \quad \mathbb{P}\text{-a.s. for every } \mathbb{P} \in \mathcal{P}, \quad (1.5)$$

where $\widehat{F}_s(y, z, a) := F(s, \omega, y, a^{1/2}z, a)$ denotes the generator F pre-composed with the pathwise volatility density $\widehat{a}$ (so that $\widehat{F}$ depends on (Y, Z) through the volatility-adjusted Brownian sensitivity), $\mathcal{P}$ is a non-dominated family of probability measures, $\widehat{a}$ is the pathwise density of the quadratic variation of X , and $K^\mathbb{P}$ is a nondecreasing process satisfying a minimality condition. The process K is the additional term that makes one process Y aggregate the family of BSDE values obtained under all measures in $\mathcal{P}$.

There are two closely related interpretations. In volatility control, the controller chooses the volatility and hence the measure. In volatility ambiguity, Nature chooses an adverse volatility scenario and the decision maker optimizes against the worst case. In both interpretations, the value is a nonlinear expectation over a non-dominated set of measures. This is why the equation is second order: in a Markovian setting, it corresponds to a fully nonlinear PDE, not merely a semilinear one. The well-posedness theory of 2BSDEs and its links with stochastic control are developed in works such as [44, 150, 94, 136, 107, 59].

This perspective is the conceptual entry point to Chapter 2. That chapter extends the 2BSDE framework to an exogenous random horizon and studies applications to erratic stochastic control, where the controller faces both uncertainty in the model and uncertainty

in the time at which the payoff is realized.

1.4.1 The aggregation principle behind 2BSDEs

Why does a single equation (1.5) capture a family of mutually singular models? The answer is the *universal aggregation* property of the canonical process X under the family $\mathcal{P}$. Let $\mathcal{P}_0$ be a class of semimartingale measures on $\Omega = \mathbb{C}([0, T]; \mathbb{R}^d)$, all making the canonical process X a continuous semimartingale with predictable quadratic variation density $\hat{a}$. Following [92, 125, 150], define $\hat{a}$ pathwise as the Karandikar limit of sums of squared increments along a fixed sequence $(\pi^n)_{n \geq 1}$ of refining partitions of $[0, T]$,

$$\hat{a}_t := \limsup_{n \rightarrow \infty} \frac{1}{|\pi_t^n|} \sum_{t_i \in \pi^n, t_i \leq t} (X_{t_{i+1} \wedge t} - X_{t_i \wedge t})(X_{t_{i+1} \wedge t} - X_{t_i \wedge t})^\top,$$

where $|\pi_t^n|$ is the mesh size of π^n on $[0, t]$. This construction is purely pathwise and does not invoke any particular $\mathbb{P}$; the resulting $\hat{a}$ exists pathwise and equals the canonical density $d\langle X \rangle_t/dt$ $\mathbb{P}$ -a.s. under every $\mathbb{P} \in \mathcal{P}_0$. Hence a single $\mathbb{F}$ -progressive process $\hat{a}$ aggregates the densities across the non-dominated family. Combined with a quasi-sure version of the martingale representation theorem [151, Theorem 6.5], this allows one to write a single process (Y, Z, K) that solves an integral identity $\mathbb{P}$ -a.s. for every $\mathbb{P} \in \mathcal{P}_0$ simultaneously.

The minimality condition $\text{ess inf}_{\mathbb{P}' \in \mathcal{P}_0(\mathbb{P}, t^+)} \mathbb{E}^{\mathbb{P}'}[K_T - K_t | \mathcal{F}_t^+] = 0$, where $\mathcal{P}_0(\mathbb{P}, t^+)$ denotes the subset of measures in $\mathcal{P}_0$ that agree with $\mathbb{P}$ on $\mathcal{F}_{t^+}$, selects, among all candidate decompositions, the one whose increasing process has the smallest possible expectation. It plays the role played by the martingale optimality principle in classical control: Y is the smallest aggregator that dominates the BSDE values under every member of $\mathcal{P}_0$, and K is the integrated penalty paid when one moves between models. The interpretation is symmetric in the volatility-control and volatility-ambiguity settings:

- Volatility control: the controller chooses $\hat{a}$ and $K^{\hat{a}}$ vanishes at the optimum;
- Volatility ambiguity: nature chooses $\hat{a}$ adversarially and $K^{\hat{a}}$ measures the misalignment between the worst-case and the actual scenario.

We refer to [150, 136, 59, 107, 108] for the well-posedness theory and to [117, 118, 133] for applications to controlled diffusions and Stackelberg games. Chapter 2 extends this aggregation principle to a setting where the canonical filtration is enlarged by an exogenous random time τ .

1.5 Random horizons and default times

Many financial contracts do not terminate at a deterministic date. A portfolio mandate can end when an investor withdraws capital, an insurance contract ends at death, a credit instrument terminates at default, and a trading strategy may be stopped by a market disruption. Let τ denote such a random time. The payoff is then received at $T \wedge \tau$, and the relevant terminal condition becomes $\xi_{T \wedge \tau}$ or a default-contingent payoff. Throughout this section, when we write $\xi_{T \wedge \tau}$ we treat $\xi = (\xi_t)_{t \in [0, T]}$ as an $\mathbb{F}$ -progressive process evaluated at the random time $T \wedge \tau$; in the Lipschitz BSDE results of Section 1.3.1 the same symbol $\xi \in L^2(\mathcal{F}_T)$ refers instead to a single random variable. We keep both usages and disambiguate by the presence or absence of a time index.

Random horizons introduce two difficulties. The first is informational: τ may not be a stopping time for the Brownian filtration generated by the market factors. One then works in an enlarged filtration in which the default event is observable. The second is analytical: the compensator or intensity of the default time may be singular near the horizon, especially when default is constrained to occur before a terminal date. BSDEs with random horizons and jumps are therefore a natural tool, but their well-posedness requires care.

Beyond the BSDE setting, similar difficulties arise for 2BSDEs, where the random horizon problem is more delicate because the equation must hold simultaneously under a non-dominated family of measures. The jump associated with τ interacts with the second-order term K and with the aggregation of conditional values across models. The first research chapter of the dissertation treats this issue directly. It develops second-order BSDEs with an uncertain, exogenous random horizon and shows how these equations connect to stochastic control problems in which drift, volatility, or Nature's volatility scenario may be optimized.

1.5.1 Enlargement of filtrations and the density hypothesis

Let $\tau : \Omega \rightarrow (0, T]$ be a strictly positive random time. The *progressive enlargement* $\mathbb{G}$ of $\mathbb{F}$ with τ (defined in the standing assumptions on p. 6, $\mathcal{G}_t = \bigcap_{s > t} (\mathcal{F}_s \vee \sigma(\tau \wedge s))$) is the smallest right-continuous filtration containing $\mathbb{F}$ such that τ is a $\mathbb{G}$ -stopping time. Two fundamental tools control how Brownian information is preserved under this enlargement.

1. *Hypothesis (H)* (immersion): every $(\mathbb{F}, \mathbb{P})$ -martingale is a $(\mathbb{G}, \mathbb{P})$ -martingale. This is the strongest possible compatibility between $\mathbb{F}$ and $\mathbb{G}$, and holds for example when τ is constructed by independent inverse-CDF sampling.
2. *Density hypothesis* [82, 96, 87]: the conditional law of τ given $\mathcal{F}_t$ admits a density $\gamma(t, u)$

(see the Unified Notation page),

$$\mathbb{P}[\tau \in du \mid \mathcal{F}_t] = \gamma(t, u) du, \quad u \geq 0,$$

in which case, defining the Azéma survival process $G_t := 1 - \int_0^t \gamma(t, s) ds$, the $\mathbb{G}$ -intensity of τ is

$$\lambda_t^\tau = \frac{\gamma(t, t)}{G_t}, \quad \text{on } \{t < \tau\}.$$

Under the density hypothesis, the $\mathbb{G}$ -compensator of $H_t = \mathbf{1}_{\{\tau \leq t\}}$ is the absolutely continuous process $\Lambda_{t \wedge \tau} := \int_0^{t \wedge \tau} \lambda_s^\tau ds$. The compensated process $M_t = H_t - \Lambda_{t \wedge \tau}$ is a $\mathbb{G}$ -martingale, and *under Hypothesis (H)* (immersion of $\mathbb{F}$ into $\mathbb{G}$) W remains a $\mathbb{G}$ -Brownian motion.

1.5.2 The before/after-default decomposition for BSDEs

The following key result of Kharroubi–Lim and Jeanblanc–Mastrolia [96, Theorem 4.1], [84, Proposition 4.4] reduces a BSDE on $[0, T \wedge \tau]$ with a single jump at τ to a Brownian-driven BSDE on $[0, T]$. We give a clean statement and a sketch of proof.

Theorem 1.5.1 (Decomposition along τ). *Assume the density hypothesis and let (ξ, F) be admissible data for a $\mathbb{G}$ -BSDE on $[0, T \wedge \tau]$ with jump term U :*

$$Y_t = \xi + \int_{t \wedge \tau}^{T \wedge \tau} F(s, Y_s, Z_s, U_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s \cdot dW_s - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s. \quad (1.6)$$

The sign of the driver integral is aligned with the convention of (1.4); in some references, including [96], the opposite sign is used — see also the sign convention used in Theorem 1.3.2. There exist measurable maps (ξ^a, ξ^b, F^b) representing ξ and F before/after τ , such that the Brownian BSDE

$$Y_t^b = \xi^b + \int_t^T F^b(s, Y_s^b, Z_s^b, \xi_s^a - Y_s^b) ds - \int_t^T Z_s^b \cdot dW_s, \quad t \in [0, T],$$

is well-posed and the unique $\mathbb{G}$ -solution of (1.6) is given by

$$Y_t = Y_t^b \mathbf{1}_{\{t < \tau\}} + \xi_\tau^a \mathbf{1}_{\{t \geq \tau\}}, \quad Z_t = Z_t^b \mathbf{1}_{\{t < \tau\}}, \quad U_t = (\xi_t^a - Y_t^b) \mathbf{1}_{\{t \leq \tau\}}.$$

Z is the predictable Brownian sensitivity and is taken on $\{t < \tau\}$; U is the jump-time predictable factor of the jump term $U_t dH_s$ and is therefore evaluated on $\{t \leq \tau\}$, matching the predictable projection of the indicator H .

Proof sketch. On $\{t \geq \tau\}$, the equation (1.6) collapses to $Y_t = \xi_\tau^a$, since the lower and upper

limits of integration both equal τ on this event, so the deterministic-time integrals reduce to the singleton $[\tau, \tau]$ and contribute zero, while H is constant equal to one (so $dH_s = 0$ on $(\tau, T]$). So $Y_t = \xi_\tau^a$ post-default. Using $\mathbb{F}$ -conditional independence under the density hypothesis and the Brownian martingale representation, one shows that Y admits, on $\{t < \tau\}$, an $\mathbb{F}$ -adapted representative Y^b , and similarly for Z, U . Substituting $U_s = \xi_s^a - Y_s^b$ into the generator of (1.6) and using the λ^τ -compensator of H produces the closed Brownian-BSDE identity stated. Existence and uniqueness of Y^b follow from Theorem 1.3.3 applied to the modified generator $F^b(s, y, z, \xi_s^a - y)$, which is Lipschitz in (y, z) as soon as F is Lipschitz in (y, z, u) . Full proof in [96, Section 4]. $\square$

This decomposition is the conceptual workhorse of Chapter 3: it lets one solve a problem on the random interval $[0, T \wedge \tau]$ by solving a Brownian BSDE on $[0, T]$, at the price of a generator that now depends explicitly on the post-default data ξ^a , an extra non-locality in the generator. The 2BSDE counterpart of Theorem 1.5.1, which is one of the main contributions of Chapter 2, requires aggregating this decomposition across a non-dominated family of measures and is the content of [67].

1.6 Portfolio delegation and principal-agent control

Chapter 3 studies a continuous-time delegation problem. A principal delegates portfolio management to an agent. The agent controls the investment strategy, but the principal designs the contract. This is a Stackelberg problem with hidden action: the principal moves first by designing the contract, the agent moves second by choosing strategy and effort, and the principal cannot directly observe the agent's action.

In continuous time, the weak formulation is especially powerful for such problems. The agent's hidden action is encoded through a change of measure. The agent's continuation utility is represented by a BSDE whose generator contains the cost of effort and the dependence of output on the trading strategy. Incentive compatibility requires that the recommended strategy maximize this generator. Once this condition is imposed, the principal optimizes over contracts, often reducing the problem to a control problem over the agent's continuation utility and its martingale sensitivity.

The random horizon is economically central. In a crypto or hedge-fund interpretation, capital may be withdrawn unexpectedly or the market may be disrupted. In an insurance interpretation, the contract may stop at death. The optimal compensation and the induced trading strategy must therefore account not only for market risk but also for horizon risk. The chapter uses this structure to derive optimal contracts and investment strategies, and

then studies numerical solutions through PDE and physics-informed neural-network methods.

Numerically, Chapter 3 solves the resulting HJB system with a physics-informed neural-network (PINN) solver in a actor critic mechanism.

1.7 Market making and signature-based approximation

A market maker supplies liquidity by posting bid and ask quotes. The dealer earns the half-spread on each filled order but accumulates inventory risk when buy and sell flow are imbalanced. The classical continuous-time formulation models the midprice, the dealer’s inventory, and the arrival intensities of market orders as controlled stochastic processes. The control is the pair of bid and ask offsets. The objective balances expected profit against terminal or running inventory penalties.

The Avellaneda–Stoikov framework and the Guéant–Lehalle–Fernandez-Tapia line of work show how this problem can be formulated as a stochastic control problem with point-process order arrivals and inventory constraints; see [9, 70]. In the simplest exponential-utility setting, the HJB equation can be reduced to a system of ordinary differential equations indexed by inventory. This gives an interpretable structure: quotes widen when inventory risk is high, become asymmetric when inventory must be mean-reverted, and depend on order-arrival intensities and volatility.

Modern limit-order-book data are path dependent. The recent order-flow history, imbalance, cancellations, and short-term price movements contain information that is not fully summarized by the current inventory and midprice. Signatures provide a systematic way to turn a path into features while preserving the ordering and iterated-integral structure of the path. The signature of a stream can be viewed as a nonparametric basis for path-dependent functionals. This makes it a natural approximation architecture for controls in high-frequency market making.

Chapter 5 as a methodological complement. The first three chapters of this dissertation work within an analytical paradigm: identify the right backward equation, prove well-posedness, derive an HJB, and use a verification theorem. That paradigm is sharp when the state is genuinely Markovian and the model class is amenable to BSDE techniques. When neither is the case, as in high-frequency market making, where path dependence is intrinsic and the data-generating process is itself uncertain, one must trade analytical control for approximation flexibility. Chapter 5 carries out this trade: it preserves the stochastic-control structure of the classical market-making literature, but replaces the low-dimensional Markovian state by a pathwise feature map (the truncated signature) and the closed-form

HJB by a likelihood-ratio gradient over a parametric policy class. The chapter is therefore not a continuation of the random-horizon programme of Chapters 2–3; it is a methodological alternative that becomes necessary when the binding constraint shifts from the time horizon to the dimensionality of the state.

Chapter 5 uses this idea to study signature-based optimal market making. Conceptually, it keeps the stochastic-control structure of the classical market-making literature but replaces a low-dimensional Markovian state representation with a pathwise feature map. The resulting method is designed for environments in which the dealer’s optimal quote depends on recent order-book trajectories, not only on the current inventory.

1.8 How the chapters fit together

The dissertation moves from theory to applications along a single axis: how to represent value in continuous-time control when the classical Markovian state description is insufficient. Chapter 2 addresses insufficiency of the probability model and the time horizon, using 2BSDEs under random horizons. Chapter 3 addresses insufficiency of observability and incentives, using weak formulation and BSDE methods for a principal-agent problem with default. Chapter 5 addresses insufficiency of finite-dimensional state variables, using signatures to approximate path-dependent controls in market making. The three chapters therefore share a common method: identify the correct dynamic value process, characterize the optimality principle it satisfies, and then adapt the representation to the uncertainty that is central in the financial problem.

Chapter 4 sits between Chapters 3 and 5 as a methodological bridge. It carries the random-horizon machinery of the first two chapters into the market-making setting that motivates the signature approach: the same exogenous default time τ that ends a delegated mandate in Chapter 3 ends the market maker’s mandate in Chapter 4, and the resulting BSDE is the limit of what the Lipschitz framework of Chapter 2 can handle in closed form. Once one has seen this limit, the move to a fully data-driven, path-dependent representation in Chapter 5 is a natural step: the signature replaces the low-dimensional Markovian state of the bridge chapter by a non-parametric encoding of the limit-order-book history, while keeping the same control objective and the same inventory penalties.

2BSDEs with uncertain horizon and stochastic control in erratic environments

This chapter is based on the paper *2BSDE with uncertain horizon and application to stochastic control in erratic environments*, joint work with Thibaut Mastrolia. The paper is currently under revision for Stochastic Processes and Applications.

2.1 Introduction

2.1.1 Backward Stochastic Differential Equations and stochastic control

Backward Stochastic Differential Equations (BSDE for short) are a class of stochastic equations characterized by a terminal condition ξ and a generator function F , which depend on time t , randomness ω and space variables y, z . On a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a Brownian motion W . A solution to a BSDE parametrized by (ξ, F) is a pair of stochastic processes (Y, Z) , adapted to the filtration of the Brownian motion satisfying the following backward equation¹

$$Y_t = \xi - \int_t^T F(s, \omega, Y_s, Z_s) ds - \int_t^T Z_s dW_s, \quad t \leq T, \quad \mathbb{P}\text{-a.s.}$$

A very special case arises when $\xi = g(X_T)$ where X is a solution to a stochastic differential equation and $F(s, \omega, Y_s, Z_s) = F(s, X_s(\omega), Y_s(\omega), Z_s(\omega))$, the corresponding BSDE is set in a Markovian framework. The well-posedness of a Markovian BSDE can thus be reduced to the existence of a solution in a PDE where Y is the solution and Z is its gradient, see, for example, [112]. This type of equation was first introduced by Bismut in [24], proving the existence and uniqueness of the solution when the generator F is assumed to be linear in Y

¹All equations in this introduction are presented informally and without full mathematical rigor. Their purpose is to motivate the study in this paper. In the main body of the work, a proper mathematical treatment is provided.

and Z . It has then been generalized in the seminal paper [129] under Lipschitz conditions in a Markovian setting, then extended to a non-Markovian framework with applications in finance (pricing and hedging of derivatives) in the celebrated article [64] or to more general assumptions to the generator F , like stochastic Lipschitz or quadratic growth [10, 29, 30] or extended to jump processes [33, 141, 139, 128]. Since these foundational studies, BSDEs have been extensively investigated, together with their links with stochastic control and applications in diverse fields of applied mathematics: finance and utility maximization, see for example [64, 143, 76, 41, 121], insurance [56, 61], risk theory [18], games theory [52, 93, 32, 105] or biology [114, 147, 89]. We refer more generally to [156, 154, 36] for more details of the BSDE theory.

In this work, we focus on a stochastic control problem, in a non-Markovian setting in which a diffusion process X is controlled both in drift and volatility:

$$dX_t = \mu(t, X, \alpha_t)dt + \sigma(t, X, \beta_t)dW_t, \quad \nu = (\alpha, \beta).$$

Echoing the analysis presented in [157, Chapter 9] and [156, Chapter 4], different types of control strategies can be considered in stochastic control problems. When the control process is open-loop - meaning it depends on time t and the underlying Brownian motion W - the problem is naturally framed in a *strong formulation*. This setting leads to the study of Forward-Backward Stochastic Differential Equations (FBSDEs for short), and the optimal control can be characterized using the stochastic maximum principle, originally formulated by Pontryagin in the 1950s; see, for instance, [156, 127]. Alternatively, in the case of *closed-loop* (or feedback) control, where the control depends on t and the controlled state process X , the problem is typically treated in a *weak formulation*. This framework often arises in contexts such as Stackelberg games or contract theory [54, Section 4]. In this setting, Girsanov's theorem is employed to shift to a family of equivalent probability measures $\mathbb{P}^\alpha$, under which a new Brownian motion is defined. This change of measure facilitates the analysis of the control problem via BSDE.

2.1.2 Second order BSDE and volatility-controlled or ambiguous optimization

Controlling the volatility becomes more challenging in the weak formulation setting, as it involves changing the probability measure associated with a diffusion process and its volatility. This requires constructing a family of non-equivalent probability measures, denoted by $\mathbb{P}^{\alpha, \beta}$ through a weak formulation and weak solution of the controlled SDE considered. The optimal

control problem then consists of maximizing an objective function over a set of probability measures parameterized by the control variables α and β . Under each such measure, the objective function is either a supermartingale or a submartingale, depending on the nature of the problem. The optimal measure corresponds to the one under which a martingale optimality principle is satisfied. As an extension of the BSDE introduced earlier, we then expect the problem to be reduced to a second-order BSDE (2BSDE for short), with terminal condition ξ and generator F given by

$$Y_t = \xi - \int_t^T \hat{F}_s(Y_s, Z_s) ds - \int_t^T \hat{a}_s^{1/2} Z_s dX_s + K_T - K_t, \quad t \leq T, \quad \mathbb{P}\text{-a.s.}, \forall \mathbb{P} \in \mathcal{P},$$

where $\mathcal{P}$ is a set of non-equivalent probability measures and $\hat{F}$ is the generator of the 2BSDE determined through the universal representation of the quadratic variation $\hat{a}_t$ of X as soon as the family $\mathcal{P}$ satisfies regularity properties, see [92, 125]. A solution to this 2BSDE consists of a triple (Y, Z, K) of stochastic processes where the K process represents the second-order representation, is non-decreasing, and satisfies the minimality condition

$$0 = \operatorname{ess\,inf}_{\mathbb{P}} \mathbb{E}^{\mathbb{P}}[K_T - K_t | \mathcal{F}_t].$$

The existence and uniqueness of solutions to 2BSDEs were first studied in [150, 44], and later extended to more general frameworks in [94, 135, 106, 97], under the assumption that both the terminal condition and the generator of the 2BSDE satisfy strong continuity conditions with respect to ω . An extension to random horizons, where the terminal time is adapted to the diffusion filtration, was developed in [107]. The connections between 2BSDEs and stochastic games or control problems involving controlled volatility or model uncertainty (ambiguity) have been explored in, for example, [117, 118, 133, 120]. Fundamental advances were made in [136], where the regularity assumptions on both the terminal condition and the generator were relaxed using a more general class of probability measures. This was further extended in [107, 78, 59]. These relaxations have been crucial in solving Stackelberg games and stochastic optimization problems involving controlled volatility or ambiguity. See, for example, [54, 115, 74, 108, 116].

2.1.3 BSDE with erratic horizon

In this work, we are interested in stochastic control problems where the horizon is not only random but also unpredictable and erratic, in the sense that it is exogenous and not adapted to the filtration of the controlled process X . To address this difficulty, we rely on the theory of filtration enlargement, which allows us to incorporate additional information into the

available filtration by accounting for the extra information provided by the termination time. This type of problem is particularly relevant in the context of credit risk, where a firm's default is modeled as a random time that is not adapted to the filtration generated by observable market variables. See, for example, [90, 1]. We refer to [87, 82, 152, 26, 7, 100, 71, 95, 35] and the monograph [23, 2] for a non-exhaustive list of works in enlargement of filtration. In a non-Markovian setting, optimization problems involving an erratic time τ have been investigated in [96, 84, 83]. It is proven that one can reduce the optimization to solve a BSDE with an erratic horizon with a single jump of the form:

$$Y_t = \xi - \int_{t \wedge \tau}^{T \wedge \tau} Z_s \cdot dW_s - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s - \int_{t \wedge \tau}^{T \wedge \tau} F(s, Y_s, Z_s, U_s) ds, \quad t \in [0, T], \quad \mathbb{P}\text{-a.s.}, \quad (2.1)$$

where $H_s = \mathbf{1}_{\tau \leq s}$ indicates if a random default occurs or not. Under different assumptions regarding the default time τ , the generator of the associated BSDE may become degenerate. This type of degeneracy has been extensively studied in the literature, particularly in the context of insurance. See, for example, [101, 132, 77, 40, 57, 8, 60, 88]. In particular, [96, Theorem 4.1] or [84, Proposition 4.4] prove that the solution (Y, Z, U) of this BSDE can be reduced to solve a Brownian driven BSDE (possibly with a degenerated generator) given by

$$Y_t^b = \xi^b - \int_t^T F^b(s, Y_s^b, Z_s^b, \xi_s^a - Y_s^b) ds - \int_t^T Z_s^b \cdot dW_s, \quad t \in [0, T],$$

and the solution to (2.1) is given by

$$Y_t = Y_t^b \mathbf{1}_{\{t < \tau\}} + \xi_\tau^a \mathbf{1}_{\{t \geq \tau\}}, \quad Z_t = Z_t^b \mathbf{1}_{\{t \leq \tau\}}, \quad U_t = (\xi_t^a - Y_t^b) \mathbf{1}_{\{t \leq \tau\}},$$

where ξ^b, ξ^a, F^b are representations of ξ and F with respect to the event $\{\tau \leq T\}$. This result is however only related to optimal control problems for either exponential portfolio utility maximization similar to [76] or drift-controlled diffusion. The goal of this study is to extend it to volatility-controlled or ambiguity models.

2.1.4 Contribution: 2BSDE with erratic horizon and erratic stochastic control

Motivation with crypto trading and cyber risk. We define **erratic stochastic control** as a class of control problems in which the decision-making process must respond to highly irregular, unpredictable stochastic environments and potentially abrupt termination events. The cryptocurrency market illustrates erratic stochastic behaviors. In this domain, asset prices are influenced by extremely volatile and unpredictable factors, including random

market noise, sudden regulatory interventions, and unforeseen news events. These dynamics result in behavior that is difficult to capture or forecast using traditional modeling techniques. Panic-driven behaviors have partially led to the cascading collapse of Terra (LUNA) and its associated stablecoin UST in May 2022, which triggered a rapid market-wide reaction [31]. Cryptocurrency markets frequently undergo sudden regime shifts— such as speculative bull runs, dramatic crashes, or fraudulent events as rug pulls. Market participants often operate under conditions of Knightian uncertainty, where the underlying probability distributions are unknown, ill-defined, or even unknowable, see [4]. Tracking 146 Proof-of-Work-based cryptocurrencies that started trading prior to the end of 2014 until December 2018, [68] shows that about 60% of those cryptocurrencies were eventually in default. In cyber-risk, Knightian uncertainty plays a crucial role, as cyber-attacks introduce greater volatility and unpredictable conditions into the system. For example, see [43, 116].

Contribution and structure. In this work, we address the problem of stochastic control in erratic environments by incorporating both control and ambiguity in the volatility of a controlled diffusion, as well as high uncertainty in the termination of the optimization problem. We begin by proposing mathematical models in Section 2.2, where we extend the so-called Density Hypothesis in the theory of enlargement of filtrations to a $\mathcal{P}$ -density assumption that holds for the default time under a set of probability measures. We then investigate the existence, uniqueness, and comparison results for solutions in Section 2.3 to a second-order backward stochastic differential equation (2BSDE) with an erratic horizon of the form:

$$\begin{aligned} Y_t = & \xi - \int_{t \wedge \tau}^{T \wedge \tau} F_s^{\mathbb{P}}(Y_s, \hat{a}_s^{1/2} Z_s, U_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s \cdot dX_s^{c, \mathbb{P}} - \int_{t \wedge \tau}^{T \wedge \tau} dK_s \\ & - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s - \int_{t \wedge \tau}^{T \wedge \tau} dM_s^{\mathbb{P}}, \quad \text{for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., for every } \mathbb{P} \in \mathcal{P}_0. \end{aligned}$$

This section extends existing results on BSDEs with random horizons and 2BSDEs with finite or random termination to the setting of 2BSDEs with an erratic horizon. Our formulation also allows for the aggregation of the U process, in contrast to existing theories on 2BSDEs with jumps; see Remark 2.3.11. We conclude this theoretical section with a connection to fully nonlinear PDE in Section 2.3.4. Finally, we investigate erratic stochastic control in Section 2.4 and explore its connection to 2BSDEs with erratic horizons. We first present a comprehensive and general structure for the erratic control problem in Section 2.4.1, unifying the volatility control and ambiguity models within our framework. We then define the 2BSDE associated with both problems in Section 2.4.2. The volatility control problem is studied in Section 2.4.3, followed by an analysis of the Knightian uncertainty problem in Section 2.4.4.

Methodological contribution. The combination of a 2BSDE with an exogenous random horizon is not, by itself, the contribution: that combination would be empty if the random-horizon BSDE theory of [96, 84, 83] and the second-order BSDE theory of [150, 136, 59] could simply be juxtaposed. They cannot. The principal technical obstruction is that the jump caused by τ at the random horizon $T \wedge \tau$ interacts with the second-order minimality term $K^{\mathbb{P}}$ across a *non-dominated* family of measures $\mathcal{P}_0$: the standard before/after-default decomposition of [96, 84] fixes a single reference measure and uses the Jacod density of τ under it, while the 2BSDE machinery requires a kernel that is consistent *simultaneously* for every $\mathbb{P} \in \mathcal{P}_0$, including those that are mutually singular. Our methodological contribution is the introduction of the $\mathcal{P}_0$ -density Hypothesis 2.2.4 (the “ $\mathcal{P}$ -density assumption” alluded to in the contribution statement above), which forces the conditional density γ of τ to be the same kernel under every measure in $\mathcal{P}_0$ and is what allows us to (i) aggregate the $\mathbb{F}$ -intensity λ across $\mathcal{P}_0$, (ii) aggregate the jump-size process U in the primal 2BSDE (2.3) (Remark 2.3.11), and (iii) carry out the auxiliary 2BSDE construction (Lemma 2.3.2–Theorem 2.3.4) that reduces the random-horizon 2BSDE to a Brownian-driven 2BSDE on $[0, T]$ with a frozen post-default boundary. The $\mathcal{P}_0$ -density assumption is strictly stronger than Jacod’s classical density hypothesis under a fixed measure (see the discussion after Remark 2.2.5), and we exhibit two structurally different families of examples (equivalent measures via Girsanov; $\mathbb{P}$ -invariant law of τ) under which it is satisfied. The connection to fully nonlinear PDE in Section 2.3.4 shows that, despite the jump and the non-Markovian structure, the random horizon does not enter the PDE itself: the dependence on τ is absorbed into a piecewise Feynman–Kac representation, with the singular value of K at τ pushed to a boundary condition. This is what we mean by “erratic stochastic control” as a methodologically distinct framework, as opposed to a relabeling of robust stochastic control: the family $\mathcal{P}_0$ is non-dominated *and* the horizon is exogenous.

2.2 Mathematical framework and Definitions

2.2.1 Probabilistic framework

Let d be a positive integer and set $\Omega := \mathbb{D}([0, T]; \mathbb{R}^d)$ defined by the set of càdlàg paths on $\mathbb{R}^d$ starting at 0, equipped with the Skorohod topology. We denote by X the canonical process on this space such that $X_t(\omega) = \omega_t$ for all $\omega \in \Omega$. We denote by $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ the filtration generated by the canonical process X , and by $\mathbb{F}^+ = (\mathcal{F}_t^+)_{0 \leq t \leq T}$ the right limit of $\mathbb{F}$, defined by

$$\mathcal{F}_t^+ := \bigcap_{s>t} \mathcal{F}_s \quad \text{for all } t \in [0, T), \quad \text{and} \quad \mathcal{F}_T^+ := \mathcal{F}_T.$$

We also define

$$\mathcal{F}_t^{\mathcal{P}} := \bigcap_{\mathbb{P} \in \mathcal{P}} \mathcal{F}_t^{\mathbb{P}}, \quad \mathcal{F}_t^{\mathcal{P}^+} := \mathcal{F}_{t+}^{\mathcal{P}}, \quad \mathbb{F}^{\mathcal{P}^+} := (\mathcal{F}_t^{\mathcal{P}^+})_{t \leq T},$$

where $\mathcal{P}$ is a set of probability measures on Ω .

Definition 2.2.1 (Semi-martingale measures). We define $\mathcal{P}_t^W$ as the set of all semi-martingale measures $\mathbb{P}$ on $(\Omega, \mathcal{F})$, such that $\mathbb{P}$ -a.s.:

1. $(X_s)_{s \in [t, T]}$ is a $(\mathbb{P}, \mathbb{F})$ -semi-martingale admitting the canonical decomposition

$$X_s = X_t + \int_t^s \mu_r^{\mathbb{P}} dr + X_s^{c, \mathbb{P}},$$

where $\mu^{\mathbb{P}}$ is an $\mathbb{F}$ -predictable $\mathbb{R}^d$ -valued process and $X^{c, \mathbb{P}}$ denotes the continuous local martingale part of X under $\mathbb{P}$;

2. the quadratic variation process $(\langle X \rangle_s)_{s \in [t, T]}$ is absolutely continuous in time s with respect to the Lebesgue measure and the process $\hat{a}$ defined by

$$\hat{a}_t := \limsup_{\varepsilon \rightarrow 0} \frac{\langle X \rangle_t - \langle X \rangle_{t-\varepsilon}}{\varepsilon},$$

takes values in $\mathcal{S}_d^{\geq 0}$, $\mathbb{P}$ -a.s., where $\mathcal{S}_d^{\geq 0}$ denotes the space of symmetric and positive semi-definite matrices in dimension d .

In this study, we consider a family $(\mathcal{P}(t, \omega))_{(t, \omega) \in [0, T] \times \Omega}$ of probability measures on $(\Omega, \mathcal{F}_T)$, where $\mathcal{P}(t, \omega) \subset \mathcal{P}_t^W$ for all $(t, \omega) \in [0, T] \times \Omega$. Denote also

$$\mathcal{P}_t := \bigcup_{\omega \in \Omega} \mathcal{P}(t, \omega), \quad \mathcal{P}_0 := \mathcal{P}_{t=0}.$$

The two symbols play distinct roles: $\mathcal{P}(t, \omega)$ is the *local* family used in the dynamic-programming construction (it is the set of measures that agree with the path ω up to time t and live in the canonical semimartingale class beyond t), whereas $\mathcal{P}_0$ is the *global* family on $(\Omega, \mathcal{F}_T)$ relative to which the global well-posedness of the 2BSDE is stated. All standing assumptions, including the density Hypothesis 2.2.4 below and the Nutz–van Handel measurability hypothesis recalled in Section 2.3.1, are imposed on $\mathcal{P}_0$ and inherited by $\mathcal{P}(t, \omega)$ via concatenation.

2.2.2 Enlargement of filtration

The default time is represented by a random variable τ taking values in $\mathbb{R}^+$. We define the default process by $H_t := \mathbf{1}_{\tau \leq t}$. Note that this process is not necessarily $\mathbb{F}$ -measurable,

since the default time is assumed to be potentially exogenous to the system and therefore independent of X . We thus enlarge the available information and so the filtration $\mathbb{F}$ taking into account the information generated by the default time occurrence.

Definition 2.2.2. Let $\mathcal{H}_u := \sigma(H_s, s \in [0, u])$ be the σ -algebra generated by H until time $u \geq 0$. Given a filtrated space $(\Omega, \mathcal{F}_T, \mathbb{F}, \mathbb{P})$, the enlarged filtration

$$\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}, \quad \mathcal{G}_t = \bigcap_{\epsilon > 0} \{\mathcal{F}_{t+\epsilon} \vee \mathcal{H}_{t+\epsilon}\},$$

is the smallest enlargement of $\mathbb{F}$ such that τ is a $\mathbb{G}$ -stopping time.

Remark 2.2.3. H is not $\mathbb{F}$ -measurable but it is $\mathbb{G}$ -measurable stochastic process.

The goal is to ensure that the inaccessible default time τ enables us to enlarge the filtration to $\mathbb{G}$ and transfer the martingale property from $\mathbb{F}$ to $\mathbb{G}$ known as the immersion property or H-hypothesis, see [28, 3]. The first fundamental hypothesis is to set the existence of a (conditional) density for the default time with a certain property.

Hypothesis 2.2.4 ($\mathcal{P}_0$ -density hypothesis). *For every $t \geq 0$, there exists a measurable kernel $\gamma(t, \cdot) : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, common to all $\mathbb{P} \in \mathcal{P}_0$, such that*

$$\mathbb{P}(\tau \geq x \mid \mathcal{F}_t) = \int_x^\infty \gamma(t, u) du, \quad x \geq 0, \mathbb{P} \in \mathcal{P}_0.$$

The kernel admits the two-regime form

$$\gamma(t, u) = \begin{cases} \gamma_{\text{pre}}(t, u), & u > t, \\ \gamma_{\text{pre}}(u, u), & u \leq t, \end{cases}$$

where $\gamma_{\text{pre}}(t, u)$ is, for each fixed $u > 0$, a $(\mathbb{P}, \mathbb{F})$ -martingale in t on $\{t < u\}$ under every $\mathbb{P} \in \mathcal{P}_0$, and the same kernel γ_{pre} represents the $\mathcal{F}_t$ -conditional density of τ simultaneously under all measures of $\mathcal{P}_0$. In particular, $\gamma(t, u) = \gamma(u, u)$ on $\{t \geq u\}$ (the freezing identity), while $\gamma(t, u) > 0$ on $\{u > t\} \cap \{\tau > t\}$.

This assumption is the robust counterpart of the density hypothesis used in the theory of progressive enlargement of filtrations. Compared with the classical formulation under a fixed probability measure, we require the same conditional density γ to work simultaneously for every $\mathbb{P} \in \mathcal{P}_0$. This stability across the model class ensures that the enlargement procedure, the immersion property, and the intensity process below are compatible with the non-dominated family of probabilities.

Remark 2.2.5. The two-regime form of Hypothesis 2.2.4 above is the standard $\mathbb{F}$ -conditional density of τ past its time of observation under Jacod's density hypothesis [87, 63]. The next lemma writes this out explicitly under the H-hypothesis and verifies, in particular, that $\mathbb{P}(\tau > t \mid \mathcal{F}_t) = e^{-\Lambda_t} \in (0, 1)$ on $\{t < \tau\}$, so the indicator notation $\mathbf{1}_{\{t \geq u\}}$ that one sometimes meets in the literature should be read as a *regime selector* and not as a multiplicative prefactor.

Lemma 2.2.6 (Explicit form of γ under H-hypothesis). *Let $G_t := \mathbb{P}(\tau > t \mid \mathcal{F}_t) = e^{-\Lambda_t}$ with $\Lambda_t = \int_0^t \lambda_s ds$ as in (2.2) below, and assume G is continuous. Then the map*

$$\gamma(t, u) := \begin{cases} \gamma_{\text{pre}}(t, u), & u > t, \\ \gamma_{\text{pre}}(u, u), & 0 \leq u \leq t, \end{cases} \quad \gamma_{\text{pre}}(t, u) := \mathbb{E}[\lambda_u e^{-\Lambda_u} \mid \mathcal{F}_t],$$

satisfies $\mathbb{P}(\tau \geq x \mid \mathcal{F}_t) = \int_x^\infty \gamma(t, u) du$ for all $x \geq 0$, and $\gamma(t, u) = \gamma(u, u)$ on $\{t \geq u\}$. In particular, $\gamma(t, u)$ is strictly positive on $\{u > t\} \cap \{\tau > t\}$ and the identity in Hypothesis 2.2.4 does not imply $\mathbb{P}(t > \tau \mid \mathcal{F}_t) = 0$ on $\{t < \tau\}$; that conditional probability is precisely $1 - G_t = 1 - e^{-\Lambda_t} > 0$.

Proof sketch. Differentiate $\mathbb{P}(\tau \geq x \mid \mathcal{F}_t)$ in x to recover the density $\gamma(t, x)$ for $x > t$; this is the “pre-default density” $\gamma_{\text{pre}}(t, x)$ which is a $(\mathbb{P}, \mathbb{F})$ -martingale in t on $\{t < x\}$ by the classical density-hypothesis argument [96, Lemma 2.3]. At $t = x$, the martingale is stopped and we set $\gamma(t, x) := \gamma_{\text{pre}}(x, x)$ for $t \geq x$, which is just the freezing of the martingale. This is the standard construction of the $\mathbb{F}$ -conditional density past the time of observation; see [63, Proposition 4.4]. $\square$

Distinction from the conditional density hypothesis. The $\mathcal{P}_0$ -density Hypothesis 2.2.4 is *stronger* than the classical density hypothesis under a fixed $\mathbb{P}$ in two respects:

- the same kernel γ is required to represent the $\mathcal{F}_t$ -conditional law of τ *simultaneously* under every $\mathbb{P} \in \mathcal{P}_0$;
- Example 1 below verifies that this is automatic when $\mathcal{P}_0$ is dominated and τ is independent of the Brownian factor under the reference measure; Example 2 below verifies the same when τ has a $\mathbb{P}$ -invariant law and is independent of $\mathbb{F}$.

This distinction is what enables the aggregation of the intensity λ and of the U -process across the non-dominated family $\mathcal{P}_0$, which is the technical ingredient making the auxiliary 2BSDE construction (Lemma 2.3.2) work in our setting.

Examples.

1. $\mathcal{P}_0$ is composed of equivalent probability measures (absolutely continuous with respect to a reference measure, admitting a Radon-Nikodym derivative). Inspired by the weak formulation of stochastic control, where only the drift is controlled and not the volatility, we define an initial probability measure $\mathbb{P}^0$ under which

$$X_t = x + \int_0^t \sigma(s, X_s) dW_s,$$

and $\mathcal{P}_0$ is composed by all the probability $\mathbb{P}^\alpha$ such that

$$\frac{d\mathbb{P}^\alpha}{d\mathbb{P}^0} = Z_T^\alpha := \mathcal{E}\left(\int_0^T \mu(t, X_t, \alpha_t) \sigma^{-1}(t, X_t) dW_t\right),$$

such that $\mathbb{E}^0[Z_T] = 1$. Thus,

$$\mathbb{P}^\alpha(\tau \geq x | \mathcal{F}_t) = \mathbb{E}^\alpha[\mathbf{1}_{\tau \geq x} | \mathcal{F}_t] = \mathbb{E}^0\left[\frac{Z_T^\alpha}{Z_t^\alpha} \mathbf{1}_{\tau \geq x} | \mathcal{F}_t\right] = \mathbb{P}^0(\tau \geq x | \mathcal{F}_t),$$

since Z_T^α/Z_t^α is independent of $\mathcal{F}_t$. The argument tacitly uses that τ is independent of $\mathbb{F}$ under the reference measure $\mathbb{P}^0$, so that Girsanov on the Brownian filtration leaves the law of τ untouched. The common conditional law then admits the explicit kernel $\gamma(t, u) \equiv f_\tau(u)$ for all $t, u \geq 0$, where f_τ is the unconditional density of τ under $\mathbb{P}^0$ (assumed to exist). In particular, the post-default condition in Hypothesis 2.2.4 reads $\gamma(t, u) = f_\tau(u) = \gamma(u, u)$ on $\{u \leq t\}$ in the two-regime form of Remark 2.2.5, while $\gamma(t, u) = f_\tau(u)$ remains a (constant, hence trivially martingale) function of t on $\{u > t\}$.

2. τ is independent of $\mathbb{F}$ with a distribution invariant with any probability in $\mathcal{P}_0$, that is for any $\mathbb{P}, \tilde{\mathbb{P}} \in \mathcal{P}_0$ and $\phi^\mathbb{P}(u) = \phi^{\tilde{\mathbb{P}}}(u)$ for any $u > 0$, where $\phi^\mathbb{P}, \phi^{\tilde{\mathbb{P}}}$ denotes the characteristic functions of τ under the probabilities $\mathbb{P}$ and $\tilde{\mathbb{P}}$ respectively. In this case $\gamma(t, u) \equiv f_\tau(u)$ where f_τ is the common density of τ under any $\mathbb{P} \in \mathcal{P}_0$, and the freezing identity in Hypothesis 2.2.4 reduces to $\gamma(t, u) = f_\tau(u) = \gamma(u, u)$ trivially.

As a consequence of this assumption, see for example [28, 63], any $(\mathbb{P}, \mathbb{F})$ -martingale is also a $(\mathbb{P}, \mathbb{G})$ -martingale. Furthermore, still under the density hypothesis, the process H admits an absolutely continuous compensator, *i.e.*, there exists a non-negative $\mathbb{G}$ -predictable process $\lambda^\mathbb{G}$, such that the compensated process $\widetilde{H}$ defined by

$$\widetilde{H}_t := H_t - \int_0^t \lambda_s^\mathbb{G} ds$$

is a $\mathbb{G}$ -martingale. The compensator vanishes after time τ (therefore $\lambda_t^{\mathbb{G}} := \lambda_t \mathbf{1}_{t \leq \tau}$) and

$$\lambda_t := \frac{\gamma(t, t)}{\mathbb{P}(\tau > t | \mathcal{F}_t)}$$

is an $\mathbb{F}$ -predictable process. For a complete and deeper discussion on the properties of enlarged filtration, we refer the reader to [71]. As a consequence of [63, Proposition 4.4]

$$\mathbb{P}(\tau > t | \mathcal{F}_t) = e^{-\Lambda_t}, \quad \Lambda_t := \int_0^t \lambda_s ds. \quad (2.2)$$

We now turn to the integrability of the process λ and the support of the default time τ . Denoting by $\mathcal{T}(\mathbb{A})$ the set of $\mathbb{A}$ -stopping times (so we will have $\mathcal{T}(\mathbb{F})$ or $\mathcal{T}(\mathbb{G})$), we assume that there exists a constant $\hat{p} > 1$ such that

$$\operatorname{ess\,sup}_{\rho \in \mathcal{T}(\mathcal{G})} \mathbb{E}^{\mathbb{P}} \left[\int_{\rho}^T |\lambda_s|^{\hat{p}} ds \mid \mathcal{G}_{\rho} \right] < +\infty.$$

As a direct consequence of the tower property and since $\mathbb{F} \subseteq \mathbb{G}$, we get

$$\operatorname{ess\,sup}_{\rho \in \mathcal{T}(\mathbb{G})} \mathbb{E}^{\mathbb{P}} \left[\int_{\rho}^T |\lambda_s|^{\hat{p}} ds \mid \mathcal{F}_{\rho} \right] < +\infty.$$

Consequently, $\mathbb{P}(\tau \in [0, T]) < 1$, so the support of τ extends strictly beyond $[0, T]$, i.e. τ has positive probability of exceeding the deterministic horizon T .

Remark 2.2.7. Alternatively, one could relax the integrability conditions for a bounded default time supported on the interval $[0, T]$ at time T . This would lead to a second-order backward stochastic differential equation with a degenerate driver. The analysis of such an extension is closely related to the work of [83, 84]. Note that the existence, uniqueness, and comparison results for 2BSDEs may benefit from the well-established theory of standard BSDEs in these references. Furthermore, one of the more recent applications of this methodology is inspired by the cryptocurrency market or cyber risk. The technical study of 2BSDEs with singular drivers is left for future research.

2.2.3 Spaces

Let $t \in [0, T]$ and $\omega \in \Omega$ be fixed. Let $\mathbb{X} := (\mathcal{X}_t)_{t \leq T}$ be an arbitrary filtration. We denote by $\mathbb{X}_{\mathbb{P}}$ the augmented filtration with respect to $\mathbb{P} \in \mathcal{P}_0$. For any $p \geq 1$ we define

- $L_{t,\omega}^p(\mathbb{X})$ as the space of random variables ξ such that

$$\|\xi\|_{L_{t,\omega}^p(\mathbb{X})}^p := \sup_{\mathbb{P} \in \mathcal{P}(t,\omega)} \mathbb{E}^{\mathbb{P}}[|\xi|^p],$$

- $L_{t,\omega}^{p,\kappa}(\mathbb{X})$ for any $p \geq \kappa \geq 1$ as the set of random variables $\xi \in L_{t,\omega}^1(\mathbb{X})$ such that

$$\|\xi\|_{L_{t,\omega}^{p,\kappa}}^p := \sup_{\mathbb{P} \in \mathcal{P}(t,\omega)} \mathbb{E}^{\mathbb{P}} \left[\text{ess sup}_{t \leq s \leq T}^{\mathbb{P}} \left(\text{ess sup}_{\mathbb{P}' \in \mathcal{P}_{(t,\omega)}(s,\mathbb{P},\mathbb{X})}^{\mathbb{P}} \mathbb{E}^{\mathbb{P}'} [|\xi|^\kappa | \mathcal{X}_s] \right)^{\frac{p}{\kappa}} \right] < \infty,$$

where

$$\mathcal{P}_{t,\omega}(s, \mathbb{P}, \mathbb{X}) := \{\mathbb{P}' \in \mathcal{P}(t, \omega), \mathbb{P} = \mathbb{P}' \text{ on } \mathcal{X}_s\}.$$

- $\mathbb{S}_{t,\omega}^p(\mathbb{X})$ as the set of $\mathbb{X}$ –progressively measurable process Y with $\mathcal{P}(t, \omega)$ quasi surely càdlàg path on $[t, T]$ such that

$$\|Y\|_{\mathbb{S}_{t,\omega}^p}^p := \sup_{\mathbb{P} \in \mathcal{P}(t,\omega)} \|Y\|_{\mathbb{S}_{t,\omega}^p(\mathbb{P})}^p < \infty,$$

where

$$\|Y\|_{\mathbb{S}_{t,\omega}^p(\mathbb{P})}^p := \mathbb{E}^{\mathbb{P}} \left[\sup_{t \leq s \leq T} |Y_s|^p \right] < \infty.$$

We similarly defined $\mathbb{S}_{t,\omega}^p(\mathbb{X}, \mathbb{P})$ for a measure $\mathbb{P}$ fixed as the set of process Y such that $\|Y\|_{\mathbb{S}_{t,\omega}^p(\mathbb{P})}^p < \infty$.

- $\mathbb{H}_{t,\omega}^p(\mathbb{X})$ as the set of $\mathbb{X}$ –predictable $\mathbb{R}^d$ –valued processes such that

$$\|Z\|_{\mathbb{H}_{t,\omega}^p}^p := \sup_{\mathbb{P} \in \mathcal{P}(t,\omega)} \|Z\|_{\mathbb{H}_{t,\omega}^p(\mathbb{P})}^p < \infty,$$

where

$$\|Z\|_{\mathbb{H}_{t,\omega}^p(\mathbb{P})}^p := \mathbb{E}^{\mathbb{P}} \left[\left(\int_t^T \|\hat{a}_s^{\frac{1}{2}} Z_s\|^2 ds \right)^{\frac{p}{2}} \right] < \infty.$$

We similarly defined $\mathbb{H}_{t,\omega}^p(\mathbb{X}, \mathbb{P})$ for a measure $\mathbb{P}$ fixed as the set of process Z such that $\|Z\|_{\mathbb{H}_{t,\omega}^p(\mathbb{P})}^p < \infty$.

- $\mathbb{I}_{t,\omega}^p(\mathbb{X}, \mathbb{P})$ as the set of $\mathbb{X}$ –predictable $\mathbb{R}$ –valued processes K with $\mathbb{P}$ -a.s. càdlàg and nondecreasing paths on $[t, T]$ with $K_t = 0$ such that

$$\|K\|_{\mathbb{I}_{t,\omega}^p(\mathbb{X}, \mathbb{P})}^p := \mathbb{E}^{\mathbb{P}}[K_T^p] < \infty.$$

We say that a family $(K^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}(t, \omega)}$ belongs to $\mathbb{I}_{t, \omega}^{\mathbb{P}}((\mathbb{X}_{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}(t, \omega)})$ if $K^{\mathbb{P}} \in \mathbb{I}_{t, \omega}^{\mathbb{P}}(\mathbb{X}, \mathbb{P})$ and

$$\sup_{\mathbb{P} \in \mathcal{P}(t, \omega)} \|K^{\mathbb{P}}\|_{\mathbb{I}_{t, \omega}^{\mathbb{P}}(\mathbb{P})} < \infty.$$

- $\mathbb{M}_{t, \omega}(\mathbb{X}, \mathbb{P})$ denotes the space of all $(\mathbb{X}, \mathbb{P})$ -optional martingales M with $\mathbb{P}$ -a.s. càdlàg paths on $[t, T]$, with $M_t = 0$, $\mathbb{P}$ -a.s. and

$$\|M\|_{\mathbb{M}_{t, \omega}^{\mathbb{P}}(\mathbb{P})} := \mathbb{E}^{\mathbb{P}}[|M|_T^{p/2}] < +\infty.$$

We say that a family $(M^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}(t, \omega)}$ belongs to $\mathbb{M}_{t, \omega}^{\mathbb{P}}((\mathbb{X}_{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}(t, \omega)})$ if $M^{\mathbb{P}} \in \mathbb{M}_{t, \omega}^{\mathbb{P}}(\mathbb{X}, \mathbb{P})$ and

$$\sup_{\mathbb{P} \in \mathcal{P}(t, \omega)} \|M^{\mathbb{P}}\|_{\mathbb{M}_{t, \omega}^{\mathbb{P}}(\mathbb{P})} < \infty.$$

- $\mathbb{J}_{t, \omega}^{\mathbb{P}}(\mathbb{X})$ as the set of $\mathbb{X}$ -predictable $\mathbb{R}$ -valued processes U such that

$$\|U\|_{\mathbb{J}_{t, \omega}^{\mathbb{P}}}^p := \sup_{\mathbb{P} \in \mathcal{P}(t, \omega)} \mathbb{E}^{\mathbb{P}} \left[\int_t^T |U_s|^p \lambda_s ds \right] < \infty.$$

For the sake of simplicity, we omit the dependence with respect to ω when $t = 0$ in the previous definitions and write, for example, $L_0^p(\mathbb{X})$ for $t = 0$.

2.2.4 Formulation and Definition

We recall that the input of a 2BSDE are a $\mathcal{G}_{\tau \wedge T}$ -measurable terminal condition ξ and a generator function f from $[0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathcal{S}_d^{\geq 0} \times \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \mathbb{R}$. We assume that the following conditions are enforced in this work.

Assumption 2.2.8. Let $C > 0$. For every $(t, \omega, y, z, u, \hat{a}, b, \lambda) \in [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathcal{S}_d^{\geq 0} \times \mathbb{R}^d \times \mathbb{R}^+$ and every $(y', z', u') \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$,

- (i) $|f(s, \omega, y, z, u, \hat{a}, b, \lambda) - f(s, \omega, y', z', u, \hat{a}, b, \lambda)| \leq C(|y - y'| + \|\hat{a}^{\frac{1}{2}}(z - z')\|)$
- (ii) $|f(s, \omega, y, z, u, \hat{a}, b, \lambda) - f(s, \omega, y, z, u', \hat{a}, b, \lambda)| \leq C\lambda|u - u'|$
- (iii) $\sup_{\mathbb{P} \in \mathcal{P}_0} \mathbb{E}^{\mathbb{P}} \left[\int_0^T |f(s, X_{\cdot \wedge s}, 0, 0, \hat{a}_s, b_s^{\mathbb{P}}, \lambda_s^{\mathbb{P}})|^p ds \right] < \infty$

We define for the sake of simplicity

$$F_s^{\mathbb{P}}(y, z, u) := f(s, X_{\cdot \wedge s}, y, z, u, \hat{a}_s, \mu_s^{\mathbb{P}}, \lambda_s).$$

We now recall the result of [83, Lemma 3.1].

Lemma 2.2.9. *Any random variable $\xi - \mathcal{G}_{T \wedge \tau}$ measurable can be decomposed as follows*

$$\xi = \xi^b \mathbf{1}_{T < \tau} + \xi^a \mathbf{1}_{T \geq \tau},$$

where ξ^b is an $\mathcal{F}_T$ -measurable random variable and ξ^a is an $\mathbb{F}$ -predictable process.

We consider the following 2BSDE with erratic horizon

$$\begin{aligned} Y_t = & \xi - \int_{t \wedge \tau}^{T \wedge \tau} F_s^\mathbb{P}(Y_s, \hat{a}_s^{\frac{1}{2}} Z_s, U_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s \cdot dX_s^{c, \mathbb{P}} - \int_{t \wedge \tau}^{T \wedge \tau} dM_s^\mathbb{P} \\ & + K_{T \wedge \tau} - K_{t \wedge \tau} - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s, \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., for every } \mathbb{P} \in \mathcal{P}_0. \end{aligned} \quad (2.3)$$

We recall from [87] that

$$F_t^\mathbb{P}(\cdot) \mathbf{1}_{t \leq \tau} = F_t^{\mathbb{P}, b}(\cdot) \mathbf{1}_{t \leq \tau},$$

where $F^{\mathbb{P}, b} : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}$ is $\mathbb{F}$ -progressively measurable.

Definition 2.2.10 (2BSDE with erratic horizon). The tuple $(Y, Z, U, (K^\mathbb{P})_{\mathbb{P} \in \mathcal{P}_0}, (M^\mathbb{P})_{\mathbb{P} \in \mathcal{P}_0})$ is solution to the 2BSDE with erratic horizon driven by $F^\mathbb{P}$ with terminal condition ξ if there exists some $p, \ell > 1$ such that

1. $(Y, Z, U, (K^\mathbb{P})_{\mathbb{P} \in \mathcal{P}_0}, (M^\mathbb{P})_{\mathbb{P} \in \mathcal{P}_0}) \in \mathbb{S}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{H}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{J}_0^\ell(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{I}_0^p((\mathbb{G}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0}) \times \mathbb{M}_0^p((\mathbb{G}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0})$ satisfies (2.3) with $Y_T = \xi$;
2. the family $(K^\mathbb{P})_{\mathbb{P} \in \mathcal{P}_0}$ satisfies the following minimality condition

$$0 = \operatorname{ess\,inf}_{\mathbb{P}' \in \mathcal{P}_0(t, \mathbb{P}, \mathbb{G}_+)} \mathbb{E}^{\mathbb{P}'}[K_T^{\mathbb{P}'} - K_t^{\mathbb{P}'} | \mathcal{G}_t^{\mathbb{P}', +}], \quad 0 \leq t \leq T \text{ } \mathbb{P}\text{-a.s., } \forall \mathbb{P} \in \mathcal{P}_0.$$

To study the existence of a solution to this second-order BSDE, we introduce an auxiliary BSDE defined by

$$\mathcal{Y}_t^\mathbb{P} = \xi^b - \int_t^T F_s^{\mathbb{P}, b}(\mathcal{Y}_s^\mathbb{P}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^\mathbb{P}, \xi_s^a - \mathcal{Y}_s^\mathbb{P}) ds - \int_t^T \mathcal{Z}_s^\mathbb{P} \cdot dX_s^{c, \mathbb{P}} - \int_t^T d\mathcal{M}_s^\mathbb{P} \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s..} \quad (2.4)$$

Assumption 2.2.11. The auxiliary BSDE (2.4) admits a unique solution $(\mathcal{Y}^\mathbb{P}, \mathcal{Z}^\mathbb{P}, \mathcal{M}^\mathbb{P}) \in \mathbb{S}_{t, \omega}^p(\mathbb{F}, \mathbb{P}) \times \mathbb{H}_{t, \omega}^p(\mathbb{F}, \mathbb{P}) \times \mathbb{M}_{t, \omega}^p(\mathbb{F}, \mathbb{P})$.

Remark 2.2.12. We provide examples under which Assumption 2.2.11 is satisfied. For example, assume that the generator $F^{b, \mathbb{P}}$ is Lipschitz in $\mathcal{Y}, \mathcal{Z}$ and ξ is an $\mathcal{G}_{T \wedge \tau}$ -measurable random

variable such that $\|\xi\|_{L^p_0(\mathbb{G})} < \infty$, [102] ensures that Assumption 2.2.11 holds for $p < 2$. Beyond the Lipschitz case for the generator of the BSDE, and when $p = 2$, [128] provides existence and uniqueness of the auxiliary BSDE (2.4) with stochastic Lipschitz condition while [114, Assumption $(sL^{p,\beta})$] ensures the existence of the solution without the additional martingale term.

2.3 Existence, uniqueness and comparison

Roadmap of this section. The reader will encounter four backward equations that fit together as follows.

- (i) Equation (2.4) is the *auxiliary BSDE*, a Brownian-driven equation on $[0, T]$ obtained by freezing the post-default piece of the data; existence and uniqueness of its solution under each fixed $\mathbb{P} \in \mathcal{P}_0$ are an external input (Assumption 2.2.11).
- (ii) Equation (2.5) is the *auxiliary 2BSDE*, also on $[0, T]$ but aggregated across $\mathcal{P}_0$; its well-posedness is Lemma 2.3.2, proved by adapting the deterministic-horizon argument of [136].
- (iii) Equation (2.3) is the *primal 2BSDE with erratic horizon*, on the random interval $[0, T \wedge \tau]$, and Theorem 2.3.4 solves it by gluing the auxiliary 2BSDE solution to a frozen post-default boundary.
- (iv) Equation (2.7) is the corresponding *erratic-horizon BSDE* under a single $\mathbb{P} \in \mathcal{P}_0$; its well-posedness is Lemma 2.3.5, and the supremum of its solutions over $\mathcal{P}_0$ recovers the 2BSDE value process via the dynamic programming Proposition 2.3.6.

The two routes are reconciled in Remark 2.3.8 and used jointly in the comparison Theorem 2.3.10.

2.3.1 Existence via auxiliary 2BSDE

Hypothesis 2.3.1 (Standing measurability). *Throughout this section, in addition to Hypothesis 2.2.4 and the integrability bound on λ , we assume that $\mathcal{P}_0$ is stable under conditioning and concatenation in the sense of [125, Theorem 4.3] (i.e., $\mathcal{P}_0$ contains all its conditional regular kernels and all its concatenations along $\mathbb{F}$ -stopping times) and that the underlying graph $\bigcup_{(t,\omega)} \{(t,\omega)\} \times \mathcal{P}(t,\omega)$ is analytic in the sense of [124, Theorem 2.1].*

These two properties guarantee that $Y_t^b = \text{ess sup}_{\mathbb{P}' \in \mathcal{P}_0(t, \mathbb{P}, \mathbb{F}^+)} \mathcal{Y}_t^{\mathbb{P}'}$, i.e. the essential-supremum value process produced by the Picard step in Lemma 2.3.2 is universally measurable, and they are exactly the conditions used in the proof of [136, Theorem 4.2] we follow. All examples of

$\mathcal{P}_0$ that appear later in Section 2.4 (in particular the volatility-control and volatility-ambiguity families) satisfy these conditions, see [136, Section 4] and [59, Section 3]. We consider the following auxiliary 2BSDE

$$\begin{aligned} Y_t^b &= \xi^b - \int_t^T F_s^{\mathbb{P},b}(Y_s^b, \hat{a}_s^{\frac{1}{2}} Z_s^b, \xi_s^a - Y_s^b) ds - \left(\int_t^T Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^{\mathbb{P}} \\ &\quad - \int_t^T dM_s^{b,\mathbb{P}} + K_T^{b,\mathbb{P}} - K_t^{b,\mathbb{P}} \quad \text{for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., } \forall \mathbb{P} \in \mathcal{P}_0. \end{aligned} \quad (2.5)$$

Lemma 2.3.2. *Let Assumption 2.2.8 and Assumption 2.2.11 be satisfied. Let ξ be an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable such that $\|\xi\|_{L_0^{p,\kappa}} < \infty$ for some $1 \leq \kappa \leq p$. The auxiliary 2BSDE (2.5) admits a unique solution $(Y^b, Z^b, (M^{b,\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0}, (K^{b,\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0}) \in \mathbb{S}_0^p(\mathbb{F}) \times \mathbb{H}_0^p(\mathbb{F}) \times \mathbb{M}_0^p((\mathbb{F}^{\mathbb{P}^+})_{\mathbb{P} \in \mathcal{P}_0}) \times \mathbb{I}_0^p((\mathbb{F}^{\mathbb{P}^+})_{\mathbb{P} \in \mathcal{P}_0})$ such that*

$$\begin{aligned} Y_t^b &= \operatorname{ess\,sup}_{\mathbb{P}'} \mathcal{Y}_t^{\mathbb{P}'}, \quad Z_t^b = \hat{a}_t^{-1} \frac{d\langle Y^b, X^c \rangle_t}{dt}, \\ M_t^{b,\mathbb{P}} - K_t^{b,\mathbb{P}} &= Y_t^b - Y_0^b + \int_0^t F_s^{\mathbb{P},b}(Y_s^b, \hat{a}_s^{\frac{1}{2}} Z_s^b, \xi_s^{\mathbb{P},a} - Y_s^b) ds + \int_0^t Z_s^b \cdot dX_s^{c,\mathbb{P}}, \end{aligned}$$

where $(\mathcal{Y}^{\mathbb{P}}, \mathcal{Z}^{\mathbb{P}}, \mathcal{M}^{\mathbb{P}})$ is the unique solution to (2.4).

Proof. Under Assumption 2.2.11 and Assumption 2.2.8 the proof follows the same lines as those of [136, Theorem 4.2]. Note that we have a slight modification of (ii) in [136]. The corresponding key quantity to compute in our case has an extra term in $|\lambda|^q$ by using Young's inequality with $\frac{1}{p} + \frac{1}{q} = 1$. Consequently, (4.4) in [136] remains unchanged. $\square$

Why the modification of [136, (ii)] is needed. The argument of [136, Theorem 4.2] establishes existence and uniqueness of the auxiliary 2BSDE (2.5) in the *deterministic-horizon* setting. Their hypothesis (ii) provides an a-priori bound on the moments of the generator at zero, of the form $\sup_{\mathbb{P} \in \mathcal{P}_0} \mathbb{E}^{\mathbb{P}}[(\int_0^T |f^{\mathbb{P}}(s, 0, 0)| ds)^p] < \infty$. In our case the auxiliary generator is

$$F_s^{\mathbb{P},b}(\mathcal{Y}_s^{\mathbb{P}}, \hat{a}_s^{1/2} \mathcal{Z}_s^{\mathbb{P}}, \xi_s^a - \mathcal{Y}_s^{\mathbb{P}}) = f\left(s, X_{\cdot \wedge s}, \mathcal{Y}_s^{\mathbb{P}}, \hat{a}_s^{1/2} \mathcal{Z}_s^{\mathbb{P}}, \xi_s^a - \mathcal{Y}_s^{\mathbb{P}}, \hat{a}_s, b_s^{\mathbb{P}}, \lambda_s^{\mathbb{P}}\right),$$

i.e. the third slot is $u_s = \xi_s^a - \mathcal{Y}_s^{\mathbb{P}}$, and Assumption 2.2.8(ii) gives $|f(\cdot, u) - f(\cdot, u')| \leq C\lambda|u - u'|$ rather than a bare Lipschitz constant. Setting $u' = 0$ and bounding $|f(s, \cdot, 0, 0, \xi_s^a, \hat{a}_s, b_s^{\mathbb{P}}, \lambda_s^{\mathbb{P}})| \leq |f(s, \cdot, 0, 0, 0, \hat{a}_s, b_s^{\mathbb{P}}, \lambda_s^{\mathbb{P}})| + C\lambda_s^{\mathbb{P}}|\xi_s^a|$ and then applying Young's inequality with conjugate exponents $(\hat{p}, \hat{q})$ where $\hat{q} := \hat{p}/(\hat{p} - 1)$ (so that $\hat{p} > 1$ matches the integrability exponent of λ fixed in the standing hypothesis $\operatorname{ess\,sup}_{\rho} \mathbb{E}^{\mathbb{P}}[\int_{\rho}^T |\lambda_s|^{\hat{p}} ds \mid \mathcal{G}_{\rho}] < \infty$) yields

$$\lambda_s^{\mathbb{P}} |\xi_s^a| \leq \frac{1}{\hat{p}} |\lambda_s^{\mathbb{P}}|^{\hat{p}} + \frac{1}{\hat{q}} |\xi_s^a|^{\hat{q}},$$

which produces the additive term $C|\lambda^\mathbb{P}|^{\hat{p}}$ inside the expectation that controls

$$\sup_{\mathbb{P}} \mathbb{E}^\mathbb{P}[(\int_0^T |F_s^{\mathbb{P},b}(0,0,0)| ds)^p].$$

The standing integrability hypothesis on λ together with Hypothesis 2.2.4 guarantees that this $|\lambda|^{\hat{p}}$ moment is finite uniformly in $\mathbb{P} \in \mathcal{P}_0$, and the $|\xi^a|^{\hat{q}}$ moment is controlled by the assumption $\|\xi\|_{L_0^{p,\kappa}} < \infty$ (taking $\kappa \geq \hat{q}$, which is compatible with the relation $\ell = p(\hat{p} - 1)/\hat{p}$ used in Theorem 2.3.4). This is exactly the required modification of [136, (ii)]: one extra $|\lambda|^{\hat{p}}$ term, which is absorbed by the standing $\hat{p}$ -integrability bound on λ . With this substitution the rest of the argument ([136, Eq. (4.4)], the dynamic programming principle for Y^b , and the universal aggregation) goes through unchanged.

Interaction of the jump term with the second-order minimality term. The jump caused by τ is absent from the auxiliary 2BSDE (2.5) *by design*: the auxiliary equation lives on $[0, T]$ and the jump is reintroduced via the before/after-default decomposition (2.6). We have to check that the non-dominated minimality condition on the primal $K^\mathbb{P}$ in Definition 2.2.10 reduces to the minimality of the auxiliary $K^{b,\mathbb{P}}$. The decomposition (2.6) is to be read as “ $K^\mathbb{P}$ tracks $K^{b,\mathbb{P}}$ up to time τ and is then frozen at its τ -value”; the unambiguous non-decreasing realization is $K_t^\mathbb{P} := K_{t \wedge \tau}^{b,\mathbb{P}}$, with which all integrals appearing in (2.3) (which are themselves stopped at τ) are unchanged, and $K_{T \wedge \tau}^\mathbb{P} - K_{t \wedge \tau}^\mathbb{P} = K_{T \wedge \tau}^{b,\mathbb{P}} - K_{t \wedge \tau}^{b,\mathbb{P}}$ pathwise. On $\{t < \tau\}$, the H-hypothesis ensures that $\mathbb{F}$ -adapted increments have the same $\mathcal{G}_t^+$ - and $\mathcal{F}_t^+$ -conditional expectations under each $\mathbb{P}' \in \mathcal{P}_0$, see, e.g., [87, 82, 28], so the minimality of $K^{b,\mathbb{P}}$ from Lemma 2.3.2 transfers verbatim:

$$\text{ess inf}_{\mathbb{P}' \in \mathcal{P}_0(t, \mathbb{P}, \mathbb{G}_+)} \mathbb{E}^{\mathbb{P}'} [K_{T \wedge \tau}^{\mathbb{P}'} - K_{t \wedge \tau}^{\mathbb{P}'} \mid \mathcal{G}_t^{\mathbb{P}',+}] = \text{ess inf}_{\mathbb{P}' \in \mathcal{P}_0(t, \mathbb{P}, \mathbb{F}_+)} \mathbb{E}^{\mathbb{P}'} [K_{T \wedge \tau}^{b,\mathbb{P}'} - K_{t \wedge \tau}^{b,\mathbb{P}'} \mid \mathcal{F}_t^{\mathbb{P}',+}], \quad \text{on } \{t < \tau\}.$$

The right-hand side vanishes because the auxiliary $K^{b,\mathbb{P}'}$ is $\mathbb{F}$ -adapted, non-decreasing, and minimal across $\mathcal{P}_0$ on $[0, T]$ by Lemma 2.3.2, and the stopping at τ only restricts the random terminal time without disturbing the essential-infimum. On the complementary event $\{t \geq \tau\}$, both $K_{T \wedge \tau}^{\mathbb{P}'} - K_{t \wedge \tau}^{\mathbb{P}'} = 0$ and $K_{T \wedge \tau}^{b,\mathbb{P}'} - K_{t \wedge \tau}^{b,\mathbb{P}'} = 0$, so the identity is trivial. This is the precise place where the random-horizon structure and the non-dominated minimality condition interact.

The following corollary enforces the aggregation of the stochastic integral with respect to $X^{c,\mathbb{P}}$ and the process $K^\mathbb{P} - M^\mathbb{P}$ can be aggregated as soon as $\mu^\mathbb{P} = 0$ (the canonical-decomposition drift of X , which appears in the generator slot named b in Assumption 2.2.8).

Corollary 2.3.3. *Assume that $\mu^\mathbb{P} = 0$ (i.e., $b_s^\mathbb{P} \equiv 0$ when the b -slot of the generator is*

evaluated under $\mathbb{P}$) then for any $\mathbb{P} \in \mathcal{P}_0$

$$\int_0^\cdot Z_s \cdot dX_s = \int_0^\cdot Z_s \cdot dX_s^{c,\mathbb{P}} \quad \mathbb{P}\text{-a.s.},$$

and there exists an $\mathbb{F}^{\mathcal{P}_0+}$ -progressively measurable process which aggregates $K^\mathbb{P} - M^\mathbb{P}$.

Proof. Under the standing assumption $\mu^\mathbb{P} \equiv 0$, $X = X^{c,\mathbb{P}}$ $\mathbb{P}$ -a.s., so the two stochastic integrals coincide. The aggregation of $K^\mathbb{P} - M^\mathbb{P}$ then follows from the $\mathcal{P}_0$ -density assumption and [136, Remark 4.1]. $\square$

Theorem 2.3.4. Assume that ξ is an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable such that $\|\xi\|_{L_0^{p,\kappa}} < \infty$ for some $1 \leq \kappa \leq p$ and (2.4) admits a unique solution $(\mathcal{Y}^\mathbb{P}, \mathcal{Z}^\mathbb{P}, \mathcal{M}^\mathbb{P}) \in \mathbb{S}_0^p(\mathbb{F}, \mathbb{P}) \times \mathbb{H}_0^p(\mathbb{F}, \mathbb{P}) \times \mathbb{M}_0^p(\mathbb{F}, \mathbb{P})$ for any $\mathbb{P} \in \mathcal{P}_0$ for some $p > 1$. Then, under Assumption 2.2.8, the 2BSDE (2.3) admits a solution $(Y, Z, U, M, K) \in \mathbb{S}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{H}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{J}_0^\ell(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{I}_0^p((\mathbb{G}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0}) \times \mathbb{M}_0^p((\mathbb{G}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0})$ where $\ell = \frac{p(\hat{p}-1)}{\hat{p}}$ given by

$$\begin{aligned} Y_t &= Y_t^b \mathbf{1}_{t < \tau} + \xi_\tau^a \mathbf{1}_{t \geq \tau}, \quad Z_t = Z_t^b \mathbf{1}_{t < \tau}, \quad M_t^\mathbb{P} = M_t^{b,\mathbb{P}} \mathbf{1}_{t < \tau}, \\ K_t^\mathbb{P} &= K_t^{b,\mathbb{P}} \mathbf{1}_{t < \tau}, \quad U_t^\mathbb{P} = (\xi_t^a - Y_t^b) \mathbf{1}_{t < \tau}, \quad \forall t < T, \quad \mathbb{P}\text{-a.s.}, \mathbb{P} \in \mathcal{P}_0 \end{aligned} \quad (2.6)$$

Proof. The proof extends [96, Theorem 4.3] to 2BSDE. We divide this proof into three main steps: first, we recall that the auxiliary 2BSDE (2.5) has a solution, then we verify the decomposition and measurability of the solution and check that (2.6) is indeed solutions to (2.3), then we show the integrability of each component.

Step 1. 2BSDE (2.5). From Lemma 2.3.2 we deduce that there exists a unique solution (2.5) such that $(Y^b, Z^b, M^b, K^b) \in \mathbb{S}_0^p(\mathbb{F}^{\mathcal{P}_0+}) \times \mathbb{H}_0^p(\mathbb{F}^{\mathcal{P}_0+}) \times \mathbb{I}_0^p((\mathbb{F}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0}) \times \mathbb{M}_0^p((\mathbb{F}^{\mathcal{P}_0+})_{\mathbb{P} \in \mathcal{P}_0})$.

Step 2. Existence, measurability, and decomposition. We prove that for any time $t \in [0, T]$ and any probability $\mathbb{P} \in \mathcal{P}_0$, the processes $(Y, Z, M^\mathbb{P}, K^\mathbb{P}, U^\mathbb{P})$ defined by (2.6) satisfies the primal 2BSDE (2.3). First, from Lemma 2.3.2, we directly deduce that Y is a càdlàg $\mathbb{G}^{\mathcal{P}_0,+}$ -adapted processes. The $\mathbb{G}^{\mathcal{P}_0,+}$ -measurability of Z, U, K, M follows from the $\mathbb{F}$ -measurability of Z^b, Y^b, K^b, M^b and the continuity of $Y^b, \xi^{a,\mathbb{P}}$. We denote by $A^+ = \{\tau > T\}$, $A = \{\tau \in (t, T]\}$ and $A_- = \{\tau \leq t\}$.

- On A^+ such that $\mathbb{P}(A^+) > 0$, we have $Y_t = Y_t^b$, $Z_t^b = Z_t$, $M_t^\mathbb{P} = M_t^{b,\mathbb{P}}$, $K_t^{b,\mathbb{P}}$ and $U_t^\mathbb{P} = \xi_t^a - Y_t^b$. Therefore, on the set A^+ , we have

$$\begin{aligned} Y_t &= Y_t^b = \xi^b + \int_t^T F_s^{\mathbb{P},b}(Y_s^b, Z_s^b, \xi_s^a - Y_s) ds - \left(\int_t^T Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} - \int_t^T dM_s^{b,\mathbb{P}} + \int_t^T dK_s^{b,\mathbb{P}}, \\ &= \xi + \int_t^T F_s^\mathbb{P}(Y_s, Z_s, U_s) ds - \left(\int_t^T Z_s \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} - \int_t^T dM_s^\mathbb{P} + \int_t^T dK_s^\mathbb{P}, \quad t \in [0, T]. \end{aligned}$$

- On A such that $\mathbb{P}(A) > 0$, we have

$$\begin{aligned}
Y_t &= Y_\tau^b + \int_t^\tau F_s^\mathbb{P}(Y_s, Z_s, \xi_s^a - Y_s) ds - \left(\int_t^\tau Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} - \int_t^\tau dM_s^{b,\mathbb{P}} + \int_t^\tau dK_s^{b,\mathbb{P}}, \\
&= \xi_\tau^a + \int_t^\tau F_s^\mathbb{P}(Y_s, Z_s, \xi_s^a - Y_s) ds - \left(\int_t^\tau Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} - \int_t^\tau dM_s^{b,\mathbb{P}} + \int_t^\tau dK_s^{b,\mathbb{P}} \\
&\quad - (\xi_\tau^a - Y_\tau), \\
&= \xi_\tau^a + \int_t^\tau F_s^\mathbb{P}(Y_s, Z_s, \xi_s^a - Y_s) ds - \left(\int_t^\tau Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} - \int_t^\tau dM_s^{b,\mathbb{P}} + \int_t^\tau dK_s^{b,\mathbb{P}} \\
&\quad - \int_t^\tau U_s^\mathbb{P} dH_s, \quad t \in [0, T].
\end{aligned}$$

- On A_- such that $\mathbb{P}(A_-) > 0$, we note that

$$\begin{aligned}
Y_t &= \xi_\tau^a \\
&= \xi_\tau^a + \int_{t \wedge \tau}^{T \wedge \tau} F_s^\mathbb{P}(Y_s, Z_s, \xi_s^a - Y_s) ds - \left(\int_{t \wedge \tau}^{T \wedge \tau} Z_s^b \cdot dX_s^{c,\mathbb{P}} \right)^\mathbb{P} \\
&\quad - \int_{t \wedge \tau}^{T \wedge \tau} dM_s^{b,\mathbb{P}} + \int_{t \wedge \tau}^{T \wedge \tau} dK_s^{b,\mathbb{P}}, \quad t \in [0, T].
\end{aligned}$$

In all cases, we have seen that for any $\mathbb{P} \in \mathcal{P}_0$, the processes $(Y, Z, M^\mathbb{P}, K^\mathbb{P}, U^\mathbb{P})$ defined by (2.6) is solution to (2.3).

Step 3. Integrability. We note that the integrability of Y, Z, M, K is inherited from the integrability of Y^b, Z^b, K^b, M^b and ξ^a . Let $\hat{q} = \frac{\hat{p}}{\hat{p}-1}$. Then, by applying Young Inequality we get

$$|U_s^\mathbb{P}|^\ell \lambda_s^\mathbb{P} \leq \frac{1}{\hat{p}} |\lambda_s^\mathbb{P}|^{\hat{p}} + \frac{1}{\hat{q}} |U_s^\mathbb{P}|^{\ell \hat{q}} \leq C(1 + |\lambda_s^\mathbb{P}|^{\hat{p}} + |Y_s^b|^{\ell \hat{q}} + |\xi_s^a|^{\ell \hat{q}}).$$

The integrability follows from the relation between $\ell, p, \hat{p}$ in the theorem. $\square$

2.3.2 Uniqueness of solution to Lipschitz 2BSDE with erratic horizon

We define for all $0 \leq t \leq T$

$$\begin{aligned}
\mathcal{Y}_t^{\mathbb{P},\tau} &= \xi - \int_{t \wedge \tau}^{T \wedge \tau} F_s^\mathbb{P}(\mathcal{Y}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) ds - \int_{t \wedge \tau}^{T \wedge \tau} \mathcal{Z}_s^{\mathbb{P},\tau} \cdot dX_s^{c,\mathbb{P}} \\
&\quad - \int_{t \wedge \tau}^{T \wedge \tau} \mathcal{U}_s^{\mathbb{P},\tau} dH_s - \int_{t \wedge \tau}^{T \wedge \tau} d\mathcal{M}_s^{\mathbb{P},\tau}, \quad \mathbb{P}\text{-a.s..}
\end{aligned} \tag{2.7}$$

Lemma 2.3.5. Assume that ξ is an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable such that $\|\xi\|_{L_0^{p,\kappa}} < \infty$ for some $1 \leq \kappa \leq p$. Fix $\mathbb{P} \in \mathcal{P}_0$. Under Assumption 2.2.11, there exists a solution $(\mathcal{Y}^{\mathbb{P},\tau}, \mathcal{Z}^{\mathbb{P},\tau}, \mathcal{U}^{\mathbb{P},\tau}, \mathcal{M}^{\mathbb{P},\tau}) \in \mathbb{S}_0^p(\mathbb{G}) \times \mathbb{H}_0^p(\mathbb{G}) \times \mathbb{J}_0^\ell(\mathbb{G}) \times \mathbb{M}_0^p(\mathbb{G})$ to the erratic horizon BSDE (2.7) where $\ell = \frac{p(\hat{p}-1)}{\hat{p}}$. If in addition, Assumption 2.2.8 is satisfied, this solution is unique.

Proof. The proof is decomposed into two parts: first, we deduce the existence of a solution to BSDE (2.7) from Assumption 2.2.11. Then, we turn to the uniqueness under Lipschitz Assumption 2.2.8.

- (i) The existence of a solution is similar to [96] or [83] extended to unbounded but integrable λ processes with a Lipschitz generator and outside a utility maximization problem. It also follows the same lines as the proof of Theorem 2.3.4 above. We only recall the decomposition of $\mathcal{Y}^{\mathbb{P},\tau}, \mathcal{Z}^{\mathbb{P},\tau}, \mathcal{U}^{\mathbb{P},\tau}, \mathcal{M}^{\mathbb{P},\tau}$ in terms of the solution $(\mathcal{Y}^{\mathbb{P}}, \mathcal{Z}^{\mathbb{P}}, \mathcal{M}^{\mathbb{P}})$ to BSDE (2.4).

$$\begin{aligned} \mathcal{Y}_t^{\mathbb{P},\tau} &= \mathcal{Y}_t^{\mathbb{P}} \mathbf{1}_{t < \tau} + \xi_\tau^a \mathbf{1}_{t \geq \tau}, \quad \mathcal{Z}_t^{\mathbb{P},\tau} = \mathcal{Z}_t^{\mathbb{P}} \mathbf{1}_{t < \tau}, \\ \mathcal{M}_t^{\mathbb{P},\tau} &= \mathcal{M}_t^{\mathbb{P}} \mathbf{1}_{t < \tau}, \quad \mathcal{U}_t^{\mathbb{P},\tau} = (\xi_t^a - \mathcal{Y}_t^{\mathbb{P}}) \mathbf{1}_{t < \tau}, \quad \forall t < T, \text{ } \mathbb{P}\text{-a.s.} \end{aligned} \quad (2.8)$$

- (ii) We now turn to the uniqueness under Assumption 2.2.8 which follows classical linearization techniques; see, for example, [64]. We consider two solutions to BSDE (2.7) denoted by $(\mathcal{Y}^{\mathbb{P},\tau}, \mathcal{Z}^{\mathbb{P},\tau}, \mathcal{U}^{\mathbb{P},\tau}, \mathcal{M}^{\mathbb{P},\tau})$ and $(\tilde{\mathcal{Y}}^{\mathbb{P},\tau}, \tilde{\mathcal{Z}}^{\mathbb{P},\tau}, \tilde{\mathcal{U}}^{\mathbb{P},\tau}, \tilde{\mathcal{M}}^{\mathbb{P},\tau})$ and we set

$$\delta_t^{\mathcal{V}} := \mathcal{V}_t^{\mathbb{P},\tau} - \tilde{\mathcal{V}}_t^{\mathbb{P},\tau}, \quad \mathcal{V} \in \{\mathcal{Y}, \mathcal{Z}, \mathcal{U}, \mathcal{M}\}.$$

Then, $(\delta\mathcal{Y}, \delta\mathcal{Z}, \delta\mathcal{U}, \delta\mathcal{M})$ is solution to

$$\delta_t^{\mathcal{Y}} = 0 - \int_{t \wedge \tau}^{T \wedge \tau} \delta_s^{F^{\mathbb{P}}} ds - \int_{t \wedge \tau}^{T \wedge \tau} \delta_s^{\mathcal{Z}} \cdot dX_s^{c,\mathbb{P}} - \int_{t \wedge \tau}^{T \wedge \tau} \delta_s^{\mathcal{U}} dH_s - \int_{t \wedge \tau}^{T \wedge \tau} d\delta_s^{\mathcal{M}}, \quad \mathbb{P}\text{-a.s.},$$

where

$$\begin{aligned} \delta_s^{F^{\mathbb{P}}} &:= F_s^{\mathbb{P}}(\mathcal{Y}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) - F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \tilde{\mathcal{Z}}_s^{\mathbb{P},\tau}, \tilde{\mathcal{U}}_s^{\mathbb{P},\tau}) \\ &= F_s^{\mathbb{P}}(\mathcal{Y}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) - F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) \\ &\quad + F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \mathcal{Z}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) - F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \tilde{\mathcal{Z}}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) \\ &\quad + F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \tilde{\mathcal{Z}}_s^{\mathbb{P},\tau}, \mathcal{U}_s^{\mathbb{P},\tau}) - F_s^{\mathbb{P}}(\tilde{\mathcal{Y}}_s^{\mathbb{P},\tau}, \hat{a}_s^{\frac{1}{2}} \tilde{\mathcal{Z}}_s^{\mathbb{P},\tau}, \tilde{\mathcal{U}}_s^{\mathbb{P},\tau}) \\ &= \theta_s \delta_s^{\mathcal{Y}} + \zeta_s \delta_s^{\mathcal{Z}} + \eta_s \lambda_s \delta_s^{\mathcal{U}}, \end{aligned}$$

where θ , ζ , η are bounded processes obtained from the mean-value theorem under Assumption 2.2.8. From Girsanov theorem, changing the probability measure for both the martingales $X^{c,\mathbb{P}}$ and $\widetilde{H}$ and by considering the process $\tilde{D}_t := e^{-\int_0^{t \wedge \tau} \theta_s ds} \delta_t^{\mathcal{Y}}$, we deduce that $\tilde{D}_t = 0$ and so $\delta_t^{\mathcal{Y}} = 0$ for any time $t \in [0, T]$, $\mathbb{P}$ -a.s. Consequently, $\mathcal{Y}_t^{\mathbb{P},\tau} = \tilde{\mathcal{Y}}_t^{\mathbb{P},\tau}$ for all $t \leq T$, $\mathbb{P}$ -a.s. The uniqueness of the solution then follows by the uniqueness of the semimartingale representation of a process. $\square$

Where the present setting differs from [96, 83]. Three steps of their construction need to be revisited under our hypotheses, and we record them here for the convenience of the reader.

- (a) *From a bounded to an integrable intensity.* [83] works in the bounded-default regime ($\text{supp}(\tau) \subset [0, T]$) under a bounded λ , while we allow $\lambda \in L^{\hat{p}}$ and $\mathbb{P}(\tau \in [0, T]) < 1$. The Picard iteration that produces $(\mathcal{Y}^{\mathbb{P}}, \mathcal{Z}^{\mathbb{P}}, \mathcal{M}^{\mathbb{P}})$ now lives in the spaces $\mathbb{S}_0^p(\mathbb{F}, \mathbb{P}) \times \mathbb{H}_0^p(\mathbb{F}, \mathbb{P}) \times \mathbb{M}_0^p(\mathbb{F}, \mathbb{P})$ defined in Section 2.2, and the only place where boundedness was used is in the contraction estimate, which goes through with the $|\lambda|^{\hat{p}}$ bound from the standing integrability hypothesis on λ and Young's inequality (the same step that appears in the proof of Lemma 2.3.2 above).
- (b) *Outside a utility-maximization framing.* [96, 83] state their representation in the language of an exponential portfolio maximization problem (CARA), which fixes the post-default value ξ^a to a specific quadratic-cost-adjusted exponential. Our formulation only assumes ξ^a is $\mathbb{F}$ -predictable with the $L_0^{p,\kappa}$ moment bound, which is what Lemma 2.3.2 above requires. The decomposition (2.8) is then read with ξ^a as a generic post-default cumulant rather than a CARA payoff.
- (c) *Verification of (2.7).* Substituting (2.8) into the right-hand side of (2.7) reduces, on the three sets $A^+ = \{\tau > T\}$, $A = \{\tau \in (t, T]\}$ and $A_- = \{\tau \leq t\}$, to the corresponding three cases of Step 2 in the proof of Theorem 2.3.4, with the auxiliary 2BSDE replaced by the auxiliary BSDE (2.4) and with the second-order term $K^{b,\mathbb{P}}$ removed. This is why the $\mathbb{F}$ -existence of $(\mathcal{Y}^{\mathbb{P}}, \mathcal{Z}^{\mathbb{P}}, \mathcal{M}^{\mathbb{P}})$ in Assumption 2.2.11 is enough.

Proposition 2.3.6 (Dynamic programming). *Assume that ξ is an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable such that $\|\xi\|_{L_0^{p,\kappa}} < \infty$ for some $1 \leq \kappa \leq p$. Then, any solution $(Y, Z, U, M, K) \in \mathbb{S}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{H}_0^p(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{J}_0^\ell(\mathbb{G}^{\mathcal{P}_0+}) \times \mathbb{I}_0^p((\mathbb{G}^{\mathbb{P}+})_{\mathbb{P} \in \mathcal{P}_0}) \times \mathbb{M}_0^p((\mathbb{G}^{\mathbb{P}+})_{\mathbb{P} \in \mathcal{P}_0})$ where $\ell = \frac{p(\hat{p}-1)}{\hat{p}}$ to the 2BSDE (2.3) satisfies the following dynamic programming principle for any $0 \leq t_1 \leq t_2 \leq T$*

$$Y_{t_1} = \text{ess sup}_{\mathbb{P}' \in \mathcal{P}(t_1, \mathbb{P}, \mathbb{G})} \mathcal{Y}_{t_1}^{\mathbb{P}', \tau}(t_2, Y_{t_2}), \quad \mathbb{P}\text{-a.s.},$$

where $(\mathcal{Y}^{\mathbb{P}',\tau}(t_2, Y_{t_2}), \mathcal{Z}^{\mathbb{P}',\tau}(t_2, Y_{t_2}), \mathcal{U}^{\mathbb{P}',\tau}(t_2, Y_{t_2}), \mathcal{M}^{\mathbb{P}',\tau}(t_2, Y_{t_2}))$ is a solution to (2.7) with horizon $T = t_2$ and terminal condition $\xi = Y_{t_2}$.

Proof. The proof follows the same lines that [136, Theorem 4.2]. We sketch the four steps of that proof adapted to our setting; the key point is that none of them is disrupted by the jump of H at τ , and we explain why for each step.

- (a) *Reduction to dyadic times.* As in [136, Step 1], it suffices to prove the identity at dyadic times $0 \leq t_1 < t_2 \leq T$ and on a $\mathbb{P}$ -null-modulo-zero set; the extension to all $t_1 \leq t_2$ uses the right-continuity of Y in $\mathbb{G}^{\mathcal{P}_0+}$ from Definition 2.2.10, which (2.6) preserves because Y^b is right-continuous in $\mathbb{F}^{\mathcal{P}_0+}$ and ξ^a is $\mathbb{F}$ -predictable.
- (b) *Concatenation of $\mathbb{P}' \in \mathcal{P}(t_1, \mathbb{P}, \mathbb{G})$.* The standing measurability hypothesis on $\mathcal{P}_0$ recalled at the beginning of this section (closure under conditioning and concatenation in the sense of [125, Theorem 4.3]) allows us to concatenate along $\mathbb{F}$ -stopping times. Since τ is exogenous and independent of $\mathbb{F}$ under each $\mathbb{P} \in \mathcal{P}_0$ in the sense of Hypothesis 2.2.4, the concatenation can be extended to $\mathbb{G}$ -stopping times $\rho \leq \tau$ without leaving $\mathcal{P}_0$; on $\{\rho > \tau\}$ the dynamics of X are deterministic past τ in the integral form (2.3) and the concatenation step is vacuous.
- (c) *$\leq$ direction (sub-solution property).* Substituting the decomposition $Y_{t_2} = Y_{t_2}^b \mathbf{1}_{\{t_2 < \tau\}} + \xi_\tau^a \mathbf{1}_{\{t_2 \geq \tau\}}$ from (2.6) into the BSDE (2.7) on $[t_1, t_2]$ and using the comparison theorem 2.3.10 for $\mathbb{P}' \in \mathcal{P}(t_1, \mathbb{P}, \mathbb{G})$ gives $Y_{t_1} \geq \mathcal{Y}_{t_1}^{\mathbb{P}',\tau}(t_2, Y_{t_2})$ $\mathbb{P}'$ -a.s., hence $Y_{t_1} \geq \text{ess sup}_{\mathbb{P}'} \mathcal{Y}_{t_1}^{\mathbb{P}',\tau}(t_2, Y_{t_2})$.
- (d) *$\geq$ direction (super-solution property).* Conversely, the family $(\mathcal{Y}_{t_1}^{\mathbb{P}',\tau}(t_2, Y_{t_2}))_{\mathbb{P}'}$ is upward directed by the concatenation argument of (b), so we may extract an increasing sequence $\mathbb{P}'_n$ such that $\mathcal{Y}_{t_1}^{\mathbb{P}'_n,\tau}(t_2, Y_{t_2}) \nearrow \text{ess sup}_{\mathbb{P}'} \mathcal{Y}_{t_1}^{\mathbb{P}',\tau}(t_2, Y_{t_2})$. The aggregation of U and λ across $\mathcal{P}_0$ from Hypothesis 2.2.4 (this is precisely the place where the $\mathcal{P}_0$ -density assumption is used in the DPP) ensures that the limiting equation has the same generator on $[t_1, t_2 \wedge \tau]$, and we conclude as in [136, Step 4].

The minimality condition on $K^{\mathbb{P}}$ enters at step (d) through the standard equivalence between the saturation of the essential-supremum and the $\text{ess inf}_{\mathbb{P}'} \mathbb{E}^{\mathbb{P}'}[K_T^{\mathbb{P}'} - K_t^{\mathbb{P}'} \mid \mathcal{G}_t^+] = 0$ condition, exactly as it does in [136], and we have already verified in the discussion after Lemma 2.3.2 that this condition transfers from $K^{b,\mathbb{P}}$ to $K^{\mathbb{P}}$ on $\{t < \tau\}$ and is trivial on $\{t \geq \tau\}$. $\square$

Theorem 2.3.7 (Lipschitz 2BSDE-uniqueness). *Assume that ξ is an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable such that $\|\xi\|_{L_0^{p,\kappa}} < \infty$ for some $1 \leq \kappa \leq p$ and Assumption 2.2.8 is enforced. Then 2BSDE (2.3) has at most one solution in the sense of Definition 2.2.10.*

Proof. The proof is a direct consequence of Proposition 2.3.6 combined with Lemma 2.3.5. $\square$

Remark 2.3.8. Note that on the one hand, Theorem 2.3.4 provides a decomposition of the solution to (2.3) under general Assumption 2.2.11, provided the existence of a solution to the auxiliary BSDE and 2BSDE (2.4) and (2.5). On the other hand, under the Lipschitz Assumption 2.2.8, the DPP gives a representation of (2.3) as a supremum of the solution to the erratic horizon BSDE (2.7). The following relation proves the consistency of the solutions given the two methods:

$$Y_t \stackrel{\text{Proposition 2.3.6}}{=} \operatorname{ess\,sup}_{\mathbb{P}} \mathcal{Y}_t^{\mathbb{P},\tau} \stackrel{\text{Lemma 2.3.5}}{=} \operatorname{ess\,sup}_{\mathbb{P}} \mathcal{Y}_t^{\mathbb{P},b} \mathbf{1}_{t<\tau} + \xi_\tau^a \mathbf{1}_{t\geq\tau} \stackrel{\text{Lemma 2.3.2}}{=} Y_t^b \mathbf{1}_{t<\tau} + \xi_\tau^a \mathbf{1}_{t\geq\tau} \stackrel{\text{Theorem 2.3.4}}{=} Y_t.$$

Remark 2.3.9. Alternatively, one could consider the following 2BSDE

$$Y_t = \xi - \int_{t\wedge\tau}^{T\wedge\tau} F_s^{\mathbb{P}}(Y_s, \hat{a}_s^{\frac{1}{2}} Z_s, U_s) ds - \int_{t\wedge\tau}^{T\wedge\tau} Z_s \cdot dX_s^{c,\mathbb{P}} - \int_{t\wedge\tau}^{T\wedge\tau} dM_s^{\mathbb{P}} \quad (2.9)$$

$$- (K_{T\wedge\tau} - K_{t\wedge\tau}) - \int_{t\wedge\tau}^{T\wedge\tau} U_s dH_s, \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., for every } \mathbb{P} \in \mathcal{P}_0,$$

where we have to consider a minus sign for the K term. In this case, the previous result still holds up to an infimum instead of a supremum and more particularly:

$$Y_t \stackrel{\text{Proposition 2.3.6}}{=} \operatorname{ess\,inf}_{\mathbb{P}} \mathcal{Y}_t^{\mathbb{P},\tau}$$

In particular, combining Proposition 2.3.6 with the sign-change above (which replaces $\operatorname{ess\,sup}$ by $\operatorname{ess\,inf}$ in the DPP), we obtain the following essential-infimum form used in the volatility-ambiguity proof of Theorem 2.4.10:

$$Y_0^\alpha = \operatorname{ess\,inf}_{\mathbb{P}' \in \mathcal{P}(0, \mathbb{P}, \mathbb{G})} \mathcal{Y}_0^{\mathbb{P}', \alpha, \beta}, \quad \mathbb{P}\text{-a.s., } \mathbb{P} \in \mathcal{P}_0.$$

2.3.3 Comparison theorem for Lipschitz 2BSDE with erratic horizon

Theorem 2.3.10 (Comparison Theorem for RH-2BSDE). *Let $(Y, Z, U, (K^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0}, (M^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0})$ and $(\tilde{Y}, \tilde{Z}, \tilde{U}, (\tilde{K}^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0}, (\tilde{M}^{\mathbb{P}})_{\mathbb{P} \in \mathcal{P}_0})$ be solutions to the 2BSDE (2.3) with respective terminal conditions ξ and $\tilde{\xi}$ and generators $F^{\mathbb{P}}$ and $\tilde{F}^{\mathbb{P}}$. Assume that $\xi^b \geq \tilde{\xi}^b$, $\xi^a = \tilde{\xi}^a$, and that $F_t^{\mathbb{P},b}(y, z, u - y) \geq \tilde{F}_t^{\mathbb{P},b}(y, z, u - y)$ for every $(y, z, u) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$, $\mathbb{P}$ -a.s. for every $\mathbb{P} \in \mathcal{P}_0$. Then, $Y_t \geq \tilde{Y}_t$, $\mathbb{P}$ -a.s., $\forall \mathbb{P} \in \mathcal{P}_0$.*

Proof. The proof is divided into two main parts: first, we show a comparison theorem for the auxiliary BSDE (2.5), then we deduce from Remark 2.3.8.

Step 1. Comparison auxiliary BSDE. Following the same argument as the proof of Lemma 2.3.5 by considering

$$\delta_t^\nu := \mathcal{V}_t^{\mathbb{P},b} - \tilde{\mathcal{V}}_t^\mathbb{P}, \quad \nu \in \{\mathcal{Y}, \mathcal{Z}, \mathcal{M}\}.$$

By using the same linearization, we deduce that a comparison theorem holds true for the solutions to the auxiliary BSDE (2.4):

$$\mathcal{Y}_t^{\mathbb{P},b} \geq \tilde{\mathcal{Y}}_t^{\mathbb{P},b}, \quad t \leq T, \quad \mathbb{P}\text{-a.s.}, \quad \mathbb{P} \in \mathcal{P}_0.$$

Step 2. Comparison and decomposition. From Remark 2.3.8, we deduce that

$$\begin{aligned} Y_t &= \operatorname{ess\,sup}_{\mathbb{P}} \mathcal{Y}_t^{\mathbb{P},b} \mathbf{1}_{t < \tau} + \xi_\tau^a \mathbf{1}_{t \geq \tau} \\ &\geq \operatorname{ess\,sup}_{\mathbb{P}} \tilde{\mathcal{Y}}_t^{\mathbb{P},b} \mathbf{1}_{t < \tau} + \xi_\tau^a \mathbf{1}_{t \geq \tau} = \operatorname{ess\,sup}_{\mathbb{P}} \tilde{\mathcal{Y}}_t^\mathbb{P} \mathbf{1}_{t < \tau} + \xi_\tau^a \mathbf{1}_{t \geq \tau} = \tilde{Y}_t. \end{aligned} \quad \square$$

Remark 2.3.11. Note that under the $\mathcal{P}_0$ -density assumption, the process γ and so λ is aggregated. We can thus aggregate the process U in the 2BSDE.

2.3.4 Markovian 2BSDE with erratic horizon and connection with fully non linear PDE

In the whole section, we denote by $\mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$ the set of maps from time and space $[0, T] \times \mathbb{R}^d$ into $\mathbb{R}$ continuously differentiable map with respect to the time component and twice continuously differentiable with respect to the space variable for any $v \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$ we denote by $\partial_t v, \nabla v, \Delta v$ its derivative in time, gradient and Laplacian respectively. In this section, we assume that the following conditions are satisfied.

Assumption 2.3.12. We set the main Markovian assumptions for this section as follows:

- There exists a function $g : [0, T] \times \mathbb{R}^d \longrightarrow \mathbb{R}$ such that $\xi = g(T, X_T)$.
- $\mu^\mathbb{P} = 0$ so that the canonical process given by X is a martingale with quadratic variation $\hat{a}$. In particular, the process $W^\mathbb{P}$ defined by

$$W_t^\mathbb{P} := \int_0^t \hat{a}_s^{-1/2} dX_s, \quad \mathbb{P}\text{-a.s.},$$

is a $\mathbb{P}$ -Brownian motion for any $\mathbb{P} \in \mathcal{P}_0$.

- There exists a generator $H^b : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times D_H \longrightarrow \mathbb{R}$ where D_H is a subset of $\mathbb{R}^{d \times d}$ containing 0 such that

$$F_t^b(x, y, z, u, a) := \sup_{\gamma \in \mathbb{R}^{d \times d}} \left\{ \frac{1}{2} \text{Tr}(a\gamma) - H^b(t, x, y, z, u, \gamma) \right\}, \quad a \in \mathcal{S}_d^{>0}.$$

The 2BSDE (2.3) thus becomes

$$Y_t = g(T, X_T) - \int_{t \wedge \tau}^{T \wedge \tau} F_s(Y_s, Z_s, U_s, \hat{a}_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s \cdot dX_s + K_{T \wedge \tau} - K_{t \wedge \tau} - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s, \\ \text{for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., for every } \mathbb{P} \in \mathcal{P}_0, \quad (2.10)$$

where F is defined by

$$F_s(y, z, u, a) := f(s, X_s, y, z, u, \hat{a}_s, \lambda_s),$$

and where f satisfied the Lipschitz Assumption 2.2.8.

Remark 2.3.13. Note that F may not be a Markovian function of t, X_t since we are not specifying any regularity assumption for the intensity process λ .

We are recalling the auxiliary 2BSDE in this Markovian framework

$$Y_t^b = g(T, X_T) - \int_t^T F_s^b(Y_s^b, Z_s^b, g(s, X_s) - Y_s^b, \hat{a}_s) ds - \int_t^T Z_s^b \cdot dX_s + K_T^{b, \mathbb{P}} - K_t^{b, \mathbb{P}}, \\ \text{for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., } \forall \mathbb{P} \in \mathcal{P}_0. \quad (2.11)$$

We define the bi-conjugate of H^b (here the Hessian slot is named Γ to avoid the clash with the density kernel $\gamma(t, u)$ of Hypothesis 2.2.4) as follows

$$\hat{H}^b(t, x, y, z, u, \Gamma) := \sup_{a \in \mathcal{S}_d^{>0}} \left\{ \frac{1}{2} \text{Tr}(a\Gamma) - F_t^b(x, y, z, u, a) \right\}.$$

We introduce the following fully nonlinear PDEs

$$(PDE) \begin{cases} \partial_t v(t, x) + \hat{H}^b(t, x, v(t, x), \nabla v(t, x), g(t, x) - v(t, x), \Delta v(t, x)) = 0, & t < T \\ v(T, x) = g(T, x), & x \in \mathbb{R}^d. \end{cases}$$

Lemma 2.3.14 (Theorem 5.3 in [150]). *Let $v \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R}^d)$ be a solution to (PDE). Then the auxiliary 2BSDE (2.11) admits a unique solution given by*

$$Y_t^b := v(t, X_t), \quad Z_t^b := \nabla v(t, x), \quad K_t^b := \int_0^t k_s^b ds,$$

where

$$k_s^b = \hat{H}^b(s, X_s, Y_s^b, Z_s^b, g(s, X_s) - Y_s^b, \Gamma_s), \Gamma_s = \Delta v(s, X_s).$$

Theorem 2.3.15 (Piecewise Feynman-Kac representation formula). *The 2BSDE (2.10) admits a unique solution given by*

$$Y_t = v(t, X_t)\mathbf{1}_{t < \tau} + g(\tau, X_\tau)\mathbf{1}_{t \geq \tau}, \quad Z_t = \nabla v(t, X_t)\mathbf{1}_{t < \tau},$$

$$K_t = K_t^b\mathbf{1}_{t < \tau}, \quad U_t = (g(t, X_t) - v(t, X_t))\mathbf{1}_{t < \tau}, \quad t < T$$

where K_t^b is given by Lemma 2.3.14 and v is a classical solution to (PDE).

Proof. This is a direct consequence of Lemma 2.3.14 together with Theorem 2.3.4. $\square$

Remark 2.3.16. Note that (PDE) is not an integropartial PDE and the jumps and the erratic time τ do not play any role in the PDE itself to describe the solution to (2.10). In particular, Y_t, Z_t, U_t, K_t are not Markovian functions of (t, X_t) since λ may be a non-Markovian process. However, we rely on the Markovianity of the auxiliary 2BSDE (2.11) to connect the solution to (2.10) with the fully nonlinear PDE (PDE).

2.4 Application to stochastic control in erratic environments

This section applies the preceding results to stochastic control problems under weak formulations. We examine two types of controlled diffusion optimization: one where both the drift and volatility are controlled by the agent, and another where only the drift is controlled in the presence of volatility uncertainty. The concept of controlled volatility and weak formulation has been well understood and extensively studied in the literature, particularly following the influential work of [54, 157], while the notion of volatility ambiguity has been investigated, for example, in [27, 126, 22]. This last case is particularly relevant for real-world applications, where modelers inevitably encounter uncertainty in parameter estimation. In this section, we extend this class of problems by adding uncertainty to the termination of the optimization problems studied.

2.4.1 A controlled model to unify volatility controlled and ambiguity

We define $\mathcal{A}$ and $\mathcal{B}$ as the sets of $\mathbb{F}$ -adapted processes taking values in two subsets $A \subset \mathbb{R}^m$ and $B \subset [\varepsilon, \bar{\beta}]^m$, for $0 < \varepsilon < \bar{\beta}$, for $m \geq 1$. We call *control process* every pair $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$.

We will also denote, to make the notation lighter, $\nu = (\alpha, \beta)$ and its respective set $\mathcal{V}$. To clarify the notations, for the rest of the paper:

- in the volatility control setting, α and β will be directly controlled by an agent, where β will affect only the volatility part and α the drift;
- in the volatility ambiguity model, β has the interpretation of a worst-case scenario, chosen by an exogenous actor, named the Nature, picking the worst volatility coefficient for the optimization problem.

We specify the drift and volatility coefficient as follows:

$\mu : [0, T] \times \Omega \times A \times B \longrightarrow \mathbb{R}^n$, such that $\mu(\cdot, \alpha, \beta)$ is $\mathbb{F}$ -optional for any $(\alpha, \beta) \in A \times B$ and bounded,

$\sigma : [0, T] \times \Omega \times B \longrightarrow \mathcal{M}_{n \times d}(\mathbb{R})$, such that $\sigma(\cdot, \beta)$ is $\mathbb{F}$ -optional for any $\beta \in B$ and bounded.

We now recall the output process definition and the control problem. In order to make the construction suitable for the volatility ambiguity, differently from the seminal work of [54] and more in line with the work of [74], we start by defining the driftless SDE driven by the d -dimensional Brownian motion W

$$X_s^{t,x,\beta} = x(t) + \int_t^s \sigma(r, X_{\cdot \wedge r}^{t,x,\beta}, \beta_r) dW_r, \quad s \in [t, T] \quad (2.12)$$

$$X_r^{t,x,\beta} = x(r), \quad x \in \Omega \quad r \in [0, t],$$

where $X_{\cdot \wedge t}$ denotes the path of X up to time t . We now define a unified control model for both volatility controlled by an agent or ambiguity controlled by an exogenous player called the Nature as a set of weak solutions as follows.

Definition 2.4.1. We say that $(\mathbb{P}, \beta)$ is a weak solution for (2.12) if the distribution of $X_t^{t,x,\beta}$ under $\mathbb{P}$ is $\delta_{x(t)}$, where δ denotes the Dirac measure and there exists a $\mathbb{P}$ -Brownian motion, denoted by $W^\mathbb{P}$, such that

$$X_s = x(t) + \int_t^s \sigma(r, X_{\cdot \wedge r}, \beta_r) dW_r^\mathbb{P}, \quad s \in [t, T], \quad \mathbb{P}\text{-a.s.}$$

We denote by

$$\mathcal{N}(t, x) := \{(\mathbb{P}, \beta) : \text{weak solutions of (2.12)}\}$$

and by $\mathcal{P}(t, x)$ the set of probability measures such that there exists an $\mathbb{F}$ -optional process β such that the probability coupled with the process β is a weak solution of (2.12).

Remark 2.4.2. Note that $\{\mathcal{P}(t, x), t \in [0, T] \times \Omega\}$ is a saturated family of probability measures, see for example [74, 54] and the references therein.

We now turn to the drift-controlled model. From now on, we denote by $(\mathcal{E}(Z_t))_{t \in [0, T]}$ the Doleans-Dade exponential of a continuous stochastic process Z defined by

$$\mathcal{E}(Z_t) := e^{Z_t - \frac{1}{2}[Z, Z]_t}.$$

Definition 2.4.3 (Admissible agent's effort). A control process α is said to be admissible, if for every $(\mathbb{P}, \beta) \in \mathcal{N}(t, x)$ the following process is an $(\mathbb{F}, \mathbb{P})$ -martingale

$$\mathcal{E} := \left(\mathcal{E} \left(\int_0^t \sigma^\top (\sigma \sigma^\top)^{-1} (s, X_{\cdot \wedge s}, \beta_s) \mu(s, X_{\cdot \wedge s}, \alpha_s, \beta_s) \cdot dW_s^\mathbb{P} \right) \right)_{t \in [0, T]}.$$

We denote by $\mathcal{A}$ the set of such admissible efforts.

Remark 2.4.4. Note that this set is not empty if we assume that σ is such that $\sigma(\sigma \sigma^\top)^{-1}$ is bounded, for example if $n = d = 1$ and $\sigma : [0, T] \times \mathbb{R} \times B \rightarrow [\underline{\sigma}, \bar{\sigma}]$ with $0 < \underline{\sigma} < \bar{\sigma}$. Alternatively, if we assume that μ can be decoupled multiplicatively as follows $\mu(s, x, \alpha, \beta) = \lambda(s, x, \alpha) \sigma(s, x, \beta)$ for some bounded function λ , the stochastic exponential in the definition of admissible effort becomes

$$\left(\mathcal{E} \left(\int_0^t \lambda(s, X_{\cdot \wedge s}, \alpha_s, \beta_s) \cdot dW_s^\mathbb{P} \right) \right)_{t \in [0, T]},$$

and is directly a martingale, similarly to [54].

Consider an admissible effort $\alpha \in \mathcal{A}$ and $(t, x) \in [0, T] \times \Omega$. For every subset $\mathcal{N} \subset \mathcal{N}(t, x)$ define

$$\mathcal{N}^\alpha := \left\{ (\mathbb{P}^\alpha, \beta), \frac{d\mathbb{P}^\alpha}{d\mathbb{P}} = \mathcal{E} \left(\int_{t \wedge \tau}^{T \wedge \tau} \sigma^\top (\sigma \sigma^\top)^{-1} (s, X_{\cdot \wedge s}, \nu_s) \mu(s, X_{\cdot \wedge s}, \alpha_s, \beta_s) \cdot dW_s^\mathbb{P} \right), (\mathbb{P}, \beta) \in \mathcal{N} \right\}.$$

By Girsanov's Theorem we then deduce that for any $\alpha \in \mathcal{A}$, and for any $(\mathbb{P}^\alpha, \beta) \in \mathcal{N}^\alpha$,

$$X_s = x(t) + \int_t^s \mu(r, X_{\cdot \wedge r}, \alpha_r, \beta_r) dr + \int_t^s \sigma(r, X_{\cdot \wedge r}, \beta_r) dW_r^{\mathbb{P}^\alpha}, \quad s \in [t, T], \quad \mathbb{P}^\alpha\text{-a.s.}$$

We point out that this construction enables us to unify the mathematical frameworks for both the controlled volatility [54] and the ambiguity [74] cases. Especially for the volatility ambiguity, the Nature fixes a diffusion coefficient first through a change of probability, selecting a family of probability measures. Then, within this family, the agent can choose its optimal effort which is changing the probability measure through a Girsanov transformation.

In the case of volatility control, this is still a valid construction, identifying the Nature with the agent, equivalent to the ones proposed in [54, 59].

We now turn to the optimization problems for both cases: volatility controlled and *sup sup* optimization or ambiguity and *sup inf* problem. We first define a discount factor by

$$\mathcal{K}_t^{\alpha, \beta} := \exp \left(- \int_0^t k(s, X_{\cdot \wedge s}, \alpha_s, \beta_s) ds \right), \quad t \in [0, T],$$

where k is a bounded and $\mathbb{F}$ -optional function. For any control $(\alpha; (\mathbb{P}^\alpha, \beta)) \in \mathcal{A} \times \mathcal{N}^\alpha$, the objective function of the agent is thus given by

$$J(\alpha; \mathbb{P}^\alpha, \beta) = \mathbb{E}^{\mathbb{P}^\alpha} [\mathcal{K}_{T \wedge \tau}^{\alpha, \beta} \Phi(\xi) - \int_0^{T \wedge \tau} \mathcal{K}_t^{\alpha, \beta} c_t(\alpha_t, \beta_t) dt]$$

where

- ξ is the liability the agent is facing as an $\mathcal{G}_{T \wedge \tau}$ -measurable random variable, Φ is a utility function assumed to be continuous and non-decreasing such that $\|\Phi(\xi)\|_{L_0^{p, \kappa}} < \infty$ for some $1 \leq \kappa \leq p$.
- $c : [0, T] \times \Omega \times A \times B \longrightarrow \mathbb{R}$ is a cost function possibly depending on the time, and the path of the process X such that

$$\sup_{\alpha \in \mathcal{A}} \sup_{(\mathbb{P}^\alpha, \beta) \in \mathcal{N}^\alpha} \mathbb{E}^{\mathbb{P}^\alpha} \left[\int_0^T |c_s(\alpha_s, \beta_s)|^p ds \right] < \infty.$$

Given this objective function, the two problems differ in their optimization routines.

Agent as a volatility controller. In the purely controlled case, we aim to maximize the expected utility over the two controls, volatility and drift simultaneously, that is to solve

$$\overline{V}_0 := \sup_{\alpha \in \mathcal{A}} \sup_{(\mathbb{P}^\alpha, \beta) \in \mathcal{N}^\alpha} J(\alpha; \mathbb{P}^\alpha, \beta) \quad (2.13)$$

Nature as the volatility controller. In a volatility ambiguity setting, we seek to optimize the worst-case scenario, tackling a problem of the form

$$\underline{V}_0 := \sup_{\alpha \in \mathcal{A}} \inf_{(\mathbb{P}^\alpha, \beta) \in \mathcal{N}^\alpha} J(\alpha; \mathbb{P}^\alpha, \beta) \quad (2.14)$$

We want to stress that this problem can be seen as the Nature being an adversarial opponent in a stochastic game, where the Nature picks its "strategy" β to make our output process as

bad as possible and we choose our strategy α to mitigate the output. This is a worst-case scenario or robust optimization framework.

We aim to prove that both problems can be linked to the solution of 2BSDE (2.3). The idea is that, in both cases, we can represent the values of both problems $\overline{V}_0, \underline{V}_0$ as the solution to a 2BSDE we studied in Section 2.3.

2.4.2 Supporting 2BSDEs and driver definitions

In this section, we introduce the 2BSDEs related to the optimization problems above. The interpretation of the two cases is different, and even if the 2BSDEs look similar, with the exception of their driver, the volatility ambiguity case will require the extra assumption of the Isaacs condition from the seminal work in stochastic differential games of [80].

We start by defining the driver f for our 2BSDEs:

$$f : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times A \times B \times \mathbb{R}^+ \rightarrow \mathbb{R}$$

defined by

$$f(t, x, y, z, u, \alpha, \beta, \lambda) := -k(t, x, \alpha, \beta)y - c_t(x, \alpha, \beta) + \mu(t, x, \alpha, \beta) \cdot z + \lambda u.$$

Define also for every $(t, x, \Sigma) \in [0, T] \times \Omega \times \mathbb{S}_d^+$ the set

$$V_t(x, \Sigma) := \left\{ \beta \in B, \sigma(t, x, \beta)\sigma^\top(t, x, \beta) = \Sigma \right\},$$

and denote by $\mathbf{B}(\hat{\sigma}^2)$ the set of controls $\beta \in \mathcal{B}$ with values in $V_t(x, \hat{\sigma}_t^2)$, $dt \otimes \mathbb{P}$ -a.e., where

$$\hat{\sigma}_t^2 := \limsup_{\varepsilon \downarrow 0} \frac{\langle X \rangle_t - \langle X \rangle_{t-\varepsilon}}{\varepsilon}, \quad t \in [0, T],$$

taking values in $\mathcal{S}_d^{>0}$ (we restrict the ambiguity family to non-degenerate volatility, so $\hat{\sigma}_t^2$ specialises $\hat{a}_t$ of Definition 2.2.1 to the strictly positive cone, which is required so that the Hamiltonians involve $(\sigma\sigma^\top)^{-1}$) and where $\langle X \rangle$ coincides with the quadratic variation of X for every probability measure $\mathbb{P} \in \mathcal{P}(t, x)$ by [92]. Throughout this section $\hat{\sigma}_t^2$ is therefore the pathwise variance density (the same object as $\hat{a}_t$ in Definition 2.2.1).

From this generator definition, we define two Hamiltonians, one associated with the

volatility-controlled problem, and the other one associated with the volatility ambiguity problem.

- **Agent as a volatility controller.** The Hamiltonian $\bar{F} : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathcal{S}_d^+ \rightarrow \mathbb{R}$ associated with the volatility controlled problem is defined by

$$\bar{F}(t, x, y, z, u, \lambda, \Sigma) := \sup_{\beta \in V_t(x, \Sigma)} \sup_{\alpha \in A} f(t, x, y, z, u, \alpha, \beta, \lambda)$$

- **Nature as an adversarial volatility controller.** The Hamiltonian $\underline{F} : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathcal{S}_d^+ \rightarrow \mathbb{R}$ associated with the volatility ambiguity problem is defined by

$$\underline{F}(t, x, y, z, u, \lambda, \Sigma) := \inf_{\beta \in V_t(x, \Sigma)} \sup_{\alpha \in A} f(t, x, y, z, u, \alpha, \beta, \lambda).$$

Given the drivers, we can define the 2BSDEs that we will use equivalently to solve our control problems. Consider the following 2BSDEs

$$\begin{aligned} \bar{Y}_t &= \Phi(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} \bar{F}(s, X_s, \bar{Y}_s, \bar{Z}_s, \bar{U}_s, \lambda_s, \hat{\sigma}_s^2) ds - \int_{t \wedge \tau}^{T \wedge \tau} \bar{Z}_s \cdot dX_s \\ &+ \int_{t \wedge \tau}^{T \wedge \tau} d\bar{K}_s - \int_{t \wedge \tau}^{T \wedge \tau} \bar{U}_s dH_s, \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., } \forall \mathbb{P} \in \mathcal{P}(t, x), \end{aligned} \quad (2.15)$$

and

$$\begin{aligned} \underline{Y}_t &= \Phi(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} \underline{F}(s, X_s, \underline{Y}_s, \underline{Z}_s, \underline{U}_s, \lambda_s, \hat{\sigma}_s^2) ds - \int_{t \wedge \tau}^{T \wedge \tau} \underline{Z}_s \cdot dX_s \\ &- \int_{t \wedge \tau}^{T \wedge \tau} d\underline{K}_s - \int_{t \wedge \tau}^{T \wedge \tau} \underline{U}_s dH_s, \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s., } \forall \mathbb{P} \in \mathcal{P}(t, x), \end{aligned} \quad (2.16)$$

Remark 2.4.5. Note that we could define the 2BSDE with an orthogonal martingale M . However, the proof of our main results below would be then taken as a conditional expectation under a saturated probability measure under which M is a martingale and therefore this part vanishes. For ease of notation and exposition, we will therefore ignore that term.

Remark 2.4.6. Given the assumption made about the intensity of the default λ , the proof below is similar to [54, 74, 136, 59]. This is because the aggregation of the process K holds thanks to the aggregation of the process U , which is only possible when the intensity is independent of the probability measure, by using [126].

Since k, σ, μ are bounded and since c does not depend on y, z, u , we have the following Lemma ensuring that the 2BSDEs (2.15) and (2.16) are well-posed in the sense of Assumption 2.2.8.

Lemma 2.4.7. *For any $(t, x, y, z, u, \lambda, \Sigma) \in [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^+ \times \mathcal{S}_d^+$, the maps $\bar{F}$ and $\underline{F}$ satisfy Assumption 2.2.8.*

2.4.3 The volatility control case: from full control to 2BSDE

In this section, we turn to the solution of the non-Markovian optimization (2.13) and its connection with the 2BSDE (2.15).

Theorem 2.4.8. *Let $(\bar{Y}, \bar{Z}, \bar{U}, \bar{K})$ be the unique solution to (2.15). Then, the value function of the agent (2.13) is given by*

$$\bar{V}_0(\xi) = \sup_{\alpha \in \mathcal{A}} \sup_{(\mathbb{P}, \beta) \in \mathcal{N}^\alpha} \mathbb{E}^\mathbb{P}[\bar{Y}_0].$$

Moreover, $(\alpha^*, \mathbb{P}^*, \beta^*)$ are optimal if and only if $(\alpha^*, \mathbb{P}^*, \beta^*) \in \mathcal{A} \times \mathcal{N}^{\alpha^*}$ such that:

- (i) (α^*, β^*) attains the supremum in the definition of $\bar{F}(\cdot, X, \bar{Y}, \bar{Z}, \bar{U}, \lambda, \hat{\sigma}^2)$, $dt \otimes \mathbb{P}^*$ -a.e.,
- (ii) $\bar{K}_{T \wedge \tau} = 0$, $\mathbb{P}^*$ -a.s.

The existence of a unique solution to the 2BSDE (2.15) is a direct consequence of the assumptions made on k, σ, μ, c and Theorem 2.3.4 and Theorem 2.3.7. We omit the proof regarding the optimizers of the problem since it follows the same lines as the one of the ambiguity case below and extend [54] to the erratic horizon $\tau \wedge T$. We refer to either an extension to [54, Proposition 5.4] or the proof of Theorem 2.4.10 below.

2.4.4 The volatility ambiguity case: from adversarial nature to 2BSDE

We now turn to the case where the agent is facing an adversarial player, the Nature, regarding the volatility control. Similarly to the previous section, we aim at connecting the robust optimization (2.14) with the 2BSDE (2.16). Note that this is a worst-case scenario framework and its link to 2BSDE has been previously studied in [74, 116]. We extend this result to the erratic horizon $T \wedge \tau$. Unlike the previous case in which the boundedness was enough to guarantee well-posedness for the optimization problem, we are facing the extra complexity of a potentially unbounded sup inf problem, where the second sub-problem becomes $-\infty$ once the worst strategy of the Nature is considered. To overcome this technical difficulty, we will add another assumption ensuring the existence of an optimizer, known as Isaacs condition, from the seminal work of [80]. This assumption has been also made in the literature regarding

zero-sum games in weak formulation, see for example [137, 74]. We assume that the following condition is satisfied.

Assumption 2.4.9. For any $(t, x, y, z, u, \lambda, \Sigma) \in [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^+ \times \mathcal{S}_d^+$

$$\inf_{\beta \in V_t(x, \Sigma)} \sup_{\alpha \in A} f(t, x, y, z, u, \alpha, \beta, \lambda) = \sup_{\alpha \in A} \inf_{\beta \in V_t(x, \Sigma)} f(t, x, y, z, u, \alpha, \beta, \lambda).$$

We refer to [74, Section 4] for explicit examples satisfying this condition.

Theorem 2.4.10. *Let Assumption 2.4.9 be true. Let $(\underline{Y}, \underline{Z}, \underline{U}, \underline{K})$ be the unique solution to (2.16). Then, the value function of the agent (2.14) is given by*

$$\underline{V}_0(\xi) = \sup_{\alpha \in \mathcal{A}} \inf_{(\mathbb{P}, \beta) \in \mathcal{N}^\alpha} \mathbb{E}^\mathbb{P}[\underline{Y}_0].$$

Moreover, $(\alpha^*, \mathbb{P}^*, \beta^*)$ are optimal if and only if $(\alpha^*, \mathbb{P}^*, \beta^*) \in \mathcal{A} \times \mathcal{N}^{\alpha^*}$ such that:

- (i) (α^*, β^*) attains the saddle point sup inf in the definition of $\underline{F}(\cdot, X, \underline{Y}, \underline{Z}, \underline{U}, \lambda, \hat{\sigma}^2)$, $dt \otimes \mathbb{P}^*$ -a.e.,
- (ii) $\underline{K}_{T \wedge \tau} = 0$, $\mathbb{P}^*$ -a.s.

Proof. From Lemma 2.4.7 together with Theorem 2.3.4 and Theorem 2.3.7 we deduce that there exists a unique solution $(\underline{Y}, \underline{Z}, \underline{U}, \underline{K})$ to (2.16). We define for each $\alpha \in \mathcal{A}$ the following 2BSDE with erratic horizon

$$\begin{aligned} Y_t^\alpha &= \Phi(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} \inf_{\beta \in V_t(X_{\cdot \wedge t}, \hat{\sigma}_t^2)} f(s, X_s, Y_s^\alpha, Z_s^\alpha, U_s^\alpha, \alpha_s, \beta_s, \lambda_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s^\alpha \cdot dX_s \\ &\quad - \int_{t \wedge \tau}^{T \wedge \tau} dK_s^\alpha - \int_{t \wedge \tau}^{T \wedge \tau} U_s^\alpha dH_s, \quad \mathbb{P}\text{-a.s.}, \quad \mathbb{P} \in \mathcal{P}(t, x). \end{aligned} \quad (2.17)$$

Similarly, we consider the 2BSDE when neither α nor β are controlled:

$$\begin{aligned} Y_t^{\alpha, \beta} &= \Phi(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} f(s, X_s, Y_s^{\alpha, \beta}, Z_s^{\alpha, \beta}, U_s^{\alpha, \beta}, \alpha_s, \beta_s, \lambda_s) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s^{\alpha, \beta} \cdot dX_s \\ &\quad - \int_{t \wedge \tau}^{T \wedge \tau} dK_s^{\alpha, \beta} - \int_{t \wedge \tau}^{T \wedge \tau} U_s^{\alpha, \beta} dH_s, \quad \mathbb{P}\text{-a.s.}, \quad \mathbb{P} \in \mathcal{P}(t, x). \end{aligned} \quad (2.18)$$

We finally denote by $(\mathcal{Y}^{\alpha, \beta}, \mathcal{Z}^{\alpha, \beta}, \mathcal{U}^{\alpha, \beta})$ the unique solutions to the Lipschitz BSDE under

some probability $\mathbb{P}' \in \mathcal{P}_0$

$$\begin{aligned} \mathcal{Y}_t^{\mathbb{P}'; \alpha, \beta} &= \Phi(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} f(s, X_s, \mathcal{Y}_s^{\mathbb{P}'; \alpha, \beta}, \mathcal{Z}_s^{\mathbb{P}'; \alpha, \beta}, U_s^{\mathbb{P}'; \alpha, \beta}, \alpha_s, \beta_s, \lambda_s) ds \\ &\quad - \int_{t \wedge \tau}^{T \wedge \tau} \mathcal{Z}_s^{\mathbb{P}'; \alpha, \beta} \cdot dX_s - \int_{t \wedge \tau}^{T \wedge \tau} \mathcal{U}_s^{\mathbb{P}'; \alpha, \beta} dH_s. \end{aligned} \quad (2.19)$$

From the dynamic programming principle Proposition 2.3.6 together with Remark 2.3.9, we get

$$Y_0^\alpha = \text{ess inf}_{\mathbb{P}' \in \mathcal{P}(0, \mathbb{P}, \mathbb{G})} \mathcal{Y}_0^{\mathbb{P}', \alpha, \beta}, \quad \mathbb{P}\text{-a.s. for all } \mathbb{P} \in \mathcal{P}_0.$$

and

$$Y_0^{\alpha, \beta} = \text{ess inf}_{\mathbb{P}' \in \mathcal{P}(0, \mathbb{P}, \mathbb{G})} \mathcal{Y}_0^{\mathbb{P}', \alpha, \beta}, \quad \mathbb{P}\text{-a.s. for all } \mathbb{P} \in \mathcal{P}_0.$$

From the comparison theorem 2.3.10, we then deduce that for any $\mathbb{P} \in \mathcal{P}_0$

$$\begin{aligned} \underline{Y}_0 &= \text{ess sup}_{\alpha \in \mathcal{A}} \mathbb{P} \text{ess inf}_{\beta \in \mathcal{B}} Y_0^{\alpha, \beta} \\ &= \text{ess sup}_{\alpha \in \mathcal{A}} \mathbb{P} \text{ess inf}_{\beta \in \mathcal{B}} \text{ess inf}_{\mathbb{P}' \in \mathcal{P}(0, \mathbb{P}, \mathbb{G})} \mathcal{Y}_0^{\mathbb{P}', \alpha, \beta} \\ &= \text{ess sup}_{\alpha \in \mathcal{A}} \mathbb{P} \text{ess inf}_{\beta \in \mathcal{B}} \text{ess inf}_{\mathbb{P}' \in \mathcal{P}(0, \mathbb{P}, \mathbb{G})} \mathbb{E}^{\mathbb{P}'_\alpha} [\mathcal{K}_{T \wedge \tau}^{\alpha, \beta} \Phi(\xi) - \int_0^{T \wedge \tau} \mathcal{K}_t^{\alpha, \beta} c_t(\alpha_t, \beta_t) dt], \end{aligned}$$

where we have used a linearization technique for the last equality with $\mathbb{P}'_\alpha$ defined by

$$\frac{d\mathbb{P}'_\alpha}{d\mathbb{P}} = \mathcal{E} \left(\int_0^T \sigma^\top (\sigma \sigma^\top)^{-1} (s, X, \beta_s) \mu(s, X, \alpha_s, \nu_s) \cdot dW_s^\mathbb{P} + \int_0^T U_s dH_s \right).$$

Consequently and taking a supremum on $(\alpha; \mathbb{P}, \beta) \in \mathcal{A} \times \mathcal{N}^\alpha$ we get

$$\underline{V}_0(\xi) = \sup_{\alpha \in \mathcal{A}} \inf_{(\mathbb{P}, \beta) \in \mathcal{N}^\alpha} \mathbb{E}^\mathbb{P}[\underline{Y}_0].$$

Finally, going back to the previous steps, $(\alpha^*, \mathbb{P}^*, \beta^*)$ is optimal if and only if all the essential suprema and infima above are attained, that is if and only if (α^*, β^*) is a saddle point and satisfies Isaacs condition, and $\mathbb{P}^*$ is such that

$$\mathbb{E}^{\mathbb{P}^*}[\underline{Y}_0] = \mathcal{Y}_0^{\mathbb{P}^*, \alpha^*, \beta^*}, \quad \mathbb{P}^*\text{-a.s.}$$

that is $\underline{K}_{T \wedge \tau} = 0$, $\mathbb{P}^*\text{-a.s.}$ □

2.5 Discussion and outlook

The results of this chapter combine three classical ingredients of stochastic analysis – second-order BSDEs in the sense of [150, 136], progressive enlargement of filtration in the sense of [87, 82, 96], and weak-formulation stochastic control with controlled volatility [54, 74] – into a single framework in which all three sources of uncertainty (drift, volatility and termination time) act simultaneously. The principal technical contribution is the $\mathcal{P}_0$ -density Hypothesis 2.2.4, which forces a single kernel γ to represent the conditional law of τ under every non-dominated probability and is what enables the aggregation of the jump-size process and the auxiliary 2BSDE construction in Lemma 2.3.2 and Theorem 2.3.4.

Three directions seem natural for follow-up work.

Singular drivers. The integrability assumption $\text{ess sup}_\rho \mathbb{E}[\int_\rho^T |\lambda_s|^{\hat{p}} ds \mid \mathcal{G}_\rho] < \infty$ implies that τ may exceed the horizon T with positive probability. The complementary regime, in which τ is supported on $[0, T]$ and λ blows up near T , leads to 2BSDEs with degenerate or singular drivers and is closely related to the analysis of [83, 84]. Extending the existence and comparison results of Section 2.3 to this singular regime is the natural next step, and the bridge Chapter 4 of this thesis points to a market-making problem in which the regime change matters economically.

Quadratic generators. Assumption 2.2.8 keeps the generator Lipschitz in the BSDE arguments. The CARA pendant of the volatility-ambiguity problem $\underline{V}_0$ is naturally formulated through a quadratic generator in $\hat{a}^{1/2}Z$, via the Hu–Imkeller–Müller exponential change of variable $V_t^A = -e^{-\eta Y_t}$ (see, e.g., [76, 64]). Extending the present 2BSDE machinery to quadratic generators in the spirit of [30, 121] would make the framework directly applicable to CARA utility maximization under volatility ambiguity and a random horizon.

Portfolio delegation under random horizon

This chapter is based on the paper *Delegated portfolio management with random default*, joint work with Thibaut Mastrolia. The paper is to appear in *Mathematical Finance*.

3.1 Introduction

Delegating portfolio management from an investor to a professional fund manager is increasingly seen as a strategic move due to the growing complexity of financial markets and their fragmentation [86, 21, 153]. The financial landscape today is marked by rapid market fluctuations, changing regulations, and an extensive range of investment opportunities, all of which necessitate not only significant time and effort but also in-depth knowledge and experience to navigate effectively. For many investors, handing over portfolio management to a professional allows them to leverage the fund manager's expertise in areas such as professional oversight, diversification strategies, and sophisticated risk management techniques, capabilities that are often difficult to achieve independently.

In this context, fund managers are expected to deliver superior performance by utilizing their specialized skills and tools. Traditionally, the compensation structure for fund managers has been a blend of a fixed fee and a performance-based component. The fixed fee provides a stable income for the manager, while the performance-based component is designed to incentivize the manager to achieve better returns by aligning their interests with those of the investor. However, this conventional compensation structure raises an important question: Is it optimally designed to align the incentives of both the investor and the fund manager? The challenge lies in ensuring that the performance-based component effectively motivates the fund manager to act in the best interests of the investor, while also accounting for the inherent uncertainties and risks associated with market fluctuations. To address this, it is crucial to evaluate whether the existing compensation structures are adequately aligned with the investor's objectives and whether alternative models could offer better alignment. This involves exploring various compensation schemes and their impact on fund performance, risk

management, and overall investor satisfaction. Ultimately, a well-designed compensation structure should incentivize fund managers not only to maximize returns but also to manage risks prudently, ensuring that both the investor's and the manager's goals are harmoniously aligned in the pursuit of financial success.

From a mathematical perspective, this issue can be framed as a variation of the principal-agent (PA) problem in wealth management, see, for example, [55, 155, 153, 58, 104, 103, 53]. The PA framework in continuous time is a game-theoretical model designed to address problems of stochastic control, where one party (the principal) delegates decision-making authority to another party (the agent), whose actions are not directly observable and evolve in a stochastic environment. The agent controls the system by choosing a strategy that influences outcomes, but because of the information asymmetry, the principal cannot directly observe the agent's actions. Instead, the principal must design a contract based on the observed outcomes, which indirectly depend on the agent's actions. The goal is to structure this contract in a way that motivates the agent to exert optimal effort, while managing the inherent trade-offs between risk-sharing and incentives. Here, the investor (the principal) has an initial capital, $X_0 = x$, and seeks an agent to invest on their behalf. The principal is willing to negotiate a compensation scheme that incentivizes the agent based on portfolio performance and risks. In the literature, the PA problem was extensively studied in a simpler setting: in the seminal work of [75], the agent controls the drift of the process, but the utility is just drawn from the terminal value of the controlled process. In [145] we still see a control based on just the drift, but the utility is now directly dependent on inter-temporal payments. In both cases, because of the modeling choice of having a single Brownian motion, there is no moral hazard with respect to the control on the volatility. A concept introduced by [53], moral hazard in the volatility control is important when multiple sources of risks arise in the control problem. What the authors proved in this work, thanks to the mathematical advancement on singular change of measure, is that in optimal contracts there are some incentives to be given with respect to the quadratic variation of the controlled process and the co-variation with respect to the risk factors. Furthermore, in their work, instead of working from a probability perspective, they cast the problem as a stochastic control problem, and they restrict the analysis to an admissible family of contracts, proving no loss of generality in doing so, paving a simpler way to the recent literature on the subject. In the subsequent paper [54], the same result was proven by using second-order BSDE, and a recent work [45] simplified even more the theoretical guarantees for volatility controlled problems, showing the same results but using the standard theory of BSDE. To summarize, in these works, they define a class of admissible contracts, they prove there is no loss of generality in considering such a form, and they find

the optimal contract in this set. In practice, they rely on the work of [129, 17, 64, 130, 99, 128] on the existence of solutions for BSDEs in the Lipschitz and quadratic cases with or without jumping terms and random horizon. This will also be the building block for this work.

Portfolio optimization has been extensively studied in the literature, beginning with the pioneering work of [113]. More recent studies have explored continuous-time versions of the framework, as well as settings that incorporate jumps, as in [122]. The PA problem captures the scenario of a fund manager investing on behalf of an investor. Despite its practical relevance, the existing literature has not fully addressed a crucial aspect of this problem: the randomness of the investment horizon. More often than not, an investment in financial markets does not have a precise duration, and duration is used as a reference for setting risk aversion for the investor.

Our contribution aims to fill this gap by examining how decision-making strategies for both the agent's investment and the principal's compensation scheme are affected by the introduction of default times, adding uncertainty to the investment horizon. Allowing for a general default time makes the problem mathematically more challenging in two ways. First, the random termination τ need not be a stopping time of the market filtration, so one must augment the available information by progressively enlarging the filtration generated by the financial market with the filtration of the default event; see [85, 23, 2, 71] and the references therein for the relevant filtration-enlargement theory. Furthermore, as highlighted by [83], there are two distinct cases to consider for default times, each with different implications:

- **Unbounded Case:** If the maximum possible default time S exceeds the investment horizon T (or is infinite), it is uncertain whether a default will occur within the investment period. In this case, the investing problem is reduced to a utility maximization under random horizon. It has been solved in, for example, [96] which proves that the solution is related to a system of BSDEs with a jump admitting a solution via a decomposition approach coming from filtration enlargement theory. Without loss of generality, we can assume that $S = +\infty$, hence the name unbounded.
- **Bounded Case:** If S is less than T , default will certainly occur before the investment horizon ends, but the exact timing is unknown. Without loss of generality, we can assume $S = T$: if the default happens almost surely before T , then we can simply rewrite the a priori horizon as S . It has been solved in, for example, [84, 83] proving that the solution is related to a system of *degenerate* BSDEs with a jump.

These two cases have not only different interpretations but also different mathematical tools required to solve them. We motivate our study with two contrasting examples inspired

by cryptocurrency trading and life-insurance portfolio management. In highly speculative markets such as crypto assets, the investment horizon is intrinsically uncertain: markets may collapse abruptly due to technological failures, regulatory bans, or market inefficiencies that effectively “terminate” the investment opportunity. Therefore, we are in a situation where the intended investment horizon can be achieved, but a black-swan event may cause the market to cease to exist. The situation is completely different for life-insurance companies. Although the investment horizon of most products is in the order of 100 years, the effective horizon remains random, as the contract expires at the death time of the insured. Furthermore, regulatory and actuarial considerations impose strict constraints on the risk profile of managed portfolios. Insurers are required to invest in a manner that does not deviate too aggressively from prescribed benchmark allocations, such as long-duration bond indices or liability-driven reference portfolios. A convenient way to encode such tracking discipline is to introduce a non-negative parameter $\varepsilon \geq 0$ that penalizes deviations between the insurer’s strategy and some benchmark, for instance through a running quadratic tracking error. In this interpretation, the crypto setting corresponds to $\varepsilon = 0$, whereas the life-insurance application is characterized by $\varepsilon > 0$, thus bridging two seemingly different investment problems within a unified principal-agent framework under random horizon.

In general, therefore, the unbounded case can be seen as a default caused by a black swan event [131], or a crash that forces authorities to close the market (the Flash crash in May 2010, see [98]), or, in the blockchain, a hacker attack or, finally, in a more structured deal with investors, when a fund sees its money withdrawn with no or little notice. This often complicates the investment strategies for funds, and for this reason certain funds (i.e. hedge funds) have very strict policies on funding withdrawal. Mathematically, the BSDE related to this case is better behaved than the other case. The bounded case instead can be representative of the well-known and studied life-insurance market: in this case, the insurance policy can have a time horizon of over 100 years, so that we can claim, with probability one, that the investor is going to pass away before the natural termination of the contract. This means that, when calibrating the contract, both the agent and the investor are aware of the fact that the horizon will not be reached, and this is taken into account in both the trading strategies and the insurance payments. Mathematically, this formulation introduces a difficulty in the family of proposed contracts, as it generates a BSDE with a singular driver (see [83] and references therein), and it also poses some difficulties in the convergence of numerical methods.

The problem’s structure also presents challenges for numerical solutions. The partial differential equation (PDE) resulting from the Hamilton-Jacobi-Bellman (HJB) control problem has a varying coefficient dependent on the solution of a maximization problem involving

the solution itself. Addressing this requires a specialized approach using an “actor–critic” iterative algorithm, where the actor solves the PDE for a fixed coefficient, and the critic updates the maximization problem based on the actor’s latest guess. While several schemes could tackle this iterative process, the most effective in our experiments have proven to be Physics-Informed Neural Networks (PINNs), a machine-learning approach that embeds the governing partial differential equations directly into the training loss so that the network respects known physical (here, stochastic-control) constraints. The surge of this methodology was made possible by automatic differentiation, one of the most useful but perhaps underused techniques in scientific computing, which allows differentiating the network with respect to both its input coordinates (yielding the derivatives appearing in the loss) and its parameters (driving the optimization); we refer to the survey [20] for a comprehensive treatment. The framework was introduced and expanded by works such as [140, 149], and is particularly effective for high-dimensional partial (integro-)differential equations arising in fluid dynamics, electromagnetism, biology, or finance (see, *e.g.*, [13, 14]), bridging the gap between traditional numerical solvers and modern machine learning techniques while reducing reliance on large datasets. Despite all these difficulties, this default-time formulation is crucial for practical applications, as it makes both contract incentives and trading strategies more robust. This study contributes significantly to the literature on principal-agent problems, extending its applicability to real-world financial scenarios by providing some insights into the effects of random investment horizons and default times.

The organization of the paper is as follows. In Section 3.2, we present the mathematical formulation of the Principal-Agent problem under time uncertainty, describing the underlying stochastic framework, the controlled wealth process governing the system dynamics, and all the necessary assumptions to make the problem tractable. We will further define some classes of admissible contracts, and the optimization problems for both the principal and the agent. In Section 3.3, we solve the problem sequentially, first focusing on the agent’s optimal strategy, which is going to be the same for both cases of bounded and unbounded default. Despite the fundamental difference in the proof of existence, this trading strategy has the same form in both cases and will be plugged into the principal’s problem. Then we will derive the Hamilton-Jacobi-Bellman (HJB) equation for the principal and we will claim, with a verification theorem, the existence of the solution to the partial differential equation (PDE) that encapsulates the principal’s optimization problem in the two proposed settings. Then, Section 3.4 provides numerical examples that demonstrate the implementation of the theoretical results in concrete scenarios, for both cases, using default times from the families of the Beta and Exponential distributions. The goals are multiple: to show the differences

arising within the same case, but also across the two different cases. Finally, we also want to highlight the sub-optimal behavior of a subset of optimal contracts that mimic real-world compensation schemes.

The code of our numerical simulations is publicly available on <https://github.com/GennaroAlberto/PrincipalAgent>.

3.2 The model and the optimization problem

3.2.1 Risky assets and portfolio dynamics

We consider a financial market represented by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with a d -dimensional Brownian motion denoted by W and a finite horizon $T > 0$. We denote by $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]}$ the natural filtration of this Brownian motion. This market consists of m risky assets with vector price S_t at time t with no risk-free rate. The risky assets have the following dynamics:

$$\frac{dS_t^i}{S_t^i} = b_t^i dt + \sigma_t^i dW_t \quad \forall i = 1 : m,$$

where σ^i, b^i are respectively $\mathbb{R}^{1 \times d}$ -valued and $\mathbb{R}$ -valued bounded $\mathbb{F}$ -predictable processes. We define the $\mathbb{R}^{m \times d}$ -valued covariance matrix σ where $\sigma^{i,j}$ is the j -th component of σ^i . We assume that $\sigma \sigma^\top$ is an invertible matrix, that is $\sigma \sigma^\top$ is $\mathbb{P} \otimes dt$ -a.e. elliptic. Note that

$$dS_t = B_t dt + \Sigma_t dW_t,$$

where $S = (S^1, \dots, S^m)$, $B = (b^i S^i)_{1 \leq i \leq m}$ and Σ is a $m \times d$ matrix with the i -th row is $\Sigma_t^i = S_t^i \sigma_t^i$.

We define $\theta_t = \sigma_t^\top (\sigma_t \sigma_t^\top)^{-1} b_t$. Let $\pi_t = (\pi_t^i)_t$ be a vector in $\mathbb{R}^{1 \times m}$ representing the fraction of money invested in every asset at time t . We refer to it as the *investment strategy* of the portfolio manager. We set $\beta_t = \pi_t \sigma_t$. Note that we can also refer to β or π as the investment strategy interchangeably up to the volatility factor σ . For every π , one can define a probability measure $\mathbb{P}^\pi$ such that the dynamics of the value of the portfolio starting with $X_0 = x$ is then given by¹

$$X_t := x + \int_0^t \pi_s \sigma_s dW_s + \int_0^t \pi_s b_s ds,$$

¹We refer to Appendix A in [14] for the rigorous formulation of the problem and the choice of the probability $\mathbb{P}^\pi$.

or equivalently

$$X_t := x + \int_0^t \beta_s dW_s + \int_0^t \beta_s \theta_s ds.$$

3.2.2 Default time and enlargement of filtration

The default time is represented by a random variable τ taking values in $\mathbb{R}^+$. We define the default process by $H_t := \mathbf{1}_{\tau \leq t}$. Note that this process is not necessarily $\mathbb{F}$ -measurable, and may require extra information or can be simply exogenous to the system and so independent of W . We thus enlarged the available information and also the filtration $\mathbb{F}$ taking into account the information generated by the default time occurrence.

Definition 3.2.1. Let $\mathcal{H}_u := \sigma(H_s, s \in [0, u])$ be the σ -algebra generated by H until time $u \geq 0$. Given a filtered space $(\Omega, \mathcal{F}_T, \mathbb{F}, \mathbb{P})$, the enlarged filtration

$$\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}, \quad \mathcal{G}_t = \bigcap_{\epsilon > 0} \{\mathcal{F}_{t+\epsilon} \vee \mathcal{H}_{t+\epsilon}\},$$

is the smallest enlargement of $\mathbb{F}$ such that τ is a $\mathbb{G}$ -stopping time.

Remark 3.2.2. H_t is not $\mathbb{F}$ -measurable but it is $\mathbb{G}$ -measurable stochastic process.

The goal is to ensure that the inaccessible default time τ enables us to enlarge the filtration to $\mathbb{G}$ and transferring the martingale property from $\mathbb{F}$ to $\mathbb{G}$ known as the immersion property or H-hypothesis. The first fundamental hypothesis is to set the existence of a (conditional) density for the default time with a certain property.

Hypothesis 3.2.3 (Density Hypothesis). *For any $t \geq 0$, there exists an $\mathcal{F}_t \otimes \mathcal{B}((0, \infty))$ measurable map $\gamma(t, \cdot)$ such that*

$$\mathbb{P}(\tau \geq x | \mathcal{F}_t) = \int_x^\infty \gamma(t, u) du \quad \forall x \geq 0,$$

and γ admits the two-regime form

$$\gamma(t, u) = \gamma_{\text{pre}}(t, u) \mathbf{1}_{u > t} + \gamma_{\text{pre}}(u, u) \mathbf{1}_{u \leq t},$$

where the pre-default density $\gamma_{\text{pre}}(t, u)$ is, for each fixed $u > 0$, a $(\mathbb{P}, \mathbb{F})$ -martingale in t on $\{t < u\}$ (so that $\gamma(t, u)$ is frozen at $\gamma(u, u)$ once the time of observation t has passed u ; the notation $\gamma(t, u) = \gamma(u, u) \mathbf{1}_{t \geq u}$ that one sometimes meets in the literature is a shorthand for this regime split, not a multiplicative prefactor that would erroneously vanish for $t < u$).

As a consequence of this assumption, see for example [28, 63], any $\mathbb{F}$ -martingale is also a $\mathbb{G}$ -martingale. Furthermore, and still under the density hypothesis, the process H admits an

absolutely continuous compensator, *i.e.*, there exists a non-negative $\mathbb{G}$ -predictable process $\lambda^{\mathbb{G}}$, such that the compensated process M defined by

$$M_t := H_t - \int_0^t \lambda_s^{\mathbb{G}} ds$$

is a $\mathbb{G}$ -martingale. The compensator vanishes after time τ (therefore $\lambda_t^{\mathbb{G}} := \lambda_t \mathbf{1}_{t \leq \tau}$) and

$$\lambda_t := \frac{\gamma(t, t)}{\mathbb{P}(\tau > t | \mathcal{F}_t)}$$

is a $\mathbb{F}$ -predictable process. For a complete and deeper discussion on the properties of enlarged filtrations, we refer the reader to [71]. We set $\Lambda_t := \int_0^t \lambda_s ds$. As a consequence of Proposition 4.4 in [63]

$$\mathbb{P}(\tau > t | \mathcal{F}_t) = e^{-\Lambda_t}.$$

We now turn to the integrability of the process λ and the support of the default time τ . Denoting by $\mathcal{T}(\mathbb{A})$ the set of $\mathbb{A}$ -stopping times (so we will have $\mathcal{T}(\mathbb{F})$ or $\mathcal{T}(\mathbb{G})$), we consider two cases

Hypothesis A - unbounded default.

$$\text{ess sup}_{\rho \in \mathcal{T}(\mathcal{G})} \mathbb{E} \left[\int_{\rho}^T \lambda_s ds \mid \mathcal{G}_{\rho} \right] < +\infty \quad (\text{HA})$$

As a direct consequence of the tower property and since $\mathbb{F} \subseteq \mathbb{G}$, Hypothesis A leads to

$$\text{ess sup}_{\rho \in \mathcal{T}(\mathbb{G})} \mathbb{E} \left[\int_{\rho}^T \lambda_s ds \mid \mathcal{F}_{\rho} \right] < +\infty$$

Consequently, $\mathbb{P}(\tau \in [0, T]) < 1$, so the support of τ extends strictly beyond $[0, T]$.

Hypothesis B - bounded default

$$\text{ess sup}_{\rho \in \mathcal{T}(\mathbb{G})} \mathbb{E} \left[\int_{\rho}^t \lambda_s ds \mid \mathcal{G}_{\rho} \right] < +\infty \quad \forall t < T, \quad \mathbb{E} \left[\Lambda_T \right] = \infty. \quad (\text{HB})$$

Consequently, $\mathbb{P}(\tau \in [0, T]) = 1$, the support of τ is included in $[0, T]$.

3.2.3 Admissible strategy and contracts

Admissible strategies π may be restricted to a closed subset of $\mathbb{R}^m$. We set the rigorous definition of an admissible strategy below, following [76, Definition 1].

Definition 3.2.4 (Admissible strategy with constraints). Let C be a closed and convex set in $\mathbb{R}^{1 \times m}$. The set of admissible strategies, denoted by $\mathcal{A}$, consists of m -dimensional $\mathbb{G}$ -predictable process π (equivalently $\mathbb{F}$ -predictable on $\{t \leq \tau\}$ under the immersion hypothesis) such that $\mathbb{E}[\int_0^T |\pi_t|^2 dt] < \infty$ and $\pi_t \in C$, $dt \otimes \mathbb{P}$ -a.e.

Note that due to the nature of the problem considered and by considering a compensation ξ , other integrability conditions are transferred to the admissibility of the contract below. Some examples of admissible strategy sets include:

- $C = [0, 1]^m$, that is π is a proportion of the total wealth X^π that is invested in the portfolio with no possibility to borrow or spend more than the actual value of X^π . It does not permit shorting stocks (i.e. selling stocks borrowed).
- $C = [-1, 1]^m$, which does not permit leveraged positions.
- $C = [-M, M]^m$ assuming that the investor can spend or borrow as much money as needed limited to a fraction M of the total wealth (possibly greater or less than 1 or -1).

For a symmetric positive definite matrix Q , we define the squared norm of a column vector x by

$$||x||_Q^2 = x^T Q x.$$

This norm is equivalent to the euclidean norm in $\mathbb{R}^n$, and the constants of equivalence are the smallest and biggest eigenvalues of the matrix Q . From this definition of Q -norm, we define the squared Q -distance of x to the set C as

$$\text{dist}_Q^2(x, C) := \inf_{y \in C} \{||x - y||_Q^2\}.$$

The contract ξ proposed by the investor follows the idea of [53]. We denote by $\eta > 0$ a risk aversion parameter for the portfolio manager with CARA exponential utility function $U_A(x) = -e^{-\eta x}$, and assume that the manager may have to track an index and stay close to a benchmark strategy α over time given by a deterministic and bounded function of time with penalty coefficient $\varepsilon \geq 0$. A more rigorous definition is given in Section 3.2.4 below.

Definition 3.2.5 (Admissible contract with contractible variables). We denote by Ξ the set of admissible contracts ξ composed of $\mathcal{G}_{\tau \wedge T}$ -measurable random variable $\xi = Y_{T \wedge \tau}^{Y_0, Z, Z^X, U, \Gamma^X, \Gamma}$ where, on $[0, T]$, the process Y is defined by the SDE below (the contract is the value of this process stopped at $T \wedge \tau$), controlled by $\mathbb{G}$ -predictable real-valued processes $U, Z =$

$(Z^i)_{1 \leq i \leq m}, Z^X, \Gamma^X, \Gamma = (\Gamma_i)_{1 \leq i \leq m}$ such that $\mathcal{I}_m - \Gamma_t^X \sigma_t \sigma_t^T$ is a positive definite matrix² and

$$\begin{aligned} Y_t^{Y_0, Z, Z^X, U, \Gamma^X, \Gamma} &= Y_0 + \int_0^t \sum_{i=1}^m \frac{Z_r^i}{S_r^i} dS_r^i + \int_0^t Z_r^X dX_r + \int_0^t U_r dH_r \\ &\quad + \frac{1}{2} \int_0^t (\Gamma_r^X + \eta(Z_r^X)^2) d\langle X, X \rangle_r + \int_0^t \sum_{i=1}^m \frac{\Gamma_r^i}{S_t^i} d\langle S^i, X \rangle_r \\ &\quad - \int_0^t F(r, Z_r, Z_r^X, \Gamma_r, \Gamma_r^X, U_r) dr, \end{aligned}$$

where

$$F(t, z, z_x, g, g_x, u) = \sup_{\nu \in C} f(t, z, z_x, g, g_x, u, \nu),$$

with $f : [0, T] \times \Omega \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R} \times C$ by (here and throughout, $z \in \mathbb{R}^{1 \times m}$ is treated as a row vector)

$$\begin{aligned} f(t, z, z_x, g, g_x, u, \nu) &= zb_t + z_x \nu \sigma_t \theta_t + \frac{1}{2} g_x \|\nu \sigma_t\|^2 - \frac{\varepsilon}{2} \|\nu - \alpha_t\|^2 + \sum_{i=1}^m g^i \nu^i \sigma_t^i (\sigma_t^i)^\top \\ &\quad - \frac{\lambda_t}{\eta} (e^{-\eta u} - 1) - \frac{\eta}{2} \|z \sigma_t\|^2, \end{aligned}$$

so that

$$F(t, z, z_x, g, g_x, u) = \frac{1}{4} q_t^T Q_t^{-1} q_t + zb_t - \frac{\varepsilon}{2} \|\alpha_t\|^2 - \frac{\lambda_t}{\eta} (e^{-\eta U_t} - 1) - \frac{\eta}{2} \|z \sigma_t\|^2 + \text{dist}_Q^2(d_t, C) \quad (3.1)$$

where

$$\begin{aligned} Q_t &= \frac{1}{2} (\mathcal{I} - g_x \sigma_t \sigma_t^T), \quad q_t = (\sigma_t \theta_t + \sigma_t \sigma_t^T g + \varepsilon \alpha_t - \eta z_x \sigma_t \sigma_t^T z) \\ d_t &= \frac{1}{2} Q_t^{-1} q_t, \end{aligned}$$

and there exists $\eta' > \eta$ such that

$$\mathbb{E} \left[\int_0^T (\|Z_s\|^2 + \|Z_s^X\|^2 + \|\Gamma_s\| + \|\Gamma_s^X\| + \|U_s\|^2 \lambda_s) ds + \sup_{0 \leq t \leq T} e^{\eta' |Y_t^{Y_0, Z, Z^X, U, \Gamma, \Gamma^X}|} \right] < \infty.$$

The set of processes $Z, Z^X, U, \Gamma, \Gamma^X$ satisfying this integrability condition is denoted by $\mathcal{U}$ while their restriction to $[0, t]$ is denoted by $\mathcal{U}_t$ for any $t < T$.

Economic interpretation.

- Y_0 is a fixed compensation determined by the reservation utility of the portfolio manager;

² $\mathcal{I}_m$ denotes the identity matrix in dimension m .

- the integrand process Z^i represents a compensation with respect to the evolution of the i -th asset S^i , if it is observable by the investor and so contractible;
- Z^X is a compensation term with respect to the portfolio dynamics, always observable by the investor;
- U is a compensation with respect to the default risk of the market;
- Γ^i is a compensation with respect to the covariation of S^i and X while the term $\Gamma^X + \eta(Z^X)^2$ is a compensation driven by the risk aversion of the manager with respect to the quadratic variation of the portfolio;
- F is the certainty equivalent utility gained by the portfolio manager when solving her optimization problem. The gain resulting from this optimization is transferred into the contract.

Remark 3.2.6. A rigorous justification of this set of contracts is given in Appendix A.2. This appendix provides a rigorous framework for defining the probability space and the weak formulation of the portfolio manager's problem together with its solution by using second-order BSDEs and the recent results in Chapter 2.

Remark 3.2.7. We can refine the set of contracts depending on the information available to the investor (see [53]).

- We denote by $\Xi^\circ \subset \Xi$ the set of random variables $\xi = Y_T^{Y_0, Z^X, U, \Gamma^X}$ with $Z^i = \Gamma^i = 0$ corresponding to the case where S is not observable by the investor and so not contractible.
- The set of linear contracts $\Xi^l \subset \Xi$ defined by

$$\begin{aligned} \xi = & Y_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{Z_s^i}{S_s^i} dS_s^i + p(X_{T \wedge \tau} - X_0) + \int_0^{T \wedge \tau} U_s dH_s + \frac{1}{2} \int_0^{T \wedge \tau} (\Gamma_s^X + \eta p^2) d\langle X, X \rangle_s \\ & - \int_0^{T \wedge \tau} F(s, Z_s, p, \Gamma_s, \Gamma_s^X, U_s) ds + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\Gamma_s^i}{S_s^i} d\langle S^i, X \rangle_s \end{aligned}$$

with contractible S . The idea around this contract is that in practice, most of fund managers ask as a compensation to their clients a fixed percentage of the terminal wealth, forcing $Z_s^X = p$ in Ξ , where p is the fixed percentage the agent receives.

In all these cases, $Z = (Z^i)_{1 \leq i \leq m}, Z^X, U, \Gamma, \Gamma^X$ are predictable processes such that $\xi \in \Xi$ and all the stochastic integrals are martingales.

Remark 3.2.8. Note that the set of contracts Ξ is stated without loss of generality as explained in [54, 45] as soon as we consider general contracts ξ with exponential moment of any order, requiring to refine the definition of $\mathcal{U}$ in order to apply the existence results in [96, 83].

3.2.4 Delegated portfolio management and bi-level stochastic programming

We are assuming that the portfolio manager is receiving the contract ξ and optimally chooses a strategy π in order to remain close to a benchmark strategy α , given by a deterministic bounded function of time, with penalty coefficient $\varepsilon \geq 0$, *à la* Almgren-Chriss, see [5], so that the objective of the manager is to solve for any contract $\xi \in \Xi$ fixed

$$V_0^A(x, \xi) = \sup_{\pi \in \mathcal{A}} J^A(\pi; x, \xi), \quad (\text{A})$$

with

$$J^A(\pi; x, \xi) := \mathbb{E}[U^A(\xi - \frac{\varepsilon}{2} \int_0^{T \wedge \tau} \|\pi_s - \alpha_s\|^2 ds)],$$

and where $\varepsilon \geq 0$ represents the weight on the penalty. Note that when $\varepsilon = 0$ we do not track any index, which is consistent with the cryptocurrency motivation. In contrast, taking $\varepsilon > 0$ introduces index-tracking power (as in an ETF) and aligns with the pension-plan motivation. In our model, the investor chooses a full delegation of the portfolio management to the manager and thus lets the manager choose the optimal strategy to optimize the terminal value of the portfolio under default. In this *second-best* case, the contracting problem is reduced to solving a bi-level optimization under constraint as follows when S is contractible

$$V_0^P(x) := \sup_{(\xi, \hat{\pi}) \in \Xi \times \mathcal{A}} \mathbb{E}[U^P(X_{T \wedge \tau} - \xi)] \quad (\text{P})$$

subject to

- (R): $V_0^A(x, \xi) \geq R_0$
- (IC): $V_0^A(x, \xi) = J^A(\hat{\pi}; x, \xi)$.

We will refer to the first problem (A) as the Problem of the Agent, and the bi-level program (P) as the Problem of the Principal. Furthermore, we want to emphasize the fact that τ does not need to be a stopping time in the natural filtration given by the assets' dynamics. Despite this, most results regarding default times are based on stopping-time theory, so we want to work with a filtration such that τ can be a stopping time.

Remark 3.2.9 (First-best solution). Alternatively, one could assume that the investor controls both the allocation strategy and the compensation proposed to the portfolio manager. In this *first-best* situation, the problem of the principal (investor) is reduced to an optimization of its projected P&L over both ξ and π subject only to the reservation constraint. That is, solving:

$$V_0^{FB} = \sup_{(\xi, \pi)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^\pi - \xi) \right],$$

subject to

$$J^A(\pi; x, \xi) \geq R_0.$$

Note that $V_0^{FB} \geq V_0^P$. We refer to Appendix A.1 for the details and the solution of this first-best benchmark.

Remark 3.2.10 (Extension to Merton's problem with consumption). A further extension of this problem considers an investor who chooses to consume a portion of the portfolio continuously over time. Here, the consumption strategy c_t is given by an $\mathbb{F}$ -adapted random process. The specific mathematical derivations for this case are provided in Appendix A.3.

3.3 Optimal contract and investment strategy

3.3.1 Optimal investment with random horizon: solving the agent problem

A common approach in the continuous stochastic optimization literature is based on the solution of Backward Stochastic Differential Equations by solving a martingale optimality principle. We refer the reader to [76] for a detailed explanation of the method in a continuous setting and to [122, 96, 83] for extensions to discontinuous processes or default times. The general idea is to generate a family of super-martingales (R^π) indexed by the control variable, in our case the investment strategy π , with terminal condition the objective function of the agent J^A at time $T \wedge \tau$. If we are able to find a specific control $\hat{\pi}$ such that $R^{\hat{\pi}}$ is a martingale, this control is optimal for (A), and the optimal value is given by the process corresponding to the optimal control at time 0.

Lemma 3.3.1 (Martingale Optimality Principle). *Let $\xi \in \Xi$ and let $(R_{t \wedge \tau}^\pi)_{t \in [0, T]}$ be a family of stochastic processes indexed by the strategy $\pi \in \mathcal{A}$ such that*

$$(i) \quad R_{T \wedge \tau}^\pi = U^A\left(\xi - \frac{\varepsilon}{2} \int_0^{T \wedge \tau} \|\pi_s - \alpha_s\|^2 ds\right),$$

(ii) R^π is a $\mathbb{G}$ -supermartingale and R_0^π is constant for all $\pi \in \mathcal{A}$,

(iii) there exists $\hat{\pi} \in \mathcal{A}$ such that $R^{\hat{\pi}}$ is a $\mathbb{G}$ -martingale.

Then, $\hat{\pi}$ is a solution to the maximization problem (A).

Proof. Take $\pi \in \mathcal{A}$. Then, we have

$$J^A(\pi; x, \xi) \stackrel{(i)}{=} \mathbb{E}[R_{T \wedge \tau}^\pi] \stackrel{(ii)}{\leq} R_0^\pi = R_0^{\hat{\pi}} \stackrel{(iii)}{=} \mathbb{E}[R_{T \wedge \tau}^{\hat{\pi}}] \stackrel{(i)}{=} J^A(\hat{\pi}; x, \xi).$$

Here, equality (i) uses the terminal condition $R_{T \wedge \tau}^\pi = U^A(\xi - \frac{\varepsilon}{2} \int_0^{T \wedge \tau} \|\pi_s - \alpha_s\|^2 ds)$ from hypothesis (i), inequality (ii) is the $\mathbb{G}$ -supermartingale property from hypothesis (ii) together with R_0^π being constant, and equality (iii) is the $\mathbb{G}$ -martingale property of $R^{\hat{\pi}}$ from hypothesis (iii). $\square$

Let $\xi \in \Xi$ be fixed. Independently of the boundedness of the default time, that is either under Hypothesis A or Hypothesis B, we define

$$R_{t \wedge \tau}^\pi := U^A\left(Y_{t \wedge \tau}^\pi - \frac{\varepsilon}{2} \int_0^{t \wedge \tau} \|\pi_s - \alpha_s\|^2 ds\right),$$

where Y^π is defined by

$$\begin{aligned} Y_{t \wedge \tau}^\pi &= Y_0 + \int_0^{t \wedge \tau} \sum_{i=1}^m \frac{Z_s^i}{S_s^i} dS_s^i + \int_0^{t \wedge \tau} Z_s^X dX_s + \int_0^{t \wedge \tau} U_s dH_s \\ &\quad + \frac{1}{2} \int_0^{t \wedge \tau} (\Gamma_s^X + \eta(Z_s^X)^2) d\langle X, X \rangle_s - \int_0^{t \wedge \tau} F(s, Z_s, Z_s^X, \Gamma_s, \Gamma_s^X, U_s) ds \\ &\quad + \int_0^{t \wedge \tau} \sum_{i=1}^m \frac{\Gamma_s^i}{S_s^i} d\langle S^i, X \rangle_s \end{aligned}$$

Theorem 3.3.2. Assume that 3.2.3 and either Hypothesis A or Hypothesis B are satisfied. For any $\xi \in \Xi$, the optimal strategy that solves (A) is

$$\hat{\pi}_t = \pi^*(Z_t, Z_t^X, \Gamma_t, \Gamma_t^X), \text{ with } \pi^*(z, z_x, g, g_x) := \text{proj}(e_t, C) \quad (3.2)$$

and the optimal value is given by $V_0^A(x, \xi) = -e^{-\eta Y_0}$, where

$$e_t = \frac{1}{2} q_t Q_t^{-1}, \quad q_t = (\sigma_t \theta_t + \sigma_t \sigma_t^T \Gamma_t + \varepsilon \alpha_t - \eta Z_t^X \sigma_t \sigma_t^T Z_t^T)$$

and

$$Q_t = \frac{1}{2} (\mathcal{I} - \Gamma_t^X \sigma_t \sigma_t^T).$$

Proof. The proof follows the same idea as [76] for the continuous case and [83, 122] for the discontinuous case extended to the multi-dimensional case and the contract ξ fixed by the principal. Note that $Z = (Z^i)_{i=1}^m$ and $\Gamma = (\Gamma^i)_{i=1}^m$ are already $\mathbb{R}^m$ stochastic row vectors and that we are going to denote by C the set of admissible trading strategies $\pi = (\pi_t)_{t \in [0, T]}$. The proof will be based on Ito's formula with Poisson jumps (as the intensity of our jump is the same as that of a simple Poisson process with varying intensity λ) and we refer to [138, 79], for more details on stochastic calculus with semi-martingales; and the notion of *Doléans-Dade exponential process* (DDE). Given a semimartingale P , we denote its DDE by

$$\mathcal{E}(P)_t = \exp\{P_t - \frac{1}{2}\langle P^c \rangle_t\} \prod_{0 \leq s \leq t} \{(1 + \Delta_s P) \exp(-\Delta_s P)\},$$

where P^c is the continuous component of the path of P while $\Delta_s P := P_s - P_{s-}$. Note that the DDE of P is a solution of the following SDE $dY_t = Y_{t-} dX_t$ and that, provided P is a *BMO*-martingale, such that its jump is greater than $-1 + \delta$, with $\delta > 0$, the resulting DDE is a martingale as well.

To find the optimal solution, we are applying the so-called martingale optimality principle. We define the following family of stochastic processes indexed by the strategy π

$$R_t^\pi = U^A(Y_t - \frac{\varepsilon}{2} \int_0^t \|\pi_s - \alpha_s\|^2 ds).$$

We set $B_t^\pi := Y_t - \frac{\varepsilon}{2} \int_0^t \|\pi_s - \alpha_s\|^2 ds$, so that $R_t^\pi = -\exp(-\eta B_t^\pi)$. Note that the Ito's decomposition of B^π is given by

$$\begin{aligned} dB_t^\pi &= \sum_{i=1}^m \frac{Z_t^i}{S_t^i} dS_t^i + Z_t^X dX_t + U_t dH_t + \frac{1}{2}(\Gamma_t^X + \eta(Z_t^X)^2) d\langle X, X \rangle_t \\ &\quad - (F(t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t) + \frac{\varepsilon}{2} \|\pi_t - \alpha_t\|^2) dt + \sum_{i=1}^m \frac{\Gamma_t^i}{S_t^i} d\langle S^i, X \rangle_t, \quad t \in [0, T \wedge \tau]. \end{aligned}$$

Hence,

$$dB_t^\pi = J_t^\pi dt + (Z_t \sigma_t + Z_t^X \pi_t \sigma_t) dW_t + U_t dH_t,$$

where

$$J_t^\pi = Z_t b_t + Z_t^X \pi_t \sigma_t \theta_t + \frac{1}{2}(\Gamma_t^X + \eta(Z_t^X)^2) \|\pi_t \sigma_t\|^2 - F(t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t) + \pi_t \sigma_t \sigma_t^T \Gamma_t - \frac{\varepsilon}{2} \|\pi_t - \alpha_t\|^2.$$

We thus deduce that

$$dR_t^\pi = -\eta R_t^\pi J_t^\pi dt + \eta R_t^\pi (Z_t \sigma_t + Z_t^X \pi \sigma_t) dW_t + R_{t-}^\pi (e^{-\eta U_t} - 1) dM_t$$

where we used the fact that $dH_t = dM_t + \lambda_t dt$. Note that $R_0 = -\exp(-\eta Y_0) < 0$ and $J^\pi \leq 0$ and is equal to zero when π is chosen by maximizing

$$-\pi_t Q_t \pi_t^T + \pi_t q^T + d_t$$

with

$$Q_t = \frac{1}{2}(\mathcal{I} - \Gamma_t^X \sigma_t \sigma_t^T),$$

$$q_t = (\sigma_t \theta_t + \sigma_t \sigma_t^T \Gamma_t + \varepsilon \alpha_t - \eta Z_t^X \sigma_t \sigma_t^T Z_t^T),$$

and

$$c_t = Z_t b_t - F - \frac{\eta}{2} \|Z_t \sigma_t\|^2 - \frac{\varepsilon}{2} \|\alpha_t\|^2 - \frac{\lambda_t}{\eta} (e^{-\eta U_t} - 1).$$

Therefore,

$$\pi^*(Z_t, Z_t^X, \Gamma_t, \Gamma_t^X) = \text{proj}(e_t, C)$$

so that π^* is optimal and is unique because our set C is closed and convex. $\square$

Remark 3.3.3. Note that if $m = d = 1$ and $C = \mathbb{R}$ and σ and b are constant we obtain

$$\begin{aligned} F(t, z, z_x, g, g_x, u) &= \frac{1}{2} \left(\frac{z_x \sigma \theta - \varepsilon \alpha_t - \eta z z_x \sigma^2 + \sigma^2 g}{\sqrt{1 - g_x \sigma^2}} \right)^2 + z b \\ &\quad - \frac{\varepsilon}{2} \alpha_t^2 - \frac{\lambda_t}{\eta} (e^{-\eta u} - 1) - \frac{\eta}{2} z^2 \sigma^2. \end{aligned}$$

Remark 3.3.4. As a consequence of the constraint (IC), we are implicitly looking for contracts such that $\hat{\pi} \in \mathcal{A}$. This requires to consider $(Z, Z^X, \Gamma, \Gamma^X, U)$ such that $\hat{\pi} = \pi^*(Z, Z^X, \Gamma, \Gamma^X) \in \mathcal{A}$, otherwise (IC) is not satisfied and there are no optimal contracts.

Remark 3.3.5 (Extension to CRRA utilities). We have shown the optimal strategy for the CARA utility. We extend the analysis to power CRRA utilities in Appendix A.4. While for general utilities, we might not be able to find an explicit formula for the driver F and the optimal strategy π , it is possible to find them at least numerically.

3.3.2 The optimal contract and verification results

In the previous chapter, we found the optimal strategy π^* for the agent under the best optimal contract with contractible S . Given that strategy, we now want to solve the problem for the

principal of (P), i.e., we want to find the compensation ξ to maximize the principal utility:

$$V_0^P = \sup_{\xi \in \Xi \text{ s.t. } V_0^A(x, \xi) \geq R} \mathbb{E}[U^P(X_{T \wedge \tau}^{\pi^*} - \xi)] \quad (3.3)$$

where we used $X_{T \wedge \tau}^{\pi^*}$ to stress that now the object of our maximization problem is dependent on the agent's best strategy π^* . Before diving into the intuition behind how to solve this problem, it is important to show how this becomes a control problem so that an HJB equation can be derived. To do so, we first want to work with ξ s.t. $V_0^A(x, \xi) \geq R$. We recall that $V_0^A(x, \xi) = -\exp(-\eta(Y_0))$, so that

$$V_0^A(x, \xi) \geq R \iff Y_0 \geq -\frac{\log(-R)}{\eta} := \hat{Y}_0.$$

Moreover,

$$\begin{aligned} \xi = Y_{T \wedge \tau}^\pi &= Y_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{Z_s^i}{S_s^i} dS_s^i + \int_0^{T \wedge \tau} Z_s^X dX_s + \int_0^{T \wedge \tau} U_s dH_s \\ &+ \frac{1}{2} \int_0^{T \wedge \tau} (\Gamma_s^X + \eta(Z_s^X)^2) d\langle X, X \rangle_s - \int_0^{T \wedge \tau} F(s, Z_s, Z_s^X, \Gamma_s, \Gamma_s^X, U_s) ds \\ &+ \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\Gamma_s^i}{S_s^i} d\langle S^i, X \rangle_s \end{aligned} \quad (3.4)$$

Therefore, we can rewrite (3.3) as

$$V_0^P = \sup_{Y_0 \geq \hat{R}_0} \sup_{(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}} \mathbb{E}[U^P(X_{T \wedge \tau}^{\pi^*} - Y_{T \wedge \tau}^{\pi^*})].$$

Since U^P is nondecreasing in Y_0 , we deduce that

$$V_0^P = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}} \mathbb{E}[U^P(X_{T \wedge \tau}^{\pi^*} - \hat{Y}_{T \wedge \tau}^{\pi^*})], \quad (3.5)$$

where $\hat{Y}^{\pi^*}$ is given by (3.4) with $Y_0 = \hat{R}_0 := -\log(-R)/\eta$.

Remark 3.3.6 (Risk neutral investor). Assume that $U^P(x) = x$, then we note that

$$V_0^P = \hat{V}_0 - \hat{R}_0,$$

where (3.5) becomes

$$\hat{V}_0 := \sup_{(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}} \mathbb{E}[U^P(X_{T \wedge \tau}^{\pi^*} - \hat{Y}_{T \wedge \tau}^{\pi^*})],$$

with

$$\left\{ \begin{array}{l} d\hat{Y}_t^{\pi^*} = \sum_{i=1}^m \frac{Z_t^i}{S_t^i} dS_t^i + Z_t^X dX_t^{\pi^*} + U_t dH_t + \frac{1}{2}(\Gamma_t^X + \eta(Z_t^X)^2) d\langle X^{\pi^*}, X^{\pi^*} \rangle_t \\ \quad - F(t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t) dt + \sum_{i=1}^m \frac{\Gamma_t^i}{S_t^i} d\langle S^i, X^{\pi^*} \rangle_t \\ \hat{Y}_0^{\pi^*} = 0. \end{array} \right.$$

In order to derive the HJB equation with control process $(Z, Z^X, \Gamma, \Gamma^X, U)$, we will use the results from [25] assuming that τ has a probability density function f and a cumulative distribution function F , whether the support of τ is bounded in $[0, T]$ or unbounded in the next two subsections. To derive the HJB equation from this problem, we set an additional assumption enforcing Markovian properties of the intensity λ , drift and volatility processes.

Assumption (M). We assume that there exist three $\mathbb{F}$ -progressive measurable functions b, σ, λ such that $b_t = b(t, S_t)$, $\sigma_t = \sigma(t, S_t)$ and $\lambda_t = \lambda(t, S_t, X_t)$.

Note that if τ is independent of $\mathbb{F}$, as we will consider for numerical simulations, the density of τ , and hence λ , becomes a deterministic function for computations.

We define the following system of coupled SDEs with jumps with solution (X^{π^*}, Y^{π^*}) controlled by $(Z, Z^X, \Gamma, \Gamma^X, U)$

$$(SDE) \left\{ \begin{array}{l} dX_t^{\pi^*} = \pi_t^* \sigma(t, S_t) dW_t + \pi_t^* b(t, S_t) dt \\ d\hat{Y}_t^{\pi^*} = \sum_{i=1}^m \frac{Z_t^i}{S_t^i} dS_t^i + Z_t^X dX_t^{\pi^*} + U_t dH_t + \frac{1}{2}(\Gamma_t^X + \eta(Z_t^X)^2) d\langle X^{\pi^*}, X^{\pi^*} \rangle_t \\ \quad - F(t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t) dt + \sum_{i=1}^m \frac{\Gamma_t^i}{S_t^i} d\langle S^i, X^{\pi^*} \rangle_t \\ X_0^{\pi^*} = x, \\ Y_0^{\pi^*} = \hat{R}_0. \end{array} \right.$$

Bounded default time

Under Hypothesis B, the support of τ is included in $[0, T]$. Recalling the results in [25], we deduce that

$$V_0^P = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}} \mathbb{E} \left[\int_0^T U^P(X_t^{\pi^*} - \hat{Y}_t^{\pi^*}) f(t) dt \right],$$

where f is the density of τ . We introduce the following Hamilton-Jacobi-Bellman integro-partial differential equation.

$$(\text{bHJB}) \quad \begin{cases} 0 = \partial_t \phi(t, s, x, y) + U^P(x - y) f(t) \\ \quad + \sup_{(z^x, z, g^x, g, u)} \left\{ \mathcal{H}^\phi(t, s, x, y, z^x, z, g^x, g, u) \right\}, \quad t < T, \\ \phi(T, s, x, y) = 0, \quad (s, x, y) \in \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}. \end{cases}$$

where terminal condition $\phi(T, \cdot) = 0$ encodes the bounded-default assumption $\mathbb{P}(\tau \leq T) = 1$ (so the survival weight $1 - F_\tau(T)$ vanishes) and where $\mathcal{H}^\phi$ is a differential operator given by

$$\begin{aligned} \mathcal{H}^\phi(t, s, x, y, z^x, z, g^x, g, u) := & \\ & \sum_{i=1}^m \phi_{s^i} s^i b^i(t, s) + \phi_x \pi^*(z, z^x, g, g^x) \sigma(t, s) \theta(t, s) \\ & + \frac{1}{2} \text{Tr}(\Sigma(t, s) \Sigma(t, s)^\top D^2 \phi) + \frac{1}{2} \phi_{xx} \|\pi^*(z, z^x, g, g^x) \sigma(t, s)\|^2 \\ & + \phi_y \left[z b(t, s) + z^x \pi^*(z, z^x, g, g^x) \sigma(t, s) \theta(t, s) + \frac{1}{2} \|\sigma(t, s) \pi^*(z, z^x, g, g^x)\|^2 (g^x + \eta |z^x|^2) \right. \\ & \left. - F(t, z, z^x, g, g^x, u) + \pi^*(z, z^x, g, g^x) \sigma(t, s) \sigma(t, s)^\top g \right] \\ & + \frac{1}{2} \phi_{yy} (\|z_t \sigma(t, s)\|^2 + (z^x)^2 \|\pi^*(z, z^x, g, g^x) \sigma(t, s)\|^2) \\ & + \phi_{xy} (z^x \|\pi^*(z, z^x, g, g^x) \sigma(t, s)\|^2 + \pi^*(z, z^x, g, g^x) \sigma(t, s) \sigma(t, s)^\top z^T) \\ & + \sum_{i=1}^m (\pi^*(z, z^x, g, g^x) \sigma(t, s) \sigma^i(t, s)^\top s^i) (\phi_{x s^i} + \phi_{y s^i} z^x) + \lambda_t (\phi(t, s, x, y + u) - \phi(t, s, x, y)) \\ & + \sum_{i=1}^m \phi_{y s^i} z^i \sigma(t, s) \sigma^i(t, s)^\top s^i, \end{aligned}$$

where $\Sigma^i(t, s) = s^i \sigma^i(t, s)$ denotes the i -th row of the matrix Σ . We define

$$\hat{Z}_t := \hat{z}(t, S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}), \quad \hat{Z}_t^X = \hat{z}_x(t, S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}), \quad \hat{U}_t = \hat{u}(t, S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}),$$

$$\hat{\Gamma}_t = \hat{g}(t, S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}), \quad \hat{\Gamma}_t^X = \hat{g}_x(t, S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}),$$

where $\hat{z}, \hat{z}_x, \hat{g}, \hat{g}_x, \hat{u}$ optimize $\mathcal{H}^\phi$.

Theorem 3.3.7 (Verification Theorem - bounded case). *Assume that there exists a function ϕ twice continuously differentiable in space and differentiable in time, such that $\phi(t, s, x, y)$*

solves **(bHJB)**. Furthermore, assume that ϕ has a quadratic growth in y and polynomial growth in s, x such that

$$|\phi(t, s, x, y)| \leq \kappa(1 + |x|^p + \|s\|^p + |y|^2), \quad p > 1, \quad \kappa > 0.$$

Then, for each $t \in [0, T]$, the strategy $(\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U})$ is an optimal strategy for the control problem

$$\phi(0, S_0, x, 0) = V_0^P = \sup_{(Z^X, Z, \Gamma^X, \Gamma, U)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^{\pi^*} - \hat{Y}_{T \wedge \tau}^{\pi^*}) \right].$$

The optimal contract is given by

$$\begin{aligned} \xi^* &= \hat{Y}_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{Z}_s^i}{S_s^i} dS_s^i + \int_0^{T \wedge \tau} \hat{Z}_s^X dX_s + \int_0^{T \wedge \tau} \hat{U}_s dH_s \\ &+ \frac{1}{2} \int_0^{T \wedge \tau} (\hat{\Gamma}_s^X + \eta(\hat{Z}_s^X)^2) d\langle X, X \rangle_s - \int_0^{T \wedge \tau} F(s, \hat{Z}_s, \hat{Z}_s^X, \hat{\Gamma}_s, \hat{\Gamma}_s^X, \hat{U}_s) ds \\ &+ \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{\Gamma}_s^i}{S_s^i} d\langle S^i, X \rangle_s \end{aligned}$$

Remark 3.3.8. If we assume that $\sigma(t, s) = \sigma \in \mathbb{R}^{m \times d}$ and $b(t, s) = b \in \mathbb{R}^m$, we note that the solution v to the HJB equations **(bHJB)** does not depend on s . It can thus be rewritten as

$$\text{(bHJB)} \quad \begin{cases} 0 = \partial_t v(t, x, y) + U^P(x - y) f(t) \\ \quad + \sup_{(z^x, z, g^x, g, u)} \left\{ \mathcal{H}^v(t, x, y, \nabla v, \Delta v, z^x, z, g^x, g, u) \right\}, \quad t < T, \\ v(T, x, y) = 0, \quad (x, y) \in \mathbb{R} \times \mathbb{R}. \end{cases}$$

where

$$\begin{aligned} \mathcal{H}^v(t, x, y, z^x, z, g^x, g, u) &:= v_x \pi^*(z, z^x, g, g^x) \sigma \theta + \frac{1}{2} v_{xx} \|\pi^*(z, z^x, g, g^x) \sigma\|^2 \\ &+ v_y \left[z b + z^x \pi^*(z, z^x, g, g^x) \sigma \theta + \frac{1}{2} \|\pi^*(z, z^x, g, g^x) \sigma\|^2 (g^x + \eta |z^x|^2) \right. \\ &\quad \left. - F(t, z, z^x, g, g^x, u) + \pi^*(z, z^x, g, g^x) \sigma \sigma^T g \right] \\ &+ \frac{1}{2} v_{yy} (\|z \sigma\|^2 + z^x \|\pi^*(z, z^x, g, g^x) \sigma\|^2) \\ &+ v_{xy} (z^x \|\pi^*(z, z^x, g, g^x) \sigma\|^2 + \pi^*(z, z^x, g, g^x) \sigma \sigma^T z) \\ &+ \lambda_t (v(t, x, y + u) - v(t, x, y)). \end{aligned}$$

Proof. The proof of the theorem relies on a localization procedure. We first assume that there exists a solution ϕ to **(bHJB)** satisfying

$$|\phi(t, s, x, y)| \leq \kappa(1 + |x|^p + \|s\|^p + |y|^2), \quad p > 1, \quad \kappa > 0.$$

Let $(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}$. We introduce the following stopping time

$$\tau_n := \inf\{t > 0, (S_t, X_t^{\pi^*}, \hat{Y}_t^{\pi^*}, Z^X, \lambda_t) \in \mathcal{B}_n(0)\} \wedge T$$

where $\mathcal{B}_n(0)$ is the ball centered at the origin with radius n in $\mathbb{R}^{m+4}$. Applying Ito's formula, we get

$$\begin{aligned} \phi(\tau_n, S_{\tau_n}, X_{\tau_n}^{\pi^*}, \hat{Y}_{\tau_n}^{\pi^*}) &= \phi(0, S_0, x, 0) + \int_0^{\tau_n} \mathcal{H}^\phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}, Z_r^X, Z_r, \Gamma_r^X, \Gamma_r, U_r) dr \\ &\quad + \int_0^{\tau_n} \sum_{i=1}^m \partial_{s^i} \phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}) S_r^i \sigma_r^i dW_r \\ &\quad + \int_0^{\tau_n} \partial_x \phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}) \pi_r^* \sigma_r dW_r \\ &\quad + \int_0^{\tau_n} \partial_y \phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}) \left(\sum_{i=1}^m Z_r^i \sigma_r^i + Z_r^X \pi_r^* \sigma_r \right) dW_r \\ &\quad + \int_0^{\tau_n} \lambda_r (\phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*} + U_r) - \phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*})) dM_r. \end{aligned}$$

By the localization procedure and considering the expectations, we get

$$\begin{aligned} &\mathbb{E}[\phi(\tau_n, S_{\tau_n}, X_{\tau_n}^{\pi^*}, \hat{Y}_{\tau_n}^{\pi^*})] \\ &= \phi(0, S_0, x, 0) - \mathbb{E}\left[\int_0^{\tau_n} U^P(X_r^{\pi^*} - \hat{Y}_r^{\pi^*}) f(r) dr\right] \\ &\quad + \mathbb{E}\left[\int_0^{\tau_n} (\mathcal{H}^\phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}, Z_r^X, Z_r, \Gamma_r^X, \Gamma_r, U_r) + \partial_t \phi(r, S_r, X_r^{\pi^*}, \hat{Y}_r^{\pi^*}) + U^P(X_r^{\pi^*} - \hat{Y}_r^{\pi^*}) f(r)) dr\right]. \end{aligned}$$

Since ϕ satisfies **(bHJB)**,

$$\mathbb{E}[\phi(\tau_n, S_{\tau_n}, X_{\tau_n}^{\pi^*}, \hat{Y}_{\tau_n}^{\pi^*}) + \int_0^{\tau_n} U^P(X_r^{\pi^*} - \hat{Y}_r^{\pi^*}) f(r) dr] \leq \phi(0, S_0, x, 0) \quad \forall (Z^X, Z, \Gamma^X, \Gamma, U) \in \mathcal{U}$$

where the equality holds for $(\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U})$. Recall that ϕ is bounded by a polynomial function in the variables (s, x, y) . We note that

$$|\phi(\tau_n, S_{\tau_n}, X_{\tau_n}^{\pi^*}, \hat{Y}_{\tau_n}^{\pi^*})| \leq C(1 + |S_{\tau_n}|^p + |X_{\tau_n}^{\pi^*}|^p + |\hat{Y}_{\tau_n}^{\pi^*}|^2) \leq C(1 + \sup_t \{|S_t|^p + |X_t^{\pi^*}|^p + |\hat{Y}_t^{\pi^*}|^2\}).$$

By applying the dominated convergence theorem, we deduce that

$$\mathbb{E}[\phi(T, S_T, X_T^{\pi^*}, \hat{Y}_T^{\pi^*}) + \int_0^T U^P(X_r^{\pi^*} - \hat{Y}_r^{\pi^*})f(r)dr] \leq \phi(0, S_0, x, 0),$$

with equality when $(Z^X, Z, \Gamma^X, \Gamma, U) = (\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U})$.

□

Viscosity solution and dynamic programming. Assuming the existence of a $\mathcal{C}^{1,2,2,2}$ solution to **(bHJB)** can be relaxed in Theorem (3.3.7) by showing that the value function of the problem is a solution in the sense of viscosity of the PDE **(bHJB)**. This relaxation of regularity for solution to PDEs has been developed by Lions and Crandall in [47, 46]. We refer to [65, 154] for more details about stochastic control with viscosity solutions. We start by recalling the dynamic programming principle and we define the continuation value objective of the principal for any control $\nu := (Z, Z^X, \Gamma, \Gamma^X, U)$, by

$$v(t, s, x, y, \nu) = \mathbb{E}\left[\int_t^T U^P(X_r^{t,x,\pi^*} - Y_r^{t,y,\pi^*})f(r)dr\right],$$

so that $v(t, s, x, y) = \sup_{\nu} v(t, s, x, y, \nu)$, where X^{t,x,π^*} and Y^{t,y,π^*} denote the flow processes starting at time t with respective initial values x and y in (SDE). The dynamic programming principle states that for any stopping time $\theta \in [t, T]$ where $t < T$ we have

$$V(t, s, x, y) = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U) \in \mathcal{U}_{[t, \theta]}} \mathbb{E}\left[V(\theta, S_{\theta}^{t,s}, X_{\theta}^{t,x,\pi^*}, Y_{\theta}^{t,y}) + \int_t^{\theta} U^P(X_r^{t,x,\pi^*} - Y_r^{t,y})f(r)dr\right]. \quad (3.6)$$

The notion of weak solutions of **(bHJB)** results from this dynamic programming principle and is defined as follows.

Definition 3.3.9 (Viscosity solution). We say that v is a lower (resp. upper) semi-continuous super-solution (resp. sub-solution) of **(bHJB)** on $[0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ if for all functions $\phi \in \mathcal{C}^{1,2,2,2}$ and $(\hat{t}, \hat{s}, \hat{x}, \hat{y}) \in [0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ satisfying

$$0 = \min_{(t,s,x,y) \in [0,T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}} (v - \phi)(t, s, x, y) = (v - \phi)(\hat{t}, \hat{s}, \hat{x}, \hat{y}),$$

$$\text{resp. } 0 = \max_{(t,s,x,y) \in [0,T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}} (v - \phi)(t, s, x, y) = (v - \phi)(\hat{t}, \hat{s}, \hat{x}, \hat{y}),$$

we have

$$-\partial_t v(\hat{t}, \hat{s}, \hat{x}, \hat{y}) - U^P(\hat{x} - \hat{y})f(\hat{t}) - \sup_{(z^x, z, g^x, g, u)} \mathcal{H}^v(\hat{t}, \hat{s}, \hat{x}, \hat{y}, \nabla v, \Delta v, z^x, z, g^x, g, u) \geq 0 \text{ (resp. } \leq 0).$$

If v is both a super-solution and a sub-solution, we say that v is a viscosity solution to **(bHJB)**.

As a consequence of the dynamic programming principle, we have the following theorem

Theorem 3.3.10. *Assume that the value function V is locally bounded on $[0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$. Then V is a viscosity solution to **(bHJB)**.*

Remark 3.3.11. Note that the infinitesimal generator $\mathcal{H}$ contains a degenerate term through the λ process. This term explodes when t approaches T in the case that the support of τ is bounded and requires to prove Theorem 3.3.10 with a localization technique for both the state variables of the problem and the λ process.

Proof of Theorem 3.3.10. We denote by V_*, V^* the lower semi-continuous and upper semi-continuous envelopes of V respectively.

Step 1. Proof of the super-solution property. Let ϕ be a $\mathcal{C}^{1,2,2,2}$ function. Let $(\hat{t}, \hat{s}, \hat{x}, \hat{y}) \in [0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ be such that

$$0 = \min(V_* - \phi)(t, s, x, y) = (V_* - \phi)(\hat{t}, \hat{s}, \hat{x}, \hat{y})$$

and a sequence $(t_n, s_n, x_n, y_n) \in [0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ such that

$$(t_n, s_n, x_n, y_n) \longrightarrow (\hat{t}, \hat{s}, \hat{x}, \hat{y}),$$

with $\lim V(t_n, s_n, x_n, y_n) = V_*(\hat{t}, \hat{s}, \hat{x}, \hat{y})$. We define

$$\varepsilon_n := V(t_n, s_n, x_n, y_n) - \phi(t_n, s_n, x_n, y_n) \geq 0.$$

Note that $\varepsilon_n \longrightarrow 0$ for large n . We also define $X^n := X^{t_n, x_n, \pi^*}, Y^n := Y^{t_n, y_n, \pi^*}$ the solution to (SDE) starting at time t_n with respective values x_n and y_n and controlled by $\nu^{(n)} \in \mathcal{U}$ such that $\nu^{(n)} = \nu \in \mathcal{U}_{\hat{t}}$, so that ν is a constant at time $t = \hat{t}$, and S^n denotes the price process starting at the price vector s_n at time t_n .

$$\delta_n := \sqrt{\varepsilon_n} \mathbf{1}_{\varepsilon_n \neq 0} + \frac{1}{n} \mathbf{1}_{\varepsilon_n = 0},$$

and the stopping time

$$\theta_n = \inf\{t > t_n, (t, S_t^n, X_t^n, Y_t^n, \lambda_t) \notin [0, t_n + \delta_n) \times \mathcal{B}_n \times [0, n]\},$$

where $\mathcal{B}_n$ is the unit ball on $\mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ centered at the point (s_n, x_n, y_n) . We note that $\theta_n \rightarrow \hat{t}$ when n goes to ∞ . According to the dynamic programming principle (3.6) we have

$$V(t_n, s_n, x_n, y_n) \geq \mathbb{E}\left[V(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) + \int_{t_n}^{\theta_n} U^P(X_r^n - Y_r^n)f(r)dr\right]$$

or equivalently

$$\begin{aligned} 0 &\leq \mathbb{E}\left[V(t_n, s_n, x_n, y_n) - V(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) - \int_{t_n}^{\theta_n} U^P(X_r^n - Y_r^n)f(r)dr\right] \\ &\leq \varepsilon_n + \mathbb{E}\left[\phi(t_n, s_n, x_n, y_n) - \phi(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) - \int_{t_n}^{\theta_n} U^P(X_r^n - Y_r^n)f(r)dr\right] \end{aligned}$$

Applying Ito's formula for the function ϕ and by the localization procedure before the stopping time θ_n we get

$$\begin{aligned} 0 &\leq \varepsilon_n + \mathbb{E}\left[\int_{t_n}^{\theta_n} \left(-\partial_t \phi(r, S_r^n, X_r^n, Y_r^n) - \mathcal{H}^\phi(r, S_r^n, X_r^n, Y_r^n, \nu) - U^P(X_r^n - Y_r^n)f(r)\right)dr\right] \\ &\leq \frac{\varepsilon_n}{\delta_n} + \mathbb{E}\left[\frac{1}{\delta_n} \int_{t_n}^{\theta_n} \left(-\partial_t \phi(r, S_r^n, X_r^n, Y_r^n) - \mathcal{H}^\phi(r, S_r^n, X_r^n, Y_r^n, \nu) - U^P(X_r^n - Y_r^n)f(r)\right)dr\right] \end{aligned}$$

Since the function ϕ is assumed to be $\mathcal{C}^{1,2,2,2}$, ϕ and its derivative are essentially locally bounded around $\mathcal{B}_n$ and uniformly in n for $t < T$. Taking the limit as n goes to ∞ , we obtain

$$-\partial_t \phi(\hat{t}, \hat{s}, \hat{x}, \hat{y}) - \mathcal{H}^\phi(\hat{t}, \hat{s}, \hat{x}, \hat{y}, \nu) - U^P(\hat{x} - \hat{y})f(\hat{t}) \geq 0, \text{ for any } \nu \in \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}.$$

We deduce that V is a viscosity super-solution in the sense of Definition 3.3.9 to **(bHJB)**.

Step 2. Proof of the sub-solution property. Let ϕ be a $\mathcal{C}^{1,2,2,2}$ function. Let $(\hat{t}, \hat{s}, \hat{x}, \hat{y}) \in [0, T) \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}$ be such that

$$0 = \max(V^* - \phi)(t, s, x, y) = (V^* - \phi)(\hat{t}, \hat{s}, \hat{x}, \hat{y}).$$

Assume by contradiction that

$$-\partial_t \phi(\hat{t}, \hat{s}, \hat{x}, \hat{y}) - \sup_{\nu} \mathcal{H}^\phi(\hat{t}, \hat{s}, \hat{x}, \hat{y}, \nu) - U^P(\hat{x} - \hat{y})f(\hat{t}) > 0$$

Since the function $\mathcal{H}^\phi$ is continuous, we deduce that there exists a ball $\mathcal{B}_\varepsilon$ around $(\hat{t}, \hat{s}, \hat{x}, \hat{y})$

with norm $\varepsilon > 0$ small enough such that

$$-\partial_t \phi(t, s, x, y) - \sup_{\nu} \mathcal{H}^{\phi}(t, s, x, y, \nu) - U^P(x - y)f(t) < 0, \quad (t, s, x, y) \in \mathcal{B}_{\varepsilon}.$$

Note that there exists $\delta_0 > 0$ independent of ν such that

$$-2\delta_0 := \max_{(t,s,x,y) \in \mathcal{B}_{\varepsilon}} V^*(t, s, x, y) - \phi(t, s, x, y).$$

Let $(t_n, s_n, x_n, y_n) \in \mathcal{B}_{\varepsilon}$ converge to $(\hat{t}, \hat{s}, \hat{x}, \hat{y})$ such that $\lim V(t_n, s_n, x_n, y_n) = V^*(\hat{t}, \hat{s}, \hat{x}, \hat{y})$ and $-\delta_0 < (V - \phi)(t_n, s_n, x_n, y_n) < 0$ for any $n \geq 1$. We define

$$\theta_n = \inf\{t > t_n, (t, S_t^n, X_t^n, Y_t^n, \lambda_t) \notin \mathcal{B}_{\varepsilon} \times [0, n]\}.$$

Applying Ito's formula, we get for any control ν

$$\begin{aligned} V(t_n, s_n, x_n, y_n) &\geq -\delta_0 + \phi(t_n, s_n, x_n, y_n) \\ &= -\delta_0 + \mathbb{E}\left[\phi(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) - \int_{t_n}^{\theta_n} [\partial_t \phi(r, S_r^n, X_r^n, Y_r^n) + \mathcal{H}^{\phi}(r, S_r^n, X_r^n, Y_r^n, \nu_r)] dr\right] \\ &\geq -\delta_0 + \mathbb{E}\left[\phi(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) + \int_{t_n}^{\theta_n} U^P(X_r^n - Y_r^n)f(r)dr\right] \\ &\geq \delta_0 + \mathbb{E}\left[V^*(\theta_n, S_{\theta_n}^n, X_{\theta_n}^n, Y_{\theta_n}^n) + \int_{t_n}^{\theta_n} U^P(X_r^n - Y_r^n)f(r)dr\right]. \end{aligned}$$

Since δ_0 is independent of the control ν , it contradicts the dynamic programming principle (3.6). We thus deduce that

$$-\partial_t \phi(\hat{t}, \hat{s}, \hat{x}, \hat{y}) - \sup_{\nu} \mathcal{H}^{\phi}(\hat{t}, \hat{s}, \hat{x}, \hat{y}, \nu) - U^P(\hat{x} - \hat{y})f(\hat{t}) \leq 0,$$

hence V is a sub-solution in the sense of Definition 3.3.9 to **(bHJB)**. $\square$

Unbounded default time

Even in the unbounded case, we start by getting rid of the $\hat{Y}_0$ we have in (3.5) as it is useless for the optimization routine: again $\hat{Y}$ and Y only differ by a constant, so they have the same governing SDE. Therefore, our starting point is the following optimization problem

$$\sup_{(Z, Z^X, \Gamma, \Gamma^X, U)} \mathbb{E}[U^P(X_{T \wedge \tau}^{\pi^*} - \hat{Y}_{T \wedge \tau}^{\pi^*})]$$

The first difference is that now, in reformulating the problem using the default density,

we will still have a terminal part $(1 - F_\tau(T))(X_T - Y_T)$. Overall, our problem is now

$$\mathbb{E}\left[\int_0^T U^P(X_t^{\pi^*} - \hat{Y}_t^{\pi^*})f(t)dt + (1 - F_\tau(T))U^P(X_T - \hat{Y}_T^{\pi^*})\right]$$

Note that this problem differs from the bounded default time case only from the terminal condition given by

$$v(T, x, y) = (1 - F_\tau(T))U^P(x - y)$$

We introduce the following integro-partial PDE and the verification theorem follows.

$$(\mathbf{uHJB}) \quad \begin{cases} 0 = \partial_t v(t, s, x, y) + U^P(x - y) f(t) \\ \quad + \sup_{(z^x, z, g^x, g, u)} \left\{ \mathcal{H}^v(t, s, x, y, \nabla v, \Delta v, z^x, z, g^x, g, u) \right\}, \quad t < T, \\ v(T, s, x, y) = (1 - F_\tau(T))U^P(x - y). \end{cases}$$

Theorem 3.3.12 (Verification Theorem - unbounded case). *Assume that there exists a function ϕ twice continuously differentiable in space and differentiable in time, such that $\phi(t, s, x, y)$ solves $(\mathbf{uHJB})$. We denote by $(\hat{Z}_t, \hat{Z}_t^X, \hat{\Gamma}_t, \hat{\Gamma}_t^X, \hat{U}_t)$ the optimizers in the supremum. Furthermore, assume that ϕ has a quadratic growth in y and polynomial growth in s, x such that*

$$|\phi(t, s, x, y)| \leq \kappa(1 + |x|^p + \|s\|^p + |y|^2), \quad p > 1, \quad \kappa > 0.$$

Then, for each $t \in [0, T]$, the strategy $(\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U})$ is an optimal strategy for the control problem and

$$\phi(0, S_0, x, 0) = V_0^P = \sup_{(Z^X, Z, \Gamma^X, \Gamma, U)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^{\pi^*} - \hat{Y}_{T \wedge \tau}^{\pi^*}) \right].$$

The optimal contract is given by

$$\begin{aligned} \xi^* &= \hat{Y}_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{Z}_s^i}{S_s^i} dS_s^i + \int_0^{T \wedge \tau} \hat{Z}_s^X dX_s + \int_0^{T \wedge \tau} \hat{U}_s dH_s \\ &\quad + \frac{1}{2} \int_0^{T \wedge \tau} (\hat{\Gamma}_s^X + \eta(\hat{Z}_s^X)^2) d\langle X, X \rangle_s - \int_0^{T \wedge \tau} F(s, \hat{Z}_s, \hat{Z}_s^X, \hat{\Gamma}_s, \hat{\Gamma}_s^X, \hat{U}_s) ds \\ &\quad + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{\Gamma}_s^i}{S_s^i} d\langle S^i, X \rangle_s \end{aligned}$$

Remark 3.3.13. The proof of this theorem follows the same lines as the proof of Theorem 3.3.7 without requiring the localization of the λ term.

Remark 3.3.14. Similarly to Theorem 3.3.10, the value function of the problem is a viscosity solution to **(uHJB)**.

A note on the uniqueness and numerical motivation. In order to carry out numerical simulations, we must first address the uniqueness of viscosity solutions to the integro-partial differential equations **(bHJB)** and **(uHJB)**. In the integrable case, *i.e.*, when $[0, T]$ is strictly contained in the support of τ so that λ is integrable on $[0, T]$, uniqueness follows directly from existing results; see, for instance, [146, Section 5]. The bounded case, in which the support of τ is contained in $[0, T]$, is considerably more delicate. The loss of integrability of λ as t approaches T introduces a degeneracy that, to the best of our knowledge, is not treated in the existing literature. We outline below a possible strategy to resolve this difficulty, but we leave a full analysis for future research, as it would require deep results in the theory of second-order BSDEs. Observe that the probabilistic formulation of **(bHJB)** corresponds to a second-order BSDE with jumps (see, for example, Chapter 2, Section 3.4), but with a degenerate (singular) generator, similar to the one studied in [84]. A first step would be to extend the analysis of [84] to the second-order setting by studying the convergence of the solutions to the second-order BSDEs with truncated intensities $\lambda^n := \lambda \wedge n$, and by proving that these solutions converge to the solution of the 2BSDE with the original singular generator. Next, by connecting the solutions of the truncated 2BSDEs to the viscosity solutions of the corresponding non-degenerate IPDEs, a comparison theorem would be obtained for the viscosity solutions of the truncated IPDEs, using, for example, [6, 123]. Finally, passing to the limit as $n \rightarrow \infty$ would yield a comparison theorem for viscosity solutions of the original IPDE with singular coefficients, thereby establishing the uniqueness of its viscosity solution.

3.4 Numerical solutions

In this section, we study numerical solutions to the coupled HJB equations **(bHJB)** and **(uHJB)**. Our goal is twofold. First, we want to contrast the bounded and unbounded default specifications and understand how they differ. Second, under a bounded default time, we are interested in extracting qualitative features of the optimal contract and trading policy. In particular, we aim to:

- understand how the random default time affects both the optimal trading strategy and the incentive compensation scheme;
- investigate how the shape of the default distribution on $[0, T]$ impacts the optimal compensation;

- compare, in the bounded default case, the behaviour of the agent under the linear specification Ξ^l and under the general functional Ξ ;
- analyse how the different sources of risk (market risk and default risk) translate into distinct patterns for the various incentive components.

From a numerical point of view, the problem is intrinsically iterative. The maximizers $(\hat{\pi}, \hat{Z}^X, \hat{\Gamma}^X, \hat{U})$ entering **(bHJB)** and **(uHJB)** depend on the value function and its derivatives, and in particular $\hat{U}$ is characterized as the argmax of a functional that explicitly involves both $v(t, x, y)$ and the shifted value $v(t, x, y + u)$. The coefficients of the PDE therefore depend nonlinearly on the unknown solution itself, which rules out straightforward one-shot schemes. To keep the numerical exploration as transparent as possible, we focus on the simpler 1-dimensional case with a single risky asset. This allows us to isolate and better interpret the dynamics of the default-linked incentive U . In this one-dimensional case the compensation for exposure to the traded asset S is switched off. The reason is twofold: such a compensation can be interpreted as a reward incentivizing the agent to invest more aggressively on the riskier assets, and in the one-dimensional setting it becomes essentially indistinguishable from a reward on portfolio growth. Therefore, the only diffusion incentive is the process Z^X attached to the wealth process X . Both principal and agent are modeled with exponential utility, with possibly different risk-aversion coefficients. An additional objective of our numerical study is to understand how the optimal default-contingent incentive U reacts to changes in the ratio between the principal's and the agent's risk aversion.

In the recent literature, neural networks showed great potential to be the state of the art for numerical schemes for PDEs, so the backbone of our algorithm will be a simple feed-forward neural network, that we will deploy in a deep learning scheme. Similar ideas have been previously developed in [13], where an actor-critic approach is used to find the optimal policy for market making together with the value function involving an integro-partial HJB equation in high dimension and in [109] in which the authors use reinforcement learning methods to solve the problem. Note that in our case we are using supervised learning for an integro-partial HJB equation which differs from the existing literature on stochastic control problems.

We recall that the value function v solves the HJB equation (specialized to the one-dimensional setting: we set $z = g = 0$ in $\mathcal{H}^v$, so that the supremum collapses to $\sup_{(z^x, g^x, u)}$)

$$\partial_t v(t, x, y) + U^P(x - y) f(t) + \sup_{(z^x, g^x, u)} \left\{ \mathcal{H}^v(t, x, y, \nabla v(t, x, y), \Delta v(t, x, y), z^x, g^x, u) \right\} = 0, \quad t < T, \quad (3.7)$$

The running term $U^P(x - y) f_T(t)$ captures the principal's instantaneous utility weighted by the density f_T of the default time, while the Hamiltonian $\mathcal{H}^v$ contains the drift, diffusion, and default terms, as well as the dependence on the controls (z^x, g^x, u) .

We equip this PDE with a terminal condition $v(T, x, y) = \Phi(x, y)$ and a boundary condition on the solvency boundary $x = 0$, which in our specification reads $v(t, 0, y) = -1$. The value -1 is the natural specialization of the principal's CARA utility: when both the portfolio and the contract collapse to zero, the principal receives $U^P(0) = -e^{-\varsigma \cdot 0} = -1$. The need for a boundary condition is given by the fact that we need to solve a PDE that lives in $[0, T] \times \mathbb{R} \times \mathbb{R}$, but practically we have to solve it in a bounded region. The economic interpretation is that if the portfolio goes to 0 the principal does not pay the agent anything, creating a fixed value of -1 .

The numerical algorithm proceeds by alternating between two steps. We start from an initial guess for the controls,

$$(z_0^{x,*}, g_0^{x,*}, u_0^*),$$

and an initial, untrained neural network representing a value function v_0 . For this fixed choice of (z^x, g^x, u) we solve the HJB equation (3.7) by training the neural network to obtain an updated value function v_1 . Following the Physics-Informed Neural Networks (PINNs) approach [140], we use automatic differentiation to compute the first and second derivatives of v_θ , construct a loss function as the squared residual of the PDE together with penalty terms enforcing the terminal and boundary conditions, and update the network parameters θ via stochastic gradient descent.

Once a new approximation v_1 is available, we can solve *numerically* the pointwise optimization problem

$$\sup_{(z^x, g^x, u)} \left\{ \mathcal{H}^v(t, x, y, \nabla v_1(t, x, y), \Delta v_1(t, x, y), z^x, g^x, u) \right\},$$

thereby obtaining updated controls $(z_1^{x,*}, g_1^{x,*}, u_1^*)$. We then freeze these new controls and solve again the HJB equation via the PINN, producing v_2 , and we iterate this actor-critic-style procedure until convergence of both the value function and the control processes. A high-level pseudo-code of the resulting alternating scheme is reported in Algorithm 1 and Algorithm 2.

Algorithm 1: Subroutines: SolvePDENN and SolveMax

```
1 Function SolvePDENN( $s_k, z_k^{x,*}, z_k^*, g_k^{x,*}, g_k^*, u_k, v_k$ ):  
    // Step 1: Evaluate the PDE residual at the sampled points  $s_k$   
    // using current variables, and evaluate boundary and terminal loss  
2    Compute PDE residual at  $s_k$  using  $v_k, z_k^{x,*}, g_k^{x,*}, u_k$ ;  
    // Step 2: One step of stochastic gradient descent for the neural  
    // network-based approximation of the PDE  
3    Use a neural network or numerical solver to minimize the residual and update  
    //  $v_{k+1}$ ;  
    // Step 3: Return the updated PDE solution  
4    return  $v_{k+1}$ ;  
5 Function SolveMax( $v_{k+1}$ ):  
    // Step 1: Compute the current approximation of the value function  
    //  $v_{k+1}$   
6    Plug in  $v_{k+1}$  into the function to be maximized;  
    // Step 2: Maximize over the variables  
7    Solve for  $z_{k+1}^{x,*}, g_{k+1}^{x,*}, u_{k+1}^*$  that maximize the objective function (grid approach);  
    // Step 3: Return the maximized variables  
8    return  $z_{k+1}^{x,*}, g_{k+1}^{x,*}, u_{k+1}^*$ ;
```

Algorithm 2: PDE-Maximizer Structure

Data: v_0 : Initial guess for the PDE solution
($z_0^{x,*}, g_0^{x,*}, u_0^*$): Initial maximization variables
Sampler: sampling of the variables
PDE inputs: parameters of the PDE
Tolerance ϵ , max_iter : Maximum iterations
 $k \leftarrow 0$: Iteration counter
Result: PDE solution v^*

```
1 while  $k \leq \text{max\_iter}$  and not converged do
    // Sample new points
2    $s_k \leftarrow \text{sampler}()$ ;
    // Solve the PDE with current guess
3    $v_{k+1} \leftarrow \text{SolvePDENN}(s_k, z_k^{x,*}, g_k^{x,*}, u_k^*, v_k)$ ;
    // Maximize over variables
4    $z_{k+1}^{x,*}, g_{k+1}^{x,*}, u_{k+1}^* \leftarrow \text{SolveMax}(v_{k+1})$ ;
    // Check for convergence
5   if  $|v_{k+1} - v_k| \leq \epsilon$  then
    // If converged, exit loop
6     break;
    // Increment iteration counter
7    $k \leftarrow k + 1$ ;
```

Main numerical challenges and remedies. The solution of the coupled HJB system (bHJB)-(uHJB) presents several numerical difficulties, both at the level of the PDE and in the iterative optimization of the controls. A first challenge is that the neural network may converge prematurely to an almost constant 0 function, given the nature of the exponential utility function: this is alleviated by the boundary condition and some pre-training on this condition to bring the initial solution away from 0. A second difficulty lies in the relative scaling of the PDE residual, boundary, and terminal losses. To mitigate this effect we employ self-aware weights: the loss has 3 components, whose weights depend on the relative magnitude of the components, so that we inform the network which part to improve the most.

A third source of complexity arises from the iterative structure of the problem. The coefficients of the Hamiltonian depend on the maximizers (z^x, Γ^x, u) , which in turn depend on the value function and its derivatives. This creates a fixed-point structure: given a value function

estimate v , one must solve three maximization problems to obtain updated controls; these controls define a new PDE, whose solution yields a new v , and the process repeats. In practice, this nested optimization may lead to oscillations or loss explosions, especially early in the training process, especially when the exponential utilities amplify errors or when the default intensity becomes large. To alleviate this instability we update the controls through a smoothed scheme: instead of using the new maximizers directly, we form an exponential moving average with the previous controls. This damping mechanism prevents violent jumps in the coefficients and results in a smoother iterative trajectory. The smoothing effect vanishes over time as we start trusting more and more the output of the optimization problems.

Regarding the optimization over (z^x, Γ^x, u) , we rely primarily on grid search, which proved robust in the early stages of training when the value function is still crude and derivative estimates are noisy. More sophisticated gradient-based maximization was found to be less reliable, as it tended to produce non-smooth or unstable controls in the first iterations. Once the controls are identified on the residual grid, a k -nearest-neighbour interpolation is applied during the forward simulation of the BSDE to recover pathwise controls at arbitrary states.

In particular, we deployed a simple feed-forward neural network (NN), with 3-d inputs (coordinates (t, x, y) ; the asset price s is absent because, as in the constant-coefficient reduction of the bounded-default HJB, the value function v does not depend on s , see the remark following **(bHJB)**), 1-d outputs, 8 hidden layers each with 32 neurons. The learning rate was piecewise linear, we were sampling 8192 points uniformly and 2048 between terminal and boundary points. Our first objective is to understand how the *shape* of the default distribution influences the optimal control structure. Since skewness plays a particularly important role in determining when default is more likely to occur, we focus on the Beta(a, b) family on $[0, T]$, which allows us to adjust the parameters (a, b) to generate different skewness profiles. In this way we can systematically test how early-default-heavy versus late-default-heavy scenarios affect the dynamic incentives encoded in U , Z^X , and Γ^X . We worked with Beta(2, 4), Beta(1, 1) (i.e., a uniform distribution), and Beta(4, 2). Of course, the exploding intensity λ was computed but capped. The market has a positive drift of 10% with a volatility of 30%. As a worked illustration, consider τ uniformly distributed on $[0, 1]$ (the symmetric Beta case). We recall that

$$\mathbb{P}(\tau > x | \mathcal{F}_t) = \int_x^{+\infty} \gamma(t, u) du.$$

Since we assume independence of τ from the filtration, we have

$$\mathbb{P}(\tau > x | \mathcal{F}_t) = \mathbb{P}(\tau > x) = (1 - x)^+ = 1 - x, \quad x \in [0, 1],$$

as $\tau \sim U(0, 1)$. Consequently,

$$\gamma(t, u) = \mathbf{1}_{1 \geq u}, \quad \lambda_t = \frac{\gamma(t, t)}{\mathbb{P}(\tau > t | \mathcal{F}_t)} = \frac{1}{1 - t}.$$

Using [63, Proposition 4.4], by defining

$$\Lambda_t = \int_0^t \lambda_s ds,$$

we have

$$\mathbb{P}(\tau > x | \mathcal{F}_t) = \exp(-\Lambda_t).$$

Therefore

$$\Lambda_t = \int_0^t \lambda_s ds = \int_0^t \frac{1}{1 - s} ds = -\ln(1 - t),$$

so that

$$\exp(-\Lambda_t) = \exp(\ln(1 - t)) = 1 - t,$$

matching the survival function of the uniform distribution. In all the experiments with an independent default time, τ has a probability density function f_τ and an associated cumulative distribution function F_τ so that

$$\lambda_t = \frac{f_\tau(t)}{1 - F_\tau(t)}.$$

Comparison of exponential and uniform random default We now undertake a comparative analysis of trading policies and incentive structures under bounded versus unbounded default horizons on $[0, 1]$. For the unbounded setting, we specify the default time as exponentially distributed with intensity $\lambda = 1.5$, and contrast it with the uniform distribution on $[0, 1]$ for the bounded benchmark. Figure 3.1 displays, for both default specifications, the resulting portfolio trajectories, optimal investment strategies, and incentive processes related to the default U and the average incentive Z^X associated with the portfolio dynamics. A salient feature emerges from the finite-horizon structure: under the exponential (unbounded) default, the optimal trading policy is initially more aggressive. This reflects the relatively high short-term hazard rate, which rationalizes front-loaded risk-taking in order to extract value before a potentially early default. The optimal policy exhibits a transition pattern: initially, investment levels are higher under the exponential default relative to the uniform case; subsequently, the allocation becomes more conservative for the unbounded case; and, finally, it declines more sharply for the bounded case as the terminal date approaches.

Near maturity, the optimal strategy becomes approximately linear, consistent with standard finite-horizon contraction effects. Intuitively, the exponential default concentrates probability mass near the origin, inducing an initially aggressive stance followed by a gradual reduction in exposure, whereas the uniform default, exhibiting a flat hazard rate, generates a smoother, more homogeneously declining investment schedule. Regarding incentives, both default-driven and portfolio-value-based compensations are systematically higher in the bounded case. This is economically consistent: when default is guaranteed to occur before T , the principal must provide stronger incentives to align the risk-averse agent's actions with the desired investment profile. Conversely, Z^X is naturally decreasing as $t \rightarrow T$ in the bounded scenario, reflecting the rising conditional probability of default and the diminishing marginal need to incentivize effort.




Figure 3.1: Random default with uniform distribution on $[0, T]$ and exponential distribution. Top left: portfolio value for bounded and unbounded distributions. Top right: average Z^X for bounded and unbounded distributions. Bottom left: average U for bounded and unbounded distributions. Bottom right: optimal investment strategy for bounded and unbounded distributions.

Sensitivities with respect to the risk aversion In Figure 3.2, we investigate the impact of the ratio $\frac{\eta}{\zeta}$ under a uniform default time on $[0, T]$ with $T = 1$ on the portfolio value, as well as on the compensations Z^X and U with respect to the default time and the optimal investment strategy. We first observe that when the agent and the principal share the same risk-aversion parameter, the portfolio value is higher for lower levels of risk-aversion as expected in risky investments. Alternatively, if the agent is more risk-averse than the principal, the resulting portfolio value decreases. In other words, delegating the portfolio to a more risk-averse manager is less efficient, and alignment of risk-aversion parameters leads to improved portfolio performance. This phenomenon is explained by the optimal investment strategy: a manager who is less risk-averse than the client invests more aggressively in the market, generating a form of risk-transfer effect. This is reflected in a higher compensation Z^X (blue versus red curves), as well as in a higher compensation U associated with the random default risk.

Impact of the skewness of the distribution A very interesting comparison, for the bounded distribution, is the impact of skewness, a good measure for understanding the shape of the distribution and where the probability mass is concentrated. We fix the risk aversion for both the agent and the principal to the higher value of 1 and obtain the plots in figure 3.3 Now, it is not surprising that, under different distributions, we can get radically different strategies and portfolio values, as the exponential utility will shrink the differences. The interesting behavior is what happens for the incentives. Let us start by analyzing the right-skewed beta (in blue): the probability mass is all concentrated in the first part of the time interval, therefore the agent knows that most of the compensation will come from the percentage on the wealth and he takes an extremely aggressive strategy. Knowing this, the principal offers a very high share of the portfolio's growth together with immediate insurance against default. Because the default is expected to occur early, there is little incentive toward the latter part of the time interval, and the incentive processes Z^X , U , and Γ^X converge to zero (the portfolio value X , on the other hand, does not).

The situation changes in the opposite case of a left-skewed Beta distribution. Here, both the agent and the principal understand that default is more likely to occur in the second half of the horizon. As a result, both have an interest in adopting aggressive strategies; however, for the same level of compensation, the agent feels less pressure to invest quickly, since he expects more time before default. Consequently, the principal must continually increase the incentive rate to motivate the agent to intensify his investment activity.

This generates a delayed response relative to what one might initially expect: the agent becomes more cautious and reduces investment just as the probability of default sharply




Figure 3.2: Random default with uniform distribution on $[0, T]$. Top left: average portfolio value for different risk-aversion ratios. Top right: average Z^X for different risk-aversion ratios. Bottom left: average U for different risk-aversion ratios. Bottom right: average optimal investment strategy for different risk-aversion ratios. The blue curve corresponds to $\eta = 0.5\varsigma$ (agent less risk averse than the principal), the red curve corresponds to $\eta = 2\varsigma$ (agent more risk averse than the principal), the orange curve corresponds to $\eta = \varsigma$ with a low level of risk aversion, and the green curve corresponds to $\eta = \varsigma$ with a high level of risk aversion.

increases, requiring the principal to offer substantially higher insurance. The reason the insurance component becomes so large is intuitive: on average, the portfolio value tends to be high due to the earlier aggressive strategy, so a late default represents a much greater lost opportunity than a default occurring near the start of the horizon.

Comparison of linear and general contracts Figure 3.4 shows the difference between the optimal solution when contracts are restricted to the class of linear contracts Ξ^l and the solution under the full class of admissible contracts Ξ . The incentive rates of 5% and 20% are selected since 5% is the best fixed percentage within the restricted class, while 20% reflects a common benchmark in the hedge fund industry. The resulting optimality gap in




Figure 3.3: Random default with 3 different beta distributions on $[0, T]$: in green, we have a left-skewed $\text{Beta}(4, 2)$, in blue a right-skewed $\text{Beta}(2, 4)$, and in orange a uniform distribution.

portfolio performance is significant (top left), and in both cases the linear strategies are noticeably more conservative (bottom right). The interaction between the growth-based incentive, the default insurance, and the investment strategy is particularly revealing. When the agent receives only 5% of the portfolio's growth, he adopts a more aggressive market stance; however, the risk of default becomes relevant earlier, since he requires sufficient time to expand the portfolio in order to earn a meaningful absolute compensation. Anticipating this, the principal provides earlier default insurance.

At the opposite end, a 20% incentive, higher than the growth share in the fully optimal contract at the initial time, induces greater risk aversion in the agent, leading to a less aggressive exposure to the risky asset. As a result, the insurance component activates later. Under the optimal growth incentive, the timing of default insurance lies between these two extremes, striking a balance between the competing incentives and supporting a more aggressive strategy when market conditions are favorable.




Figure 3.4: Linear incentive on portfolio growth under bounded uniform default on $[0, T]$ (so that the effective horizon is $T \wedge \tau$). Top left: portfolio value for different percentages of the incentive. Top right: average Z^X for different percentages of the incentive. Bottom left: average U for different percentages of the incentive. Bottom right: optimal investment strategy for different percentages of the incentive.

3.5 From portfolio delegation to market making under default

The principal–agent analysis above shows how a single random horizon τ couples the contracting incentives, the trading strategy, and the stochastic control of a controlled wealth process. The three building blocks that emerged are: (i) a weak formulation in which the agent’s hidden choice changes the law of the canonical wealth process; (ii) a BSDE/2BSDE representation of the agent’s continuation utility on the random interval $[0, T \wedge \tau]$, with a jump compensating the default event $H_t = \mathbf{1}_{\tau \leq t}$; and (iii) a Hamilton–Jacobi–Bellman equation for the principal whose driver and terminal data degenerate as $t \uparrow T$ in the bounded case.

Each of these three ingredients re-appears in the next chapter, but in a different financial setting. Instead of a fund manager investing on behalf of a client, we consider a market maker who continuously posts bid and ask quotes on a single risky asset. The wealth process is now the dealer's mark-to-market P&L; the unobserved state of the principal is replaced by a controlled inventory; and the analogue of the random horizon τ is a default-time-style market disruption (a venue shutdown, a regulatory halt, or a counterparty failure) that forcibly liquidates the dealer's position. The control is no longer the portfolio weight π but the pair of half-spreads (δ^a, δ^b) posted to the order book, and the noise driving the wealth dynamics is a pair of point processes rather than a Brownian motion. As we show in the next chapter, the same combination of weak formulation, default-time enlargement, and random-horizon HJB equation gives a tractable benchmark for market making under uncertain termination, and this benchmark turns out to be the natural baseline against which the signature-based approximations of Chapter 5 should be compared.

Market making under a random terminal horizon

This is a thesis-only bridging chapter; it is not based on a published paper, and is not intended to stand on its own. Its role is to connect the analytical random-horizon machinery of Chapters 2–3 with the data-driven, signature-based methodology of Chapter 5. The chapter recalls the classical inventory-based market-making problem of Avellaneda and Stoikov [9], records the bounded-inventory linearisation of Guéant, Lehalle, and Fernandez-Tapia [70] (the *GLFT reduction*), and extends the construction to a random terminal horizon $T \wedge \tau$ via a Hopf–Cole substitution that mimics GLFT’s. The reader should treat the chapter as motivational scaffolding for Chapter 5; the open analytical points are listed in S4.3.

4.1 The market-making problem

4.1.1 Stylised facts and the dealer’s setup

A market maker is a liquidity provider who continuously posts buy (*bid*) and sell (*ask*) quotes on a risky asset. Two sources of risk shape the dealer’s problem:

- *Inventory risk*: the exposure to the asset’s price moves between the moment a position is opened and the moment it is unwound;
- *Adverse-selection risk*: the exposure to toxic order flow, in which counterparties trade against the dealer’s quotes only when they have private information about the asset’s near-future direction.

We adopt the canonical continuous-time formulation of Avellaneda and Stoikov [9], Guéant, Lehalle, and Fernandez-Tapia [70], and Cartea, Jaimungal, and Penalva [39]. Fix a finite deterministic horizon $T > 0$ and a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ carrying a Brownian motion W and two independent counting processes N^a, N^b that count arrivals of liquidity-taking buy and sell market orders. The mid-price follows

$$dS_t = \sigma dW_t, \quad S_0 = s_0,$$

with constant volatility $\sigma > 0$. The dealer's controls are non-negative half-spreads $\delta_t^a, \delta_t^b \geq 0$, which drive the controlled order-arrival intensities

$$\lambda^a(\delta^a) = A e^{-k\delta^a}, \quad \lambda^b(\delta^b) = A e^{-k\delta^b},$$

where $A > 0$ is an order-flow prefactor and $k > 0$ is the spread-elasticity coefficient.

The dealer's cash C_t and inventory Q_t evolve as

$$dC_t = (S_t + \delta_t^a) dN_t^a - (S_t - \delta_t^b) dN_t^b, \quad dQ_t = dN_t^b - dN_t^a,$$

with mark-to-market wealth $X_t := C_t + Q_t S_t$. By Itô,

$$dX_t = \delta_t^a dN_t^a + \delta_t^b dN_t^b + Q_t dS_t,$$

displaying the two sources of P&L: realised half-spreads on fills, and unrealised P&L through the inventory's exposure to mid-price moves. Throughout this chapter the inventory is constrained to the risk-management range $|Q_t| \leq Q_{\max}$, with reflecting boundary conditions imposed by forcing $\delta_t^a = +\infty$ when $Q_t = -Q_{\max}$ and $\delta_t^b = +\infty$ when $Q_t = Q_{\max}$. We write $\mathcal{Q} := \{-Q_{\max}, \dots, Q_{\max}\}$.

The dealer is risk-averse with CARA utility $U(x) = -e^{-\eta x}$, $\eta > 0$.

4.1.2 The Avellaneda–Stoikov benchmark

In the original GLFT setup (which we follow here), Avellaneda and Stoikov [9] and Guéant, Lehalle, and Fernandez-Tapia [70] take the dealer's terminal payoff to be the mark-to-market wealth itself,

$$X_T = C_T + Q_T S_T,$$

with *no* inventory-penalty term. The dealer's value function is

$$v(t, x, s, q) = \sup_{(\delta^a, \delta^b)} \mathbb{E} \left[-e^{-\eta(X_T)} \mid C_t = x, S_t = s, Q_t = q \right],$$

with terminal condition $v(T, x, s, q) = -e^{-\eta(x+qs)}$. The HJB equation reads, for $t < T$,

$$\begin{aligned} 0 = & \partial_t v + \frac{1}{2} \sigma^2 \partial_{ss} v \\ & + \sup_{\delta^a \geq 0} \lambda^a(\delta^a) \left[v(t, x + s + \delta^a, s, q - 1) - v(t, x, s, q) \right] \\ & + \sup_{\delta^b \geq 0} \lambda^b(\delta^b) \left[v(t, x - s + \delta^b, s, q + 1) - v(t, x, s, q) \right]. \end{aligned} \tag{4.1}$$

The *Avellaneda–Stoikov ansatz*

$$v(t, x, s, q) = -e^{-\eta(x+qs)} w(t, q) \quad (4.2)$$

separates the price-dependence from the inventory dependence. Substituting (4.2) into (4.1) and dividing through by $e^{-\eta(x+qs)} > 0$ reduces the problem to the inventory-indexed ODE

$$\partial_t w(t, q) + \frac{\eta^2 \sigma^2}{2} q^2 w(t, q) - H^a(w_q, w_{q-1}) - H^b(w_q, w_{q+1}) = 0, \quad t < T, \quad (4.3)$$

with terminal condition $w(T, q) = 1$ and Hamiltonians

$$H^a(w_q, w_{q-1}) := \sup_{\delta^a \geq 0} A e^{-k\delta^a} [w_q - e^{-\eta\delta^a} w_{q-1}], \quad H^b(w_q, w_{q+1}) := \sup_{\delta^b \geq 0} A e^{-k\delta^b} [w_q - e^{-\eta\delta^b} w_{q+1}].$$

For the exponential intensity $\lambda^\bullet(\delta) = A e^{-k\delta}$ the inner suprema are explicit. The first-order condition for δ^a gives

$$e^{-\eta\delta^{a,*}} = \frac{k}{k + \eta} \cdot \frac{w(t, q)}{w(t, q - 1)},$$

and analogously for δ^b . Taking logs:

$$\delta_t^{a,*} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{\eta} \log \frac{w(t, q-1)}{w(t, q)}, \quad \delta_t^{b,*} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{\eta} \log \frac{w(t, q+1)}{w(t, q)}. \quad (4.4)$$

Substituting the optimisers back gives

$$H^a(w_q, w_{q-1}) = K \frac{w_q^{1+k/\eta}}{w_{q-1}^{k/\eta}}, \quad H^b(w_q, w_{q+1}) = K \frac{w_q^{1+k/\eta}}{w_{q+1}^{k/\eta}}, \quad (4.5)$$

with $K := A \frac{\eta}{k+\eta} \left(\frac{k}{k+\eta} \right)^{k/\eta}$. Equation (4.3) is then the non-linear ODE system in $\{w(t, q)\}_{q \in \mathcal{Q}}$ studied by Avellaneda and Stoikov [9].

4.1.3 Bounded inventory: the GLFT linearisation

Guéant, Lehalle, and Fernandez-Tapia [70] observed that the Hopf–Cole substitution

$$u(t, q) := w(t, q)^{-k/\eta}, \quad q \in \mathcal{Q}, \quad (4.6)$$

linearises (4.3). Writing $w = u^{-\eta/k}$,

$$\partial_t w(t, q) = -\frac{\eta}{k} u^{-\eta/k-1} \partial_t u, \quad \frac{w_q^{1+k/\eta}}{w_{q\pm 1}^{k/\eta}} = u_q^{-\eta/k-1} u_{q\pm 1}.$$

Substituting into (4.3) and dividing through by $-\frac{\eta}{k} u_q^{-\eta/k-1}$ collapses the system to

$$\partial_t u(t, q) - \frac{k\eta\sigma^2}{2} q^2 u(t, q) + B \left[u(t, q-1) \mathbf{1}_{q > -Q_{\max}} + u(t, q+1) \mathbf{1}_{q < Q_{\max}} \right] = 0, \quad (4.7)$$

$t < T$, with terminal $u(T, q) = 1$ and

$$B := A \left(\frac{k}{k+\eta} \right)^{1+k/\eta}.$$

The reflecting boundary at $q = \pm Q_{\max}$ is encoded by indicators dropping the off-grid neighbour. Equation (4.7) is a linear, tridiagonal, autonomous ODE system on $\mathbb{R}^{2Q_{\max}+1}$. Stacking $\mathbf{u}(t) := (u(t, -Q_{\max}), \dots, u(t, Q_{\max}))^\top$ we have $\dot{\mathbf{u}}(t) = \mathbf{M} \mathbf{u}(t)$ with $\mathbf{M} := \frac{k\eta\sigma^2}{2} \text{diag}(q^2)_{q \in \mathcal{Q}} - B(S_+ + S_-)$ symmetric tridiagonal, so that

$$\mathbf{u}(t) = e^{(t-T)\mathbf{M}} \mathbf{u}(T), \quad u(T, q) = 1. \quad (4.8)$$

A one-time eigendecomposition $\mathbf{M} = O\Lambda O^\top$ makes the evaluation $O(Q_{\max})$. The optimal half-spreads, expressed directly in u , read

$$\delta_t^{a,\star} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{k} \log \frac{u(t, q)}{u(t, q-1)}, \quad \delta_t^{b,\star} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{k} \log \frac{u(t, q)}{u(t, q+1)}. \quad (4.9)$$

Remark 4.1.1 (Soft inventory penalty). The GLFT base setup carries no soft inventory penalty: the dealer's terminal payoff is just X_T , and the terminal $u(T, q) = 1$ for all q . A common extension, popular in Cartea, Jaimungal, and Penalva [39] and subsequent work, augments the terminal payoff by a quadratic penalty $-\frac{\phi}{2} Q_T^2$ with $\phi \geq 0$. This only shifts the terminal of (4.8) to

$$u(T, q) = e^{-k\phi q^2/2}, \quad w(T, q) = e^{+\eta\phi q^2/2},$$

leaving the ODE (4.7) unchanged. We treat the GLFT base case $\phi = 0$ as the default throughout this chapter; the soft-penalty extension is recovered by the obvious terminal swap.

4.2 Random terminal horizon

We now allow the dealer's trading mandate to be cut short by an exogenous random time τ that is not a stopping time of the order-flow filtration. The interpretation is that τ models a default-style market disruption—an exchange halt, a regulator-imposed trading curfew, a counterparty failure, or, in cryptoasset markets, a venue collapse—which forcibly liquidates the dealer's inventory at the prevailing mid-price.

The enlargement-of-filtrations apparatus is the same as in Chapter 3. Let $\mathbb{F} = (\mathcal{F}_t)$ be

the natural filtration of (W, N^a, N^b) , let $H_t := \mathbf{1}_{\{\tau \leq t\}}$ be the default indicator, and let $\mathbb{G}$ be the smallest right-continuous enlargement of $\mathbb{F}$ such that τ is a $\mathbb{G}$ -stopping time. We work under the density hypothesis: the conditional law of τ given $\mathcal{F}_t$ admits a density $\gamma(t, u)$ in u , so that the Azéma supermartingale

$$g(t) := \mathbb{P}(\tau > t \mid \mathcal{F}_t) = 1 - \int_0^t \gamma(t, s) ds$$

is well-defined, and the $\mathbb{G}$ -intensity of τ on $\{t < \tau\}$ is $\lambda_t^\tau = \gamma(t, t)/g(t)$. (For τ independent of $\mathbb{F}$, which is our running assumption, $g(t) = 1 - F_\tau(t)$ and $\gamma(t, t) = f_\tau(t)$.)

In the GLFT-aligned setup, no extra penalty is charged at default: the dealer simply receives the mark-to-market wealth $X_\tau = C_\tau + Q_\tau S_\tau$ upon forced liquidation. The dealer's overall objective is therefore

$$J(\delta^a, \delta^b) = \mathbb{E} \left[-e^{-\eta X_{T \wedge \tau}} \right].$$

We derive the random-horizon HJB and the resulting Bernoulli ODE in two equivalent ways.

4.2.1 Method 1 —running-density formulation with the GLFT ansatz

The simplest path is to integrate the random horizon out using the density of τ . Conditioning on $\mathcal{F}_t$,

$$v(t, x, s, q) := \mathbb{E} \left[- \int_t^T e^{-\eta X_u} f_\tau(u) du - e^{-\eta X_T} g(T) \mid \mathcal{F}_t \right],$$

which collects the contribution of τ landing in du (with probability $f_\tau(u) du$) and of the survival event $\{\tau > T\}$ (with probability $g(T)$). Standard stochastic-control arguments give the HJB

$$\partial_t v + \frac{\sigma^2}{2} \partial_{ss} v + \sup_{\delta^a \geq 0} [\cdot] + \sup_{\delta^b \geq 0} [\cdot] - e^{-\eta(x+qs)} f_\tau(t) = 0, \quad (4.10)$$

with terminal $v(T, x, s, q) = -g(T) e^{-\eta(x+qs)}$. The running-cost term $-e^{-\eta(x+qs)} f_\tau(t)$ comes from the density-weighted integration.

We now mimic GLFT's bounded-inventory Hopf–Cole linearisation (4.6), combined with the survival weight that handles the random-horizon discounting. Set

$$v(t, x, s, q) = -e^{-\eta(x+qs)} h(t, q)^{-\eta/k} g(t), \quad (4.11)$$

so that $w(t, q) := h(t, q)^{-\eta/k} g(t)$ plays the role of the AS factor of S4.1.2. The terminal of (4.11) matches the terminal of (4.10) provided $h(T, q) = 1$ for all $q \in \mathcal{Q}$ (just as in GLFT).

Substituting (4.11) into (4.10) requires the building blocks

$$\partial_t v = -e^{-\eta(x+qs)} \left[-\frac{\eta}{k} h^{-\eta/k-1} \partial_t h \cdot g + h^{-\eta/k} \dot{g} \right], \quad \dot{g} = -f_\tau,$$

and the same explicit Hamiltonian computation as in S4.1.2, applied here to the rescaled variable $w(t, q) = h(t, q)^{-\eta/k} g(t)$. The supremum at $\delta^{a,\star}$ yields

$$\sup_{\delta^a} A e^{-k\delta^a} [w_q - e^{-\eta\delta^a} w_{q-1}] = K \frac{w_q^{1+k/\eta}}{w_{q-1}^{k/\eta}} = K h_q^{-\eta/k-1} h_{q-1} g(t),$$

since the $g(t)$ factors collect cleanly under the substitution. The running cost $-e^{-\eta(x+qs)} f_\tau$ and the $h^{-\eta/k} \dot{g} = -f_\tau h^{-\eta/k}$ contribution from $\partial_t v$ combine into the crucial pair $f_\tau(h^{-\eta/k} - 1)$.

Collecting everything, dividing through by $e^{-\eta(x+qs)}$, dividing again by $g(t) h(t, q)^{-\eta/k-1}$, and rearranging, the equation collapses to

$$\partial_t h(t, q) = \frac{k\eta\sigma^2}{2} q^2 h(t, q) - B [h(t, q-1) + h(t, q+1)] - \frac{k\lambda_t^\tau}{\eta} h(t, q) + \frac{k\lambda_t^\tau}{\eta} h(t, q)^{1+\eta/k}, \quad (4.12)$$

with reflecting boundary conditions at $q = \pm Q_{\max}$ and terminal $h(T, q) = 1$. This is the multidimensional Bernoulli ODE with exponent

$$\alpha := 1 + \eta/k \notin \{0, 1\}.$$

In the deterministic-horizon limit $\lambda^\tau \equiv 0$, the last two terms vanish and (4.12) reduces exactly to the GLFT linear ODE (4.7) in $u = h$.

Stacking $\mathbf{h}(t) := (h(t, -Q_{\max}), \dots, h(t, Q_{\max}))^\top$, equation (4.12) reads

$$\dot{\mathbf{h}}(t) = \left(\mathbf{M} - \frac{k\lambda_t^\tau}{\eta} \mathbf{I} \right) \mathbf{h}(t) + \frac{k\lambda_t^\tau}{\eta} \mathbf{h}(t)^\alpha, \quad (4.13)$$

where $\mathbf{h}^\alpha$ is the componentwise power and $\mathbf{M}$ the same GLFT tridiagonal matrix of (4.8).

The optimal half-spreads expressed in h are

$$\delta_t^{a,\star} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{k} \log \frac{h(t, q)}{h(t, q-1)}, \quad \delta_t^{b,\star} = \frac{1}{\eta} \log(1 + \eta/k) + \frac{1}{k} \log \frac{h(t, q)}{h(t, q+1)}, \quad (4.14)$$

identical in form to GLFT's (4.9).

4.2.2 Method 2 —direct AS HJB and an equivalent change of variable

An equivalent route is to write the random-horizon HJB directly in the before-default value $v^b(t, x, s, q) := \mathbb{E}[-e^{-\eta X_{T \wedge \tau}} \mid \mathcal{F}_t, \tau > t]$, which satisfies

$$\partial_t v^b + \frac{\sigma^2}{2} \partial_{ss} v^b + \sup_a [\cdot] + \sup_b [\cdot] + \lambda_t^\tau (v^a - v^b) = 0, \quad v^b(T, \cdot) = -e^{-\eta(x+qs)}, \quad (4.15)$$

where the post-default value (no penalty in the GLFT base) is $v^a(t, x, s, q) = -e^{-\eta(x+qs)}$. The link between the two formulations is $v(t, x, s, q) = g(t) v^b(t, x, s, q)$. Plugging the AS ansatz $v^b = -e^{-\eta(x+qs)} w^b(t, q)$ into (4.15) (with $w^a(q) = 1$ in the GLFT base) gives

$$\partial_t w^b + \frac{\eta^2 \sigma^2}{2} q^2 w^b - K \frac{w_q^{b, 1+k/\eta}}{w_{q-1}^{b, k/\eta}} - K \frac{w_q^{b, 1+k/\eta}}{w_{q+1}^{b, k/\eta}} - \lambda_t^\tau (w^b - 1) = 0, \quad (4.16)$$

with $w^b(T, q) = 1$. Numerically, (4.16) can be integrated directly without any further substitution, and the optimal half-spreads are recovered via (4.4) with w replaced by w^b .

Equivalence between the two methods. From $v = g v^b$ and the two ansätze $v = -e^{-\eta(x+qs)} h^{-\eta/k} g$ (Method 1) and $v^b = -e^{-\eta(x+qs)} w^b$ (Method 2) one reads off

$$h(t, q) = w^b(t, q)^{-k/\eta} \quad (4.17)$$

exactly. That is, h is the same Hopf–Cole transform $u = w^{-k/\eta}$ applied to the random-horizon AS factor w^b , with no extra survival or compensator factor. The two methods solve the same problem in different coordinates; the multidimensional Bernoulli ODE (4.12) and the direct ODE (4.16) are equivalent and produce identical optimal half-spreads (4.14).

4.2.3 Scalar warmup: zero inventory, decoupled from the system

To make the structural effect of a random horizon visible in closed form, restrict to $q = 0$ and pretend $w^b(t, \pm 1) = w^b(t, 0)$, i.e. the dealer earns the symmetric spread but does not respond dynamically to inventory excursions. The Hamiltonians at $q = 0$ then collapse to the constant per-side spread-earning rate

$$\mu_{\text{spr}} := A \frac{\eta}{k + \eta} \left(\frac{k}{k + \eta} \right)^{k/\eta}, \quad (4.18)$$

and the scalar $w_0(t) := w^b(t, 0)$ satisfies the linear inhomogeneous ODE

$$\dot{w}_0(t) - (2\mu_{\text{spr}} + \lambda_t^\tau) w_0(t) = -\lambda_t^\tau, \quad w_0(T) = 1. \quad (4.19)$$

The reciprocal $u(t) := 1/w_0(t)$ converts (4.19) into the scalar Bernoulli ODE $\dot{u} + (2\mu_{\text{spr}} + \lambda^\tau)u = \lambda^\tau u^2$, which is the source of the “Bernoulli” label in the multidimensional version. Equation (4.19) is a toy: it does not describe the value at non-zero inventory and not even at $q = 0$ in the genuinely coupled system. We record it only as a motivating warmup.

Remark 4.2.1 (Closed form under uniform default). For $\tau \sim \text{Unif}[0, T]$, $\lambda_t^\tau = 1/(T - t)$, and an integrating-factor calculation gives

$$w_0(t) = \frac{1 - e^{-2\mu_{\text{spr}}(T-t)}}{2\mu_{\text{spr}}(T-t)}. \quad (4.20)$$

Sanity check: $w_0(T) = 1$ and $w_0(t) \sim 1/[2\mu_{\text{spr}}(T - t)]$ as $T - t \rightarrow \infty$, a polynomial-in- $(T - t)$ decay slower than the deterministic-horizon exponential decay $w_0^{\text{det}}(t) = e^{-2\mu_{\text{spr}}(T-t)}$. The qualitative reading is that a uniformly distributed default *slows* the decay of the value function from exponential to polynomial in the time-to-go.

4.2.4 Numerical verification

Both formulations were implemented as standalone ODE integrators (Radau with $\text{rtol} = 10^{-9}$, $\text{atol} = 10^{-12}$) in the supplementary package `MM_RH`. Three sanity checks were performed:

1. *Deterministic-horizon limit.* Setting $\lambda^\tau \equiv 0$, the Bernoulli ODE (4.12) reduces to the GLFT linear system. The h-solver and the matrix-exponential closed form (4.8) agree to $|h - u_{\text{GLFT}}| < 1.3 \times 10^{-9}$ on the parameter set $A = 1$, $k = 1.5$, $\eta = 0.1$, $\sigma = 0.3$, $T = 1$, $Q_{\text{max}} = 10$.
2. *Equivalence of Method 1 and Method 2.* For each of $\lambda^\tau \in \{\text{none}, \text{Exp}(\mu = 1), \text{Exp}(\mu = 0.1), \text{Unif}[0, T]\}$, the relationship $h(t, q) = w^b(t, q)^{-k/\eta}$ of (4.17) holds to better than 10^{-9} pointwise, and the optimal half-spreads computed via (4.14) from h agree with those computed from w^b to better than 6×10^{-11} . The two methods solve identical problems.
3. *Toy $q = 0$ closed form.* Equation (4.20) is reproduced by a scalar RK4 integration of (4.19) to within 2.4×10^{-5} (step-size limited).

The recovery of GLFT in the vanishing-hazard limit is the characteristic numerical signature of the framework. The L^∞ norm of the deviation between the random-horizon

optimal half-spreads at $t = 0$ and the GLFT half-spreads, over $q \in \mathcal{Q}$ and both sides, decays linearly in μ for $\tau \sim \text{Exp}(\mu)$; Figure 4.1 displays the log-log convergence for both solvers.




Figure 4.1: Recovery of GLFT optimal half-spreads as the hazard rate μ of the exponential default time vanishes. The two solvers (Method 1 in h , Method 2 in w^b) are run independently and produce identical deviations from GLFT, both scaling linearly in μ . The dotted slope-1 reference confirms the order. Parameters: $A = 1$, $k = 1.5$, $\eta = 0.1$, $\sigma = 0.3$, $T = 1$, $Q_{\max} = 10$, $\phi = 0$.

4.3 Open points and connection with the next chapter

The closed-form treatment above leaves three precise items for follow-up work; together they motivate the move to the signature-based, data-driven framework of Chapter 5.

- *Well-posedness of the singular bounded case.* The random-horizon AS equation (4.16) has a singular driver when $\text{supp}(\tau) \subset [0, T]$, since $\lambda_t^\tau \uparrow \infty$ at $t \uparrow T$. The Bernoulli substitution (4.12) regularises the equation on $[0, T - \varepsilon]$ for any $\varepsilon > 0$, but the behaviour as $\varepsilon \downarrow 0$ requires the singular-driver tools of [84], extended to the inventory-indexed system. This is the natural next step beyond the Lipschitz Hypothesis 3.2.3 of Chapter 3.
- *Comparison with deterministic GLFT.* A precise comparison theorem between the deterministic-horizon optimal quotes (4.9) and the random-horizon quotes (4.14) would quantify the spread distortion as a function of the hazard rate. The pieces are in place:

the random-horizon h converges to the GLFT u as the hazard rate vanishes (Figure 4.1), and a careful expansion in λ^τ should yield the leading-order correction.

- *Path-dependent extensions.* Once the order-flow model ceases to be memoryless Poisson (e.g. Hawkes-type self-exciting arrivals or limit-order-book imbalance signals), closed-form analysis is no longer available, the AS ansatz (4.2) no longer separates the price and inventory dependence, and the control problem requires either a high-dimensional PDE solver or a data-driven approximation. The signature-based approach of Chapter 5 is the natural successor: the rich path-dependent state is encoded by truncated signatures, and the optimal control is sought directly in a signature-linear policy class.

Signature methods in optimal market making

This chapter is based on the paper *Signature Methods for Optimal Market Making* (in preparation), joint work with Francesca Primavera and Thibaut Mastrolia. It connects classical inventory-based market making with path-dependent approximation through signatures.

Note to the reader. The market-making problem is introduced from scratch in the bridging Chapter 4, together with a first-draft analysis of the optimal quotes under a default-style random horizon. The introduction below is the introduction of the underlying paper and focuses on the signature-based methodology.

How this chapter relates to the bridge chapter. Chapter 4 solves the market-making problem under a Markovian state (t, X, S, Q) and an exponential CARA utility, with risk-aversion parameter η and a random termination τ , by an Avellaneda–Stoikov ansatz. The present chapter takes the same financial setup but replaces the Markovian state by a *path-dependent* feature map – the signature of an augmented market path – and the closed-form ODE-based reduction by a *policy-gradient* numerical method. The two chapters therefore sit at opposite ends of a methodological spectrum: closed-form control under a tractable Markovian state on one side, and non-parametric, path-dependent control trained by score-function gradients on the other.

5.1 Introduction

Market makers provide immediacy by posting bid and ask quotes and absorbing the random imbalance between buy and sell market orders. Their profitability depends on the spread they earn when trades arrive, while their risk is driven by the inventory accumulated through these trades and by subsequent asset-price moves. This trade-off has made market making a central example of continuous-time stochastic control in financial mathematics. A large part of the existing literature obtains tractable policies by imposing a low-dimensional Markovian state, typically consisting of the current inventory, price, and possibly a small number of order-flow or volatility factors. Such reductions are analytically powerful, but they are also

restrictive: in high-frequency markets, the recent path of order arrivals, quote revisions, price moves and inventory excursions may contain predictive information that is not summarized by the current state alone.

This paper proposes a path-dependent approximation method for market making based on signatures. The signature of a path is the collection of its iterated integrals. It provides a canonical, non-parametric feature map for paths and plays a role for path functionals analogous to the role played by monomials for functions on Euclidean space. In particular, linear functionals of the signature form a rich approximation class for non-anticipative path functionals under suitable topological conditions. This suggests a natural policy class for market making: at each time t , the ask and bid half-spreads are chosen as linear functionals of the truncated signature of the observed market path up to time t .

Our starting point is a market maker who trades a risky asset over a finite horizon and posts ask and bid offsets δ^a and δ^b . Buy market orders, which hit the ask and decrement inventory, are counted by N^a ; sell market orders, which hit the bid and increment inventory, are counted by N^b . Thus the inventory is $Q = N^b - N^a$. We consider an augmented path containing time, price information, quadratic-variation covariates and the order-arrival processes. The feedback controls are parametrized as

$$\delta_t^a = \langle \boldsymbol{\alpha}^a, \mathbb{Y}_t \rangle, \quad \delta_t^b = \langle \boldsymbol{\alpha}^b, \mathbb{Y}_t \rangle,$$

where $\mathbb{Y}_t$ denotes the signature of the augmented path stopped at time t , and α^a, α^b are trainable tensors supported on words up to a fixed truncation level. This class allows quotes to depend on the whole observed history while retaining an explicit finite-dimensional parametrization after truncation.

The main structural observation is that the signature algebra makes the resulting objective tractable. For semimartingale paths with jumps, stochastic integrals of signature-linear functionals can be rewritten as linear functionals of the terminal signature after an explicit correction, which we call the tilde transform. This gives an algebraic expression for the execution revenue. Moreover, for a mean-variance objective, the quadratic inventory term is handled by the shuffle identity, so that the realized payoff is again a linear functional of the terminal signature, with coefficients depending polynomially on the control parameters. The only non-algebraic part of the optimization is the dependence of the controlled law on the policy through the order-arrival intensities. We handle this dependence by a likelihood-ratio, or score-function, gradient.

The contributions of the paper are as follows. First, we formulate a signature-linear class of path-dependent market-making controls for semimartingale price and point-process

order-flow dynamics. Second, we derive an explicit tilde transformation which rewrites the relevant stochastic integrals against price and order-arrival components as linear functionals of the signature of an augmented path. Third, we show how the mean–variance objective can be reduced, via the shuffle product, to an algebraic functional of the expected terminal signature under the controlled law. Fourth, we combine this representation with a likelihood-ratio gradient for spread-dependent point-process intensities. Finally, we provide numerical evidence in Poisson and Hawkes order-flow environments. In the Poisson model, where the expert benchmark is tailored to the memoryless structure, the signature policy recovers the qualitative behavior of the benchmark with a modest loss in mean reward. In the Hawkes model, where order flow is self-exciting and the control problem is more genuinely path-dependent, the signature policy improves the observed mean–variance trade-off.

5.2 Related literature

The stochastic-control approach to market making goes back to classical models of dealer inventory risk and has been developed extensively in modern high-frequency settings. The influential model of Avellaneda and Stoikov [9] derives optimal bid and ask quotes under inventory risk and Poisson order arrivals, and Guéant, Lehalle, and Fernandez-Tapia [69] provide a tractable treatment of inventory-risk management with exponential intensities. General introductions and models for algorithmic and high-frequency trading can be found in Cartea, Jaimungal, and Penalva [39]. Related developments include market making with strategic traders, make–take fees, multiple market makers, option market making and partial-information order-flow models; see, among others, Carmona and Webster [37], El Aoud and Abergel [62], Baldacci, Possamai⁺, and Rosenbaum [15], Baldacci, Bergault, and Possamai⁺ [11], Baldacci, Jusselin, and Rosenbaum [12], Chávez-Casillas et al. [42], Rosenbaum and Zhang [142] and Barucci, Mathieu, and Sánchez-Betancourt [19]. Our contribution is complementary to this literature: rather than solving a particular Markovian model in closed form, we introduce a flexible path-dependent policy class and an associated training procedure.

Signatures originated in rough path theory, following the foundational work of Lyons [110]; see also the uniqueness result of Hambly and Lyons [72]. Their use in finance has expanded rapidly because they provide universal and numerically tractable features for path-dependent problems. Signature methods have been used for optimal execution by Kalsi, Lyons, and Arribas [91], double-execution strategies by Cartea, Arribas, and Sánchez-Betancourt [38], non-parametric pricing and hedging by Lyons, Nejad, and Arribas [111], signature trading by Futter, Horvath, and Wiese [66], and signature-based models and calibration by Cuchiero,

Gazzani, and Svaluto-Ferro [49] and Cuchiero et al. [48]. Closest in spirit to the present paper are the works of Cuchiero and Möller [50], who use signature portfolios in stochastic portfolio theory, and Bank et al. [16], who study stochastic control with signature-parametrized controls. The latter gives theoretical support for signature control classes in general stochastic-control problems. Our focus is different and more specific: we exploit the algebra of signatures to obtain explicit objective and gradient formulae for market making with point-process order flow.

The numerical part of the paper also connects to the literature on point processes in market microstructure. Poisson models with spread-dependent intensities are the standard benchmark in market-making theory. Hawkes processes, introduced by Hawkes [73], provide a natural extension for clustered and self-exciting order arrivals. In such models, past events affect future intensities, making the optimal quoting rule inherently path-dependent unless the excitation variables are included explicitly in the state. This is precisely the type of setting in which signature features are expected to be useful.

5.3 Preliminaries

5.3.1 Notation

Fix $d \in \mathbb{N}$ and let $\mathbb{R}^d$ be the Euclidean space. The extended tensor algebra over $\mathbb{R}^d$ is defined by

$$T((\mathbb{R}^d)) := \left\{ \mathbf{u} = (\mathbf{u}^{(0)}, \mathbf{u}^{(1)}, \dots, \mathbf{u}^{(n)}, \dots) \mid \mathbf{u}^{(n)} \in (\mathbb{R}^d)^{\otimes n} \right\},$$

where $(\mathbb{R}^d)^{\otimes n}$ denotes the n -fold tensor product of $\mathbb{R}^d$, with the convention $(\mathbb{R}^d)^{\otimes 0} := \mathbb{R}$. We equip $T((\mathbb{R}^d))$ with the standard addition $+$, tensor multiplication $\otimes$, and scalar multiplication. For $N \in \mathbb{N}$, the truncated tensor algebra is defined by

$$T^N(\mathbb{R}^d) := \left\{ \mathbf{u} = (\mathbf{u}^{(0)}, \mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}) \mid \mathbf{u}^{(n)} \in (\mathbb{R}^d)^{\otimes n} \text{ for } n \leq N \right\}.$$

Let $\pi_n : T((\mathbb{R}^d)) \rightarrow (\mathbb{R}^d)^{\otimes n}$ be the map such that, for $\mathbf{u} \in T((\mathbb{R}^d))$, $\pi_n(\mathbf{u}) = \mathbf{u}^{(n)}$, and let $\pi_{\leq N} : T((\mathbb{R}^d)) \rightarrow T^N(\mathbb{R}^d)$ be such that, for $\mathbf{u} \in T((\mathbb{R}^d))$,

$$\pi_{\leq N}(\mathbf{u}) = \mathbf{u}^N := (\mathbf{u}^{(n)})_{n=0}^N.$$

Let $I = (i_1, \dots, i_n)$ be a multi-index with entries in $\{1, \dots, d\}$. Denoting by $\epsilon_1, \dots, \epsilon_d$ the

canonical basis of $\mathbb{R}^d$, we use the notation $|I| := n$ and

$$\epsilon_I := \epsilon_{i_1} \otimes \epsilon_{i_2} \otimes \cdots \otimes \epsilon_{i_n}.$$

Observe that $(\epsilon_I)_I$ is the canonical orthonormal basis of $(\mathbb{R}^d)^{\otimes n}$. Furthermore, we denote by $\epsilon_\emptyset$ the basis element of $(\mathbb{R}^d)^{\otimes 0}$ and set $|\emptyset| := 0$. We also set $I' := (i_1, \dots, i_{n-1})$ for $n > 1$, $I' := \emptyset$ for $n = 1$, and $I'' := (I')'$ for $n > 1$, with the convention $\epsilon_{I''} = 0$ for $n = 1$.

For two multi-indices $I \in \{1, \dots, d\}^{|I|}$ and $J \in \{1, \dots, d\}^{|J|}$, and two indices of length 1, $a, b \in \{1, \dots, d\}$, the shuffle product $\sqcup$ is defined recursively by

$$\begin{aligned} I \sqcup \emptyset &= \emptyset \sqcup I = I, \\ (I, a) \sqcup (J, b) &= ((I \sqcup (J, b)), a) + (((I, a) \sqcup J), b), \end{aligned} \tag{5.1}$$

where (I, a) denotes the concatenation of multi-indices.

Given $\mathbf{x} \in T((\mathbb{R}^d))$, we write $\mathbf{x}_I := \langle \mathbf{x}, \epsilon_I \rangle$, and for each $\mathbf{u} \in T(\mathbb{R}^d)$, we set

$$\langle \mathbf{u}, \mathbf{x} \rangle := \sum_{|I| \geq 0} \mathbf{u}_I \mathbf{x}_I \in \mathbb{R}. \tag{5.2}$$

Notice that the sum in (5.2) is finite by definition of $T(\mathbb{R}^d)$.

Define

$$G((\mathbb{R}^d)) := \left\{ \mathbf{x} \in T((\mathbb{R}^d)) : \langle \epsilon_\emptyset, \mathbf{x} \rangle = 1, \langle \epsilon_I \sqcup \epsilon_J, \mathbf{x} \rangle = \langle \epsilon_I, \mathbf{x} \rangle \langle \epsilon_J, \mathbf{x} \rangle, \text{ for all } I, J \right\}. \tag{5.3}$$

5.3.2 Signature of a càdlàg semimartingale

Definition 5.3.1. For a semimartingale $X = (X_t)_{t \in [0, T]}$, we call the $G((\mathbb{R}^d))$ -valued process $\mathbb{X} = (\mathbb{X}_t)_{t \in [0, T]}$ the *signature of X* if $\langle \epsilon_\emptyset, \mathbb{X}_t \rangle = 1$ and

$$\begin{aligned} \langle \epsilon_I, \mathbb{X}_t \rangle &= \int_0^t \langle \epsilon_{I'}, \mathbb{X}_{s-} \rangle dX_s^{i_{|I|}} 1_{\{|I| > 0\}} + \frac{1}{2} \int_0^t \langle \epsilon_{I''}, \mathbb{X}_{s-} \rangle d[X^{i_{|I|-1}, X^{i_{|I|}}]_s^c 1_{\{|I| > 1\}} \\ &\quad + \sum_{0 < s \leq t} \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \frac{1}{|I_2|!} 1_{\{|I_2| > 1\}} \langle \epsilon_{I_1}, \mathbb{X}_{s-} \rangle \langle \epsilon_{I_2}, (\Delta X_s)^{\otimes |I_2|} \rangle, \end{aligned} \tag{5.4}$$

for each $|I| > 0$.

5.4 Signature market-making optimization

5.4.1 Financial modeling and optimal quotes

We consider a market maker trading a financial product with mid-price process $S = (S_t)_{t \in [0, T]}$ over an intraday horizon $[0, T]$, so that interest rates can be neglected. We also allow for a hedging instrument $P = (P_t)_{t \in [0, T]}$, which may represent, for instance, the price of a European option written on S , or simply coincide with S when the market maker only trades the underlying asset. More generally, we consider a self-financing hedging portfolio with predictable strategy Θ_t and dynamics

$$dP_t = \Theta_t dS_t.$$

At each time t , the market maker posts an ask price S_t^a and a bid price S_t^b . Equivalently, she controls the ask and bid offsets

$$\delta_t^a := S_t^a - S_t, \quad \delta_t^b := S_t - S_t^b.$$

In addition, she may continuously allocate wealth to the hedging instrument through a control π_t .

Let N^a and N^b be point processes representing the arrivals of market orders hitting the ask and the bid, respectively. The inventory process is

$$Q_t = N_t^b - N_t^a.$$

The profit-and-loss process is written as

$$dX_t^{\pi, \delta} = dZ_t^\delta + dI_t^\pi, \quad X_0 > 0,$$

where Z^δ denotes the execution revenue and I^π denotes the inventory and hedging component, namely

$$\begin{aligned} dZ_t^\delta &= \delta_t^a dN_t^a + \delta_t^b dN_t^b, \\ dI_t^\pi &= Q_{t-} dS_t + \pi_{t-} \frac{dP_t}{P_{t-}}, \\ X_t^{\pi, \delta} &= X_0 + Z_t^\delta + I_t^\pi. \end{aligned}$$

Under risk-neutral pricing, both S and P are $\mathbb{P}$ -martingales, where $\mathbb{P}$ denotes the risk-

neutral measure, and therefore $\mathbb{E}[I_t^\pi] = 0$. However, the general framework below does not require S or P to be martingales.

The objective of the market maker is to solve

$$V_0^{MM} = \sup_{\pi, \delta=(\delta^a, \delta^b)} \mathbb{E}\left[U\left(Z_T^\delta, I_T^\pi, (I_T^\pi)^2\right)\right], \quad (5.5)$$

for either

1. $U(\zeta, \iota, \nu) = \bar{U}(\zeta + \iota)$, where $\bar{U} : \mathbb{R} \rightarrow \mathbb{R}$ is increasing and concave, for example

$$\bar{U}(x) = -e^{-\eta x}, \quad \eta > 0,$$

2. or a mean–variance-type criterion

$$U(\zeta, \iota, \nu) = \zeta - w \nu, \quad w > 0.$$

Here $\eta > 0$ is the CARA risk-aversion parameter (case 1) and $w > 0$ is the inventory-variance penalty (case 2).

Remark 5.4.1. This framework is general: at this stage we do not specify the dynamics of the asset price S , the hedging instrument P , or the order-arrival processes N^a, N^b .

Remark 5.4.2. This framework unifies market-making optimization with pricing and hedging via self-financing portfolios, and will cover the examples considered below.

5.4.2 Signature methods

Assume that $\Delta N_t^a \Delta N_t^b = 0$ for all t , and consider the process

$$Y_t := \left(t, S_t, P_t, [S]_t, [S, P]_t, [P]_t, N_t^a, \sum_{s \leq t} (\Delta N_s^a)^2, \dots, \sum_{s \leq t} (\Delta N_s^a)^M, N_t^b, \sum_{s \leq t} (\Delta N_s^b)^2, \dots, \sum_{s \leq t} (\Delta N_s^b)^M \right),$$

for some $M \geq 2$. To simplify notation, we write the components of Y as follows:

$$\begin{aligned} Y_t^0 &:= t, & Y_t^s &:= S_t, & Y_t^p &:= P_t, & Y_t^{qs} &:= [S]_t, & Y_t^{qsp} &:= [S, P]_t, \\ Y_t^{qp} &:= [P]_t, & Y_t^a &:= N_t^a, & Y_t^b &:= N_t^b, \end{aligned} \quad (5.6)$$

and, for all $2 \leq m \leq M$,

$$Y_t^{a^m} := \sum_{s \leq t} (\Delta N_s^a)^m, \quad Y_t^{b^m} := \sum_{s \leq t} (\Delta N_s^b)^m.$$

We consider multi-indices $I = (i_1, \dots, i_n)$ with entries in

$$\mathcal{A} := \{0, s, p, qs, qp, qsp, a, a^2, \dots, a^M, b, b^2, \dots, b^M\}.$$

Note that $|\mathcal{A}| = 6 + 2M$ and that, up to a bijection

$$b_{\mathcal{A}} : \mathcal{A} \rightarrow \{1, \dots, 6 + 2M\},$$

we are considering multi-indices with entries in $\{1, \dots, 6 + 2M\}$. For concreteness, we fix the bijection $b_{\mathcal{A}}$ as the lexicographic order $0, s, p, qs, qp, qsp, a, a^2, \dots, a^M, b, b^2, \dots, b^M \leftrightarrow 1, 2, \dots, 6 + 2M$. Therefore, the notation introduced in Section 5.3.1 applies directly in this setting.

Definition 5.4.3. For each multi-index I with entries in $\{0, s, p, a, b\}$ and each $j \in \{0, s, p, a, b\}$, we define the following transformed tensor:

$$\begin{aligned} (\epsilon_I; \epsilon_j)^\sim &:= \epsilon_I \otimes \epsilon_j - \frac{1}{2} \epsilon_{I'} \otimes \epsilon_{qs} 1_{\{i_{|I|}=j=s\}} - \frac{1}{2} \epsilon_{I'} \otimes \epsilon_{qp} 1_{\{i_{|I|}=j=p\}} \\ &\quad - \frac{1}{2} \epsilon_{I'} \otimes \epsilon_{qsp} 1_{\{(i_{|I|}, j) \in \{(s,p), (p,s)\}\}} \\ &\quad + \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_{a^{|I_2|+1}}) 1_{\{j=a\}} 1_{\{I_2=(a, \dots, a)\}} \\ &\quad + \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_{b^{|I_2|+1}}) 1_{\{j=b\}} 1_{\{I_2=(b, \dots, b)\}}, \end{aligned}$$

where

$$\alpha(r) := \sum_{k=1}^r (-1)^k \alpha(r, k),$$

and, denoting by $\mathcal{I}(r, k)$ the set of all multi-indices $J \in \mathbb{N}^k$ with $j_i \geq 1$ for every i such that $j_1 + \dots + j_k = r$,

$$\alpha(r, k) := \sum_{J \in \mathcal{I}(r, k)} \prod_{i=1}^k \frac{1}{(j_i + 1)!}.$$

Remark 5.4.4. In the setting of [9], one has

$$dP_t = dS_t = \sigma dW_t,$$

hence

$$[P]_t = [S]_t = [S, P]_t = \sigma^2 t.$$

In that case, the transformation introduced in Definition 5.4.3 reduces to

$$\begin{aligned}
(\epsilon_I; \epsilon_j)^\sim &:= \epsilon_I \otimes \epsilon_j - \sigma^2 \frac{1}{2} \epsilon_{I'} \otimes \epsilon_0 \mathbf{1}_{\{i_{|I|}=j=s\}} \\
&+ \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_{a^{|I_2|+1}}) \mathbf{1}_{\{j=a\}} \mathbf{1}_{\{I_2=(a,\dots,a)\}} \\
&+ \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_{b^{|I_2|+1}}) \mathbf{1}_{\{j=b\}} \mathbf{1}_{\{I_2=(b,\dots,b)\}},
\end{aligned}$$

for each multi-index I with entries in $\{0, s, a, b\}$ and each $j \in \{0, s, a, b\}$.

If, furthermore, all jumps are of unit size, i.e.

$$\Delta N_t^a = \Delta N_t^b = 1 \quad \text{at every jump time } t,$$

then the transformation simplifies further to

$$\begin{aligned}
(\epsilon_I; \epsilon_j)^\sim &:= \epsilon_I \otimes \epsilon_j - \sigma^2 \frac{1}{2} \epsilon_{I'} \otimes \epsilon_0 \mathbf{1}_{\{i_{|I|}=j=s\}} \\
&+ \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_a) \mathbf{1}_{\{j=a\}} \mathbf{1}_{\{I_2=(a,\dots,a)\}} \\
&+ \sum_{\epsilon_{I_1} \otimes \epsilon_{I_2} = \epsilon_I} \alpha(|I_2|) (\epsilon_{I_1} \otimes \epsilon_b) \mathbf{1}_{\{j=b\}} \mathbf{1}_{\{I_2=(b,\dots,b)\}},
\end{aligned}$$

for each multi-index I with entries in $\{0, s, a, b\}$ and each $j \in \{0, s, a, b\}$.

Example 5.4.5. • If I is a multi-index with entries in $\{0, s, p, a, b\}$ and $j = 0$, then

$$(\epsilon_I; \epsilon_0)^\sim = \epsilon_I \otimes \epsilon_0 = \epsilon_{(I,0)}.$$

- If I is a multi-index with entries in $\{0, s, p, a, b\}$ such that $I = (I', s)$, then

$$(\epsilon_I; \epsilon_s)^\sim = \epsilon_I \otimes \epsilon_s - \frac{1}{2} \epsilon_{I'} \otimes \epsilon_{qs} = \epsilon_{(I,s)} - \frac{1}{2} \epsilon_{(I',qs)}.$$

- If $I = (0, s, p, a, a)$, then

$$\begin{aligned}
(\epsilon_I; \epsilon_a)^\sim &= \epsilon_I \otimes \epsilon_a + \alpha(1) \epsilon_{I'} \otimes \epsilon_{a^2} + \alpha(2) \epsilon_{I''} \otimes \epsilon_{a^3} \\
&= \epsilon_{(I,a)} + \alpha(1) \epsilon_{(I',a^2)} + \alpha(2) \epsilon_{(I'',a^3)}, \\
(\epsilon_I; \epsilon_b)^\sim &= \epsilon_I \otimes \epsilon_b = \epsilon_{(I,b)}.
\end{aligned}$$

The proof of the following proposition is a straightforward adaptation of Proposition 4.10 in [51].

Proposition 5.4.6. Fix $n \in \mathbb{N}$, $I \in \{0, s, p, a, b\}^n$, and $j \in \{0, s, p, a, b\}$. Assume that each component of Y is a càdlàg semimartingale with locally bounded coordinate processes, so that the stochastic integrals below are well-defined. If $M \geq n + 1$, then

$$\int_0^t \langle \epsilon_I, \mathbb{Y}_{s-} \rangle dY_s^j = \langle (\epsilon_I; \epsilon_j)^\sim, \mathbb{Y}_t \rangle.$$

Remark 5.4.7. 1. By linearity of the stochastic integral, for every $\mathbf{u} \in T(\mathbb{R}^5)$ (here $\{0, s, p, a, b\}$ is identified with $\{1, \dots, 5\}$) and every $j \in \{0, s, p, a, b\}$,

$$\int_0^t \langle \mathbf{u}, \mathbb{Y}_s \rangle dY_s^j = \langle (\mathbf{u}; \epsilon_j)^\sim, \mathbb{Y}_t \rangle, \quad (5.7)$$

where

$$(\mathbf{u}; \epsilon_j)^\sim := \sum_{|I| \geq 0} u_I (\epsilon_I; \epsilon_j)^\sim.$$

Thus, Itô integrals of linear functionals of $\mathbb{Y}$ against selected components of Y are again linear functionals of $\mathbb{Y}$.

2. For every $j \in \{0, s, p, a, b\}$, the map

$$\mathbf{u} \mapsto (\mathbf{u}; \epsilon_j)^\sim$$

is linear from $T(\mathbb{R}^5)$ to $T(\mathbb{R}^{6+2M})$.

We are now ready to reformulate the problem (5.5) in terms of the signature. To this end, we assume that

$$\delta_t^a = \langle \boldsymbol{\alpha}^a, \mathbb{Y}_t \rangle := \sum_{|I| \leq K} \alpha_I^a \langle \epsilon_I, \mathbb{Y}_t \rangle, \quad (5.8)$$

$$\delta_t^b = \langle \boldsymbol{\alpha}^b, \mathbb{Y}_t \rangle := \sum_{|I| \leq K} \alpha_I^b \langle \epsilon_I, \mathbb{Y}_t \rangle, \quad (5.9)$$

$$\pi_t^P := \frac{\pi_t}{P_t} = \langle \boldsymbol{\beta}, \mathbb{Y}_t \rangle := \sum_{|I| \leq K} \beta_I \langle \epsilon_I, \mathbb{Y}_t \rangle, \quad (5.10)$$

for some $K \in \mathbb{N}$, where

$$\boldsymbol{\alpha}^a := (\alpha_I^a \epsilon_I)_I, \quad \boldsymbol{\alpha}^b := (\alpha_I^b \epsilon_I)_I, \quad \boldsymbol{\beta} := (\beta_I \epsilon_I)_I,$$

and $\alpha_I^a, \alpha_I^b, \beta_I \in \mathbb{R}$. We also assume that

$$\alpha_I^a = \alpha_I^b = \beta_I = 0 \quad \text{whenever some index in } I \text{ belongs to } \mathcal{A} \setminus \{0, s, p, a, b\}.$$

Recalling that

$$dZ_t^\delta = \delta_t^a dN_t^a + \delta_t^b dN_t^b, \quad dI_t^\pi = Q_{t-} dS_t + \pi_{t-} \frac{dP_t}{P_{t-}},$$

we obtain

$$\begin{aligned} Z_t^\delta &= \langle (\boldsymbol{\alpha}^a; \epsilon_a)^\sim + (\boldsymbol{\alpha}^b; \epsilon_b)^\sim, \mathbb{Y}_t \rangle, \\ I_t^\pi &= \langle (\epsilon_b - \epsilon_a; \epsilon_s)^\sim + (\boldsymbol{\beta}; \epsilon_p)^\sim, \mathbb{Y}_t \rangle. \end{aligned} \quad (5.11)$$

Notice that the identities in (5.11) are linear in $\boldsymbol{\alpha}^a$, $\boldsymbol{\alpha}^b$, and $\boldsymbol{\beta}$.

If the utility function U in (5.5) is polynomial in (ζ, ι, ν) , or can be approximated by a polynomial, then the problem (5.5) can be rewritten as

$$\sup_{\boldsymbol{\alpha}_I^a, \boldsymbol{\alpha}_I^b, \boldsymbol{\beta}_I \in \mathbb{R}} \mathbb{E} \left[\tilde{U} \left(\langle (\boldsymbol{\alpha}^a; \epsilon_a)^\sim + (\boldsymbol{\alpha}^b; \epsilon_b)^\sim, \mathbb{Y}_T \rangle, \langle (\epsilon_b - \epsilon_a; \epsilon_s)^\sim + (\boldsymbol{\beta}; \epsilon_p)^\sim, \mathbb{Y}_T \rangle \right) \right], \quad (5.12)$$

for some polynomial $\tilde{U}$ defined by $\tilde{U}(\zeta, \iota) := U(\zeta, \iota, \iota^2)$. Finally, by the shuffle property of the signature, see for instance (5.3), the problem (5.12) reduces to an optimization problem that is polynomial in the coefficients $\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b, \boldsymbol{\beta}$ and linear in the expected signature.

We provide more details for the specific utility functions 1 and 2 in the next sections.

5.5 Mean–variance objective and controlled dynamics

We now specialize (5.5) to the mean–variance criterion, i.e. case 2. We first keep the asset and order-flow models general, and only later introduce a simple benchmark specification to explain the optimization algorithm.

We consider

$$J(\pi, \delta) := \mathbb{E} \left[Z_T^\delta \right] - w \mathbb{E} \left[\left(I_T^\pi \right)^2 \right], \quad w > 0, \quad (5.13)$$

where

$$dZ_t^\delta = \delta_t^a dN_t^a + \delta_t^b dN_t^b, \quad dI_t^\pi = Q_{t-} dS_t + \pi_{t-} \frac{dP_t}{P_{t-}}, \quad Q_t = N_t^b - N_t^a.$$

When S and P are $\mathbb{P}$ -martingales, (5.13) captures the usual trade-off between expected execution revenue and inventory/hedging risk. However, the formulation remains meaningful when prices are not martingales.

For the optimization procedure discussed below, we will focus on the case $\pi = 0$, so that

only the spread coefficients are trained. In particular, the trainable parameter will be

$$\theta = \alpha := (\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b), \quad \boldsymbol{\beta} = 0.$$

5.5.1 Signature-linear feedback and shuffle reduction

We adopt the signature-linear feedback parametrization

$$\delta_t^a = \langle \boldsymbol{\alpha}^a, \mathbb{Y}_t \rangle, \quad \delta_t^b = \langle \boldsymbol{\alpha}^b, \mathbb{Y}_t \rangle, \quad \pi_t^P = \frac{\pi_t}{P_t} = \langle \boldsymbol{\beta}, \mathbb{Y}_t \rangle, \quad (5.14)$$

where $\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b, \boldsymbol{\beta} \in T(\mathbb{R}^5)$ are supported on words over $\{0, s, p, a, b\}$ up to some truncation level K .

By Proposition 5.4.6, stochastic integrals of linear functionals of the signature against components of Y remain linear functionals of the signature. In particular, (5.11) yields

$$Z_T^\delta = \left\langle (\boldsymbol{\alpha}^a; \epsilon_a)^\sim + (\boldsymbol{\alpha}^b; \epsilon_b)^\sim, \mathbb{Y}_T \right\rangle, \quad (5.15)$$

$$I_T^\pi = \left\langle (\epsilon_b - \epsilon_a; \epsilon_s)^\sim + (\boldsymbol{\beta}; \epsilon_p)^\sim, \mathbb{Y}_T \right\rangle. \quad (5.16)$$

Define the linear functionals

$$\ell_Z(\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b) := (\boldsymbol{\alpha}^a; \epsilon_a)^\sim + (\boldsymbol{\alpha}^b; \epsilon_b)^\sim, \quad \ell_I(\boldsymbol{\beta}) := (\epsilon_b - \epsilon_a; \epsilon_s)^\sim + (\boldsymbol{\beta}; \epsilon_p)^\sim. \quad (5.17)$$

Then

$$Z_T^\delta = \langle \ell_Z, \mathbb{Y}_T \rangle, \quad I_T^\pi = \langle \ell_I, \mathbb{Y}_T \rangle.$$

The quadratic term in (5.13) can be reduced using the shuffle identity:

$$\left(I_T^\pi \right)^2 = \left\langle \ell_I(\boldsymbol{\beta}) \sqcup \ell_I(\boldsymbol{\beta}), \mathbb{Y}_T \right\rangle. \quad (5.18)$$

Hence

$$J(\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b, \boldsymbol{\beta}) = \left\langle \ell_Z(\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b) - w(\ell_I(\boldsymbol{\beta}) \sqcup \ell_I(\boldsymbol{\beta})), \mathbb{E}[\mathbb{Y}_T] \right\rangle. \quad (5.19)$$

Equation (5.19) shows that the mean-variance objective is an inner product between an explicit algebraic functional, built from the tilde and shuffle operations, and the expected terminal signature.

5.5.2 Semi-linear structure under the controlled law

The representation (5.19) is not closed form because $\mathbb{E}[\mathbb{Y}_T]$ depends on the control through the law of (N^a, N^b) , and more generally through the controlled trajectory distribution.

Collect the trainable coefficients into a parameter vector θ and write

$$J(\theta) = \left\langle L(\theta), \mathbb{E}_\theta[\mathbb{Y}_T] \right\rangle, \quad L(\theta) := \ell_Z(\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b) - w(\ell_I(\boldsymbol{\beta}) \sqcup \ell_I(\boldsymbol{\beta})). \quad (5.20)$$

This highlights a semi-linear structure: $L(\theta)$ is explicit and polynomial in the coefficients, hence easy to differentiate, whereas $\mathbb{E}_\theta[\mathbb{Y}_T]$ is implicit and depends on θ through the controlled law.

Equivalently, one may write

$$J(\theta) = q(\theta)^\top \theta + c(\theta),$$

where $q(\theta)$ and $c(\theta)$ are expectations of explicit functionals under the law $\mathbb{P}_\theta$.

5.5.3 Likelihood-ratio gradient and the role of the intensity model

We now explain the optimization procedure in the simplified case $\pi = 0$, so that $\theta = \alpha = (\boldsymbol{\alpha}^a, \boldsymbol{\alpha}^b)$ and $\boldsymbol{\beta} = 0$. In that case,

$$I_T = \langle \ell_I, \mathbb{Y}_T \rangle, \quad \ell_I = (\epsilon_b - \epsilon_a; \epsilon_s)^\sim,$$

which does not depend explicitly on α , while

$$Z_T^\delta = \langle \ell_Z(\alpha), \mathbb{Y}_T \rangle$$

does depend explicitly on α .

Accordingly, the realized return can be written as

$$R(\omega; \alpha) := Z_T^\delta(\omega; \alpha) - w I_T(\omega)^2 = \left\langle \ell_Z(\alpha), \mathbb{Y}_T(\omega) \right\rangle - w \left\langle \ell_I \sqcup \ell_I, \mathbb{Y}_T(\omega) \right\rangle. \quad (5.21)$$

The key point is that the dependence of $R(\omega; \alpha)$ on α is only partially explicit. Indeed,

$$\nabla_\alpha R(\omega; \alpha) = \left\langle \nabla_\alpha \ell_Z(\alpha), \mathbb{Y}_T(\omega) \right\rangle,$$

which is an explicit algebraic derivative obtained from the tilde calculus. By contrast, the dependence of the law of $\mathbb{Y}_T$ on α must be handled through a likelihood-ratio correction.

For any integrable functional $F_\alpha(\omega)$ that may depend on α both through the trajectory law and through an explicit α -dependence, the likelihood-ratio identity gives

$$\nabla_\alpha \mathbb{E}_\alpha[F_\alpha] = \mathbb{E}_\alpha[\nabla_\alpha F_\alpha(\omega) + F_\alpha(\omega) \nabla_\alpha \log p_\alpha(\omega)], \quad (5.22)$$

where $p_\alpha(\omega)$ denotes the likelihood of ω under parameter α . Applying (5.22) to (5.21), we obtain

$$\nabla_\alpha J(\alpha) = \mathbb{E}_\alpha[\nabla_\alpha R(\omega; \alpha) + R(\omega; \alpha) \nabla_\alpha \log p_\alpha(\omega)]. \quad (5.23)$$

Thus, the gradient splits into:

1. an *explicit algebraic term*, coming from the direct dependence of the execution revenue on α through the tilde representation;
2. a *likelihood-ratio term*, accounting for the dependence of the controlled law on α .

To make the algorithm fully explicit, one needs a specification of the map $\delta \mapsto \lambda(\delta)$ linking the quoted spreads to the order-arrival intensities. For the sake of exposition, and in the spirit of [69], we assume in this section that:

1. the price process follows a Bachelier model,

$$dS_t = \sigma dW_t, \quad (5.24)$$

2. N^a and N^b are conditionally Poisson processes,
3. the arrival intensities depend exponentially on the quoted spreads:

$$\lambda_t^a(\delta_t^a) = A e^{-k\delta_t^a}, \quad \lambda_t^b(\delta_t^b) = A e^{-k\delta_t^b}, \quad (5.25)$$

where $A, k > 0$.

This specification is used only to derive an explicit score. The algebraic signature linearization above does *not* require the Bachelier model, Poisson arrivals, martingale prices, or exponential intensities. In particular, the same decomposition (5.23) remains valid for more general price dynamics and more complex point-process models, provided the score of the controlled law can be computed or estimated.

Under the Poisson specification (5.25), the conditional log-likelihood of the counting paths

is

$$\begin{aligned} \log p_\alpha(N^a, N^b \mid \text{cont.}) &= \int_0^T \log \lambda_t^a(\alpha) dN_t^a - \int_0^T \lambda_t^a(\alpha) dt \\ &\quad + \int_0^T \log \lambda_t^b(\alpha) dN_t^b - \int_0^T \lambda_t^b(\alpha) dt + (\text{terms independent of } \alpha). \end{aligned} \quad (5.26)$$

Therefore,

$$\nabla_\alpha \log p_\alpha = \int_0^T \nabla_\alpha \log \lambda_t^a(\alpha) (dN_t^a - \lambda_t^a(\alpha) dt) + \int_0^T \nabla_\alpha \log \lambda_t^b(\alpha) (dN_t^b - \lambda_t^b(\alpha) dt). \quad (5.27)$$

Since

$$\delta_t^a = \langle \boldsymbol{\alpha}^a, \mathbb{Y}_t \rangle, \quad \delta_t^b = \langle \boldsymbol{\alpha}^b, \mathbb{Y}_t \rangle,$$

and

$$\lambda_t^{a/b} = A e^{-k \delta_t^{a/b}}, \quad \nabla_\alpha \log \lambda_t^{a/b} = -k \nabla_\alpha \delta_t^{a/b},$$

we obtain, for the spread parameters,

$$\nabla_{\alpha^a} \log p_\alpha = -k \int_0^T \mathbb{Y}_t (dN_t^a - \lambda_t^a dt), \quad \nabla_{\alpha^b} \log p_\alpha = -k \int_0^T \mathbb{Y}_t (dN_t^b - \lambda_t^b dt). \quad (5.28)$$

Discrete-time form. On a grid $t_n = n dt$, with

$$\Delta N_{n+1}^{a/b} := N_{t_{n+1}}^{a/b} - N_{t_n}^{a/b},$$

and prefix features

$$\phi_n = \pi_{\leq K}(\mathbb{Y}_{t_n}),$$

or their flattened representation, (5.28) becomes

$$\nabla_{\alpha^a} \log p_\alpha = -k \sum_{n=0}^{N-1} \phi_n (\Delta N_{n+1}^a - \lambda_n^a dt), \quad \nabla_{\alpha^b} \log p_\alpha = -k \sum_{n=0}^{N-1} \phi_n (\Delta N_{n+1}^b - \lambda_n^b dt), \quad (5.29)$$

where $\lambda_n^{a/b} := \lambda_{t_n}^{a/b}(\delta_{t_n}^{a/b})$.

Remark 5.5.1. The tilde linearization (5.15)–(5.16) and the shuffle identity are purely algebraic and hold for any semimartingale model of (S, P, N^a, N^b) satisfying the standing assumptions. The Poisson/Bachelier specification is used only to provide an explicit and simple likelihood-ratio term. In particular, one may replace (5.24)–(5.25) by non-martingale price dynamics or more general point-process models, at the cost of replacing the score (5.27)–(5.29) accordingly.

5.5.4 Pseudo-algorithm: likelihood-ratio optimization with tilde/shuffle payoff

We summarize a generic gradient-ascent scheme for maximizing (5.19) in the case $\pi = 0$, $\theta = \alpha = (\alpha^a, \alpha^b)$.

Algorithm 3: Signature market making via likelihood-ratio gradient (mean–variance, $\pi = 0$).

Input: Signature truncation level K ; coefficients $\alpha^{(0)} = (\alpha^{a,(0)}, \alpha^{b,(0)})$; batch size B ; learning-rate schedule $(\rho_m)_{m \geq 0}$; inventory-penalty weight $w > 0$; maximum iterations $M_{\max}$; convergence tolerance $\varepsilon > 0$, iteration counter $m \leftarrow 0$.

1 **repeat**

 /* Sample paths under the current policy law. */

2 Simulate B trajectories $\omega^{(1)}, \dots, \omega^{(B)}$ under $\mathbb{P}_{\alpha^{(m)}}$;

3 For each trajectory, compute the terminal signature $\mathbb{Y}_T^{(b)}$ and the prefix features $(\phi_n^{(b)})_{n=0}^{N-1}$;

 /* Evaluate the realized reward via the algebraic identity. */

4 Compute the realized returns

$$R^{(b)}(\alpha^{(m)}) = \langle \ell_Z(\alpha^{(m)}), \mathbb{Y}_T^{(b)} \rangle - w \langle \ell_I \sqcup \ell_I, \mathbb{Y}_T^{(b)} \rangle, \quad b = 1, \dots, B,$$

 where $\ell_I = (\epsilon_b - \epsilon_a; \epsilon_s)^\sim$;

 /* Algebraic gradient of the (signature-linear) reward in α . */

5 Compute

$$\nabla_\alpha R^{(b)}(\alpha^{(m)}) = \langle \nabla_\alpha \ell_Z(\alpha^{(m)}), \mathbb{Y}_T^{(b)} \rangle;$$

 Compute the score terms using (5.29);

 /* Monte-Carlo policy-gradient estimator. */

6

$$\widehat{\nabla_\alpha J}^{(m)} = \frac{1}{B} \sum_{b=1}^B \left(\nabla_\alpha R^{(b)}(\alpha^{(m)}) + \left(R^{(b)}(\alpha^{(m)}) - \bar{R}^{(m)} \right) \nabla_\alpha \log p_{\alpha^{(m)}}(\omega^{(b)}) \right);$$

 /* Gradient ascent with per-level step sizes. */

7 Update $\alpha^{(m+1)} \leftarrow \alpha^{(m)} + \rho_m \widehat{\nabla_\alpha J}^{(m)}$;

8 $m \leftarrow m + 1$;

9 **until** $\|\widehat{\nabla_\alpha J}^{(m)}\| < \varepsilon$ or $m = M_{\max}$;

Return: Trained policy parameters $\alpha^\star = \alpha^{(m)}$.

In Algorithm 3, the signature algebra is used twice: first, to write the payoff in the linear/polynomial form (5.21); second, to compute the explicit derivative $\nabla_{\alpha}R(\omega; \alpha)$ without differentiating through the signature map itself. The simplified Poisson/Bachelier specification is used only to obtain the explicit score term. The same optimization principle extends to more general intensity models and more general price dynamics whenever the score of the controlled law is available.

5.6 Numerical experiments

5.6.1 Setup and benchmarks

We now illustrate the performance of the signature-linear market-making policy. The purpose of the experiments is to test whether the policy class and the likelihood-ratio training procedure are able to reproduce an expert strategy in a classical benchmark and to exploit additional path dependence when the order flow has memory.

We compare two policies. The first is the learned signature policy, in which the ask and bid half-spreads are linear functionals of the truncated signature of the observed market path. The second is an expert benchmark computed from the reference model. In the pathwise comparisons, the two policies are evaluated on matched random-number paths whenever possible, so that the differences reflect the quoting rule rather than exogenous simulation noise.

5.6.2 Policy optimization versus expert approximation

Why this particular comparison structure? A natural objection is that the expert benchmarks used below (LGFT in the Poisson experiment; a Hawkes-calibrated LGFT in the self-exciting experiment) are not themselves path-dependent learners, and one might expect a fairer test against, say, a deep-learning baseline trained on the same objective. We deliberately do not do this.

The expert is the natural ablation, not the natural competitor. The purpose of the Poisson experiment is not to claim that the signature policy outperforms a state-of-the-art deep RL agent; it is to test whether the signature parametrization combined with the likelihood-ratio gradient *recovers* the analytical optimum in a setting where the analytical optimum is known. The expert is the right benchmark for that question. The Hawkes experiment then asks the complementary question: when the analytical optimum is not available because the order-flow law is path-dependent, does the signature parametrization exploit the additional information?

The right benchmark is then the same Avellaneda–Stoikov-type expert evaluated under the misspecified (Poisson-projected) intensity, and the comparison gap measures the gain from path-dependence.

A principled comparison against a deep-learning baseline will require constructing a reward decomposition that is consistent in the relevant continuous-time limit; we leave this to future work. Our experiments compare the learned signature policy with an expert benchmark at the level of realized performance. A complementary question is whether the learning procedure is merely recovering a linear approximation of the expert quotes, or whether it is genuinely optimizing the market-making objective under the controlled dynamics. This distinction is important. Since signature features form a rich basis for path-dependent functionals, one could also use them in a purely supervised way: simulate expert trajectories, regress the expert ask and bid quotes on the truncated signature of the observed path, and then deploy the resulting signature-linear approximation as a quoting rule. Such an approach tests the expressiveness of the signature parametrization, but it does not solve the control problem directly.

To separate these two effects, we consider the following additional benchmark. Given trajectories generated under the expert policy, we compute the expert quotes $(\delta_t^{a,\text{exp}}, \delta_t^{b,\text{exp}})$ and fit coefficients $(\hat{\alpha}_{\text{sup}}^a, \hat{\alpha}_{\text{sup}}^b)$ by least squares:

$$\min_{\alpha^a, \alpha^b} \sum_{m=1}^{M_{\text{train}}} \sum_{n=0}^{N-1} \left| \delta_{t_n}^{a,\text{exp},m} - \langle \alpha^a, \mathbb{Y}_{t_n}^m \rangle \right|^2 + \left| \delta_{t_n}^{b,\text{exp},m} - \langle \alpha^b, \mathbb{Y}_{t_n}^m \rangle \right|^2.$$

The fitted policy is then evaluated out of sample in the same controlled environment as the policy-gradient solution. This provides a natural ablation: the supervised policy uses the same signature features and the same linear-in-signature parametrization, but it is trained to imitate the expert locally rather than to maximize the global mean–variance objective.

The comparison is informative for two reasons. First, if the supervised signature policy performs close to the expert, then the chosen truncation level is sufficiently expressive to represent the benchmark quotes. Second, if the policy-gradient solution differs from the supervised approximation while achieving a higher objective, then the learning procedure is not simply imitating the expert; it is exploiting the feedback between quotes, execution intensities, inventory, and terminal performance.

In the Poisson benchmark, where the expert policy is well matched to the simulation model, we expect the supervised signature approximation and the policy-gradient solution to be close. This is consistent with the purpose of the Poisson experiment: the order flow is conditionally memoryless, and therefore there is limited additional value in learning a richer path-dependent policy. The relevant test is whether the learned policy recovers a

stable quoting rule with performance close to the expert and to its supervised signature approximation.

In the Hawkes experiment, the same comparison is more revealing. Because order arrivals are self-exciting, the expert benchmark fails to exploit all the path-dependent information contained in the recent order-flow history unless the state is explicitly augmented with the relevant excitation variables. A signature-linear policy, by contrast, can condition on iterated path features of the observed process. The figures and tables collected in Section 5.6.5 report both quote-level and performance-level diagnostics. The quote-level diagnostics are the inventory–spread correlations of Table 5.2 and the matched-path comparisons of Figures 5.7, 5.10 and 5.11; the performance-level diagnostics are the mean reward, terminal profit-and-loss, inventory penalty, and terminal-inventory statistics shown in Figures 5.8 and 5.9, together with the regime-by-regime gap table 5.1. Read together, they separate approximation quality from control quality: in every Hawkes regime considered, the policy-gradient solution departs from the near-affine expert in quote space (Table 5.2) while delivering a better mean–variance objective (Table 5.1), which is precisely the regime in which the signature parametrization is supposed to add value.

The numerical evidence therefore distinguishes three regimes. In the first, the policy-gradient solution is close to both the expert and the supervised signature approximation; this indicates that the expert is essentially optimal within the tested environment and that the signature parametrization recovers it. In the second, the policy-gradient solution is close in performance but not in quote space; this indicates that several quoting rules may be nearly equivalent for the objective. In the third, most relevant for path-dependent order flow, the policy-gradient solution departs from the expert and improves the objective; this suggests that the signature policy is using historical information that is not fully captured by the benchmark strategy.

5.6.3 Choice of the signature truncation level

Before turning to the optimization experiments, we briefly justify the truncation level used throughout the rest of this section. Recall that the signature-linear policy of Section 5.5.1 is parametrized by a vector of coefficients indexed by words of length at most K on the alphabet of the augmented path. Increasing K enlarges the policy class but also the dimension of the parameter to be estimated, and in our experience the variance of the likelihood-ratio gradient grows accordingly. Figure 5.1 only addresses the first, expressivity, side of this trade-off; the second side is reflected indirectly in the choice $K = 3$ for the experiments below.

We assess the expressive power of each truncation level by a model-dependent oracle

test. For each setting (Poisson and Hawkes), we generate paths under the expert policy, compute the truncated signature of the augmented path, and run a ridge regression of the expert’s quoted spread on the signature features. We then evaluate the in-sample R^2 and the out-of-sample reward obtained by replaying the fitted linear-in-signature rule. The result is a one-dimensional diagnostic of how well a signature of given level can *represent* the expert before any optimization is involved.




Figure 5.1: R^2 of the ridge fit to the expert policy and out-of-sample reward of the fitted linear-in-signature rule, as a function of the signature truncation level K . Top panel: Poisson order flow. Bottom panel: Hawkes order flow.

In both settings, the marginal gain from $K = 2$ to $K = 3$ is substantial, while moving to $K = 4$ contributes only a small improvement at the cost of a much larger parameter space. We therefore fix $K = 3$ as the default truncation level in the optimization experiments below.

5.6.4 Poisson order flow

The first experiment uses conditionally Poisson ask and bid arrivals with exponential dependence on the quoted half-spreads. This is the natural benchmark for the method because it is close to the classical Avellaneda–Stoikov and Guéant–Lehalle–Fernandez-Tapia specifications. Since the model is memoryless conditional on the current quotes, a well-designed expert policy should already capture most of the relevant structure. We therefore view the Poisson experiment as a stability and approximation test rather than as a setting where one should expect a large outperformance from path-dependent features.

Figure 5.2 reports the training history. The average reward increases sharply during the initial iterations and then stabilizes. The execution-revenue component and the inventory-penalty component also settle into a stable regime. This suggests that the likelihood-

ratio gradient identifies profitable spread coefficients without generating unstable inventory dynamics.




Figure 5.2: Training history in the Poisson experiment. The top panel reports the average reward, while the lower panels decompose the objective into the execution-revenue component and the inventory-penalty component.

The pathwise comparison in Figure 5.3 shows that the learned policy produces dynamic spreads in a reasonable range and keeps inventory centered around zero. In the displayed matched paths, the learned policy tends to quote wider spreads than the expert. This conservativeness is consistent with the distributional results below: the policy sacrifices some execution revenue, but maintains controlled inventory risk.

Figures 5.4 and 5.5 summarize the out-of-sample distributions. In the displayed run, the expert has a modest advantage in the mean objective value, with mean reward approximately 436.6 compared with 424.9 for the learned policy. The expert also has a higher mean terminal execution component, approximately 519.7 compared with 503.3. At the same time, the learned policy attains smaller average inventory-risk statistics in the reported run, with lower mean values of both I_T^2 and Q_T^2 . Overall, the learned policy approximates the expert closely in a setting where the expert is well matched to the data-generating model.




Figure 5.3: Matched-path comparison in the Poisson experiment. The left panels compare the total spread $\delta^a + \delta^b$ under the learned policy and the expert. The right panels compare the corresponding inventory paths.




Figure 5.4: Poisson experiment: distribution of the mean-variance reward, left, and terminal execution component, right, for the learned policy and the expert.

5.6.5 Hawkes order flow

The second experiment replaces the memoryless Poisson order flow by a Hawkes-type specification. This introduces self-excitation: recent arrivals increase future arrival intensities. The resulting control problem is more path-dependent, because the recent order-flow history carries information about the near-future arrival rate. This is the environment in which signature features should be most valuable.

Hawkes specification used in this experiment. We use a self-exciting model with exponential kernel ϕ and intensities

$$\lambda_t^{a/b}(\delta_t^{a/b}) = A e^{-k\delta_t^{a/b}} \left(\mu + \sum_{\tau_i^{a/b} < t} \phi(t - \tau_i^{a/b}) \right),$$




Figure 5.5: Poisson experiment: inventory-risk statistics for the learned policy and the expert. The learned policy is slightly more conservative in this run.

with $\phi \geq 0$ and $\|\phi\|_{L^1} < 1$ (sub-criticality), so that the intensity is finite a.s. and the conditional log-likelihood retains the compensated-jump form used in (5.27). The “self-excitation strength” k used in the regime tables below is the spread sensitivity of (5.25); the kernel ϕ is held fixed across regimes.

The training history in Figure 5.6 again shows fast stabilization. Compared with the Poisson experiment, the inventory-penalty component is less erratic after the initial transient, and the average reward remains stable over the later iterations. The behavior suggests that the same likelihood-ratio approach can be used beyond the memoryless benchmark, provided that the score of the controlled law is available or can be estimated.

Figure 5.7 displays the learned spreads and inventories along sample paths. The policy reacts to the realized order-flow history, producing substantial variation across paths while keeping inventory fluctuations controlled. The clustered nature of the Hawkes arrivals makes this pathwise adaptation economically meaningful: past arrivals contain information about future intensities, and the signature features encode this information without imposing a fixed finite-dimensional Markovian representation by hand.

The distributional comparison is more favorable to the signature policy than in the Poisson case. In the displayed Hawkes run, the learned policy achieves a higher mean objective value, approximately 284.2 compared with 277.3 for the expert. The improvement does not come from higher terminal execution revenue: the expert has a slightly higher mean terminal execution component, approximately 376.3 compared with 368.7. Instead, the learned policy improves the mean–variance trade-off by reducing inventory risk. The reported average values of I_T^2 and Q_T^2 are lower for the learned policy than for the expert.




Figure 5.6: Training history in the Hawkes experiment. The learned policy stabilizes quickly despite the self-exciting order-flow dynamics.

Robustness across Hawkes regimes

To probe the qualitative claim of the previous paragraphs, we repeat the Hawkes experiment under several combinations of horizon T and self-excitation strength k , and compare the learned signature policy to the expert benchmark using paired random-number-generator out-of-sample paths at $B = 16,384$ (inventory-penalty weight $w = 0.05$). Table 5.1 reports, for each regime, the signed difference (signature policy minus expert) in average mean–variance reward $\Delta\bar{R}$, mean terminal execution component $\Delta\bar{Z}$, mean squared inventory $\Delta\bar{I}^2$, and mean number of arrivals Δarr .

Table 5.1: Differences (learned policy – expert) in execution PnL, inventory variance, and arrival count across Hawkes regimes. Paired-RNG out-of-sample evaluation at $B = 16,384$, inventory-penalty weight $w = 0.05$.

T	k	$\Delta\bar{R}$	$\Delta\bar{R}$ (%)	$\Delta\bar{Z}$	$\Delta\bar{I}^2$	Δarr
300	0.2	+6.86	+2.47%	−12.48	−373.3	−8.0
300	1.0	+2.29	+9.20%	−14.39	−339.3	−12.1
100	1.0	+1.09	+13.09%	−5.74	−131.7	−6.4
60	0.4	+1.38	+6.08%	−2.82	−78.7	−1.1




Figure 5.7: Hawkes experiment: sample paths of ask and bid half-spreads together with inventory paths. Solid curves are quotes and dotted curves are inventories.

Across all four regimes the learned policy obtains a strictly positive mean–variance gap. The signs of the auxiliary diagnostics are consistent with the economic mechanism described above: the learned policy quotes more conservatively, executes fewer trades ($\Delta_{\text{arr}} < 0$), and pays a small cost in execution revenue ($\Delta \bar{Z} < 0$), but achieves a sizable reduction in inventory variance ($\Delta \bar{I}^2 < 0$). The trade is therefore a controlled exchange of pure execution PnL for inventory-risk control, which is precisely the trade-off that a Hawkes-aware path-dependent policy is supposed to exploit.

For visual intuition, Figure 5.10 shows a representative pathwise diagnostic in the main configuration $(T, k) = (300, 0.2)$, and Figure 5.11 compares the policies in a shorter-horizon, more strongly self-exciting regime, where the non-linearity of the learned quotes becomes more apparent.




Figure 5.8: Hawkes experiment: distribution of the mean–variance reward, left, and terminal execution component, right, for the learned policy and the expert.




Figure 5.9: Hawkes experiment: inventory-risk statistics for the learned policy and the expert. In this path-dependent setting, the learned policy achieves a better risk profile.

To quantify how the structural relationship between quotes and inventory differs from the expert, Table 5.2 reports the linear correlations $\rho(\delta^a, Q)$ and $\rho(\delta^b, Q)$ for the learned policy (π) and the expert (exp) in the same four Hawkes regimes. The expert is, by construction, almost perfectly affine in the inventory: its correlations are uniformly close to ± 1 . The learned policy preserves the correct sign and the qualitative skew-symmetry between the ask and bid legs, but exhibits visibly weaker linear correlations whenever T or k make the problem more strongly path-dependent. Read together with Figure 5.11, this is consistent with the learned policy genuinely using non-linear features of the augmented path rather than collapsing onto an inventory-only rule.




Figure 5.10: Hawkes experiment, main configuration $(T, k) = (300, 0.2)$: pathwise comparison of the learned signature policy and the expert benchmark on a representative out-of-sample trajectory.

5.6.6 Discussion

The two experiments convey a consistent message. In the Poisson benchmark, the expert policy is naturally adapted to the memoryless structure of the model. The signature policy learns a stable quoting rule and approximates the expert, but it does not produce a large improvement. In the Hawkes setting, the order flow contains additional path dependence. The signature policy can use this information through the history of the augmented path, and the resulting strategy improves the observed mean–variance objective by reducing inventory risk.

The robustness check across (T, k) regimes (Table 5.1) shows that the mean–variance improvement is not specific to one parameter combination: in every regime considered, the learned policy obtains a positive $\Delta\bar{R}$ together with strictly negative $\Delta\bar{Z}$, $\Delta\bar{I}^2$, and Δarr . The same regimes that produce the largest $\Delta\bar{R}$ are also the ones in which the inventory–spread correlation of the learned policy departs the most from the near-affine expert (Table 5.2), supporting the interpretation that the additional value of the signature parametrization is delivered through genuinely non-linear, path-dependent quoting rather than through a




Figure 5.11: Hawkes experiment, shorter horizon and stronger self-excitation: the learned policy departs more visibly from a linear-in-inventory rule, reflecting the additional value of path-dependent features in this regime.

re-scaling of the expert.

5.7 Conclusion

This paper develops a signature-based approach to optimal market making with path-dependent order flow. The central idea is to parametrize the market maker's quotes as linear functionals of the signature of an augmented market path. This produces a flexible class of feedback controls that can depend on the entire observed history of prices, inventories, and order arrivals, while preserving an explicit algebraic structure. In particular, stochastic integrals of signature-linear controls can be rewritten as linear functionals of the terminal signature, and quadratic inventory penalties can be handled through the shuffle identity. For a mean-variance objective, this yields a tractable representation of the realized payoff and its direct derivative with respect to the policy coefficients.

A key feature of the method is that the learning structure is largely model-agnostic. The signature algebra does not rely on a specific price model, nor on Poisson order arrivals. The price process may be Brownian, jump-driven, rough, stochastic-volatility-based, empirically

Table 5.2: Linear correlation between the half-spreads and the inventory Q for the learned policy (π) and the expert (exp) across Hawkes regimes.

T	k	$\rho(\delta^a, Q)_\pi$	$\rho(\delta^a, Q)_{\text{exp}}$	$\rho(\delta^b, Q)_\pi$	$\rho(\delta^b, Q)_{\text{exp}}$
300	0.2	-0.980	-0.989	+0.973	+0.989
300	1.0	-0.756	-0.994	+0.752	+0.994
100	1.0	-0.689	-0.986	+0.671	+0.986
60	0.4	-0.945	-0.967	+0.939	+0.967

simulated, or generated by a black-box model. What is required for the policy-gradient step is not a closed form solution of the control problem, but the ability to differentiate the controlled order-flow likelihood, or an appropriate surrogate for it, with respect to the quoted spreads. More precisely, whenever the arrival intensity or conditional execution mechanism is differentiable with respect to the control δ , the likelihood-ratio term can be computed or estimated. The same training architecture can therefore be applied across a broad class of models, including Poisson, Hawkes, state-dependent, neural, or otherwise black-box intensity specifications.

This separation between the price dynamics and the execution mechanism is particularly useful in market making. The market maker’s quotes directly affect the probability of execution, while the mid-price dynamics may be exogenous, partially observed, or learned from data. The proposed approach accommodates this separation naturally. The payoff representation uses the realized path, whereas the score term only requires sensitivity of the execution law with respect to the quotes. As a result, the method can be used in environments where the price model is too complex to yield a tractable Hamilton–Jacobi–Bellman equation, but where one can still simulate trajectories and evaluate the derivative of the intensity model. The numerical experiments provide preliminary evidence for this perspective. In the Poisson benchmark, the signature policy learns a stable strategy whose performance is close to that of an expert policy specifically designed for the model. In the Hawkes setting, where order flow is self-exciting and more strongly path-dependent, the learned policy can improve the mean–variance trade-off relative to the expert. This suggests that the value of the signature parametrization increases when the control problem contains genuine memory effects. Several directions remain open. First, the numerical study should be extended to richer market simulators and to empirical limit-order-book data, where the intensity model may be estimated from observations or represented by a neural point process. Second, it would be useful to investigate variance-reduction techniques for the likelihood-ratio estimator, including control variates, baselines, actor–critic variants, and pathwise derivatives when available. Third, one should study the interaction between truncation level, feature selection,

and economic interpretability of the learned signature coefficients. Finally, the theoretical approximation results for signature controls suggest a route toward convergence guarantees for market-making policies trained over increasing signature levels. Developing such results in the presence of controlled point-process intensities is a promising topic for future research.

Extra results from the portfolio-delegation paper

This appendix gathers the supplementary material from the paper *Delegated portfolio management with random default* (joint work with Thibaut Mastrolia, to appear in *Mathematical Finance*). All sections in what follows correspond verbatim, up to renumbering, to the appendix of that paper.

A.1 First-best optimal delegated contract

As a benchmark case, we consider here the situation where the principal proposes both a contract and an investment strategy to the agent. We rewrite the principal problem in the first-best case as

$$V_0^{FB} = \sup_{(\xi, \pi)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^\pi - \xi) \right],$$

subject to

$$J^A(\pi; x, \xi) \geq R.$$

We consider exponential utilities for both the agent and the principal so that $U^P(x) = -e^{-\varsigma x}$ and $U^A(x) = -e^{-\eta x}$, for risk-aversion parameters $\varsigma, \eta > 0$ and no penalty on the trading strategy $\varepsilon = 0$. Introducing a Lagrange multiplier $\rho \geq 0$, the problem is reduced to

$$\begin{aligned} V_0^{FB} &= \inf_{\rho \geq 0} \sup_{\pi} \sup_{\xi} \mathbb{E} \left[U^P(X_{T \wedge \tau}^\pi - \xi) + \rho(J^A(\pi; x, \xi) - R) \right], \\ &= \inf_{\rho \geq 0} \sup_{\pi} \sup_{\xi} \mathbb{E} \left[U^P(X_{T \wedge \tau}^\pi - \xi) - \rho e^{-\eta \xi} - \rho R \right]. \end{aligned}$$

Using Gâteaux derivatives, see for example [115], the first-order condition on an optimizer ξ^* gives

$$e^{-\varsigma(X_{T \wedge \tau} - \xi^*)} = \frac{\rho \eta}{\varsigma} e^{-\eta \xi^*}. \quad (\text{A.1})$$

Consequently,

$$\xi^* = \frac{1}{\varsigma + \eta} \left[\log\left(\frac{\eta \rho}{\varsigma}\right) + \varsigma X_{T \wedge \tau} \right].$$

Note that the optimal contract is linear with respect to $X_{T \wedge \tau}$ and similar to the contract in [53, Section 2.1] as a function of the Lagrangian ρ . Substituting this optimizer into the investor's first-best problem and using the first-order condition (A.1) we get the following:

$$\begin{aligned}
V_0^{FB} &= \inf_{\rho \geq 0} \sup_{\pi} \mathbb{E} \left[U^P(X_{T \wedge \tau}^{\pi} - \xi^*) + \rho(J^A(\pi; x, \xi^*) - R) \right], \\
&= \inf_{\rho \geq 0} \sup_{\pi} \mathbb{E} \left[-\rho \left(\frac{\varsigma + \eta}{\varsigma} \right) e^{-\eta \xi^*} \right] - \rho R \\
&= \inf_{\rho \geq 0} \sup_{\pi} \mathbb{E} \left[-\rho \left(\frac{\varsigma + \eta}{\varsigma} \right) \left(\frac{\varsigma}{\eta \rho} \right)^{\frac{\eta}{\varsigma + \eta}} e^{-\frac{\varsigma \eta}{\varsigma + \eta} X_{T \wedge \tau}} \right] - \rho R \\
&= \inf_{\rho \geq 0} \left\{ \rho \left(\frac{\varsigma + \eta}{\varsigma} \right) \left(\frac{\varsigma}{\eta \rho} \right)^{\frac{\eta}{\varsigma + \eta}} \tilde{V}_0^{FB} - \rho R \right\},
\end{aligned}$$

where

$$\tilde{V}_0^{FB} = \sup_{\pi} \mathbb{E} \left[-e^{-\frac{\varsigma \eta}{\varsigma + \eta} X_{T \wedge \tau}} \right].$$

We introduce the following HJB equations depending on whether the support of τ is bounded in $[0, T]$ or not.

$$(\mathbf{bHJB-FB}) \quad \begin{cases} \partial_t \tilde{v}(t, s, x) - e^{-\frac{\varsigma \eta}{\varsigma + \eta} x} f(t) + \sup_{\nu} \{ \tilde{\mathcal{H}}^{\tilde{v}}(t, s, x; \nu) \} = 0, & t < T \\ \tilde{v}(T, s, x) = 0, & (s, x) \in \mathbb{R}^m \times \mathbb{R}, \end{cases}$$

and

$$(\mathbf{uHJB-FB}) \quad \begin{cases} \partial_t \tilde{v}(t, s, x) - e^{-\frac{\varsigma \eta}{\varsigma + \eta} x} f(t) + \sup_{\nu} \{ \tilde{\mathcal{H}}^{\tilde{v}}(t, s, x; \nu) \} = 0, & t < T \\ \tilde{v}(T, s, x) = -(1 - F_{\tau}(T)) e^{-\frac{\varsigma \eta}{\varsigma + \eta} x}; & (s, x) \in \mathbb{R}^m \times \mathbb{R}, \end{cases}$$

where

$$\begin{aligned}
\tilde{\mathcal{H}}^{\tilde{v}}(t, s, x; \nu) &= \nu b(t, s) \partial_x \tilde{v}(t, s, x) + \frac{1}{2} \|\nu \sigma(t, s)\|^2 \partial_{xx} \tilde{v}(t, s, x) \\
&\quad + \sum_{i=1}^m \partial_i \tilde{v}(t, s, x) b^i(t, s) s^i + \frac{1}{2} \text{Tr}(\Sigma(t, s) \Sigma(t, s)^{\top} D^2 \tilde{v}(t, s, x)) \\
&\quad + \sum_{i=1}^m \partial_{x s^i} \tilde{v}(t, s, x) \nu \sigma(t, s) (\sigma^i(t, s))^{\top} s^i.
\end{aligned}$$

We then get the following theorem as a direct application of Ito's formula with localization procedure similar to the proof of Theorem 3.3.7.

Theorem A.1.1. *Assume that Hypothesis B (resp. Hypothesis A) holds and there exists a solution of **(bHJB-FB)** (resp. **(uHJB-FB)**) denoted by $\tilde{v}$ with polynomial growth in s, x .*

Then, the optimal contract is given by

$$\xi^* = \frac{1}{\varsigma + \eta} [\log(\frac{\eta \hat{\rho}}{\varsigma}) + \varsigma X_{T \wedge \tau}],$$

where $\hat{\rho}$ is the optimizer of

$$\inf_{\rho \geq 0} \left\{ \rho \left(\frac{\varsigma + \eta}{\varsigma} \right) \left(\frac{\varsigma}{\eta \rho} \right)^{\frac{\eta}{\varsigma + \eta}} \tilde{V}_0^{FB} - \rho R \right\},$$

where $\tilde{V}_0^{FB} = \tilde{v}(0, S_0, x)$ and the optimal trading strategy is given by

$$\hat{\pi}_t = \nu^*(t, S_t, X_t),$$

where $\nu^*(t, s, x)$ is an optimizer of $\tilde{\mathcal{H}}^{\bar{v}}(t, s, x; \nu)$.

Remark A.1.2. Unlike moral hazard, the optimal contract is a linear function of $X_{T \wedge \tau}$ which is often observed in first-best scenario in contract theory, see for example [115]. Note that the random default only differs from the standard first-best contract through the modified terminal-value distribution induced by $T \wedge \tau$.

A.2 Portfolio manager's problem and second-order BSDE

In this appendix, we provide a rigorous mathematical framework and introduce the weak formulation of the portfolio optimization problem, extending [53, 54] to the case of an erratic horizon. It is now well established that the agent's problem under moral hazard in contract theory can be reduced to solving a second-order BSDE when the volatility of the asset is controlled, as first explained in [54].¹

The results used in this section rely on several new developments in the theory of second-order BSDEs with jumps. Second-order BSDEs were initially introduced in [44, 150], then reformulated for stochastic control in [136], and later extended to random horizons in [107, 108, 67]. The theory has more recently been generalized to jumps and general semimartingales in [59], with explicit links to stochastic control problems in [116]. The work [67] is particularly relevant for our analysis, as it directly connects 2BSDEs with jumps and *erratic termination* to the associated stochastic control problem.

We begin by establishing a rigorous probabilistic setting under which the portfolio

¹Note that the very recent results of Chiusolo and Hubert [45] circumvent the need for second-order BSDEs by taking a detour through a first-best reformulation. However, since we now have access to a general theory of 2BSDEs with jumps and random horizons, we prefer to rely directly on the 2BSDE approach rather than to employ the technical shortcut of Chiusolo and Hubert.

manager's optimization problem is well-posed for any fixed compensation scheme ξ , under general integrability assumptions. We then prove the correspondence between the 2BSDE with erratic horizon and the value of the optimization problem. Throughout this appendix, we focus on the case $\epsilon = 0$, in order to stay consistent with the model of Cadenillas, Cvitanic and Zapatero [34], thereby extending [54, Section 6.3] to the case of portfolio optimization with erratic termination. For the sake of simplicity, and in line with a more realistic contractual structure, we assume in this appendix that the price process S is not contractible for the investor (client, principal). Consequently, we take X as the only contractible variable and thus as the main canonical process in the probabilistic weak formulation developed below.

A.2.1 Probabilistic framework and weak formulation

Let $\Omega := C([0, T], \mathbb{R})$ be the canonical space endowed with the canonical process X and canonical filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$. Let $\tau : \Omega \rightarrow [0, T]$ be a finite $\mathbb{F}$ -measurable stopping time. We denote by $\mathbb{G} = (\mathcal{G}_t)_{0 \leq t \leq T}$ the progressive enlargement of $\mathbb{F}$ with τ , defined by

$$\mathcal{G}_t := \bigcap_{s > t} (\mathcal{F}_s \vee \sigma(\tau \wedge s)), \quad 0 \leq t \leq T,$$

which satisfies the usual conditions. A control process is a $\mathbb{G}$ -progressively measurable process $\pi = (\pi_t)_{0 \leq t \leq T}$ taking values in a set $A \subset \mathbb{R}$ and satisfying Definition 3.2.4. For such a control π , we say that a probability measure $\mathbb{P}$ on $(\Omega, \mathcal{G}_T)$ belongs to the set $\mathcal{P}(\pi)$ if there exists a $\mathbb{G}$ -Brownian motion $W^\mathbb{P}$ under $\mathbb{P}$ such that

$$X_t = x + \int_0^{t \wedge \tau} \pi_s \sigma_s dW_s^\mathbb{P} + \int_0^{t \wedge \tau} \pi_s b_s ds, \quad 0 \leq t \leq T, \quad \mathbb{P}\text{-a.s.}$$

The global set of admissible probability measures on the enlarged space is

$$\mathcal{P} := \bigcup_{\pi \in \mathcal{A}} \mathcal{P}(\pi),$$

where $\mathcal{A}$ denotes the class of all $\mathbb{G}$ -progressively measurable admissible controls. The problem of the portfolio manager is then given by

$$V_0^A(x, \xi) = \sup_{(\pi, \mathbb{P}^\pi) \in \mathcal{A} \times \mathcal{P}} J^A(\pi; x, \xi), \quad (\text{A})$$

where

$$J^A(\pi; x, \xi) := \mathbb{E}^{\mathbb{P}^\pi}[U^A(\xi)].$$

Remark A.2.1. In order to apply the result in [67] we need to enforce the $\mathcal{P}$ -density hypothesis,

see [67, Hypothesis 2.4 ($\mathcal{P}_0$ -Density Hypothesis)]. It is in particular satisfied as soon as τ is independent of $\mathbb{F}$ with a distribution invariant under any probability in $\mathcal{P}$.

A.2.2 The portfolio manager's optimization: 2BSDE with erratic termination

We consider the dynamic version of the value function. Following the same lines as [53, Proof of Theorem 4.2] and assuming that the family of probabilities in $\mathcal{P}$ satisfies [136, 2.1.3. Conditioning and concatenation of probability measures and Assumption 2.1 (iii)-(iv)-(v)], we deduce that $V_0 = Y_0$ where (Y, Z, U, K) is a solution to the following second-order BSDE

$$Y_t = U^A(\xi) + \int_{t \wedge \tau}^{T \wedge \tau} F_s(Z_s^X, \hat{\sigma}_s^2) ds - \int_{t \wedge \tau}^{T \wedge \tau} Z_s^X dX_s - \int_{t \wedge \tau}^{T \wedge \tau} U_s dH_s + K_T - K_t, \quad (\text{A.2})$$

where for any $\Sigma > 0$

$$F_s(z, \Sigma) := \sup_{\nu \in B_s(\Sigma)} \{z_x \nu \sigma_s \theta_s\}, \quad B_s(\Sigma) := \{\nu \in C, \|\nu \sigma_s\|^2 = \Sigma\},$$

and where $\hat{\sigma}$ is the quadratic variation of the process X under each $\mathbb{P} \in \mathcal{P}$ (see [92, 54]).

Proposition A.2.2 ([67]). *Assume that $U^A(\xi) \in \mathbb{L}^{p,k}$ for some $1 \leq k \leq p$ and σ, b are bounded.² Then there exists a unique solution (Y, Z^X, U, K) to the 2BSDE with erratic horizon (A.2).*

Following the same steps as [54, section 5.4], without loss of generality, there exists an $\mathbb{G}$ -predictable process Γ^X such that

$$dK_t = \left(H_t(Z_t^X, \Gamma_t^X) - F_t(Z_t^X, \hat{\sigma}_t^2) - \frac{1}{2} \hat{\sigma}_t^2 \Gamma_t^X \right) dt,$$

where

$$H_s(z, g) = \sup_{\Sigma} \left\{ \frac{1}{2} \Sigma g + F_s(z, \Sigma) \right\}, \quad \Sigma \in \{\|\nu \sigma_s\|^2, \nu \in C\}.$$

We thus get the decomposition

$$U^A(\xi) = Y_0 + \int_0^{T \wedge \tau} Z_s^X dX_s + \int_0^{T \wedge \tau} U_s dH_s - \int_0^{T \wedge \tau} \left(H_s(Z_s^X, \Gamma_s^X) - \frac{1}{2} \hat{\sigma}_s^2 \Gamma_s^X \right) ds.$$

²We refer to [67, Section 2.3] for the exact integrability condition on $U^A(\xi)$. Note that it is satisfied for exponential utilities as soon as ξ has exponential moments of any order greater than the risk aversion parameter for any probability $\mathbb{P} \in \mathcal{P}$.

Applying finally Ito's formula with the function $(U^A)^{-1}(x) = -\log(-x)/\eta$ we get the decomposition $\xi = Y_T^{Y_0, Z^X, U, \Gamma^X}$ from Definition 3.2.5.

A.3 Extension to Merton's portfolio problem with consumption

We propose to extend the study by assuming that the investor decides to consume part of the portfolio continuously over time with consumption strategy c_t given by an $\mathbb{F}$ -adapted random process. This problem was initially developed by Merton in [119] and this extension will follow the same type of problem extended to random termination. The value of the portfolio becomes

$$X_t := x + \int_0^t \pi_s \sigma_s dW_s + \int_0^t \pi_s b_s ds - \int_0^t c_s ds,$$

or equivalently

$$X_t := x + \int_0^t \beta_s dW_s + \int_0^t (\beta_s \theta_s - c_s) ds.$$

We denote by $\mathfrak{C}$ the set of admissible consumption strategies defined as the set of $\mathbb{F}$ -adapted processes c with values in $[0, \infty)$ such that

$$\mathbb{E}[\int_0^T c_t^2 dt] < \infty.$$

Definition 3.2.5 of admissible contract thus becomes the following:

Definition A.3.1 (Admissible contract with contractible variables and consumption). We denote by Ξ the set of admissible contracts ξ composed of $\mathcal{G}_{\tau \wedge T}$ -measurable random variable $\xi = Y_{T \wedge \tau}^{Y_0, Z, Z^X, U, \Gamma^X, \Gamma}$, controlled by $\mathbb{G}$ -predictable real-valued processes $U, Z = (Z^i)_{1 \leq i \leq m}, Z^X, \Gamma^X, \Gamma = (\Gamma_i)_{1 \leq i \leq m}$ such that $\mathcal{I}_m - \Gamma_t^X \sigma_t \sigma_t^T$ is a positive definite matrix³ and

$$\begin{aligned} Y_t^{Y_0, Z, Z^X, U, \Gamma^X, \Gamma} &= Y_0 + \int_0^t \sum_{i=1}^m \frac{Z_r^i}{S_r^i} dS_r^i + \int_0^t Z_r^X dX_r + \int_0^t U_r dH_r \\ &+ \frac{1}{2} \int_0^t (\Gamma_r^X + \eta(Z_r^X)^2) d\langle X, X \rangle_r + \int_0^t \sum_{i=1}^m \frac{\Gamma_r^i}{S_r^i} d\langle S^i, X \rangle_r \\ &- \int_0^t F(r, Z_r, Z_r^X, \Gamma_r, \Gamma_r^X, U_r, c_r) dr, \end{aligned}$$

³ $\mathcal{I}_m$ denotes the identity matrix in dimension m .

for any $c \in \mathfrak{C}$ where

$$F(t, z, z_x, g, g_x, u, c) = \sup_{\nu \in C} f(t, z, z_x, g, g_x, u, \nu, c),$$

with $f : [0, T] \times \Omega \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R} \times C \times [0, \infty) \rightarrow \mathbb{R}$ given by

$$\begin{aligned} f(t, z, z_x, g, g_x, u, \nu, c) = & z b_t + z_x(\nu \sigma_t \theta_t - c) + \frac{1}{2} g_x \|\nu \sigma_t\|^2 - \frac{1}{2} \|\nu - \alpha_t\|^2 + \sum_{i=1}^m g^i \nu^i \sigma_t^i (\sigma_t^i)^\top \\ & - \frac{\lambda_t}{\eta} (\exp(-\eta u) - 1) - \frac{\eta}{2} \|z \sigma_t\|^2, \end{aligned}$$

and there exists $\eta' > \eta$ such that

$$\mathbb{E} \left[\int_0^T (\|Z_s\|^2 + \|Z_s^X\|^2 + \|\Gamma_s\| + \|\Gamma_s^X\| + \|U_s\|^2 \lambda_s) ds + \sup_{0 \leq t \leq T} e^{\eta' |Y_t^{Y_0, Z, Z^X, U, \Gamma, \Gamma^X}|} \right] < \infty.$$

We are assuming that the portfolio manager is receiving the contract ξ and optimally chooses a strategy π in order to stay close to a benchmark strategy α *à la* Almgren-Chriss, see [5] so that the objective of the manager is to solve for any contract $\xi \in \Xi$ and consumption process $c \in \mathfrak{C}$ fixed by the investor

$$V_0^A(x, \xi, c) = \sup_{\pi \in \mathcal{A}} J^A(\pi; x, \xi, c), \quad (\text{A})$$

where

$$J^A(\pi; x, \xi, c) := \mathbb{E} \left[U^A \left(\xi - \frac{\varepsilon}{2} \int_0^{T \wedge \tau} \|\pi_s - \alpha_s\|^2 ds \right) \right].$$

The bilevel optimization corresponding to this portfolio problem with consumption becomes

$$V_0^P(x) := \sup_{(\xi, \hat{\pi}, c) \in \Xi \times \mathcal{A} \times \mathfrak{C}} \mathbb{E} [U^P(X_{T \wedge \tau} - \xi) + \int_0^{T \wedge \tau} u_P(c_t) dt], \quad (\text{P})$$

where $u_P : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a utility function, non-convex and non-decreasing in the consumption control variable,

subject to

- (R): $V_0^A(x, \xi, c) \geq R_0$
- (IC): $V_0^A(x, \xi, c) = J^A(\hat{\pi}; x, \xi, c)$.

Note that Theorem 3.3.2 is unchanged since the only impact of c for the agent (portfolio manager) will be given through the compensation Z^X . This is consistent with Merton's

approach where the optimal portfolio strategy $\hat{\pi}$ and consumption can be optimized independently in a first-best case model. We thus define the following system of coupled SDEs with jumps with solution $(X^{\pi^*,c}, \hat{Y})$ controlled by $(Z, Z^X, \Gamma, \Gamma^X, U, c)$

$$(SDE) \left\{ \begin{array}{l} dX_t^{\pi^*,c} = \pi_t^* \sigma(t, S_t) dW_t + (\pi_t^* b(t, S_t) - c_t) dt \\ d\hat{Y}_t = \sum_{i=1}^m \frac{Z_t^i}{S_t^i} dS_t^i + Z_t^X dX_t^{\pi^*,c} + U_t dH_t + \frac{1}{2}(\Gamma_t^X + \eta(Z_t^X)^2) d\langle X^{\pi^*,c}, X^{\pi^*,c} \rangle_t \\ \quad - F(t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t, c_t) dt + \sum_{i=1}^m \frac{\Gamma_t^i}{S_t^i} d\langle S^i, X^{\pi^*,c} \rangle_t \\ X_0^{\pi^*,c} = x, \\ \hat{Y}_0 = 0. \end{array} \right.$$

The problem of the investor becomes

$$V_0^P(x) = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U, c) \in \mathcal{U} \times \mathfrak{C}} \mathbb{E}[U^P(X_{T \wedge \tau} - \hat{Y}_{T \wedge \tau}) + \int_0^{T \wedge \tau} u_P(c_t) dt].$$

Note that this class of problems does not exactly fit the framework [25]. It requires to introduce a new state variable

$$I_t := \int_0^t u_P(c_s) ds \iff dI_t = u_P(c_t) dt, \quad I_0 = 0.$$

The problem of the investor thus becomes

$$V_0^P(x) = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U, c) \in \mathcal{U} \times \mathfrak{C}} \mathbb{E}[U^P(X_{T \wedge \tau} - \hat{Y}_{T \wedge \tau}) + I_{T \wedge \tau}].$$

- For bounded random time τ with support in $[0, T]$:

$$\hat{V}_0 = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U, c) \in \mathcal{U} \times \mathfrak{C}} \mathbb{E}\left[\int_0^T \left(U^P(X_t^{\pi^*,c} - \hat{Y}_t) + I_t\right) f(t) dt\right],$$

with corresponding HJB PDE

$$\left\{ \begin{array}{l} 0 = \partial_t v(t, s, x, y, \iota) + \left(U^P(x - y) + \iota \right) f(t) \\ \quad + \sup_{(z^x, z, g^x, g, u, c)} \mathcal{H}^v(t, s, x, y, \iota, \nabla v, \Delta v, z^x, z, g^x, g, u, c), \quad t < T, \\ v(T, s, x, y, \iota) = 0, \quad (s, x, y, \iota) \in \mathbb{R}^m \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}. \end{array} \right. \quad (\text{A.3})$$

where $\mathcal{H}^\phi$ is a differential operator given by

$$\begin{aligned}
\mathcal{H}^\phi(t, s, x, y, \iota, z^x, z, g^x, g, u, c) := & \\
& \sum_{i=1}^m \phi_{s^i} s^i b^i(t, s) + \phi_x \pi^*(z, z^x, g, g^x) \sigma(t, s) \theta(t, s) + \mathcal{L}^\phi(t, s, x, y, \iota, c) \\
& + \frac{1}{2} \text{Tr}(s \sigma(t, s) \sigma^T(t, s) s D^2 \phi) + \frac{1}{2} \phi_{xx} \|\pi^*(z, z^x, g, g^x) \sigma(t, s)\|^2 \\
& + \phi_y \left[z b(t, s) + z^x \pi^*(z, z^x, g, g^x) \sigma(t, s) \theta(t, s) + \frac{1}{2} \|\sigma(t, s) \pi^*(z, z^x, g, g^x)\|^2 (g^x + \eta |z^x|^2) \right. \\
& \left. - F(t, z, z^x, g, g^x, u, c) + \pi^*(z, z^x, g, g^x) \sigma(t, s) \sigma(t, s)^T g \right] \\
& + \frac{1}{2} \phi_{yy} (\|z_t \sigma(t, s)\|^2 + (z^x)^2 \|\sigma(t, s) \pi^*(z, z^x, g, g^x)\|^2) \\
& + \phi_{xy} (z^x \|\pi^*(z, z^x, g, g^x) \sigma(t, s)\|^2 + \pi^*(z, z^x, g, g^x) \sigma(t, s) (z \sigma)^T) \\
& + \sum_{i=1}^m (\pi^*(z, z^x, g, g^x) \sigma(t, s) \sigma^i(t, s)^T s^i) (\phi_{x s^i} + \phi_{y s^i} z^x) + \lambda_t (\phi(t, s, x, y + u) - \phi(t, s, x, y)) \\
& + \sum_{i=1}^m \phi_{y s^i} z^i \sigma(t, s) \sigma^i(t, s)^T s^i,
\end{aligned}$$

and with

$$\mathcal{L}^\phi(t, s, x, y, \iota, c) := \phi_\iota u_P(c) - \phi_x c.$$

Theorem A.3.2 (Verification Theorem - bounded case). *Assume that there exists a function ϕ twice continuously differentiable in space and differentiable in time, such that $\phi(t, s, x, y, \iota)$ solves (A.3). Furthermore, assume that ϕ has a quadratic growth in y and polynomial growth in s, x, ι such that*

$$|\phi(t, s, x, y, \iota)| \leq \kappa(1 + |x|^p + \|s\|^p + |\iota|^p + |y|^2), \quad p > 1, \quad \kappa > 0.$$

Then, for each $t \in [0, T]$, $(\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U}, \hat{c})$ is an optimal control for the investor's bilevel problem, and

$$\phi(0, S_0, x, 0, 0) = V_0 = \sup_{(Z^X, Z, \Gamma^X, \Gamma, U, c)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^{\pi^*, c} - \hat{Y}_{T \wedge \tau}^{\pi^*, c}) + I_{T \wedge \tau} \right].$$

Let $c^*(t, s, x, y, \iota)$ be an optimizer of $\mathcal{L}^\phi(t, s, x, y, \iota, c)$ such that (SDE) admits a unique

solution with this choice of c . The optimal contract is given by

$$\begin{aligned}\xi^* &= \hat{Y}_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{Z}_t^i}{\hat{S}_t^i} dS_t^i + \int_0^{T \wedge \tau} \hat{Z}_t^X dX_t + \int_0^{T \wedge \tau} \hat{U}_s dH_s \\ &+ \frac{1}{2} \int_0^{T \wedge \tau} (\hat{\Gamma}_t^X + \eta(\hat{Z}_t^X)^2) d\langle X, X \rangle_t - \int_0^{T \wedge \tau} F(t, \hat{Z}_t, \hat{Z}_t^X, \hat{\Gamma}_t, \hat{\Gamma}_t^X, \pi_t^*, \hat{c}_t) dt \\ &+ \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{\Gamma}_t^i}{\hat{S}_t^i} d\langle S^i, X \rangle_t,\end{aligned}$$

and the optimal consumption is given by $\hat{c}_t := c^*(t, S_t, X_t^{\pi^*, \hat{c}}, \hat{Y}_t, I_t)$.

- For unbounded random time with $[0, T] \subsetneq \text{supp}(\tau)$

$$\hat{V}_0 = \sup_{(Z, Z^X, \Gamma, \Gamma^X, U, c) \in \mathcal{U} \times \mathfrak{C}} \mathbb{E} \left[\int_0^T \left(U^P(X_t^{\pi^*, c} - \hat{Y}_t) + I_t \right) f(t) dt + (1 - F_\tau(T)) \left(U^P(X_T - \hat{Y}_T) + I_T \right) \right],$$

with corresponding HJB PDE is

$$\begin{cases} 0 = \partial_t v(t, s, x, y, \iota) + \left(U^P(x - y) + \iota \right) f(t) \\ \quad + \sup_{(z^x, z, g^x, g, u, c)} \mathcal{H}^v(t, s, x, y, \iota, \nabla v, \Delta v, z^x, z, g^x, g, u, c), \quad t < T, \\ v(T, s, x, y, \iota) = (1 - F_\tau(T)) \left(U^P(x - y) + \iota \right), \quad (s, x, y, \iota) \in \mathbb{R}^m \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}. \end{cases} \quad (\text{A.4})$$

with the same $\mathcal{H}$ operator as in the bounded case.

Theorem A.3.3 (Verification Theorem - unbounded case). *Assume that there exists a function ϕ twice continuously differentiable in space and differentiable in time, such that $\phi(t, s, x, y, \iota)$ solves (A.4). Furthermore, assume that ϕ has a quadratic growth in y and polynomial growth in s, x, ι such that*

$$|\phi(t, s, x, y, \iota)| \leq \kappa(1 + |x|^p + \|s\|^p + |\iota|^p + |y|^2), \quad p > 1, \quad \kappa > 0.$$

Then, for each $t \in [0, T]$, $(\hat{Z}^X, \hat{Z}, \hat{\Gamma}^X, \hat{\Gamma}, \hat{U}, \hat{c})$ is an optimal control for the investor's bilevel problem, and

$$\phi(0, S_0, x, 0, 0) = V_0 = \sup_{(Z^X, Z, \Gamma^X, \Gamma, U, c)} \mathbb{E} \left[U^P(X_{T \wedge \tau}^{\pi^*, c} - \hat{Y}_{T \wedge \tau}^{\pi^*, c}) + I_{T \wedge \tau} \right].$$

Let $c^*(t, s, x, y, \iota)$ be an optimizer of $\mathcal{L}^\phi(t, s, x, y, \iota, c)$ such that (SDE) admits a unique

solution with this choice of c . The optimal contract is given by

$$\begin{aligned}\xi^* &= \hat{Y}_0 + \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{Z}_t^i}{\hat{S}_t^i} dS_t^i + \int_0^{T \wedge \tau} \hat{Z}_t^X dX_t + \int_0^{T \wedge \tau} \hat{U}_s dH_s \\ &+ \frac{1}{2} \int_0^{T \wedge \tau} (\hat{\Gamma}_t^X + \eta(\hat{Z}_t^X)^2) d\langle X, X \rangle_t - \int_0^{T \wedge \tau} F(t, \hat{Z}_t, \hat{Z}_t^X, \hat{\Gamma}_t, \hat{\Gamma}_t^X, \pi_t^*, \hat{c}_t) dt \\ &+ \int_0^{T \wedge \tau} \sum_{i=1}^m \frac{\hat{\Gamma}_t^i}{\hat{S}_t^i} d\langle S^i, X \rangle_t,\end{aligned}$$

and the optimal consumption is given by $\hat{c}_t := c^*(t, S_t, X_t^{\pi^*, \hat{c}}, \hat{Y}_t, I_t)$.

Remark A.3.4. Note that, if we instead assume that the investor uses singular controls for the consumption strategy c , modeled as a non-decreasing process with bounded variation (see, for example, [148]), this would lead to

$$X_t := x + \int_0^t \beta_s dW_s + \int_0^t \beta_s \theta_s ds - \int_0^t dc_s.$$

This case involves several difficulties, including

1. second-order BSDE with jumps driven by general martingales to solve the portfolio manager problem.
2. an obstacle PDE with gradient constraint for the investor's optimization with random horizon.

The first point is critical and requires to link the manager problem to a very new type of 2BSDE. Recent advances in this direction have been made in [134]. However, the link with stochastic control problem with singular control remains an open problem beyond the scope of our paper that we will leave for future research.

A.4 The agent optimal strategy with power utility

Let us now consider, under the same setting of a Geometric Brownian Motion for the risky assets, the same portfolio evolution and therefore the same trading strategy π as in Section 3.3. The agent's problem is the same, but for the sake of simplicity let us assume that there is no running penalty, so that the agent problem becomes

$$V_0^A(x, \xi) = \sup_{\pi \in \mathcal{A}} J^A(\pi; x, \xi), \quad (\text{Power-A})$$

where

$$J^A(\pi; x, \xi) := \mathbb{E}[U^A(\xi)].$$

with

$$U^A(x) = \frac{1}{1-\eta} x^{1-\eta}, \quad \eta \in [0, \infty) \setminus \{1\}.$$

For $\eta > 1$ the utility is negative and the supermartingale argument below carries over with a sign flip; the verification still proceeds unchanged. We then have the following result.

Theorem A.4.1. *Assume that 3.2.3 and Hypothesis A or Hypothesis B are satisfied. For any $\xi \in \Xi$, assume in addition that $\Gamma_t^X \sigma_t \sigma_t^\top$ is $dt \otimes \mathbb{P}$ -a.e. positive definite (in particular $\Gamma_t^X \geq 0$), and that there exists a $(C, Z, Z^X, \Gamma, \Gamma^X, U)$ solution to*

$$dC_t = C_t - dY_t$$

with $C_{T \wedge \tau} = \xi$ and Y a semi-martingale with the following decomposition

$$\begin{aligned} dY_t &= Z_t(b_t dt + \sigma_t dW_t) + Z_t^X \pi_t(\sigma_t dW_t + b_t dt) \\ &\quad - \frac{1}{2}(\Gamma_t^X - \eta(Z_t^X)^2) \pi_t \sigma_t \sigma_t^\top \pi_t^\top dt + \pi_t \sigma_t \sigma_t^\top \Gamma_t dt \\ &\quad - F(t, Y_t, Z_t, Z_t^X, \Gamma_t, \Gamma_t^X, U_t) dt + U_t dH_t \end{aligned} \tag{A.5}$$

Then, the optimal strategy solving (Power-A) is

$$\hat{\pi}_t = \pi^*(Z_t, Z_t^X, \Gamma_t, \Gamma_t^X), \quad \text{with } \pi^*(z, z_x, g, g_x) := \text{proj}(e_t, C) \tag{A.6}$$

and the optimal value is given by $V_0^A(x, \xi) = \frac{1}{1-\eta} C_0^{1-\eta}$, where

$$q_t = \frac{1}{\Gamma_t^X} (Z_t^X (\sigma_t \sigma_t^\top)^{-1} b_t + \Gamma_t)$$

and

$$Q_t = \Gamma_t^X \sigma_t \sigma_t^\top$$

Proof. The proof follows a similar idea to the one in the main body of the paper: we want to find a process related to our optimization problem and ensure that the driver of a related BSDE is such that the process is a super-martingale for each strategy and a martingale for the optimal strategy. In this case, in order to achieve this, we are further assuming that the auxiliary process solves the auxiliary exponential BSDE. As before, we define the family of

processes indexed by the strategy π

$$R_t^\pi = U^A(C_t).$$

This time we set $C = (C_t)_{t \in [0, T]}$ as the solution of the following simple BSDE, $dC_t = C_{t-} dY_t$, coupled with $C_{T \wedge \tau} = \xi$, where the semi-martingale Y follows the dynamics:

$$\begin{aligned} Y_t^{Y_0, Z, Z^X, U, \Gamma^X, \Gamma} &= Y_0 + \int_0^t \sum_{i=1}^m \frac{Z_r^i}{S_r^i} dS_r^i + \int_0^t Z_r^X dX_r + \int_0^t U_r dH_r \\ &\quad - \frac{1}{2} \int_0^t (\Gamma_r^X - \eta(Z_r^X)^2) d\langle X, X \rangle_r + \int_0^t \sum_{i=1}^m \frac{\Gamma_r^i}{S_r^i} d\langle S^i, X \rangle_r \\ &\quad - \int_0^t F(r, Z_r, Z_r^X, \Gamma_r, \Gamma_r^X, U_r) dr, \end{aligned}$$

or equivalently the infinitesimal decomposition in (A.5). Therefore,

$$R_t^\pi = \psi(C_t)$$

where $\psi(x) = \frac{1}{1-\eta} x^{1-\eta}$. Using a simple Ito's expansion, we get that

$$\begin{aligned} dR_t &= \partial_t \psi dt + \partial_x \psi dC_t^c + \frac{1}{2} \partial_{xx} \psi d\langle C_t^c, C_t^c \rangle + (\psi(C_t) - \psi(C_{t-})) dH_t \\ &= (1-\eta) \frac{1}{1-\eta} C_t^{-\eta} C_t dY_t - \frac{1}{2} (1-\eta) \eta \frac{1}{1-\eta} C_t^{-\eta-1} C_t^2 d\langle Y_t^c, Y_t^c \rangle + (\psi(C_t) - \psi(C_{t-})) dH_t \end{aligned}$$

In this way, we observe that we can write the first two terms as R_t times some other quantities. At the same time, since $C_t = C_{t-}(1 + U_t)$ when H jumps,

$$\psi(C_t) - \psi(C_{t-}) = \frac{1}{1-\eta} (C_{t-}(1 + U_t))^{1-\eta} - \frac{1}{1-\eta} C_{t-}^{1-\eta} = \frac{1}{1-\eta} C_{t-}^{1-\eta} ((1 + U_t)^{1-\eta} - 1).$$

Therefore, we can write the increment of R at time t as

$$dR_t = (1-\eta) R_{t-} dY_t^c - \frac{1}{2} (1-\eta) \eta R_{t-} d\langle Y_t^c, Y_t^c \rangle_t + R_{t-} ((1 + U_t)^{1-\eta} - 1) dH_t.$$

We now expand dY_t , whose continuous part satisfies:

$$\begin{aligned} dY_t^c &= Z_t b_t dt + Z_t \theta_t dW_t + Z_t^X \pi_t \sigma_t dW_t + Z_t^X \pi_t \sigma_t \theta_t dt + \\ &\quad - \frac{1}{2} (\Gamma_t^X - \eta(Z_t^X)^2) \pi_t \sigma_t \sigma_t^\top \pi_t^\top dt - F_t dt + \pi_t \sigma_t \sigma_t^\top \Gamma_t dt. \end{aligned}$$

Its quadratic variation is

$$d\langle Y^c, Y^c \rangle_t = Z_t \sigma_t \sigma_t^\top Z_t^\top dt + (Z_t^x)^2 \pi_t \sigma_t \sigma_t^\top \pi_t^\top dt.$$

Substituting into the expression for dR_t , we get that

$$\begin{aligned} dR_t = & R_t (1 - \eta) \left(Z_t b_t + Z_t^X \pi_t \sigma_t \theta_t - \frac{1}{2} \Gamma_t^X \pi_t \sigma_t \sigma_t^\top \pi_t^\top - F_t + \pi_t \sigma_t \sigma_t^\top \Gamma_t - \frac{\eta}{2} Z_t \sigma_t \sigma_t^\top Z_t^\top \right) dt \\ & + R_t (1 - \eta) \left(Z_t \sigma_t + Z_t^X \pi_t \sigma_t \right) dW_t + R_t \left((1 + U_t)^{1-\eta} - 1 \right) dH_t. \end{aligned}$$

Now, because the terminal condition on C is the reward for the agent, and the agent has a reserved utility, we know that the DDE of Y is positive and therefore $C_0 \geq 0$. Now, the utility is a power function, we have $R_0 \geq 0$ and

$$dR_t = (1 - \eta) R_{t-} a_t dt + R_{t-} d\tilde{M}_t$$

where $\tilde{M}$ is a $\mathbb{G}$ -martingale and

$$a_t = Z_t b_t + Z_t^X \pi_t \sigma_t \theta_t + \frac{1}{2} \Gamma_t^X \pi_t \sigma_t \sigma_t^\top \pi_t^\top - F_t + \pi_t \sigma_t \sigma_t^\top \Gamma_t - \frac{\eta}{2} Z_t \sigma_t \sigma_t^\top Z_t^\top + \lambda_t \left((1 + U_t)^{1-\eta} - 1 \right).$$

Now, because of the positivity of the utility, we have that, in order for R to be a super-martingale, we need to ensure that $a_t \leq 0$ and optimality is achieved when the process is equal to 0. Completing again the square, we have that

$$a_t = -\frac{1}{2} (\pi_t - q_t) Q_t (\pi_t - q_t)^\top$$

where $Q_t = \Gamma_t^X \sigma_t \sigma_t^\top$ and

$$q_t = \frac{1}{\Gamma_t^X} (Z_t^X (\sigma_t \sigma_t^\top)^{-1} b_t + \Gamma_t)$$

which is possible when F has the following form:

$$F = \frac{1}{2} \left(q_t - \frac{\eta}{2\Gamma_t^X} Z_t^\top \right)^\top Q_t \left(q_t - \frac{\eta}{2\Gamma_t^X} Z_t^\top \right) + Z_t b_t + \frac{\lambda_t}{1-\eta} \left[(1 + U_t)^{1-\eta} - 1 \right] + \text{dist}_{Q_t}^2(d_t, C),$$

where $d_t := q_t - \frac{\eta}{2\Gamma_t^X} Z_t^\top$ is the un-projected optimum (with $e_t := \text{proj}_C(d_t)$ as in Chapter 3, line of definition of e_t). $\square$

Bibliography

- [1] Beatrice Acciaio, Claudio Fontana, and Constantinos Kardaras. “Arbitrage of the first kind and filtration enlargements in semimartingale financial models”. In: *Stochastic Processes and their Applications* 126.6 (2016), pp. 1761–1784.
- [2] Anna Aksamit and Monique Jeanblanc. “Enlargement of filtration with finance in view”. In: (2017).
- [3] Anna Aksamit and Libo Li. “Projections, pseudo-stopping times and the immersion property”. In: *Séminaire de Probabilités XLVIII*. Springer, 2016, pp. 459–467.
- [4] Dora Almeida, Andreia Dionísio, Isabel Vieira, and Paulo Ferreira. “Uncertainty and risk in the cryptocurrency market”. In: *Journal of Risk and Financial Management* 15.11 (2022), p. 532.
- [5] Robert Almgren and Neil Chriss. “Optimal execution of portfolio transactions”. In: *Journal of Risk* 3 (2001), pp. 5–40.
- [6] Olivier Alvarez and Agnes Tourin. “Viscosity solutions of nonlinear integro-differential equations”. In: *Annales de l’Institut Henri Poincaré C, Analyse non linéaire*. Vol. 13. 3. Elsevier. 1996, pp. 293–317.
- [7] Stefan Ankirchner. “On filtration enlargements and purely discontinuous martingales”. In: *Stochastic Processes and their Applications* 118.9 (2008), pp. 1662–1678.
- [8] Joshua Aurand and Yu-Jui Huang. “Epstein-Zin utility maximization on a random horizon”. In: *Mathematical Finance* 33.4 (2023), pp. 1370–1411.
- [9] Marco Avellaneda and Sasha Stoikov. “High-frequency trading in a limit order book”. In: *Quantitative Finance* 8.3 (2008), pp. 217–224.
- [10] Khaled Bahlali. “Backward stochastic differential equations with locally Lipschitz coefficient”. In: *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics* 333.5 (2001), pp. 481–486.
- [11] Bastien Baldacci, Philippe Bergault, and Dylan Possamai. “A mean-field game of market-making against strategic traders”. In: *SIAM Journal on Financial Mathematics* 14.4 (2023), pp. 1080–1112.
- [12] Bastien Baldacci, Paul Jusselin, and Mathieu Rosenbaum. “How to design a derivatives market?” In: *Peter Carr Gedenkschrift: Research Advances in Mathematical Finance*. World Scientific, 2024, pp. 657–699.

- [13] Bastien Baldacci, Iuliia Manziuk, Thibaut Mastrolia, and Mathieu Rosenbaum. “Market making and incentives design in the presence of a dark pool: a deep reinforcement learning approach”. In: *arXiv preprint arXiv:1912.01129* (2019).
- [14] Bastien Baldacci and Dylan Possamai. “Governmental incentives for green bonds investment”. In: *Mathematics and Financial Economics* 16.3 (2022), pp. 539–585.
- [15] Bastien Baldacci, Dylan Possamai, and Mathieu Rosenbaum. “Optimal make-take fees in a multi market-maker environment”. In: *SIAM Journal on Financial Mathematics* 12.1 (2021), pp. 446–486.
- [16] P. Bank, C. Bayer, P. P. Hager, S. Riedel, and T. Nauen. “Stochastic Control with Signatures”. In: *Preprint, arXiv:2406.01585* (2024).
- [17] Guy Barles, Rainer Buckdahn, and Etienne Pardoux. “Backward stochastic differential equations and integral-partial differential equations”. In: *Stochastics: An International Journal of Probability and Stochastic Processes* 60.1-2 (1997), pp. 57–83.
- [18] Pauline Barrieu and Nicole El Karoui. “Pricing, hedging, and designing derivatives with risk measures”. In: *Indifference pricing: Theory and applications* (2009), pp. 77–146.
- [19] Emilio Barucci, Adrien Mathieu, and Leandro Sánchez-Betancourt. “Market Making with Fads, Informed, and Uninformed Traders”. In: *arXiv preprint arXiv:2501.03658* (2025).
- [20] Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. “Automatic differentiation in machine learning: a survey”. In: *Journal of machine learning research* 18.153 (2018), pp. 1–43.
- [21] Peter L Bernstein and Aswath Damodaran. *Investment management*. J. Wiley, 1998.
- [22] Tomasz R Bielecki, Tao Chen, Igor Cialenco, Areski Cousin, and Monique Jeanblanc. “Adaptive robust control under model uncertainty”. In: *SIAM Journal on Control and Optimization* 57.2 (2019), pp. 925–946.
- [23] Tomasz R Bielecki and Marek Rutkowski. *Credit risk: modeling, valuation and hedging*. Springer Science & Business Media, 2013.
- [24] Jean-Michel Bismut. “Conjugate convex functions in optimal stochastic control”. In: *Journal of mathematical analysis and applications* 44.2 (1973), pp. 384–404.
- [25] Christopette Blanchet-Scalliet, Nicole El Karoui, Monique Jeanblanc, and Lionel Martellini. “Optimal investment decisions when time-horizon is uncertain”. In: *Journal of Mathematical Economics* 44.11 (2008), pp. 1100–1113.

- [26] Christopette Blanchet-Scalliet, Caroline Hillairet, and Ying Jiao. “Successive enlargement of filtrations and application to insider information”. In: *Advances in Applied Probability* 49.3 (2017), pp. 653–685.
- [27] Giuliana Bordigoni, Anis Matoussi, and Martin Schweizer. “A stochastic control approach to a robust utility maximization problem”. In: *Stochastic Analysis and Applications: The Abel Symposium 2005*. Springer. 2007, pp. 125–151.
- [28] Pierre Brémaud and Marc Yor. “Changes of filtrations and of probability measures”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 45.4 (1978), pp. 269–295.
- [29] Philippe Briand and Fulvia Confortola. “BSDEs with stochastic Lipschitz condition and quadratic PDEs in Hilbert spaces”. In: *Stochastic Processes and their Applications* 118.5 (2008), pp. 818–838.
- [30] Philippe Briand and Ying Hu. “Quadratic BSDEs with convex generators and unbounded terminal conditions”. In: *Probability Theory and Related Fields* 141 (2008), pp. 543–567.
- [31] Antonio Briola, David Vidal-Tomás, Yuanrong Wang, and Tomaso Aste. “Anatomy of a Stablecoin’s failure: The Terra-Luna case”. In: *Finance Research Letters* 51 (2023), p. 103358.
- [32] Rainer Buckdahn, Pierre Cardaliaguet, and Marc Quincampoix. “Some recent aspects of differential game theory”. In: *Dynamic Games and Applications* 1 (2011), pp. 74–114.
- [33] Rainer Buckdahn and Etienne Pardoux. “BSDEs with jumps and associated integro-partial differential equations”. In: *preprint* 79 (1994).
- [34] Abel Cadenillas, Jakša Cvitanić, and Fernando Zapatero. “Optimal risk-sharing with effort and project choice”. In: *Journal of Economic Theory* 133.1 (2007), pp. 403–440.
- [35] Giorgia Callegaro, Monique Jeanblanc, and Behnaz Zargari. “Carthaginian enlargement of filtrations”. In: *ESAIM: Probability and Statistics* 17 (2013), pp. 550–566.
- [36] René Carmona. *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*. SIAM, 2016.
- [37] Rene Carmona and Kevin Webster. “High frequency market making”. In: *arXiv preprint arXiv:1210.5781* (2012).

- [38] Álvaro Cartea, Imanol Pérez Arribas, and Leandro Sánchez-Betancourt. “Double-Execution Strategies Using Path Signatures”. In: *SIAM Journal on Financial Mathematics* 13.4 (2022), pp. 1379–1417.
- [39] Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and high-frequency trading*. Cambridge University Press, 2015.
- [40] Claudia Ceci, Katia Colaneri, and Alessandra Cretarola. “Indifference pricing of pure endowments via BSDEs under partial information”. In: *Scandinavian Actuarial Journal* 2020.10 (2020), pp. 904–933.
- [41] Claudia Ceci, Alessandra Cretarola, and Francesco Russo. “BSDEs under partial information and financial applications”. In: *Stochastic Processes and their Applications* 124.8 (2014), pp. 2628–2653.
- [42] Jonathan Chávez-Casillas, José E Figueroa-López, Chuyi Yu, and Yi Zhang. “Adaptive optimal market making strategies with inventory liquidation cost”. In: *SIAM Journal on Financial Mathematics* 15.3 (2024), pp. 653–699.
- [43] Hamid Cheraghali, Peter Molnár, Mattis Storsveen, and Florent Veliqui. “The impact of cryptocurrency-related cyberattacks on return, volatility, and trading volume of cryptocurrencies and traditional financial assets”. In: *International Review of Financial Analysis* 95 (2024), p. 103439.
- [44] Patrick Cheridito, H Mete Soner, Nizar Touzi, and Nicolas Victoir. “Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs”. In: *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* 60.7 (2007), pp. 1081–1110.
- [45] Alessandro Chiusolo and Emma Hubert. “A new approach to principal-agent problems with volatility control”. In: *arXiv preprint arXiv:2407.09471* (2024).
- [46] Michael G Crandall, Hitoshi Ishii, and Pierre-Louis Lions. “User’s guide to viscosity solutions of second order partial differential equations”. In: *Bulletin of the American mathematical society* 27.1 (1992), pp. 1–67.
- [47] Michael G Crandall and Pierre-Louis Lions. “Viscosity solutions of Hamilton-Jacobi equations”. In: *Transactions of the American mathematical society* 277.1 (1983), pp. 1–42.
- [48] C. Cuchiero, G. Gazzani, J. Möller, and S. Svaluto-Ferro. “Joint calibration to SPX and VIX options with signature-based models”. In: *Mathematical Finance* (2024), pp. 1–53.

- [49] Christa Cuchiero, Guido Gazzani, and Sara Svaluto-Ferro. “Signature-Based Models: Theory and Calibration”. In: *SIAM Journal on Financial Mathematics* 14.3 (2023), pp. 910–957.
- [50] Christa Cuchiero and Janka Möller. *Signature Methods in Stochastic Portfolio Theory*. 2024. arXiv: 2310.02322 [q-fin.MF]. URL: <https://arxiv.org/abs/2310.02322>.
- [51] Christa Cuchiero, Francesca Primavera, and Sara Svaluto-Ferro. “Universal approximation theorems for continuous functions of càdlàg paths and Lévy-type signature models”. In: *Forthcoming in Finance and Stochastics* (2024).
- [52] Jakša Cvitanić and Ioannis Karatzas. “Backward stochastic differential equations with reflection and Dynkin games”. In: *The Annals of Probability* (1996), pp. 2024–2056.
- [53] Jakša Cvitanić, Dylan Possamai, and Nizar Touzi. “Moral hazard in dynamic risk management”. In: *Management Science* 63.10 (2017), pp. 3328–3346.
- [54] Jakša Cvitanić, Dylan Possamai, and Nizar Touzi. “Dynamic programming approach to principal–agent problems”. In: *Finance and Stochastics* 22 (2018), pp. 1–37.
- [55] Flávia Zóboli Dalmácio, Valcemiro Nossa, et al. “The agency theory applied to the investment funds”. In: *Brazilian Business Review* 1.1 (2004), pp. 31–44.
- [56] Łukasz Delong. *Backward stochastic differential equations with jumps and their actuarial and financial applications*. Springer, 2013.
- [57] Łukasz Delong. “Optimal investment for insurance company with exponential utility and wealth-dependent risk aversion coefficient”. In: *Mathematical Methods of Operations Research* 89.1 (2019), pp. 73–113.
- [58] Peter M DeMarzo and Yuliy Sannikov. “Optimal security design and dynamic capital structure in a continuous-time agency model”. In: *The journal of Finance* 61.6 (2006), pp. 2681–2724.
- [59] Laurent Denis, Anis Matoussi, and Chao Zhou. “Second order BSDEs with jumps by measurable selection argument”. In: (2024).
- [60] Paolo Di Tella. “On the weak representation property in progressively enlarged filtrations with an application in exponential utility maximization”. In: *Stochastic Processes and their Applications* 130.2 (2020), pp. 760–784.
- [61] Gonçalo Dos Reis. “On some properties of solutions of quadratic growth BSDE and applications in finance and insurance”. PhD thesis. Verlag nicht ermittelbar, 2010.
- [62] Sofiene El Aoud and Frédéric Abergel. “A stochastic control approach to option market making”. In: *Market microstructure and liquidity* 1.01 (2015), p. 1550006.

- [63] Nicole El Karoui, Monique Jeanblanc, and Ying Jiao. “What happens after a default: the conditional density approach”. In: *Stochastic processes and their applications* 120.7 (2010), pp. 1011–1032.
- [64] Nicole El Karoui, Shige Peng, and Marie Claire Quenez. “Backward stochastic differential equations in finance”. In: *Mathematical finance* 7.1 (1997), pp. 1–71.
- [65] Wendell H Fleming and Halil Mete Soner. *Controlled Markov processes and viscosity solutions*. Vol. 25. Springer Science & Business Media, 2006.
- [66] Owen Futter, Blanka Horvath, and Magnus Wiese. “Signature Trading: A Path-Dependent Extension of the Mean-Variance Framework with Exogenous Signals”. In: *arXiv/2308.15135* (2023).
- [67] Alberto Gennaro and Thibaut Mastrolia. “2BSDE with uncertain horizon and application to stochastic control in erratic environments”. In: *arXiv preprint arXiv:2506.15037* (2025).
- [68] Klaus Grobys and Niranjana Sapkota. “Predicting cryptocurrency defaults”. In: *Applied Economics* 52.46 (2020), pp. 5060–5076.
- [69] Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. “Dealing with the inventory risk: a solution to the market making problem”. In: *Mathematics and Financial Economics* 7.4 (Sept. 2012), pp. 477–507. ISSN: 1862-9660. DOI: 10.1007/s11579-012-0087-0. URL: <http://dx.doi.org/10.1007/s11579-012-0087-0>.
- [70] Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. “Dealing with the inventory risk: a solution to the market making problem”. In: *Mathematics and financial economics* 7 (2013), pp. 477–507.
- [71] Xin Guo and Yan Zeng. “Intensity process and compensator: A new filtration expansion approach and the Jeulin–Yor theorem”. In: (2008).
- [72] Ben M. Hambly and Terry Lyons. “Uniqueness for the signature of a path of bounded variation and the reduced path group”. In: *Annals of Mathematics* 171 (2005), pp. 109–167.
- [73] Alan G. Hawkes. “Spectra of Some Self-Exciting and Mutually Exciting Point Processes”. In: *Biometrika* 58.1 (1971), pp. 83–90.
- [74] Nicolas Hernandez-Santibanez and Thibaut Mastrolia. “Contract theory in a VUCA world”. In: *SIAM Journal on Control and Optimization* 57.4 (2019), pp. 3072–3100.

- [75] Bengt Holmstrom and Paul Milgrom. “Aggregation and linearity in the provision of intertemporal incentives”. In: *Econometrica: Journal of the Econometric Society* (1987), pp. 303–328.
- [76] Ying Hu, Peter Imkeller, and Matthias Müller. “Utility maximization in incomplete markets”. In: (2005).
- [77] Zongyuan Huang, Haiyang Wang, and Zhen Wu. “A kind of optimal investment problem under inflation and uncertain time horizon”. In: *Applied Mathematics and Computation* 375 (2020), p. 125084.
- [78] O Hun, Mun-Chol Kim, and Kon-Gun Kim. “Wellposedness of second order reflected BSDEs: A new formulation”. In: *ESAIM: Probability and Statistics* 28 (2024), pp. 1–21.
- [79] Nobuyuki Ikeda and Shinzo Watanabe. *Stochastic differential equations and diffusion processes*. Elsevier, 2014.
- [80] Rufus Isaacs. *Differential games: a mathematical theory with applications to warfare and pursuit, control and optimization*. Courier Corporation, 1999.
- [81] Jean Jacod and Albert N. Shiryaev. *Limit Theorems for Stochastic Processes*. 2nd. Vol. 288. Grundlehren der mathematischen Wissenschaften. Springer, 2003.
- [82] Monique Jeanblanc and Yann Le Cam. “Progressive enlargement of filtrations with initial times”. In: *Stochastic Processes and their Applications* 119.8 (2009), pp. 2523–2543.
- [83] Monique Jeanblanc, Thibaut Mastrolia, Dylan Possamai, and Anthony Réveillac. “Utility maximization with random horizon: a BSDE approach”. In: *International Journal of Theoretical and Applied Finance* 18.07 (2015), p. 1550045.
- [84] Monique Jeanblanc and Anthony Réveillac. “A note on BSDEs with singular driver coefficients”. In: *Arbitrage, credit and informational risks*. World Scientific, 2014, pp. 207–224.
- [85] Monique Jeanblanc, Marc Yor, and Marc Chesney. *Mathematical methods for financial markets*. Springer Science & Business Media, 2009.
- [86] Michael C Jensen. “The performance of mutual funds in the period 1945-1964”. In: *The Journal of finance* 23.2 (1968), pp. 389–416.
- [87] Thierry Jeulin and Marc Yor. *Grossissements de filtrations: exemples et applications: Séminaire de Calcul Stochastique 1982/83 Université Paris VI*. Vol. 1118. Springer, 2006.

- [88] Ying Jiao and Idris Kharroubi. “Information uncertainty related to marked random times and optimal investment”. In: *Probability, Uncertainty and Quantitative Risk* 3 (2018), pp. 1–24.
- [89] Paul Jusselin and Thibaut Mastrolia. “Scaling limit for stochastic control problems in population dynamics”. In: *Applied Mathematics & Optimization* 88.1 (2023), p. 14.
- [90] Yuri Kabanov, Robert Liptser, Jordan Stoyanov, Dario Gasbarra, Esko Valkeila, and Lioudmila Vostrikova. “Enlargement of filtration and additional information in pricing models: Bayesian approach”. In: *From Stochastic Calculus to Mathematical Finance: The Shiryaev Festschrift* (2006), pp. 257–285.
- [91] Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. “Optimal Execution with Rough Path Signatures”. In: *SIAM Journal on Financial Mathematics* 11.2 (2020), pp. 470–493. DOI: 10.1137/19M1259778. eprint: <https://doi.org/10.1137/19M1259778>. URL: <https://doi.org/10.1137/19M1259778>.
- [92] Rajeeva L Karandikar. “On pathwise stochastic integration”. In: *Stochastic Processes and their applications* 57.1 (1995), pp. 11–18.
- [93] Ioannis Karatzas and Qinghua Li. “BSDE approach to non-zero-sum stochastic differential games of control and stopping”. In: *Stochastic Processes, Finance and Control: A Festschrift in Honor of Robert J Elliott*. World Scientific, 2012, pp. 105–153.
- [94] Nabil Kazi-Tani, Dylan Possamai, and Chao Zhou. “Second-order BSDEs with jumps: formulation and uniqueness”. In: (2015).
- [95] Idris Kharroubi and Thomas Lim. “Progressive enlargement of filtrations and backward stochastic differential equations with jumps”. In: *Journal of Theoretical Probability* 27 (2014), pp. 683–724.
- [96] Idris Kharroubi, Thomas Lim, and Armand Ngupeyou. “Mean-variance hedging on uncertain time horizon in a market with a jump”. In: *Applied Mathematics & Optimization* 68 (2013), pp. 413–444.
- [97] Kon-Gun Kim, Mun-Chol Kim, and Ho-Jin Hwang. “Representation of solutions to quadratic 2BSDEs with unbounded terminal values”. In: *Statistics & Probability Letters* 213 (2024), p. 110191.
- [98] Andrei Kirilenko, Albert S Kyle, Mehrdad Samadi, and Tugkan Tuzun. “The flash crash: High-frequency trading in an electronic market”. In: *The Journal of Finance* 72.3 (2017), pp. 967–998.

- [99] Magdalena Kobylanski. “Backward stochastic differential equations and partial differential equations with quadratic growth”. In: *The annals of probability* 28.2 (2000), pp. 558–602.
- [100] Arturo Kohatsu-Higa and Makoto Yamazato. “Enlargement of filtrations with random times for processes with jumps”. In: *Stochastic Processes and their applications* 118.7 (2008), pp. 1136–1158.
- [101] Thomas Kruse and Alexandre Popier. “Minimal supersolutions for BSDEs with singular terminal condition and application to optimal position targeting”. In: *Stochastic Processes and their Applications* 126.9 (2016), pp. 2554–2592.
- [102] Thomas Kruse and Alexandre Popier. “Lp-solution for BSDEs with jumps in the case $p < 2$: corrections to the paper ‘BSDEs with monotone generator driven by Brownian and Poisson noises in a general filtration’”. In: *Stochastics* 89.8 (2017), pp. 1201–1227.
- [103] Raymond CW Leung. “Continuous-time principal-agent problem with drift and stochastic volatility control: with applications to delegated portfolio management”. In: *Available at SSRN* (2014).
- [104] C Wei Li and Ashish Tiwari. “Incentive contracts in delegated portfolio management”. In: *The Review of Financial Studies* 22.11 (2009), pp. 4681–4714.
- [105] Qian Lin. “A BSDE approach to Nash equilibrium payoffs for stochastic differential games with nonlinear cost functionals”. In: *Stochastic Processes and their Applications* 122.1 (2012), pp. 357–385.
- [106] Yiqing Lin. “A new existence result for second-order BSDEs with quadratic growth and their applications”. In: *Stochastics* 88.1 (2016), pp. 128–146.
- [107] Yiqing Lin, Zhenjie Ren, Nizar Touzi, and Junjian Yang. “Second order backward SDE with random terminal time”. In: *Electron. J. Probab.* 25 (2020).
- [108] Yiqing Lin, Zhenjie Ren, Nizar Touzi, and Junjian Yang. “Random horizon principal-agent problems”. In: *SIAM Journal on Control and Optimization* 60.1 (2022), pp. 355–384.
- [109] Liwei Lu, Ruimeng Hu, Xu Yang, and Yi Zhu. “Multiagent Relative Investment Games in a Jump Diffusion Market with Deep Reinforcement Learning Algorithm”. In: *SIAM Journal on Financial Mathematics* 16.2 (2025), pp. 707–746.
- [110] Terry Lyons. “Differential equations driven by rough signals”. In: *Revista Matemática Iberoamericana* 14.2 (1998), pp. 215–310.

- [111] Terry Lyons, Sina Nejad, and Imanol Perez Arribas. *Nonparametric pricing and hedging of exotic derivatives*. Papers 1905.00711. arXiv.org, May 2019. URL: <https://ideas.repec.org/p/arx/papers/1905.00711.html>.
- [112] Jin Ma and Jianfeng Zhang. “Representation theorems for backward stochastic differential equations”. In: *The annals of applied probability* 12.4 (2002), pp. 1390–1418.
- [113] Harry M Markowitz. “Foundations of portfolio theory”. In: *The journal of finance* 46.2 (1991), pp. 469–477.
- [114] Thibaut Mastrolia. “Density analysis of non-Markovian BSDEs and applications to biology and finance”. In: *Stochastic Processes and their Applications* 128.3 (2018), pp. 897–938.
- [115] Thibaut Mastrolia and Dylan Possamai. “Moral hazard under ambiguity”. In: *Journal of Optimization Theory and Applications* 179 (2018), pp. 452–500.
- [116] Thibaut Mastrolia and William Yan. “Agency Problems and Adversarial Bilevel Optimization under Uncertainty and Cyber Threats”. In: *arXiv preprint arXiv:2505.08989* (2025).
- [117] Anis Matoussi, Lambert Piozin, and Dylan Possamai. “Second-order BSDEs with general reflection and game options under uncertainty”. In: *Stochastic Processes and their Applications* 124.7 (2014), pp. 2281–2321.
- [118] Anis Matoussi, Dylan Possamai, and Chao Zhou. “Robust utility maximization in nondominated models with 2BSDE: the uncertain volatility model”. In: *Mathematical Finance* 25.2 (2015), pp. 258–287.
- [119] Robert C Merton. “Optimum consumption and portfolio rules in a continuous-time model”. In: *Stochastic optimization models in finance*. Elsevier, 1975, pp. 621–661.
- [120] Zaineb Mezhdoud, Carsten Hartmann, Mohamed Riad Remita, and Omar Kebiri. “ α -Hypergeometric Uncertain Volatility Models and their Connection to 2BSDEs”. In: *arXiv preprint arXiv:2108.06965* (2021).
- [121] Marie-Amélie Morlais. “Quadratic BSDEs driven by a continuous martingale and applications to the utility maximization problem”. In: *Finance and Stochastics* 13 (2009), pp. 121–150.
- [122] Marie-Amélie Morlais. “Utility maximization in a jump market model”. In: *Stochastics: An International Journal of Probability and Stochastic Processes* 81.1 (2009), pp. 1–27.

- [123] Chenchen Mou and Andrzej Świąch. “Uniqueness of viscosity solutions for a class of integro-differential equations”. In: *Nonlinear Differential Equations and Applications NoDEA* 22.6 (2015), pp. 1851–1882.
- [124] Ariel Neufeld and Marcel Nutz. “Measurability of semimartingale characteristics with respect to the probability law”. In: *Stochastic Processes and their Applications* 124.11 (2014), pp. 3819–3845.
- [125] Marcel Nutz. “Robust superhedging with jumps and diffusion”. In: *Stochastic Processes and their Applications* 125.12 (2015), pp. 4543–4555.
- [126] Marcel Nutz and H Mete Soner. “Superhedging and dynamic risk measures under volatility uncertainty”. In: *SIAM Journal on Control and Optimization* 50.4 (2012), pp. 2065–2089.
- [127] Bernt Øksendal and Agnes Sulem. “Maximum principles for optimal control of forward-backward stochastic differential equations with jumps”. In: *SIAM Journal on Control and Optimization* 48.5 (2010), pp. 2945–2976.
- [128] Antonis Papapantoleon, Dylan Possamai, and Alexandros Saplaouras. “Existence and uniqueness results for BSDE with jumps: the whole nine yards”. In: *Electron. J. Probab.* 23 (2018).
- [129] Etienne Pardoux and Shige Peng. “Adapted solution of a backward stochastic differential equation”. In: *Systems & control letters* 14.1 (1990), pp. 55–61.
- [130] Etienne Pardoux and Shige Peng. “Backward stochastic differential equations and quasilinear parabolic partial differential equations”. In: *Stochastic Partial Differential Equations and Their Applications: Proceedings of IFIP WG 7/1 International Conference University of North Carolina at Charlotte, NC June 6–8, 1991*. Springer. 2005, pp. 200–217.
- [131] Elisabeth Paté-Cornell. “On “black swans” and “perfect storms”: Risk analysis and management when statistics are not enough”. In: *Risk Analysis: An International Journal* 32.11 (2012), pp. 1823–1833.
- [132] Alexandre Popier. “Limit behaviour of BSDE with jumps and with singular terminal condition”. In: *ESAIM: Probability and Statistics* 20 (2016), pp. 480–509.
- [133] Alexandre Popier and Chao Zhou. “Second-order BSDE under monotonicity condition and liquidation problem under uncertainty”. In: *Ann. Appl. Probab.* 29 (3 2019).
- [134] Dylan Possamai, Marco Rodrigues, and Alexandros Saplaouras. “Mind the jumps: when 2BSDEs meet semi-martingales”. In: *arXiv preprint arXiv:2507.01767* (2025).

- [135] Dylan Possamai and Xiaolu Tan. “Weak approximation of second-order BSDEs”. In: *Ann. Appl. Probab.* 25 (25 2015).
- [136] Dylan Possamai, Xiaolu Tan, and Chao Zhou. “Stochastic control for a class of nonlinear kernels and applications”. In: *The Annals of Probability* 46.1 (2018), pp. 551–603.
- [137] Dylan Possamai, Nizar Touzi, and Jianfeng Zhang. “Zero-sum path-dependent stochastic differential games in weak formulation”. In: *The Annals of Applied Probability* 30.3 (2020), pp. 1415–1457.
- [138] Nicolas Privault. *Introduction to stochastic finance with market examples*. Chapman and Hall/CRC, 2022.
- [139] Marie-Claire Quenez and Agnès Sulem. “BSDEs with jumps, optimization and applications to dynamic risk measures”. In: *Stochastic Processes and their Applications* 123.8 (2013), pp. 3328–3357.
- [140] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. “Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations”. In: *arXiv preprint arXiv:1711.10561* (2017).
- [141] Situ Rong. “On solutions of backward stochastic differential equations with jumps and applications”. In: *Stochastic Processes and their Applications* 66.2 (1997), pp. 209–236.
- [142] Mathieu Rosenbaum and Jianfei Zhang. *Multi-asset market making under the quadratic rough Heston*. 2022. arXiv: 2212.10164 [q-fin.MF]. URL: <https://arxiv.org/abs/2212.10164>.
- [143] Richard Rouge and Nicole El Karoui. “Pricing via utility maximization and entropy”. In: *Mathematical Finance* 10.2 (2000), pp. 259–276.
- [144] Manuela Royer. “BSDEs with a random terminal time driven by a monotone generator and their links with PDEs”. In: *Stochastics and Stochastic Reports* 76.4 (2004), pp. 281–307. DOI: 10.1080/10451120410001696270. eprint: <https://doi.org/10.1080/10451120410001696270>. URL: <https://doi.org/10.1080/10451120410001696270>.
- [145] Yuliy Sannikov. “A continuous-time version of the principal-agent problem”. In: *The Review of Economic Studies* 75.3 (2008), pp. 957–984.
- [146] Roland C Seydel. “Existence and uniqueness of viscosity solutions for QVI associated with impulse control of jump-diffusions”. In: *Stochastic Processes and their Applications* 119.10 (2009), pp. 3719–3748.

- [147] Evelina Shamarova, Roman Chertovskih, Alexandre F Ramos, and Paulo Aguiar. “Backward-stochastic-differential-equation approach to modeling of gene expression”. In: *Physical Review E* 95.3 (2017), p. 032418.
- [148] Steven E Shreve and H Mete Soner. “Optimal investment and consumption with transaction costs”. In: *The Annals of Applied Probability* (1994), pp. 609–692.
- [149] Justin Sirignano and Konstantinos Spiliopoulos. “DGM: A deep learning algorithm for solving partial differential equations”. In: *Journal of computational physics* 375 (2018), pp. 1339–1364.
- [150] H Mete Soner, Nizar Touzi, and Jianfeng Zhang. “Wellposedness of second order backward SDEs”. In: *Probability Theory and Related Fields* 153.1 (2012), pp. 149–190.
- [151] H. Mete Soner, Nizar Touzi, and Jianfeng Zhang. “Dual formulation of second order target problems”. In: *Annals of Applied Probability* 23.1 (2013), pp. 308–347.
- [152] Shiqi Song. “Optional splitting formula in a progressively enlarged filtration”. In: *ESAIM: Probability and Statistics* 18 (2014), pp. 829–853.
- [153] Livio Stracca. “Delegated portfolio management: A survey of the theoretical literature”. In: *Journal of Economic surveys* 20.5 (2006), pp. 823–848.
- [154] Nizar Touzi. *Optimal stochastic control, stochastic target problems, and backward SDE*. Vol. 29. Springer Science & Business Media, 2012.
- [155] Hui Ou-Yang. “Optimal contracts in a continuous-time delegated portfolio management problem”. In: *The Review of Financial Studies* 16.1 (2003), pp. 173–208.
- [156] Jiongmin Yong, Xun Yu Zhou, Jiongmin Yong, and Xun Yu Zhou. “Backward stochastic differential equations”. In: *Stochastic Controls: Hamiltonian Systems and HJB Equations* (1999), pp. 345–400.
- [157] Jianfeng Zhang. “Stochastic Calculus Under Weak Formulation”. In: *Backward Stochastic Differential Equations: From Linear to Fully Nonlinear Theory* (2017), pp. 205–244.