

DISS. ETH NO. 22018

NUMÉRAIRE-INDEPENDENT MODELLING
OF FINANCIAL MARKETS

A dissertation submitted to
ETH ZURICH

for the degree of
Doctor of Sciences

presented by

MARTIN PETER GERHARD HERDEGEN

MSc in Mathematical Sciences
University of Bath, United Kingdom
born April 10, 1986
citizen of Germany

accepted on the recommendation of

Prof. Dr. Martin Schweizer	examiner
Prof. Dr. Josef Teichmann	co-examiner
Prof. Dr. Konstantinos Kardaras	co-examiner

2014

DOI: 10.3929/ethz-a-010315552
ISBN: 978-3-906031-76-7

SOLI DEO GLORIA

MINUS ENIM TE AMAT
QUI TECUM ALIQUID AMAT
QUOD NON PROPTER TE AMAT*
Augustinus, *Confessiones. Liber X*

*HE LOVES THEE TOO LITTLE/WHO LOVES ANYTHING TOGETHER WITH THEE/WHICH HE LOVES NOT FOR THY SAKE. This phrase is to be understood in the sense that I love mathematics because I love God, and that I do mathematics in order to glorify God. Any other (historical) (mis)interpretation of this quote by any denomination or person is *theirs*, not mine.

Abstract

In this thesis, we develop a new framework for modelling financial markets which does not depend on an ex-ante choice of a currency unit and a numéraire. In this framework, we then study central concepts of mathematical finance like (no-)arbitrage or financial bubbles. The motivation for doing this is to overcome several deficiencies of the standard modelling approach in mathematical finance, first and foremost the initial choice of currency unit and numéraire, which is not justified from an economic point of view. In addition, our framework can be seen as a great unifying paradigm encompassing various other recent approaches to modelling financial markets and studying (no-)arbitrage.

Numéraire-independent modelling starts by describing a financial market not by a single process but by an equivalence class of processes. As a consequence, familiar concepts such as admissible trading strategies, contingent claims or superreplication prices have to be revisited and redefined in a numéraire-independent way. In particular—even though this sounds absurd—we have to give a numéraire-independent definition of the key notion of “numéraire” itself.

After this groundwork has been finished, we translate in our framework the catchphrase that arbitrage means “making a profit out of nothing without risk” into a rigorous mathematical definition as literally as possible by saying that a market satisfies *numéraire-independent no-arbitrage* (NINA) if and only if there does not exist a nonzero contingent claim (“the profit”) which can be superreplicated for free (“out of nothing”). We study this notion in detail and compare it to classic no-arbitrage concepts from the literature.

If a market fails NINA, a natural question is whether there is a stopping time σ such that strictly before σ , we can never “make a profit out of nothing without risk”, and immediately after σ , we always can. Another interesting question in the same setup is whether it is possible to find a probability measure \mathbb{Q} that is absolutely continuous to the physical measure \mathbb{P} such that the market satisfies NINA *under* \mathbb{Q} . We answer both questions in the affirmative and in full generality.

On the other hand, if a market satisfies NINA, an important mathematical problem is to find a dual characterisation of that property. We show that the corresponding dual objects are *pairs* (η, \mathbb{Q}) of numéraire strategies and equivalent σ -martingale measures. Of course, we also compare this numéraire-independent version of the fundamental theorem of asset pricing to the classic result by Delbaen and Schachermayer.

Last but not least, we propose a numéraire-independent approach to modelling financial bubbles. Unlike most papers in the recent literature we do not define

bubbles by a dual object, usually a *strict* local martingale measure, but start from primary concepts that are economically motivated. To this end, we introduce the notions of static and dynamic viability and efficiency of a market and derive their dual characterisations. In particular, we show that strict local martingale measures arise naturally in the context of modelling financial bubbles.

Kurzfassung

In dieser Arbeit entwickeln wir einen neuen Modellierungsansatz für Finanzmärkte, welcher nicht von der Vorab-Wahl einer Währungseinheit und eines Numéraire abhängt. Innerhalb dieses Ansatzes studieren wir dann zentrale Konzepte der Finanzmathematik wie Arbitrage(freiheit) oder Spekulationsblasen. Die Motivation dafür ist, mehrere Schwächen des Standardmodellierungsansatzes der Finanzmathematik zu überwinden, vor allem die Vorab-Wahl von Währungseinheit und Numéraire, welche aus ökonomischer Sicht nicht gerechtfertigt ist. Unseren Modellierungsansatz kann man zudem als ein grosses einheitliches Paradigma betrachten, in das sich zahlreiche andere neuere Modellierungsansätze für Finanzmärkte und Arbitrage(freiheit) einordnen lassen.

Numéraire-unabhängige Modellierung beginnt damit, dass man Finanzmärkte nicht durch einen einzigen Prozess, sondern eine Äquivalenzklasse von Prozessen beschreibt. Folglich muss man bekannte Konzepte wie zulässige Handelsstrategien, Eventualforderungen oder Superreplikationspreise überdenken und in Numéraire-unabhängiger Weise neu definieren. Insbesondere – auch wenn das absurd klingt – muss man den Schlüsselbegriff „Numéraire“ selbst Numéraire-unabhängig definieren.

Nach Abschluss dieser Vorarbeit, übersetzen wir innerhalb unseres Modellierungsansatzes den Allgemeinplatz, das Arbitrage bedeutet „aus Nichts einen Gewinn ohne Risiko zu erzielen“, so wörtlich wie möglich in eine präzise mathematische Definition, indem wir sagen, dass ein Markt genau dann *Numéraire-unabhängige Arbitragefreiheit* (NINA¹) erfüllt, wenn es keine Eventualforderung („den Gewinn“) gibt, der kostenlos („aus nichts“) superrepliziert werden kann. Wir studieren dann diesen Begriff sorgfältig und vergleichen ihn mit klassischen Konzepten der Arbitrage(freiheit) aus der Literatur.

Wenn ein Markt NINA nicht erfüllt, stellt sich ganz natürlich die Frage, ob eine Stoppzeit σ existiert, so dass man strikt vor σ in keinem Fall „aus Nichts einen Gewinn ohne Risiko erzielen“ kann, aber unmittelbar nach σ dies immer kann. Eine weitere interessante Frage unter der gleichen Voraussetzung ist, ob es möglich ist, ein Wahrscheinlichkeitsmass \mathbb{Q} zu finden, welches absolut stetig zum physischen Mass \mathbb{P} ist, so dass der Markt NINA *unter* \mathbb{Q} erfüllt. Wir geben auf beide Fragen eine positive Antwort und dies in voller Allgemeinheit.

Im anderen Fall, dass der Markt NINA erfüllt, besteht ein wichtiges mathematisches Problem darin, eine duale Charakterisierung dieser Eigenschaft zu finden.

¹Dies ist die englische Abkürzung des Begriffs, die wir im Folgenden aus Konsistenzgründen verwenden.

Wir zeigen, dass die geeigneten dualen Objekte *Paare* (η, \mathbb{Q}) von Numéraire-Strategien und äquivalenten σ -Martingalmassen sind. Selbstverständlich vergleichen wir diese Numéraire-unabhängige Version des Fundamentalsatzes der Anlagenbewertung mit dem klassischen Resultat von Delbaen und Schachermayer.

Zum Schluss schlagen wir einen Numéraire-unabhängigen Modellierungsansatz für Spekulationsblasen vor. Anders als die meisten neueren Artikel in der Literatur definieren wir dabei eine Spekulationsblase nicht mittels eines dualen Objekts, normalerweise eines *strikt* lokalen Martingalmasses, sondern starten von primären Konzepten, die ökonomisch motiviert sind. Dazu führen wir die Begriffe der statischen und dynamischen Viabilität und Effizienz ein und leiten deren duale Charakterisierung her. Insbesondere zeigen wir, dass strikt lokale Martingalmasse auf ganz natürliche Weise im Zusammenhang mit Spekulationsblasen auftauchen.

Acknowledgements

First and foremost, I wish to express my sincere and deep gratitude to Martin Schweizer, my doctoral advisor. In numerous discussions, he helped me not only to advance my understanding of my research topic but also to sharpen my thinking about mathematical problems in general. His steadfast quest for clarity, elegance and rigor in mathematical writing has improved my own style immeasurably—even outside mathematics. But above all, I would like to thank him for his great and constant encouragement, in particular in times of drawback and failure. Next, I am much indebted to Josef Teichmann and Konstantinos Kardaras for readily accepting to act as co-examiners and for sharing with me their view on my research topic. I am also profoundly grateful to Andreas Kyprianou, the supervisor of my Master’s thesis, for introducing me to mathematical research and for strongly motivating me to pursue a PhD. Moreover, I want to thank my colleagues in Gruppe 3 at ETH for creating such a relaxed and stimulating working atmosphere. I am especially grateful for the friendship of Blanka Horvarth and Sebastian Herrmann. Last but not least, I would like to thank my parents for their unconditional love and tremendous support during all the years of my life. Without them, I would not be where I am.

Financial support by the National Centre of Competence in Research “Financial Valuation and Risk Management” (NCCR FINRISK), Project D1 (Mathematical Methods in Financial Risk Management) for parts of this thesis is gratefully acknowledged.

Contents

I Introduction	1
1 Numéraire-independent modelling in a nutshell	3
2 Overview of the thesis	4
3 Probabilistic setup and general notation	6
II Key concepts of numéraire-independent modelling	9
1 Exchange rate processes and markets	9
2 Self-financing strategies	11
3 Numéraire strategies and representatives	14
4 Strategy cones and undefaultable strategies	18
5 Contingent claims	19
6 Superreplication prices	21
7 Robustness of limit quantile superreplication prices	28
III Numéraire-independent no-arbitrage (NINA)	33
1 Gratis events	33
2 Numéraire-independent no-arbitrage (NINA)	37
3 Maximal strategies	39
IV Separating stopping times for markets failing NINA	51
1 Conceptual prélude	51
2 The lower separating stopping time	53
3 The upper separating stopping time	54
4 Examples	57
V Absolutely continuous measures for markets failing NINA	59
1 Technical preliminaries	59
2 Continuous markets	63
3 General markets	65
VI Dual characterisation of markets satisfying NINA	71
1 Numéraire-independent FTAP	71
2 Martingale properties	75
3 Dual characterisation of superreplication prices	78
4 Appendix	82

VII Comparison to other modelling frameworks	85
1 The standard modelling framework	85
2 The fair market framework	86
3 The numéraire portfolio and the benchmark approach	87
4 The equivalent local martingale deflator framework	88
VIII Bubbles from a numéraire-independent perspective	89
1 Main concepts	90
2 First examples	101
3 Dual characterisation of nontrivial bubbly markets.	103
4 Further examples	116
5 Comparison to the literature	123
6 Appendix	130

Chapter I

Introduction

The standard modelling approach in mathematical finance consists of four steps. First, one chooses a *currency unit*. This is a *label* attaching numbers over time to real commodities and services. Examples are an ounce of gold, a Euro, or its precursor, the European currency unit (ECU). Second, one describes in that unit the evolution of N financial assets by a stochastic process $\tilde{S} = (\tilde{S}_t^1, \dots, \tilde{S}_t^N)_{t \in [0, T]}$. Third, one chooses in that unit a *numéraire*, which for the moment may be understood to mean the *positive* price process \tilde{S}^k of a basic asset. (A bit more precisely, we need that $\mathbb{P}[\inf_{t \in [0, T]} \tilde{S}_t^k > 0] = 1$.) Fourth, one expresses all other assets in units of the numéraire asset \tilde{S}^k by defining $X^i := \tilde{S}^i / \tilde{S}^k$. After relabelling the assets, the numéraire asset k is called “asset 0” or *bank account*; the other $d := N - 1$ assets X^1, \dots, X^d are called *risky assets*, and it is often said that they are *discounted by* or *in units of* the bank account.

The vast majority of papers in mathematical finance—with an obvious exception in the literature on interest rate modelling—starts only after the final step of the above procedure. These papers consider from the beginning an \mathbb{R}^d -valued process X and call this the (discounted) price process of d (risky) assets. Almost always, they also assume (but very often do not mention explicitly) that there is in addition to X a (riskless) bank account whose price is identically 1. Moreover, it is sometimes said and probably widely believed throughout the community that starting right away with X can be done without loss of generality, i.e., that doing this is a purely mathematical convenience.

However, starting with X is a loss of generality in two respects. First, this implicitly assumes that the price process of at least one of the basic assets is positive. To illustrate this point, consider a market with two assets \tilde{S}^1 and \tilde{S}^2 which both may default with positive probability, but on disjoint events, i.e., $\inf_{0 \leq t \leq T} |\tilde{S}_t^i| = 0$ on A_i with $0 < P[A_i] < 1$ and $\mathbb{P}[A_1 \cap A_2] = 0$. Then neither \tilde{S}^1 nor \tilde{S}^2 can be used for discounting. Second, and this is the crucial point, when passing from \tilde{S} to X , two important pieces of information are lost: one forgets the original currency unit of \tilde{S} and the price process \tilde{S}^k of the numéraire asset in the original currency unit. From an economic perspective, it makes a difference whether the currency unit of \tilde{S} is EUR or USD, say, or whether \tilde{S}^k is a deterministic function or a geometric Brownian motion, say.

If we ignore the first point for the moment, one might well be led to think the second point does not really matter when studying only *qualitative* and *preference-independent* properties of the market, such as absence of arbitrage. We think that it is in this sense that people generally understand the statement that starting with X can be done without loss of generality. However, it is an unfortunate fact that this intuitive understanding does not rest on solid foundations. The most basic definition in the standard framework, the notion of *admissible strategies* [9], crucially depends on the original currency unit of \tilde{S} and on the price process of the numéraire asset \tilde{S}^k . This is because it imposes a (strategy-dependent) credit line, which is expressed in quantities of the numéraire asset \tilde{S}^k . As a consequence, all the classic no-arbitrage concepts like *no-arbitrage (NA)* [9], *no free lunch with vanishing risk (NFLVR)* [9] and *no unbounded profit with bounded risk (NUPBR)* [46, 73] depend *by their very definition*, at least formally, on the currency unit and on the numéraire chosen in the four steps above. The same holds for the notion of *maximal strategies* in the sense of Delbaen and Schachermayer [11, 12] because their comparison class consists of admissible strategies. It would be more appropriate to call such strategies (asset 0)-admissible and to speak of (asset 0)-NA, (asset 0)-NFLVR, etc., to emphasise the dependence on (the choice of) asset 0 also in the notation.

In view of the above dependence, calling failure of NFLVR simply “arbitrage” and thereby suggesting a preference-independent concept is misleading. This has led to some results in the recent literature which might be surprising at a first glance. For example, in the benchmark approach [65], a market may have “arbitrage” (violate NFLVR); but in units of the numéraire portfolio, the theory works as if there was no arbitrage. For stochastic portfolio theory and relative arbitrage [22, 70], a market may have “arbitrage”; but studying portfolio choice still makes sense, and Delta hedging still works. Finally, financial bubbles in the sense of Protter et al. [40, 41, 67] and “arbitrage” seem to be two sides of the same coin in the sense that if $(1, X)$ is a “bubble model”, then $(1/X, 1)$ is an “arbitrage model”; see also [27]. On closer inspection, however, none of the above situations is really astonishing: Failure of NFLVR is a preference-dependent concept—the preference is encoded in the original currency unit of \tilde{S} and the price process of the numéraire asset \tilde{S}^k —and so the use of the (preference-independent) word “arbitrage” is simply misleading in each case.

To some extent, specialists are aware of the dependence of the standard framework on the choice of a currency unit and a numéraire. (Let us stress that the notion of “numéraire” in the standard framework implicitly assumes that a currency unit has been chosen. A numéraire is always a numéraire in some fixed currency unit; cf. the discussion after Remark II.3.2 below.) Delbaen and Schachermayer [11, 12] fix a currency unit and an initial “good” numéraire (in the sense that X satisfies NFLVR), and then study how properties transfer from one numéraire to another. They analyse for example how NFLVR for X is related to NFLVR for X' , where $X'_i = \tilde{S}^i / \tilde{S}^{k'}$. The experts in the field also know that the no-arbitrage conditions of *no-arbitrage of the first kind (NA₁)* [48], *no cheap thrills* [56] or NUPBR are equivalent, and it is folklore that they do not depend on the choice

of a numéraire (in a *given* fixed currency unit)—even though (in the case of NUPBR) this is not trivial and does not appear to have been *rigorously* proved.¹ But these results are based on initial choices of a currency unit and a numéraire to even *define* the concepts involved; an example where all the assets can default is not included. This is discussed in more detail in Chapter III.3.3.

The goal of this thesis is to tackle the root of the problem explained above—we want to define and study, in a systematic way, concepts like no-arbitrage or financial bubbles *without first choosing a currency unit and a numéraire*. To this end, we develop a new modelling framework and reconsider from scratch how to describe markets, trading strategies, etc. in that setting. This needs care, and sometimes apparently familiar concepts must be re-introduced to explain the subtle differences. We call our approach *numéraire-independent* because it is based upon concepts that do not fix a currency unit or a numéraire.²

1 Numéraire-independent modelling in a nutshell

Before providing a short overview over the material presented in this thesis, let us briefly sketch some of the key ideas of numéraire-independent modelling.

The starting point of our work is to recall that mathematical finance can only describe *relative* prices of traded assets (e.g. stocks, derivatives or interest rate products) in terms of a given currency unit, or the relative price of one currency unit in terms of another, i.e., an exchange rate. For this reason, we model a financial market not by a *single* price process S , but by an *equivalence class* \mathcal{S} of processes S . Each $S \in \mathcal{S}$ describes the evolution of asset prices in one of the (infinitely many) possible currency units. Economic sense then dictates that all definitions and results must be formulated in such a way that they hold for *some* representative $S \in \mathcal{S}$ if and only if they hold for *every* $S \in \mathcal{S}$. This gives a truly numéraire-independent framework.

Trading in the market \mathcal{S} is “as usual” restricted to self-financing strategies, and we show that the self-financing concept is numéraire-independent. Again “as usual”, we need some additional conditions to exclude doubling phenomena when there are infinitely many trading dates. The familiar notion of (asset 0)-admissibility cannot be used since it is based on initial choices of a currency unit and a numéraire. Instead, we consider self-financing strategies with nonnegative value processes. If the value process of an undefaultable strategy η is positive, we call η a *numéraire strategy*. For each such η , there is a unique *numéraire representative* $S^{(\eta)} \in \mathcal{S}$ in whose units the value process of η is identically 1.

¹For instance, it is stated in Schweizer and Takaoka [73, before Proposition 2.7, italics added] that NUPBR is a “numéraire-free property in a *certain sense*”, but this is not made precise. Moreover, in [73, Proposition 2.7] there is a change of dimension, and so this result is strictly speaking *not* numéraire-free.

²The more precise terminology *currency-unit- and numéraire-independent* is too long to be useful. Moreover, we note in passing that the paper by Yan [83] on a numéraire-free framework is a misnomer; it uses as initial numéraire (in the original currency unit) instead of one basic asset \tilde{S}^k their sum $\sum_{i=1}^N \tilde{S}^i$.

This simple observation has remarkably far-reaching consequences. Instead of imposing on \mathcal{S} that one asset is positive in some (and hence in every) currency unit, we only need to require that there exists a (dynamic) numéraire strategy for \mathcal{S} . This allows us to deal with very general situations, including the above example where all assets can default.

We next want to introduce a no-arbitrage condition which is not based on initial choices of a currency unit and a numéraire (strategy). Simply modifying the classic concepts of (asset 0)-NA and (asset 0)-NFLVR by replacing (asset 0)-admissible strategies in the definitions by undefaultable strategies starting at 0 does not work; see the discussion after Proposition III.3.21 below. Instead, we go back to the fundamental economic catchphrase that arbitrage means “making a profit out of nothing without risk”. This can be rephrased as “anything good and riskless must cost something”, and so we say that a market \mathcal{S} satisfies *numéraire-independent no-arbitrage (NINA)* if every nonzero (nonnegative) contingent claim has a positive superreplication price. Making this rigorous also involves defining, like the market \mathcal{S} , the notions of a contingent claim and a superreplication price in a numéraire-independent way.

In the standard framework, the dual characterisation of no-arbitrage is given by the classic fundamental theorem of asset pricing (FTAP), in its most general form due to Delbaen and Schachermayer [9, 13]. It says that the (asset 0)-denominated price process $(1, X)$ satisfies (asset 0)-NFLVR if and only if there exists an equivalent σ -martingale measure (E σ MM) \mathbb{Q} for X . One of our main results is a numéraire-independent version of the FTAP. We show in Theorem VI.1.10 below that a market \mathcal{S} satisfies NINA (which is strictly weaker than (asset 0)-NFLVR) if and only if there exists a pair (η, \mathbb{Q}) , where η is a “good” numéraire strategy and \mathbb{Q} an E σ MM for the corresponding numéraire representative $S^{(\eta)}$. So our dual characterisation includes the *existence* of a suitable “good” numéraire, which can be different from any of the basic assets, and which is neither given nor chosen ex ante. For this reason, we think that our version of the FTAP is more natural from an economic perspective.

2 Overview of the thesis

The concepts and results obtained in this thesis are presented in seven chapters. The last chapter, which is joint work with Martin Schweizer, is unlike the others largely self-contained, and so there are some redundancies there. The material for Chapter II is mainly taken from [30], Chapters III, VI and VII are based on [31] and [32], and Chapters IV, V and VIII are unpublished material.

II Key concepts of numéraire-independent modelling. In this chapter, we motivate and develop the key concepts of numéraire-independent modelling. After explaining in Section II.1 what we understand under a market in our framework, we discuss in Section II.2 how trading can be described in a numéraire-independent way. In Section II.3, we define the central concepts of numéraire strategies and numéraire representatives and compare them to the classic notion

of a numéraire (in a fixed currency unit). In Section II.4, we introduce strategy cones and show in which sense the prime example of undefaultable strategies is the numéraire-independent counterpart of the classic notion of admissible strategies. In Section II.5, we look at contingent claims from a numéraire-independent perspective and discuss the notion of numéraire-independent derivative securities. In Section II.6, we study in our framework two notions of superreplication prices, which we call *ordinary* and *limit quantile* superreplication prices. The first one is the numéraire-independent counterpart of superreplication prices in the standard framework; the second one is new and captures a slightly relaxed version of the concept of superreplication. It has nicer (continuity) properties than the first one and is the cornerstone of our concept of (no-)arbitrage in Chapter III. Finally, we show in Section II.7, that limit quantile superreplication prices enjoy a robustness property that ordinary superreplication prices lack. Theorem II.7.3, which is mathematically the most sophisticated result in this chapter, is also the technical anchor of the results in Chapter V.

III Numéraire-independent no-arbitrage (NINA). In this chapter, we study the concept of (no-)arbitrage in our numéraire-independent framework. In Section III.1, we introduce a very general and *quantitative* notion of arbitrage, called *gratis events*, based on limit quantile superreplication prices. The results in Section III.1 also lay the ground for Chapters IV and V. For the most relevant case of undefaultable strategies, we derive an equivalent and simpler characterisation of the absence of gratis events based on ordinary superreplication prices in Section III.2. We call the corresponding notion *numéraire-independent no-arbitrage (NINA)*. In Section III.3, we compare NINA to classic notions of no-arbitrage including no-arbitrage (NA), no free lunch with vanishing risk (NFLVR) and no unbounded profit with bounded risk (NUPBR) by presenting a new unifying characterisation of those concepts in terms of *maximal strategies*. The results in Section III.3 are also foundational for Chapter VI.

IV Separating stopping times for markets failing NINA. In this chapter, we study markets failing NINA. We seek to find a stopping time σ such that (1) strictly before σ , we can never “make a profit out of nothing without risk”, and (2) immediately after σ , we always can. It is natural to call σ a *separating stopping time* for the market \mathcal{S} . After making the above concept mathematically precise and establishing a preliminary lemma in Section IV.1, we show the existence of a smallest separating stopping time in Section IV.2 and prove the existence of a largest separating stopping time in Section IV.3. We conclude the chapter by illustrating our results by several examples in Section IV.4.

V Absolutely continuous measures for markets failing NINA. In this chapter, we study markets which fail NINA under the physical measure \mathbb{P} . We seek to answer the question whether it is then possible to find an *absolutely continuous* probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T such that the market satisfies NINA *under* \mathbb{Q} . After addressing some technical issues related to an absolutely

continuous change of measure in Section V.1, we explain in Section V.2 how to construct \mathbb{Q} in the case of *continuous* markets. The more difficult general case is fully solved in Section V.3.

VI Dual characterisation of markets satisfying NINA. In this chapter, we study markets satisfying NINA. After proving the existence of nonzero strongly maximal (numéraire) strategies, we derive a numéraire-independent version of the fundamental theorem of asset pricing (FTAP) in Section VI.1. In Section VI.2, we provide a dual characterisation of (weakly and strongly) maximal strategies and provide conditions for the existence of (*true*) *martingale representatives*. In Section VI.3, we derive a numéraire-independent dual characterisation of super-replication prices and discuss the notion of (strongly) *maximal* and (strongly) *attainable* contingent claims.

VII Comparison to other modelling frameworks. In this short chapter, we compare our numéraire-independent approach of modelling financial markets and studying no-arbitrage to the standard and other recent approaches to these issues.

VIII Bubbles from a numéraire-independent perspective. In this final chapter, we develop a new approach for modelling financial bubbles using our numéraire-independent paradigm. Unlike most papers in the recent literature, e.g. [56, 8, 40, 41, 67], we do not define bubbles by a dual object, usually a *strict* local martingale measure, but start from primary notions that are economically motivated. After explaining the main concepts of static and dynamic *viability* and *efficiency* in Section VIII.1, we illustrate them by several examples in Section VIII.2 before deriving their dual characterisations in Section VIII.3. In particular, we show that strict local martingale measures *arise naturally* in the context of modelling financial bubbles; see Theorem VIII.3.22. After providing some further examples of what we call *nontrivial bubbly markets* in Section VIII.4, we compare our definitions and results to the existing literature on bubbles in Section VIII.5.

3 Probabilistic setup and general notation

We always work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ satisfying the usual conditions of right-continuity and completeness; $T > 0$ is a finite time horizon. We assume that \mathcal{F}_0 is \mathbb{P} -trivial. We denote the collection of all \mathbb{P} -nullsets in \mathcal{F}_T by \mathcal{N} and the set of all stopping times with values in $[0, T]$ by $\mathcal{T}_{[0, T]}$. We also consider stopping times with values in $[0, T] \cup \{+\infty\}$ and agree as usual that $\inf \emptyset = +\infty$.

We set $\mathbb{R}_+ := [0, \infty)$, $\mathbb{R}_{++} := (0, \infty)$, $\overline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{+\infty\}$, $\underline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{-\infty\}$ and $\overline{\underline{\mathbb{R}}}_+ := \mathbb{R}_+ \cup \{+\infty\} \cup \{-\infty\}$. All these sets are convex cones, when we agree that $c + \infty = \infty + c = \infty$ for $c \in \overline{\mathbb{R}}_+$, $-\infty + c = c - \infty = -\infty$ for $c \in \underline{\mathbb{R}}_+$, $c \times (\pm\infty) = \pm\infty$ for $c \in \mathbb{R}_{++}$, and $0 \times (\pm\infty) = 0$. For $\tau \in \mathcal{T}_{[0, T]}$, we denote

by $\mathbf{L}_+^0(\mathcal{F}_\tau)$, $\mathbf{L}_{++}^0(\mathcal{F}_\tau)$, $\overline{\mathbf{L}}_+^0(\mathcal{F}_\tau)$, $\underline{\mathbf{L}}_+^0(\mathcal{F}_\tau)$ and $\overline{\underline{\mathbf{L}}}_+^0(\mathcal{F}_\tau)$ the sets of all \mathcal{F}_τ -measurable random variables taking values in \mathbb{R}_+ , \mathbb{R}_{++} , $\overline{\mathbb{R}}_+$, $\underline{\mathbb{R}}_+$, and $\overline{\underline{\mathbb{R}}}_+$, respectively.

An \mathbb{R}^N -valued predictable process $\zeta = (\zeta_t^1, \dots, \zeta_t^N)_{t \in [0, T]}$ is called *bounded* if there exists a constant $K > 0$ such that $\sup_{t \in [0, T]} \|\zeta_t\| \leq K$ \mathbb{P} -a.s., where $\|\cdot\|$ denotes any norm in \mathbb{R}^N . It is called *simple predictable* if it is bounded and if there exist stopping times $0 = \tau_0 \leq \dots \leq \tau_n = T$ in $\mathcal{T}_{[0, T]}$ and \mathbb{R}^N -valued $\mathcal{F}_{\tau_{k-1}}$ -measurable³ random vectors ξ_k such that

$$\zeta = \xi_0 \mathbf{1}_{[0]} + \sum_{k=1}^n \xi_k \mathbf{1}_{] \tau_{k-1}, \tau_k]} \quad \mathbb{P}\text{-a.s.}$$

For an \mathbb{R}^N -valued semimartingale $X = (X_t^1, \dots, X_t^N)_{t \in [0, T]}$, we denote by $L(X)$ the set of all \mathbb{R}^N -valued predictable processes $\zeta = (\zeta_t^1, \dots, \zeta_t^N)_{t \in [0, T]}$ that are integrable with respect to X in the sense of N -dimensional (vector) stochastic integration (consult Jacod and Shiryaev [37] for details). For $\zeta \in L(X)$ and $0 \leq t \leq T$, we write $\zeta \bullet X_t$ for the stochastic integral $\int_{(0, t]} \zeta_u dX_u$ and $\zeta_t \cdot X_t$ for the inner product $\sum_{k=1}^N \zeta_t^k X_t^k$.

³with $\tau_{-1} := 0$.

Chapter II

Key concepts of numéraire-independent modelling

In this chapter, we motivate and develop the key concepts of numéraire-independent modelling. After explaining in Section 1 what we understand under a market in our framework, we discuss in Section 2 how trading can be described in a numéraire-independent way. In Section 3, we define the central concepts of numéraire strategies and numéraire representatives and compare them to the classic notion of a numéraire (in a fixed currency unit). In Section 4, we introduce strategy cones and show in which sense the prime example of undefaultable strategies is the numéraire-independent counterpart of the classic notion of admissible strategies. In Section 5, we look at contingent claims from a numéraire-independent perspective and discuss the notion of numéraire-independent derivative securities. In Section 6, we study in our framework two notions of superreplication prices, which we call *ordinary* and *limit quantile* superreplication prices. The first one is the numéraire-independent counterpart of superreplication prices in the standard framework; the second one is new and captures a slightly relaxed version of the concept of superreplication. It has nicer (continuity) properties than the first one and is the cornerstone of our concept of (no-)arbitrage in Chapter III. Finally, we show in Section 7, that limit quantile superreplication prices enjoy a robustness property that ordinary superreplication prices lack. Theorem 7.3, which is mathematically the most sophisticated result in this chapter, is also the technical anchor of the results in Chapter V. The material for this chapter closely follows [30].

1 Exchange rate processes and markets

One primal economic entity in financial markets is a *currency unit*. But since mathematical finance can only model *relative* prices of one currency unit in terms of another,¹ the corresponding primal mathematical entity is an *exchange rate process*.

¹For an excellent discussion of this fact, we refer to the introduction of Vecer [80].

Definition 1.1. An *exchange rate process* is a positive semimartingale D satisfying

$$\inf_{0 \leq t \leq T} D_t > 0 \text{ P-a.s.} \quad (1.1)$$

We denote the set of all exchange rate processes by \mathcal{D} .

If $S = (S^1, \dots, S^N)$ is an \mathbb{R}^N -valued semimartingale describing the evolution of N assets in EUR and D is the exchange rate process of EUR against USD, i.e., 1 EUR corresponds to D_t USD at time t , then $S' = DS$ describes the evolution of the assets in USD. Clearly, S and S' describe the *same* financial market, and so we should treat them as one object, not two.

Definition 1.2. Two \mathbb{R}^N -valued semimartingales S and S' are called *economically equivalent* if $S' = DS$ P-a.s. for some $D \in \mathcal{D}$.

It is easy to check that economic equivalence is indeed a mathematical equivalence relation, this uses (1.1).

Let us repeat: If S describes N assets in some currency unit and S' is economically equivalent to S , then S' describes the same assets in another currency unit. So instead of just considering S , we can consider the set \mathcal{S} of all semimartingales which are economically equivalent to S —or to S' . Since it does not matter, whether \mathcal{S} was derived from S or from S' , we obtain a numéraire-independent (i.e., a currency-unit- and numéraire-independent) description of the assets.

Definition 1.3. A set $\mathcal{S} \neq \emptyset$ of \mathbb{R}^N -valued semimartingales is called an (N -dimensional) *market* if $\mathcal{S} = \{DS : D \in \mathcal{D}\}$ for some (and hence every) $S \in \mathcal{S}$ and if some (and hence every) $S \in \mathcal{S}$ satisfies

$$\inf_{0 \leq t \leq T} \sum_{i=1}^N |S_t^i| > 0 \text{ P-a.s.} \quad (1.2)$$

Each $S \in \mathcal{S}$ is called a *representative* of \mathcal{S} . \mathcal{S} is called *nonnegative* if some (and hence every) $S \in \mathcal{S}$ is P-a.s. componentwise nonnegative; it is called *continuous* if there exists a representative $S \in \mathcal{S}$ which has P-a.s. continuous trajectories.

To define a market \mathcal{S} , one usually starts with an (\mathbb{R}^N -valued) semimartingale S satisfying (1.2) describing the evolution of the market in some currency unit, and then sets $\mathcal{S} := \{DS : D \in \mathcal{D}\}$, which we also call the *market generated by S* . A special case is

$$S = (1, X), \quad (1.3)$$

where $X = (X^1, \dots, X^d)$ denotes $d = N - 1$ “risky” asset and 1 a “riskless” asset, in “discounted” units. We refer to (1.3) as a *classic model* in the sequel; note the index shift $S^{i+1} = X^i$.

Remark 1.4. Note that (1.2) is strictly weaker than $\inf_{0 \leq t \leq T} S_t^i > 0$ P-a.s. for some $i = 1, \dots, N$; see the discussion in Chapter I. Mathematically, (1.2) implies

that for each pair $(S, S') \in \mathcal{S} \times \mathcal{S}$, there exists a \mathbb{P} -a.s. *unique* exchange rate process D satisfying $S' = DS$ \mathbb{P} -a.s.; this follows from the identity

$$D = \frac{\sum_{i=1}^N |S'^i|}{\sum_{i=1}^N |S^i|} \mathbb{P}\text{-a.s.}$$

Economically, (1.2) is a very weak *nondegeneracy condition* on the market; it requires in particular that there do not exist a stopping time $\tau \in \mathcal{T}_{[0, T]}$ and an event $A_\tau \in \mathcal{F}_\tau$ with $\mathbb{P}[A_\tau] > 0$ on which all assets have no value whatsoever and where one could exchange *any* position in the assets against *any* other position.²

Since our approach is numéraire-independent, we must be careful with interpretations. If we pick an arbitrary representative $S \in \mathcal{S}$, an expression like $S_1^1 = 100$ does not have any financial meaning whatsoever *unless* we also have precise information about the currency unit corresponding to S . Similarly, an expression like $S_1^1 = 100$, $S_1^2 = 300$ in general only conveys the information that at time 1, we can exchange $3 = 300/100$ quantities (think of stock certificates) of asset 1 against 1 quantity of asset 2.

Remark 1.5. If one generalises the notion of exchange rate processes to non-negative semimartingales which satisfy (1.1) only strictly prior to a (predictable) stopping time $\tau \in \mathcal{T}_{[0, T]}$ and says that S and S' are economically equivalent if $S' = DS$ \mathbb{P} -a.s. on $\llbracket 0, \tau \llbracket$ for some generalised exchange rate process D , one can extend the notion of a market correspondingly. Economically, this captures the notion of defaulting/exploding exchange rates. For a setup with two representatives $S^\$$ and S^ϵ and one exchange rate process X , we refer to Carr et al. [7].

2 Position processes, investment processes and self-financing strategies

We proceed to describe trading in our numéraire-independent framework. We always assume that an (N -dimensional) market \mathcal{S} is *exogenously* given and that there are *no market frictions* such as transaction costs or liquidity premia. Nevertheless, this is not a completely trivial exercise. On the one hand, the standard definition of self-financing strategies found in most textbooks on continuous-time mathematical finance assumes the classic framework, and so we cannot use it. On the other hand, unlike in the classic framework, we also have to consider strategies which are not self-financing; see the discussion after Corollary 3.8.

2.1 Position processes

First, trading can be seen *statically*: At each stopping time $\tau \in \mathcal{T}_{[0, T]}$, one looks at the *position* $\varphi_\tau = (\varphi_\tau^1, \dots, \varphi_\tau^N)$ held in the assets. This can be interpreted as the point of view of tax or regulatory authorities, which check at some point in time

²For a precise definition of a *position* in a market, we refer to the following section.

the inventory of some market participant. Note that this static view does not say anything about how *changes* in positions are made. Therefore, mathematically, we only have to require that for each stopping time $\tau \in \mathcal{T}_{[0,T]}$, the quantity φ_τ only depends upon information available up to time τ , i.e., is \mathcal{F}_τ -measurable.

Definition 2.1. A *position process for the market \mathcal{S}* is an \mathbb{R}^N -valued progressive process $\varphi = (\varphi_t^1, \dots, \varphi_t^N)_{t \in [0,1]}$. For each $S \in \mathcal{S}$, the process $V(\varphi)(S)$ defined by

$$V_t(\varphi)(S) := \varphi_t \cdot S_t, \quad t \in [0, T],$$

is called the *position value process of φ in the currency unit corresponding to S* .

If $S \in \mathcal{S}$ describes the evolution of the assets in EUR, then $V_t(\varphi)(S)$ is the value in EUR of the position φ_t held at t . Formally, $V(\varphi)$ is a *map* from \mathcal{S} to the space of progressive processes, which for all $S \in \mathcal{S}$ and all $D \in \mathcal{D}$ satisfies the *exchange rate consistency condition*

$$V(\varphi)(DS) = \varphi \cdot (DS) = D(\varphi \cdot S) = DV(\varphi)(S) \text{ P-a.s.} \quad (2.1)$$

Thus, if one knows $V(\varphi)(S)$ for *some* $S \in \mathcal{S}$, one also knows it via (2.1) for *every* $S \in \mathcal{S}$, without any need for explicit knowledge of φ .

Remark 2.2. We often write $V(\varphi) \geq 0$ P-a.s., or $\mathbb{P}[V_\tau(\varphi) > 0] > 0$, etc. as a shorthand for $V(\varphi)(S) \geq 0$ P-a.s. for all $S \in \mathcal{S}$, or $\mathbb{P}[V_\tau(\varphi)(S) > 0] > 0$ for all $S \in \mathcal{S}$, etc. We use this notation only if the validity of the expression for *some* $S \in \mathcal{S}$ implies its validity for *every* $S \in \mathcal{S}$. This is for instance not the case in expressions like $\mathbb{P}[V_\tau(\varphi)(S) > c] > 0$ for $c \neq 0$.

2.2 Investment processes

Positions in our terminology are for static views. Now we take a *dynamic* perspective: A position process φ for the market \mathcal{S} is not just held, but actually *invested* into the assets, and thereby generates cash flows which have to be measured in some currency unit. Replacing φ by ζ to emphasise this dynamic view in the notation as well, we obtain for each $S \in \mathcal{S}$ a *cumulative cash flow process* $\zeta \bullet S$ and, together with an initial position ζ_0 , an *investment value process* $\zeta_0 \cdot S_0 + \zeta \bullet S$, both expressed in the currency unit corresponding to S . One can interpret this dynamic perspective as the point of view of a broker, who carries out the transactions of an investor and reports back how the positions have performed over time. Note that this dynamic view still does not say anything about how the investment is *financed*, i.e., how changes from ζ_s to ζ_t for $s < t$ are paid for. Mathematically, two additional assumptions on ζ (as opposed to φ) have to be made. First, since an investor at time $s < t$ does not know the market situation at time t , ζ must be predictable. Second, due to the continuous-time setup, we have to assume that $\zeta \in L(S)$ in order to make sense of the stochastic integral $\zeta \bullet S$.

Definition 2.3. An *investment process* for the representative $S \in \mathcal{S}$ is a position process $\zeta = (\zeta_t^1, \dots, \zeta_t^N)_{t \in [0, T]}$ that is predictable and satisfies $\zeta \in L(S)$. The process $\tilde{V}(\zeta)(S)$ defined by

$$\tilde{V}_t(\zeta)(S) := \zeta_0 \cdot S_0 + \zeta \bullet S_t, \quad t \in [0, T],$$

is called the *investment value process* of ζ in the currency unit corresponding to S .

Investment processes are—unlike position processes—in general not numéraire-independent objects. This is because the condition $\zeta \in L(S)$ refers to a *specific* representative $S \in \mathcal{S}$. But even if ζ is such that $\tilde{V}(\zeta)(S)$ is defined for all $S \in \mathcal{S}$, $\tilde{V}(\zeta)$ satisfies an exchange rate consistency condition similar to (2.1) only if ζ is *self-financing*; see Lemma 2.5 below.

2.3 Self-financing strategies

There is no general link between the static and the dynamic perspective of trading. Notwithstanding, if an investor neither *consumes* nor receives any further *endowment*, then his value and investment value process must *coincide*. We use yet another notation ϑ to reflect this below.

Definition 2.4. A *self-financing strategy* for the market \mathcal{S} is a position process $\vartheta = (\vartheta_t^1, \dots, \vartheta_t^N)_{t \in [0, T]}$ that is an investment process for *every* $S \in \mathcal{S}$ and satisfies for every $S \in \mathcal{S}$,

$$\vartheta \cdot S = V(\vartheta)(S) = \tilde{V}(\vartheta)(S) = \vartheta_0 \cdot S_0 + \vartheta \bullet S \quad \mathbb{P}\text{-a.s.} \quad (2.2)$$

We call $V(\vartheta)(S) = \tilde{V}(\vartheta)(S)$ the *value process* of ϑ in the currency unit corresponding to S . We denote the vector space of all self-financing strategies for \mathcal{S} by $L^{\text{sf}}(\mathcal{S})$ or just L^{sf} .

The next result provides two equivalent but formally weaker characterisations of self-financing strategies. The first one, which seems to be new, shows that an investment process is self-financing if and only if its investment value process satisfies an exchange rate consistency condition like (2.1). The second one, which is probably due to Xia and Yan [81, Theorem 2.1],³ shows that the self-financing condition (2.2) is numéraire-independent. A similar calculation as in our proof also appears in Schweizer and Takaoka [73, proof of Proposition 2.7 (ii)].

Lemma 2.5. *Let ϑ be a position process for \mathcal{S} . Then the following are equivalent:*

- (a) ϑ is a self-financing strategy for \mathcal{S} .
- (b) ϑ is an investment process for all $S \in \mathcal{S}$, and the investment value process $\tilde{V}(\vartheta)$ satisfies for all $S \in \mathcal{S}$ and all $D \in \mathcal{D}$ the exchange rate consistency condition

$$\tilde{V}(\vartheta)(DS) = D\tilde{V}(\vartheta)(S) \quad \mathbb{P}\text{-a.s.} \quad (2.3)$$

³The economic statement as such is of course much older. For continuous price processes, the proof can be traced back at least to El Karoui et al. [17, Proposition 1].

(c) ϑ is an investment process for some $S \in \mathcal{S}$, and $V(\vartheta)(S) = \tilde{V}(\vartheta)(S)$ \mathbb{P} -a.s.

Proof. “(a) \Rightarrow (b)”. This follows immediately from (2.2) and (2.1).

“(b) \Rightarrow (c)”. Fix $S \in \mathcal{S}$. Then for all $D \in \mathcal{D}$, (2.3), the product rule for semimartingales and associativity and linearity of the stochastic integral give

$$\begin{aligned}
0 &= D\tilde{V}(\vartheta)(S) - \tilde{V}(\vartheta)(DS) \\
&= D(\vartheta_0 \cdot S_0 + \vartheta \bullet S) - (\vartheta_0 \cdot (D_0 S_0) + \vartheta \bullet (DS)) \\
&= D_0(\vartheta_0 \cdot S_0) + D_- \bullet (\vartheta \bullet S) + (\vartheta_0 \cdot S_0 + \vartheta \bullet S)_- \bullet D + [D, \vartheta \bullet S] \\
&\quad - D_0(\vartheta_0 \cdot S_0) - \vartheta \bullet (D_- \bullet S + S_- \bullet D + [D, S]) \\
&= (\vartheta_0 \cdot S_0 + \vartheta \bullet S_- - \vartheta \cdot S_-) \bullet D \quad \mathbb{P}\text{-a.s.}
\end{aligned} \tag{2.4}$$

Since (2.4) holds simultaneously for all $D \in \mathcal{D}$, an easy exercise in stochastic analysis shows that $\vartheta_0 \cdot S_0 + \vartheta \bullet S_- - \vartheta \cdot S_- \equiv 0$ \mathbb{P} -a.s., and hence,

$$\begin{aligned}
V(\vartheta)(S) &= \vartheta \cdot S = \vartheta \cdot S_- + \vartheta \cdot \Delta S = \vartheta_0 \cdot S_0 + \vartheta \bullet S_- + \vartheta \cdot \Delta S \\
&= \tilde{V}(\vartheta)(S) \quad \mathbb{P}\text{-a.s.}
\end{aligned} \tag{2.5}$$

“(c) \Rightarrow (a)”. Let $S \in \mathcal{S}$ be such that $\vartheta \in L(S)$ and $V(\vartheta)(S) = \tilde{V}(\vartheta)(S)$ \mathbb{P} -a.s. Then using that $\vartheta \cdot \Delta S = \Delta(\vartheta \bullet S)$, the same calculation as in (2.5) gives

$$\vartheta_0 \cdot S_0 + (\vartheta \bullet S)_- = \vartheta \cdot S_- \quad \mathbb{P}\text{-a.s.} \tag{2.6}$$

Let $S' \in \mathcal{S}$ be arbitrary and $D \in \mathcal{D}$ such that $S' = DS$. Now the claim follows from the product rule for semimartingales, (2.1) and (2.6). \square

3 Numéraire strategies and representatives

Self-financing strategies are conceptually the good objects to describe trading in our framework. But how does one check in practice that an investment process ζ for some $S \in \mathcal{S}$ is a self-financing strategy? Using Definition 2.4 or Lemma 2.5, this is a very hard task in general since we have to check (2.2) for all stopping times $\tau \in \mathcal{T}_{[0,T]}$. In this section, we provide a satisfactory solution to this problem—the main ingredient being the central concepts of numéraire strategies and numéraire representatives.

Definition 3.1. A *numéraire strategy* for the market \mathcal{S} is a self-financing strategy $\eta \in L^{\text{sf}}$ satisfying $\inf_{t \in [0,T]} V_t(\eta)(S) > 0$ for some (and hence every) $S \in \mathcal{S}$, i.e., $V_t(\eta)(S)$ is an exchange rate process for some (and hence every) $S \in \mathcal{S}$. If a numéraire strategy for \mathcal{S} exists, \mathcal{S} is called a *numéraire market*.

Remark 3.2. For a nonnegative market \mathcal{S} , the *market portfolio* $\eta^{\mathcal{S}} := (1, \dots, 1)$ is always a numéraire strategy. This follows immediately from (1.2). In general, however, there need not exist a numéraire strategy. Indeed, consider the one-period market \mathcal{S} generated by $S = (S_k^1, S_k^2)_{k \in \{0,1\}}$, where $S_0^1 = S_0^2 = 1$ and

$S_1^1 = N^1$ and $S_1^2 = N^2$ are independent standard normal random variables.⁴ Then any $\vartheta \in L^{\text{sf}}$ can be identified with a vector $(\vartheta^1, \vartheta^2) \in \mathbb{R}^2$, and the corresponding value process $V(\vartheta)$ satisfies

$$V_1(\vartheta)(S) = \vartheta^1 N^1 + \vartheta^2 N^2.$$

Since $\vartheta^1 N^1 + \vartheta^2 N^2$ is normally distributed with mean 0 and variance $(\vartheta^1)^2 + (\vartheta^2)^2$, it cannot be positive.

A numéraire strategy is the currency-unit-independent counterpart of a “numéraire” in the standard framework. Indeed, in the classic setup, one first fixes a currency unit and then calls any positive value process (in that unit) of a self-financing strategy a “numéraire”. So in our terminology, a classic “numéraire” is the value process of a numéraire strategy in a fixed currency unit. So oddly enough, the classic notion of “numéraire” is not numéraire-independent (in the sense of currency-unit-independent) because it is based on the initial choice of a currency unit. By contrast, our notion of numéraire strategy clearly is currency-unit-independent.

To each numéraire strategy η , there corresponds a unique *numéraire representative* $S^{(\eta)} \in \mathcal{S}$ in which units the value process of η is identically 1.

Proposition 3.3. *Let \mathcal{S} be numéraire market and η a numéraire strategy. Then there exists a \mathbb{P} -a.s. unique representative $S^{(\eta)} \in \mathcal{S}$ satisfying*

$$V(\eta)(S^{(\eta)}) \equiv 1 \text{ } \mathbb{P}\text{-a.s.}$$

Starting from any $S \in \mathcal{S}$, it can be computed by

$$S^{(\eta)} = \frac{S}{V(\eta)(S)} \text{ } \mathbb{P}\text{-a.s.} \quad (3.1)$$

Moreover, if \mathcal{S} is continuous, $S^{(\eta)}$ has \mathbb{P} -a.s. continuous paths.

Proof. Existence and uniqueness of $S^{(\eta)}$ follow from (3.1) and the exchange rate consistency (2.1) of $V(\eta)$. If \mathcal{S} is continuous, there exists a representative $S \in \mathcal{S}$ which has \mathbb{P} -a.s. continuous paths. Then $V(\eta)(S)$ has \mathbb{P} -a.s. continuous paths, and by (3.1), $S^{(\eta)}$ has so, too. \square

For some results, it is important that a numéraire market does not only contain some arbitrary numéraire strategy but one that is *bounded* or *simple predictable* and whose numéraire representative is bounded, too.

Definition 3.4. A market \mathcal{S} is called a *bounded numéraire market* if there exists a bounded numéraire strategy η such that also $S^{(\eta)}$ is bounded. It is called a *simple predictable numéraire market* if in addition η can be chosen simple predictable.

⁴This can of course be easily embedded into our general setup.

A nonnegative market \mathcal{S} is always a simple predictable numéraire market. Indeed, the market portfolio $\eta^{\mathcal{S}} = (1, \dots, 1)$ is a simple predictable numéraire strategy by (1.2) and its (nonnegative) numéraire representative $S^{(\eta)}$ satisfies $\sum_{i=1}^N S^{(\eta),i} = 1$, where N is the dimension of the market. Hence, $\|S^{(\eta)}\| \leq 1$ \mathbb{P} -a.s., where $\|\cdot\|$ denotes the maximum norm in \mathbb{R}^N .

Remark 3.5. Even if there exists a simple predictable or bounded numéraire strategy η , \mathcal{S} need not be a simple predictable or bounded numéraire market because its numéraire representative $S^{(\eta)}$ may fail to be bounded. Indeed, consider the one-period market \mathcal{S} generated by a classic model $S = (1, X_k)_{k \in \{0,1\}}$, where $X_0 = 1$ and $X_1 = N$ is a standard normal random variable. Then any $\vartheta \in L^{\text{sf}}$ can be identified with a random vector $(\vartheta^1, \vartheta^2) \in \mathbb{R}^2$, and its corresponding value process $V(\vartheta)$ satisfies

$$V_1(\vartheta)(S) = \vartheta^1 + \vartheta^2 N.$$

Since $\vartheta^1 + \vartheta^2 N$ is normally distributed with mean ϑ^1 and variance $(\vartheta^2)^2$, the value process $V(\vartheta)(S)$ is a exchange rate process if and only if $\vartheta^1 > 0$ and $\vartheta^2 = 0$. Hence, every numéraire strategy is of the form $\eta = (\eta^1, 0)$ and in particular simple predictable, but by (3.1), its numéraire representative $S^{(\eta)}$ satisfies

$$S_1^{(\eta),2} = \frac{S_1^2}{V_1(\eta)(S)} = \frac{X_1}{\eta^1} = \frac{N}{\eta^1} \text{ P-a.s.},$$

and is therefore unbounded.

The following easy but important example shows that a classic model (1.3) is nothing else than a special case of a numéraire representative.

Example 3.6. Let \mathcal{S} be the market generated by a classic model $S = (1, X)$. Then $e_1 := (1, 0, \dots, 0)$, the buy-and-hold strategy of the “bank account”, is a numéraire strategy for \mathcal{S} since $V(e_1)(S) = 1$ \mathbb{P} -a.s. Moreover, by (3.1),

$$S^{(e_1)} = S = (1, X).$$

We proceed to establish the main result of this section that given *one* numéraire strategy η , *all* self-financing strategies for \mathcal{S} can be easily described. This kind of result is well known in the standard framework, i.e., for markets \mathcal{S} generated by a classic model $S = (1, X)$ and the numéraire strategy $e_1 = (1, 0, \dots, 0)$; see Corollary 3.8 (b).

Theorem 3.7. *Let \mathcal{S} be a numéraire market and η a numéraire strategy. Then for all investment processes ζ for $S^{(\eta)}$, there exists $\vartheta \in L^{\text{sf}}$ such that*

$$\tilde{V}(\vartheta)(S^{(\eta)}) = \vartheta_0 \cdot S_0^{(\eta)} + \vartheta \bullet S^{(\eta)} = \zeta_0 \cdot S_0^{(\eta)} + \zeta \bullet S^{(\eta)} = \tilde{V}(\zeta)(S^{(\eta)}) \text{ P-a.s.}$$

Moreover, if η , $S^{(\eta)}$, ζ and $\zeta \bullet S^{(\eta)}$ are bounded, ϑ can be chosen bounded, and if η and ζ are simple predictable and $S^{(\eta)}$ is bounded, ϑ can be chosen simple predictable.

Proof. Fix an investment process ζ for $S^{(\eta)}$. Writing $\zeta = \zeta' + \zeta''$ with $\zeta' := \zeta_0$ and $\zeta'' := \zeta - \zeta_0$ and noting that constant processes are trivially self-financing shows that we may assume without loss of generality that $\zeta_0 = 0$. For $0 \leq t \leq T$, set $\zeta_t^0 := (\zeta \bullet S^{(\eta)})_t - \zeta_t \cdot S_t^{(\eta)}$. Then ζ^0 is predictable since

$$\zeta^0 = (\zeta \bullet S^{(\eta)})_- + \Delta(\zeta \bullet S^{(\eta)}) - \zeta \cdot S^{(\eta)} = (\zeta \bullet S^{(\eta)})_- - \zeta \cdot S_-^{(\eta)} \text{ P-a.s.}$$

Moreover, $\eta \bullet S^{(\eta)} = V(\eta)(S^{(\eta)}) - \eta_0 \cdot S_0^{(\eta)} \equiv 1 - \eta_0 \cdot S_0^{(\eta)}$ P-a.s. trivially gives $\zeta^0 \in L(\eta \bullet S^{(\eta)})$, and so $\zeta^0 \eta \in L(S^{(\eta)})$ by associativity of the stochastic integral. Set $\vartheta := \zeta + \zeta^0 \eta$. Then ϑ is an investment process for $S^{(\eta)}$ with $\vartheta_0 = 0$ and satisfies

$$\begin{aligned} \tilde{V}(\vartheta)(S^{(\eta)}) &= \vartheta \bullet S^{(\eta)} = \zeta \bullet S^{(\eta)} + (\zeta^0 \eta) \bullet S^{(\eta)} = \zeta \bullet S^{(\eta)} + \zeta^0 \bullet (1 - \eta_0 \cdot S_0^{(\eta)}) \\ &= \zeta \bullet S^{(\eta)} = \zeta^0 + \zeta \cdot S^{(\eta)} = \zeta^0 (\eta \cdot S^{(\eta)}) + \zeta \cdot S^{(\eta)} = \vartheta \cdot S^{(\eta)} \\ &= V(\vartheta)(S^{(\eta)}) \text{ P-a.s.} \end{aligned}$$

By Lemma 2.5, ϑ is a self-financing strategy for \mathcal{S} . Moreover, if η , $S^{(\eta)}$, ζ and $\zeta \bullet S^{(\eta)}$ are bounded, then by construction, ζ^0 and ϑ are bounded, and if η and ζ are simple predictable processes and $S^{(\eta)}$ is bounded, it is an easy exercise to show that ζ^0 and ϑ are simple predictable. \square

The following corollary links the above result to the standard framework. To this end, recall from Example 3.6 that a classic model $S = (1, X)$ is nothing else than the numéraire representative $S^{(e_1)}$, where $e_1 = (1, 0, \dots, 0)$ is the buy-and-hold strategy of the “bank account”.

Corollary 3.8. *Let \mathcal{S} be the market generated by a classic model $S = (1, X)$.*

- (a) *For each investment process ζ for $S = S^{(e_1)}$, there exists a self-financing strategy ϑ such that*

$$\tilde{V}(\vartheta)(S) = \tilde{V}(\zeta)(S) \text{ P-a.s.} \quad \text{and} \quad \vartheta^i = \zeta^i, \quad i = 2, \dots, N = d + 1.$$

- (b) *For each $v_0 \in \mathbb{R}$ and each \mathbb{R}^d -valued predictable process $\zeta \in L(X)$, there exists a real-valued predictable process ϑ^1 such that $\vartheta = (\vartheta^1, \zeta)$ is a self-financing strategy for \mathcal{S} and satisfies*

$$V(\vartheta)(S) = v_0 + \zeta \bullet X \text{ P-a.s.}$$

Corollary 3.8 (b) shows that for a market \mathcal{S} that is generated by a classic model $S = (1, X)$, one can identify each self-financing strategy ϑ for \mathcal{S} with a pair (v_0, ζ) consisting of an “initial capital” $v_0 \in \mathbb{R}$ in the currency unit corresponding to $S = S^{(e_1)}$ and a predictable process ζ which is integrable with respect to the “discounted risky” assets X . This identification is tacitly done throughout most of the literature on mathematical finance, and sometimes—in an abuse of notation—the d -dimensional process ζ itself is called “self-financing”. But this hides the fact that one needs apart from the d “risky” assets X a “riskless” asset 1 to implement trading in a self-financing manner.

Remark 3.9. The concept of numéraire strategies and numéraire representatives can be slightly generalised. First, we say that a real-valued semimartingale \overline{D} is a *signed exchange rate process* if $\inf_{t \in [0, T]} |\overline{D}_t| > 0$. Second, if \mathcal{S} is an N -dimensional market, we say that an N -dimensional semimartingale \overline{S} is a *signed representative* of \mathcal{S} if there exist $S \in \mathcal{S}$ and a signed exchange rate process \overline{D} such that $\overline{S} = \overline{D}S$ \mathbb{P} -a.s. We denote the set of all signed representatives of \mathcal{S} by $\overline{\mathcal{S}}$ and call this a *signed market*. It is not difficult to show that the exchange rate consistency of value processes (2.1) and the self-financing property (2.2) extend to signed markets. Third, we say that a self-financing strategy $\overline{\eta} \in L^{\text{sf}}$ is a *signed numéraire strategy* if for some (and hence every) signed representative $\overline{S} \in \overline{\mathcal{S}}$, the value process $V(\overline{\eta})(\overline{S})$ is a signed exchange rate process, and call $\overline{\mathcal{S}}$ a *signed numéraire market* if such an $\overline{\eta}$ exists. It can then be shown that there exist a \mathbb{P} -a.s. unique signed representative $\overline{S}^{(\overline{\eta})} \in \overline{\mathcal{S}}$ such that $V(\overline{\eta})(\overline{S}^{(\overline{\eta})}) \equiv 1$ \mathbb{P} -a.s. and that Theorem 3.7 carries over to signed numéraire markets.

4 Strategy cones, undefaultable strategies and admissible investment processes

As in the standard framework, to exclude doubling phenomena if there are infinitely many trading dates, we have to take proper subsets of self-financing strategies to describe “allowed” trading.

Definition 4.1. A *strategy cone* for the market \mathcal{S} is a convex cone $\Gamma \subset L^{\text{sf}}$ containing 0. If Γ is a strategy cone, we denote by $\mathbf{b}\Gamma$ and $\mathbf{s}\Gamma$ the subcone of all bounded and simple predictable strategies in Γ , respectively.

Basic examples of strategy cones include buy-and-hold strategies or self-financing strategies with some short-selling constraint, e.g. all $\vartheta \in L^{\text{sf}}$ satisfying $\vartheta^i \geq 0$ for $1 \leq i \leq N$. Another prime example is given in the following definition.

Definition 4.2. An *undefaultable strategy* is a self-financing strategies $\vartheta \in L^{\text{sf}}$ satisfying $V(\vartheta) \geq 0$ \mathbb{P} -a.s. The strategy cone of all such strategies is denoted by $\mathcal{U}(\mathcal{S})$ or just \mathcal{U} .

\mathcal{U} is the “natural” strategy cone for studying no-arbitrage in a numéraire-independent setup; see Chapter III.3.3. This has already been noted by El Karoui et al. [17].

Next, we introduce $S^{(n)}$ -*admissible investment processes*, a generalisation of the classic key concept of *admissible strategies*. This is clearly not a numéraire-independent notion; see also Proposition 4.4 below.

Definition 4.3. Let \mathcal{S} be a numéraire market and η a numéraire strategy. An investment process ζ for $S^{(n)}$ with $\zeta_0 = 0$ is called an a - $S^{(n)}$ -*admissible investment process*, where $a \geq 0$, if

$$\zeta \bullet S^{(n)} \geq -a \text{ } \mathbb{P}\text{-a.s.} \quad (4.1)$$

We write $\zeta \in L^{\text{ad}}(S^{(n)}, a)$. The union of all $L^{\text{ad}}(S^{(n)}, a)$ over all $a \geq 0$ is denoted by $L^{\text{ad}}(S^{(n)})$, and its elements are called $S^{(n)}$ -admissible investment processes.

If $S = (1, X)$ is a classic model, then a (classic) *admissible strategy* is a d -dimensional predictable process $\zeta \in L(X)$ with $\zeta_0 = 0$ such that $\zeta \bullet X \geq -a$ \mathbb{P} -a.s. for some $a \geq 0$; see e.g. Delbaen and Schachermayer [14, Section 8.1]. From a purely mathematical point of view, if \mathcal{S} is a d -dimensional (sic!) market, $S^{(n)} \in \mathcal{S}$ a numéraire representative and ζ an $S^{(n)}$ -admissible investment process, then ζ is a classic admissible strategy for the $(d+1)$ -dimensional classic model $S = (1, X)$ with $(X^1, \dots, X^d) = (S^{(n),1}, \dots, S^{(n),d})$. Thus, all results for classic admissible strategies carry over to $S^{(n)}$ -admissible investment processes. From an economic perspective, however, there is a fundamental difference: In the standard framework, there is the “riskless” *traded* asset 1 in the background, and so every classic admissible strategy ζ can be extended to a $(d+1)$ -dimensional *self-financing* strategy (ϑ^1, ζ) ; see Corollary 3.8. By contrast, for $S^{(n)}$ -admissible investment processes, there is no traded constant asset in the background, and so a direct extension to self-financing strategies is not possible—an indirect identification, however, is.

Proposition 4.4. *Let \mathcal{S} be a numéraire market and η a numéraire strategy. Then for each $\zeta \in L^{ad}(S^{(n)})$, there exists $\vartheta \in L^{sf}$ (with $\vartheta_0 \cdot S_0^{(n)} = 0$) such that*

$$V(\vartheta)(S^{(n)}) = \vartheta \bullet S^{(n)} = \zeta \bullet S^{(n)} \quad \mathbb{P}\text{-a.s.} \quad (4.2)$$

Moreover, for all $a \geq 0$ with $\zeta \in L^{ad}(S^{(n)}, a)$, there exists $\tilde{\vartheta} \in \mathcal{U}$ such that $\vartheta = \tilde{\vartheta} - a\eta$.

We interpret $\tilde{\vartheta}$ as the “numéraire-independent” and $a\eta$ as the “numéraire-dependent” part of ζ .

Proof. Fix $\zeta \in L^{ad}(S^{(n)})$. Then $\zeta \in L(S^{(n)})$ and Theorem 3.7 gives $\vartheta \in L^{sf}$ satisfying (4.2). Let $a \geq 0$ be such that $\zeta \bullet S^{(n)} \geq -a$ \mathbb{P} -a.s. and set $\tilde{\vartheta} := \vartheta + a\eta$. Then $\tilde{\vartheta} \in L^{sf}$. Since $V(\eta)(S^{(n)}) \equiv 1$, the claim follows from (4.2) and the choice of a via

$$V(\tilde{\vartheta})(S^{(n)}) = V(\vartheta)(S^{(n)}) + aV(\eta)(S^{(n)}) = \zeta \bullet S^{(n)} + a \geq 0 \quad \mathbb{P}\text{-a.s.} \quad \square$$

The following corollary shows that undefaultable strategies implicitly appear rather frequently in the literature on mathematical finance, e.g. in the context of utility maximisation on \mathbb{R}_+ .

Corollary 4.5. *Let \mathcal{S} be a numéraire market, η a numéraire strategy, $v_0 \geq 0$ and $\zeta \in L^{ad}(S^{(n)})$. Suppose that $v_0 + \zeta \bullet S^{(n)} \geq 0$ \mathbb{P} -a.s. Then there is $\vartheta \in \mathcal{U}$ such that*

$$V(\vartheta)(S^{(n)}) = \vartheta_0 \cdot S_0^{(n)} + \vartheta \bullet S^{(n)} = v_0 + \zeta \bullet S^{(n)} \quad \mathbb{P}\text{-a.s.}$$

5 Contingent claims

In this section, we introduce a numéraire-independent notion of *contingent claims* for a market \mathcal{S} . Moreover, we look at the concept of *numéraire-independent*

derivative securities and show that even plain vanilla call options are in general not in that class.

A *contingent claim* due at time τ pays “something” at time τ depending on the state of the world at time τ . Therefore, a contingent claim is usually modelled as a nonnegative \mathcal{F}_τ -measurable random variable. This, however, implicitly assumes that the units in which “something” is measured are known, i.e., that some representative $S \in \mathcal{S}$ is fixed. But as our paradigm consists in not fixing any representative, the definition of a contingent claim in our framework requires some thought.

Recall that $\bar{\mathbf{L}}_+^0$, $\underline{\mathbf{L}}_+^0$ and $\bar{\mathbf{L}}_+^0$ denote the set of all random variables taking values in $\bar{\mathbf{R}}_+ = \mathbf{R}_+ \cup \{+\infty\}$, $\underline{\mathbf{R}}_+ = \mathbf{R}_+ \cup \{-\infty\}$ and $\bar{\mathbf{R}}_+ = \mathbf{R}_+ \cup \{+\infty\} \cup \{-\infty\}$, respectively, and that we have the conventions $c + \infty = \infty + c = +\infty$ for $c \in \bar{\mathbf{R}}_+$, $-\infty + c = c - \infty = -\infty$ for $c \in \bar{\mathbf{R}}_+$, $c \times (\pm\infty) = \pm\infty$ for $c \in \mathbf{R}_{++}$, and $0 \times (\pm\infty) = 0$.

Definition 5.1. A *generalised contingent claim* at time $\tau \in \mathcal{T}_{[0,T]}$ for the market \mathcal{S} is a map $F : \mathcal{S} \rightarrow \bar{\mathbf{L}}_+^0(\mathcal{F}_\tau)$ satisfying for all $S \in \mathcal{S}$ and all $D \in \mathcal{D}$ the *exchange rate consistency condition*

$$F(DS) = D_\tau F(S) \quad \mathbb{P}\text{-a.s.} \quad (5.1)$$

F is called an *improper contingent claim* if it is valued in $\bar{\mathbf{L}}_+^0(\mathcal{F}_\tau)$, it is called a *defaultable contingent claim* if it is valued in $\underline{\mathbf{L}}_+^0(\mathcal{F}_\tau)$, it is called a *contingent claim* if it is valued in $\mathbf{L}_+^0(\mathcal{F}_\tau)$, and it is called a *positive contingent claim* if it is valued in $\mathbf{L}_{++}^0(\mathcal{F}_\tau)$.

Remark 5.2. (a) The notions of generalised, improper and defaultable contingent claims are introduced for technical reasons; from an economic perspective, only (positive) contingent claims are relevant.

(b) It is easy to check that the set of all (generalised, improper, defaultable) contingent claims at time τ is a convex cone containing 0.

(c) The idea that a contingent claim is a map satisfying an exchange rate consistency condition is implicitly also behind the definition of a contingent claim in Carr et al. [7]; cf. Remark 1.5.

(d) We often write $F \geq 0$ \mathbb{P} -a.s. or $\mathbb{P}[F > 0] > 0$, etc., as a shorthand for $F(S) \geq 0$ \mathbb{P} -a.s. for all $S \in \mathcal{S}$, or $\mathbb{P}[F(S) > 0] > 0$ for all $S \in \mathcal{S}$, etc.; cf. Remark 2.2.

A generalised contingent claim is uniquely characterised by a representative $S \in \mathcal{S}$ (fixing a currency unit) and a random variable g (describing the payoff structure in that unit). This is easy.

Proposition 5.3. Let \mathcal{S} be a market and $\tau \in \mathcal{T}_{[0,T]}$. For any pair (S, g) , where $S \in \mathcal{S}$ and $g \in \bar{\mathbf{L}}_+^0(\mathcal{F}_\tau)$, there exists a \mathbb{P} -a.s.-unique generalised contingent claim F at time τ satisfying $F(S) = g$ \mathbb{P} -a.s. It is given by $F(DS) := D_\tau g$, where $D \in \mathcal{D}$ is arbitrary.

Directly linked to the notion of a contingent claim is the notion of a *derivative security*. Loosely speaking, a derivative security is a contingent claim which can be described by some *payoff function* $h : \mathbb{R}^N \rightarrow \mathbb{R}^+$ only.

Definition 5.4. A contingent claim F at time $\tau \in \mathcal{T}_{[0,T]}$ for the (N -dimensional) market \mathcal{S} is called a *numéraire-independent derivative security* if there exists a measurable function $h : \mathbb{R}^N \rightarrow \mathbb{R}_+$, called the *payoff function* of F , such that

$$F(S) = h(S_\tau) \quad \text{for all } S \in \mathcal{S}.$$

It follows immediately from the exchange rate consistency condition (5.1) and Proposition 5.3 that F is a numéraire-independent derivative security if and only if h is *positively homogeneous of degree 1*, i.e., $h(\lambda x) = \lambda h(x)$ for all $\lambda > 0$, $x \in \mathbb{R}^N$. Note that F and h are mathematically completely different objects— F goes from processes to random variables while h goes from \mathbb{R}^N to \mathbb{R}_+ .

Even plain vanilla call options are only in special cases numéraire-independent derivative securities. Indeed, let \mathcal{S} be the market generated by $S = (S_t^1, S_t^2)_{t \in [0,T]}$, where $S^1 := \exp(\int_0^\cdot r_s ds)$ describes the evolution of a bank account in EUR with some (possibly stochastic) short rate process $(r_t)_{t \in [0,T]}$ and S^2 models the evolution of a stock in EUR. A *European call option* written on S^2 with strike $K > 0$ (denoted in EUR) and maturity T is the contingent claim F at time T from Proposition 5.3 satisfying

$$F(S) = (S_T^2 - K)^+ = (S_T^2 - \tilde{K} S_T^1)^+,$$

where $\tilde{K} := K \exp(-\int_0^T r_s ds)$. Clearly, F is a numéraire-independent derivative security if and only if \tilde{K} is *deterministic*, e.g. if the short rate process r is deterministic. In that case, the call option can be interpreted as an *exchange option*, where one can exchange \tilde{K} units of S^1 against 1 unit of S^2 at time T if this gives a profit. This insight is already apparent in the paper on exchange options by Margrabe [61]. The key point here, however, is that in general we cannot describe the call option by only specifying the strike K as a *number*—we also have to specify the *units* of that number.

6 Superreplication prices

When considering a contingent claim F at time $\tau \in \mathcal{T}_{[0,T]}$, one fundamental question is how to attach to F a *price* today.⁵ Loosely speaking, a “price” is a number expressed in some currency unit. If we change the currency unit, then the number will change, notwithstanding that both numbers describe the same economic entity (“price of F at time 0”). If we consider this economic entity

⁵The term *price* suggests that it is objective or preference-independent. In incomplete markets, however, it is well known that an objective “price” of a contingent claim does not exist in general. Therefore, one should better speak of a *value* instead of *the price* of a contingent claim. Nevertheless, in order to be consistent with the existing literature, we have chosen to use the somewhat misleading word “price” instead of the correct term “value”.

as “price”, then it is natural to define “price” as a mapping from representatives (fixing a currency unit) to nonnegative real numbers satisfying an exchange rate consistency condition like (5.1) at time 0. Put differently, a (time 0) “price” is nothing else than a contingent claim at time 0.

We study two notions of superreplication prices in our framework: *ordinary* and *limit quantile* superreplication prices—quantile superreplication prices are mainly an intermediate concept, introduced for technical convenience. Ordinary superreplication prices are the numéraire-independent counterpart of superreplication prices in the standard framework, whereas limit quantile superreplication prices are based on a slightly relaxed version of the concept of superreplication, which has nicer continuity properties than the ordinary concept of superreplication. Limit quantile superreplication prices are also the cornerstone of our notion of (no-)arbitrage in Chapter III.

6.1 Ordinary superreplication prices

Definition 6.1. Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. The *ordinary superreplication price* of F for Γ is the map $\Pi(F | \Gamma) : \mathcal{S} \rightarrow [0, \infty]$ given by

$$\begin{aligned} \Pi(F | \Gamma)(S) = \inf \{ v \geq 0 : \text{there exists } \vartheta \in \Gamma \text{ with} \\ V_0(\vartheta)(S) = v \text{ and } V_\tau(\vartheta)(S) \geq F(S) \text{ P-a.s.} \} \end{aligned}$$

If $S \in \mathcal{S}$ describes the evolution of asset prices in EUR, then $\Pi(F | \Gamma)(S)$ is the classic superreplication price in EUR of the classic contingent claim $F(S)$ in EUR.

Remark 6.2. It is easy to check that $\Pi(F | \Gamma)$ is an improper contingent claim at time 0. So in view of Proposition 5.3, it suffices to compute $\Pi(F | \Gamma)(S)$ for *one* $S \in \mathcal{S}$. In particular, $\Pi(F | \Gamma)(S) = \Pi(F | \Gamma)(\tilde{S})$ if $S_0 = \tilde{S}_0$, which suggests that when looking for a *dual* characterisation of $\Pi(F | \Gamma)(S)$ for *some* fixed $S \in \mathcal{S}$, we should consider *all* $\tilde{S} \in \mathcal{S}$ with $\tilde{S}_0 = S_0$; see Chapter VI.3.

The next two results provide a useful characterisation of ordinary superreplication prices in terms of strategies and collect their basic properties. The proofs are easy and hence omitted.

Proposition 6.3. *Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ with $\Pi(F | \Gamma) < \infty$. Then for each $\delta > 0$ and each positive contingent claim C at time 0, there is $\vartheta \in \Gamma$ satisfying*

$$V_0(\vartheta) \leq \Pi(F | \Gamma) + \delta C \quad \text{and} \quad V_\tau(\vartheta) \geq F \text{ P-a.s.}$$

Proposition 6.4. *Let \mathcal{S} be a market, Γ a strategy cone, F, F_1, F_2, G generalised contingent claims at time $\tau \in \mathcal{T}_{[0,T]}$ with $F \leq G$ P-a.s. and $\lambda \geq 0$. Then*

$$\begin{aligned} \Pi(F | \Gamma) &\leq \Pi(G | \Gamma) && \text{(monotonicity),} \\ \Pi(\lambda F | \Gamma) &= \lambda \Pi(F | \Gamma) && \text{(positive homogeneity),} \\ \Pi(F_1 + F_2 | \Gamma) &\leq \Pi(F_1 | \Gamma) + \Pi(F_2 | \Gamma) && \text{(subadditivity).} \end{aligned}$$

Remark 6.5. The above properties of ordinary superreplication prices are reminiscent of coherent risk measures; only cash-invariance is missing. Indeed, for fixed $S \in \mathcal{S}$, $\tau \in \mathcal{T}_{[0,T]}$ and Γ , consider the map $\rho : \mathbf{L}_-^0(\mathcal{F}_\tau) \rightarrow [0, \infty]$ defined by $\rho(X) = \Pi_\tau(F | \Gamma)(S)$, where F is the contingent claim at time τ from Proposition 5.3 satisfying $F(S) = -X$ \mathbb{P} -a.s. Then ρ is a coherent *capital requirement*, a generalisation of risk measures introduced by Frittelli and Scandolo [23].

6.2 Quantile superreplication prices

The ordinary superreplication price $\Pi(F | \Gamma)$ of a contingent claim F is the cheapest initial “capital”—to be understood in a numéraire-independent sense—that is needed to find a strategy ϑ in a given class Γ whose value process dominates F at maturity τ *almost surely*. If we relax the latter condition and only require $V(\vartheta)$ to dominate F with a high probability, we are led to the concept of *quantile superreplication prices*.

Definition 6.6. Let \mathcal{S} be a market, Γ a strategy cone, F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ and $\varepsilon \in (0, 1)$. The ε -*quantile superreplication price* of F for Γ is the map $\Pi^\varepsilon(F | \Gamma) : \mathcal{S} \rightarrow [0, \infty]$ given by

$$\Pi^\varepsilon(F | \Gamma)(S) = \inf\{v \geq 0 : \text{there exists } \vartheta \in \Gamma \text{ such that} \\ V_0(\vartheta)(S) = v \text{ and } \mathbb{P}[V_\tau(\vartheta)(S) \geq F(S)] \geq 1 - \varepsilon\}.$$

If $S \in \mathcal{S}$ describes the evolution of asset prices in EUR, then $\Pi^\varepsilon(F | \Gamma)(S)$ may be interpreted as the price in EUR of a *quantile superhedge* for $F(S)$ with *shortfall risk* ε . Regarding the notion of *quantile hedging*, we refer to Föllmer and Schied [25, Section 8.1].

Remark 6.7. (a) It is easy to check that $\Pi^\varepsilon(F | \Gamma)$ is an improper contingent claim at time 0.

(b) One conceptual drawback of the notion of quantile superreplication prices is that only the *probability* that F is superreplicated but not the *size* of the *shortfall* $F - V_\tau(\vartheta)$ in case that F is not superreplicated is taken into account.

The following result links quantile superreplication prices to ordinary superreplication prices. The proof is easy and hence omitted.

Proposition 6.8. *Let \mathcal{S} be a market, Γ a strategy cone, F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ and $\varepsilon \in (0, 1)$.*

(a) *For all $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon$,*

$$\Pi^\varepsilon(F | \Gamma) \leq \Pi(F - \infty \mathbb{1}_A | \Gamma).$$

(b) *For each $\delta > 0$ and each positive contingent claim C at time 0, there exists $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon$ such that*

$$\Pi(F - \infty \mathbb{1}_A | \Gamma) \leq \Pi^\varepsilon(F | \Gamma) + \delta C.$$

Moreover, if $\Gamma \subset \mathcal{U}$, $\Pi(F - \infty \mathbf{1}_A | \Gamma)$ can be replaced by $\Pi(F \mathbf{1}_{A^c} | \Gamma)$ in both assertions.

For a fixed $\varepsilon \in (0, 1)$, quantile superreplication prices are clearly *monotonic* and *positively homogeneous*. But unlike classic superreplication prices they are not *subadditive*. However, if we do not insist on fixing ε , we get some sort of subadditivity.

Proposition 6.9. *Let \mathcal{S} be a market, Γ a strategy cone, F, F_1, F_2, G generalised contingent claims at time $\tau \in \mathcal{T}_{[0, T]}$ with $F \leq G$ \mathbb{P} -a.s. and $\lambda \geq 0$. Then for all $\varepsilon \in (0, 1)$,*

$$\begin{aligned} \Pi^\varepsilon(F | \Gamma) &\leq \Pi^\varepsilon(G | \Gamma) && \text{(monotonicity),} \\ \Pi^\varepsilon(\lambda F | \Gamma) &= \lambda \Pi^\varepsilon(F | \Gamma) && \text{(positive homogeneity),} \\ \Pi^\varepsilon(F_1 + F_2 | \Gamma) &\leq \Pi^{\varepsilon/2}(F_1 | \Gamma) + \Pi^{\varepsilon/2}(F_2 | \Gamma) && \text{(semi-subadditivity).} \end{aligned}$$

Proof. We only prove semi-subadditivity. To this end, we may assume without loss of generality that $\Pi^{\varepsilon/2}(F_1 | \Gamma), \Pi^{\varepsilon/2}(F_2 | \Gamma) < \infty$. Let C be a positive contingent claim at time 0 and $\delta > 0$. By Proposition 6.8 (b), there exist $A_1, A_2 \in \mathcal{F}_\tau$ with $\mathbb{P}[A_i] \leq \varepsilon/2$ such that

$$\Pi(F_i - \infty \mathbf{1}_{A_i} | \Gamma) \leq \Pi^{\varepsilon/2}(F_i | \Gamma) + \frac{\delta}{2} C.$$

Set $A := A_1 \cup A_2$. Then $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon$. By Propositions 6.8 (a) and 6.4,

$$\begin{aligned} \Pi^\varepsilon(F_1 + F_2 | \Gamma) &\leq \Pi((F_1 + F_2) - \infty \mathbf{1}_A | \Gamma) \\ &= \Pi((F_1 - \infty \mathbf{1}_{A_1}) + (F_2 - \infty \mathbf{1}_{A_2}) | \Gamma) \\ &\leq \Pi(F_1 - \infty \mathbf{1}_{A_1} | \Gamma) + \Pi(F_2 - \infty \mathbf{1}_{A_2} | \Gamma) \\ &\leq \Pi^{\varepsilon/2}(F_1 | \Gamma) + \Pi^{\varepsilon/2}(F_2 | \Gamma) + \delta C. \end{aligned}$$

Letting $\delta \searrow 0$ establishes the claim. \square

6.3 Limit quantile superreplication prices

Quantile superreplication prices are clearly *decreasing* in ε . So it is natural to consider the limit $\lim_{\varepsilon \searrow 0} \Pi^\varepsilon(F | \Gamma)$. From an economic perspective, this captures the idea of superreplication with an “arbitrary high probability” (cf. Corollary 6.15).

Definition 6.10. Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0, T]}$. The *limit quantile superreplication price* of F for Γ is the map $\Pi^*(F | \Gamma) : \mathcal{S} \rightarrow [0, \infty]$ given by

$$\Pi^*(F | \Gamma) = \lim_{\varepsilon \searrow 0} \Pi^\varepsilon(F | \Gamma). \quad (6.1)$$

Remark 6.11. It is not difficult to check that $\Pi^*(F | \Gamma)$ is an improper contingent claim at time 0.

Since quantile superreplication prices are dominated by ordinary superreplication prices, for each contingent F claim at time $\tau \in \mathcal{T}_{[0,T]}$ and each strategy cone Γ , we have the inequality

$$\Pi^*(F | \Gamma) \leq \Pi(F | \Gamma). \quad (6.2)$$

The following example shows that the inequality (6.2) may be *strict* even in the case of a one-period model with countable Ω .

Example 6.12. Let $\Omega := \{\omega_1, \omega_2, \dots\}$, $\mathcal{F} := 2^\Omega$, \mathbb{P} be a probability measure on (Ω, \mathcal{F}) charging every singleton $\{\omega_n\}$, $\mathcal{F}_0 := \{\emptyset, \Omega\}$ and $\mathcal{F}_1 := \mathcal{F}$. Let \mathcal{S} be the market generated by the classic model $S = (1, X_k)_{k \in \{0,1\}}$ where $X_0 = 1$ and $X_1(\omega_n) := 1 + 2^{-n}$. Let $F := V(e_1)$, where $e_1 = (1, 0)$. We claim that $\Pi(F | \mathcal{U}) = V_0(e_1) > 0$ and $\Pi^*(F | \mathcal{U}) = 0$.

For the first claim, note that \mathcal{U} can be identified with the set

$$\left\{ \vartheta \in \mathbb{R}^2 : \vartheta^2 \geq 0 \text{ and } \vartheta^1 \geq -\vartheta^2, \text{ or } \vartheta^2 < 0 \text{ and } \vartheta^1 > -\frac{3}{2}\vartheta^2 \right\}.$$

Let $\vartheta \in \mathcal{U}$ be such that $V_1(\vartheta)(S) \geq F(S)$ \mathbb{P} -a.s. Then

$$\vartheta^1 + \vartheta^2(1 + 2^{-n}) \geq 1, \quad n \in \mathbb{N}.$$

Letting $n \rightarrow \infty$, yields $\vartheta^1 + \vartheta^2 \geq 1$. Thus, $V_0(\vartheta)(S) \geq 1$ and so $\Pi(F | \mathcal{U}) = V_0(e_1)$.

For the second claim, let $\varepsilon \in (0, 1)$. Choose $n \in \mathbb{N}$ large enough that $\mathbb{P}[\{\omega_1, \dots, \omega_n\}] \geq 1 - \varepsilon$ and set $\vartheta^{(n)} := (-2^n, 2^n) \in \mathcal{U}$. Then

$$V_1(\vartheta^{(n)})(S)(\omega_k) = 2^{n-k} \geq 1 = F(S)(\omega_k), \quad k \leq n,$$

and so $\Pi^\varepsilon(F | \mathcal{U}) = 0$. Letting $\varepsilon \searrow 0$ yields $\Pi^*(F | \mathcal{U}) = 0$.

Remark 6.13. The market in Example 6.12 clearly admits *arbitrage* in a very strong sense. Indeed, one can show that in markets which satisfy *numéraire-independent no-arbitrage (NINA)* (see Chapter III.2), ordinary and limit quantile superreplication prices for \mathcal{U} coincide (see [31, Theorem 7.28] and note that NINA is called NGE there).

Despite the last example, ordinary and limit quantile superreplication prices are “close”, in the following sense.

Proposition 6.14. *Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. Then for each $\varepsilon \in (0, 1)$, there exists $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon$ such that*

$$\Pi(F - \infty \mathbf{1}_A | \Gamma) \leq \Pi^*(F | \Gamma).$$

Moreover, if $\Gamma \subset \mathcal{U}$, $\Pi(F - \infty \mathbf{1}_A | \Gamma)$ can be replaced by $\Pi(F \mathbf{1}_{A^c} | \Gamma)$.

Proof. The second claim follows immediately from the first one by noting that $\Pi(F - \infty \mathbb{1}_A | \Gamma) = \Pi(F \mathbb{1}_{A^c} | \Gamma)$ if $\Gamma \subset \mathcal{U}$. For the first claim, we may assume without loss of generality that $\Pi^*(F | \Gamma) < \infty$. Let $\varepsilon \in (0, 1)$ and C be a positive contingent claim at time 0. Then by Proposition 6.8 (b), there exists a sequence $(A_n)_{n \in \mathbb{N}}$ in \mathcal{F}_τ with $\mathbb{P}[A_n] \leq \varepsilon 2^{-n}$ such that

$$\Pi(F - \infty \mathbb{1}_{A_n} | \Gamma) \leq \Pi^{\varepsilon 2^{-n}}(F | \Gamma) + 2^{-n}C \leq \Pi^*(F | \Gamma) + 2^{-n}C.$$

Set $A = \bigcup_{n \in \mathbb{N}} A_n$. Then $A \in \mathcal{F}_\tau$ and $\mathbb{P}[A] \leq \varepsilon$. Moreover, monotonicity of ordinary superreplication prices and the above construction give

$$\Pi(F - \mathbb{1}_A | \Gamma) \leq \Pi(F - \mathbb{1}_{A_n} | \Gamma) \leq \Pi^*(F | \Gamma) + 2^{-n}C.$$

Letting $n \rightarrow \infty$ establishes the claim. \square

The following corollary provides an economic interpretation of limit quantile superreplication prices by showing that if F can be superreplicated in the limit quantile sense, in any currency unit $S \in \mathcal{S}$, with arbitrarily little more initial capital than $\Pi^*(F | \Gamma)(S)$ and with a probability arbitrarily close to 1, there exists a strategy $\vartheta \in \Gamma$ which superreplicates $F(S)$ at time τ with that probability.

Corollary 6.15. *Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0, T]}$ with $\Pi^*(F | \Gamma) < \infty$. Then for each $\varepsilon \in (0, 1)$, each $\delta > 0$ and each positive contingent claim C at time 0, there exists $\vartheta \in \Gamma$ satisfying*

$$V_0(\vartheta) \leq \Pi^*(F | \Gamma) + \delta C \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta) \geq F] \geq 1 - \varepsilon.$$

Proof. Fix ε, δ and C as above. By Proposition 6.14, there exists $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon$ such that $\Pi(F - \infty \mathbb{1}_A | \Gamma) \leq \Pi^*(F | \Gamma)$. Moreover, by Proposition 6.3, there exists $\vartheta \in \Gamma$ satisfying

$$V_0(\vartheta) \leq \Pi(F - \infty \mathbb{1}_A | \Gamma) + \delta C \quad \text{and} \quad V_\tau(\vartheta) \geq F - \infty \mathbb{1}_A \text{ P-a.s.}$$

Since $\mathbb{P}[V_\tau(\vartheta) \geq F] \geq \mathbb{P}[A^c] \geq 1 - \varepsilon$, this establishes the claim. \square

Like ordinary superreplication prices, limit quantile superreplication prices are *monotonic, positively homogeneous and subadditive*.

Proposition 6.16. *Let \mathcal{S} be a market, Γ a strategy cone, F, F_1, F_2, G generalised contingent claims at time $\tau \in \mathcal{T}_{[0, T]}$ with $F \leq G$ P-a.s. and $\lambda \geq 0$. Then*

$$\begin{aligned} \Pi^*(F | \Gamma) &\leq \Pi^*(G | \Gamma) && \text{(monotonicity),} \\ \Pi^*(\lambda F | \Gamma) &= \lambda \Pi^*(F | \Gamma) && \text{(positive homogeneity),} \\ \Pi^*(F_1 + F_2 | \Gamma) &\leq \Pi^*(F_1 | \Gamma) + \Pi^*(F_2 | \Gamma) && \text{(subadditivity).} \end{aligned}$$

Proof. All three assertions follow from Proposition 6.9 by letting $\varepsilon \searrow 0$. \square

Remark 6.17. In the same way as ordinary superreplication prices, limit quantile superreplication prices can be identified with (generalised) coherent risk measures. Indeed, for fixed $S \in \mathcal{S}$, $\tau \in \mathcal{T}_{[0,T]}$ and Γ , the map $\rho : \mathbf{L}^0_{-}(\mathcal{F}_{\tau}) \rightarrow [0, \infty]$ defined by $\rho(X) = \Pi_{\tau}^{*}(F | \Gamma)(S)$, where F is the contingent claim at time τ from Proposition 5.3 satisfying $F(S) = -X$ \mathbb{P} -a.s., is a coherent *capital requirement*; see [23].

Limit quantile superreplication prices enjoy in addition a *monotone convergence property*. This makes them mathematically much nicer than ordinary superreplication prices, which lack this property.

Lemma 6.18. *Let \mathcal{S} be a market, Γ a strategy cone, $(F_n)_{n \in \mathbb{N}}$ an nondecreasing sequence of generalised contingent claims at time $\tau \in \mathcal{T}_{[0,T]}$, and assume that $F := \lim_{n \rightarrow \infty} F_n < \infty$ \mathbb{P} -a.s. Then*

$$\lim_{n \rightarrow \infty} \Pi^{*}(F_n | \Gamma) = \Pi^{*}(F | \Gamma).$$

Proof. By monotonicity of limit quantile superreplication prices, it suffices to show that

$$\liminf_{n \rightarrow \infty} \Pi^{*}(F_n | \Gamma) \geq \Pi^{*}(F | \Gamma). \quad (6.3)$$

Let $\varepsilon \in (0, 1)$ be given. Since each F_n and F are valued in $\overline{\mathbb{R}}_{+}$, it is straightforward to check that

$$\lim_{n \rightarrow \infty} \mathbb{P}[(1 + \varepsilon)F_n \geq F] = 1.$$

Choose $N \in \mathbb{N}$ large enough that $\mathbb{P}[(1 + \varepsilon)F_N < F] \leq \frac{\varepsilon}{2}$, and let B_N in \mathcal{F}_{τ} be such that $B_N = \{(1 + \varepsilon)F_N < F\}$ \mathbb{P} -a.s. By Proposition 6.14, there exists $A_N \in \mathcal{F}_{\tau}$ with $\mathbb{P}[A_N] \leq \frac{\varepsilon}{2}$ such that

$$\Pi^{*}(F_N | \Gamma) \geq \Pi(F_N - \infty \mathbf{1}_{A_N} | \Gamma).$$

Note that $\mathbb{P}[B_N \cup A_N] \leq \varepsilon$. Now, by Propositions 6.16, 6.4 and 6.8 (a),

$$\begin{aligned} \liminf_{n \rightarrow \infty} \Pi^{*}(F_n | \Gamma) &\geq \Pi^{*}(F_N | \Gamma) \geq \Pi(F_N - \infty \mathbf{1}_{A_N} | \Gamma) \\ &= \frac{1}{1 + \varepsilon} \Pi((1 + \varepsilon)F_N - \infty \mathbf{1}_{A_N} | \Gamma) \\ &\geq \frac{1}{1 + \varepsilon} \Pi((1 + \varepsilon)F_N - \infty \mathbf{1}_{A_N \cup B_N} | \Gamma) \\ &\geq \frac{1}{1 + \varepsilon} \Pi(F - \infty \mathbf{1}_{A_N \cup B_N} | \Gamma) \\ &\geq \frac{1}{1 + \varepsilon} \Pi^{\varepsilon}(F | \Gamma). \end{aligned}$$

Letting $\varepsilon \searrow 0$ establishes (6.3). □

Remark 6.19. If we identify limit quantile superreplication prices with generalised coherent risk measures (cf. Remark 6.17), Lemma 6.18 shows that these risk measures are *continuous from above* or equivalently satisfy the *Fatou property*; see [25, Section 4.2] for the definition and implications of the above properties in terms of the *robust representation* of risk measures.

All objects that are defined almost surely are trivially invariant under an *equivalent change of measure*. We proceed to establish the nontrivial fact that limit quantile superreplication prices are invariant under an equivalent change of measure, even though quantile superreplication prices are clearly not invariant under an equivalent change of measure.

In the following result we indicate by a left superscript the measure under which limit quantile and quantile superreplication prices are taken.

Proposition 6.20. *Let \mathcal{S} be a market, Γ a strategy cone, F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T an equivalent probability measure. Then*

$${}^{\mathbb{P}}\Pi^*(F | \Gamma) = {}^{\mathbb{Q}}\Pi^*(F | \Gamma).$$

Proof. By the symmetry of the claim in \mathbb{P} and \mathbb{Q} and the definition of limit quantile superreplication prices, it suffices to show that for all $\varepsilon \in (0, 1)$,

$${}^{\mathbb{P}}\Pi^*(F | \Gamma) \geq {}^{\mathbb{Q}}\Pi^\varepsilon(F | \Gamma).$$

Let $\varepsilon > 0$ be given. Since $\mathbb{Q} \ll \mathbb{P}$ there exists $\delta \in (0, 1)$ such that $\mathbb{Q}[A] \leq \varepsilon$ for all $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \delta$ (see Kallenberg [50, Theorem 7.37]). Moreover, by Proposition 6.14, there exists $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \delta$ such that

$${}^{\mathbb{P}}\Pi(F - \infty \mathbf{1}_A | \Gamma) \leq {}^{\mathbb{P}}\Pi^*(F | \Gamma).$$

Then $\mathbb{Q}[A] \leq \varepsilon$, and by Proposition 6.9 (a), we may conclude that

$${}^{\mathbb{P}}\Pi^*(F | \Gamma) \geq {}^{\mathbb{P}}\Pi(F - \infty \mathbf{1}_A | \Gamma) = {}^{\mathbb{Q}}\Pi(F - \infty \mathbf{1}_A | \Gamma) \geq {}^{\mathbb{Q}}\Pi^\varepsilon(F | \Gamma) \mathbb{P}\text{-a.s.} \quad \square$$

7 Robustness of limit quantile superreplication prices

In this section, we show that under a mild technical condition, limit quantile superreplication prices for *bounded* and general undefaultable strategies coincide. For continuous markets, we prove that even limit quantile superreplication prices for *simple predictable* and general undefaultable strategies coincide.

This *robustness* of limit quantile superreplication is significant in two respects: From an economic perspective, it means that properties of the market that are (or can be) formulated in terms of limit quantile superreplication prices for general undefaultable strategies—the key example being the notion of *numéraire-independent no-arbitrage* (NINA) (see Chapter III.2)—can be reformulated in terms of bounded or even simple predictable undefaultable strategies, which are somewhat closer to reality. From a mathematical perspective, considering bounded undefaultable strategies only has the technical advantage that they can be integrated with respect to *any* semimartingale under *any* measure that is absolutely continuous with respect to the physical measure \mathbb{P} ; see Chapter V.

Before proving this robustness property of limit quantile superreplication prices, we present a counter-example showing that ordinary superreplication lack this property.

Example 7.1. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting a $(0, 1)$ -uniformly distributed random variable U , $\mathcal{F}_t^0 := \sigma(\{U \leq u\} : u \leq t)$, and $\mathcal{F}_t := \sigma(\mathcal{F}_t^0, \mathcal{N})$, $t \in [0, 1]$, where \mathcal{N} denotes the \mathbb{P} -null sets in \mathcal{F}_1^0 . It is not difficult to check that $(\mathcal{F}_t)_{t \in [0, 1]}$ satisfies the usual condition and that U is an $(\mathcal{F}_t)_{t \in [0, 1]}$ -stopping time. Let \mathcal{S} be the market generated by the classic model $S = (1, X_t)_{t \in [0, 1]}$, where $X_t := (1 - U)\mathbb{1}_{[U, 1]}$. Set $F := V_1(e_1)$, where $e_1 = (1, 0)$. We claim that $\Pi(F | \mathbf{b}\mathcal{U}) = V_0(\eta) > 0$ and $\Pi(F | \mathcal{U}) = 0$.

To establish both claims, note that for each $\vartheta \in \mathcal{U}$,

$$V_1(\vartheta)(S) = V_0(\vartheta)(S) + \vartheta_V^2(1 - U) \text{ P-a.s.}$$

Thus, if ϑ is bounded, $V_1(\vartheta)(S) \geq 1$ P-a.s. if and only if $V_0(\vartheta)(S) \geq 1$, which gives the first claim. The second claim follows from observing that the strategy $\vartheta^* := (\mathbb{1}_{[U, 1]}, \frac{1}{1-t}\mathbb{1}_{[0, U]})$ is in \mathcal{U} and satisfies $V_0(\vartheta^*) = 0$ and $V_1(\vartheta^*) \geq V_1(e_1)$ P-a.s.

The following technical lemma is the key ingredient for the proof of Theorem 7.3 below. It shows that if one can superreplicate a contingent claim F with a *general* undefaultable strategy, one can, with a little more initial capital and with probability almost 1, even superreplicate F with a *bounded* or—if the market is continuous—with a *simple predictable* undefaultable strategy.

Recall from Definition 3.4 that a numéraire market is called *bounded* if there exists a bounded numéraire strategy η such that $S^{(\eta)}$ is bounded, and *simple predictable* if in addition η can be chosen simple predictable. Moreover, recall from Proposition 3.3 that if a market is continuous each numéraire representative has P-a.s. continuous paths.

Lemma 7.2. *Let \mathcal{S} be a bounded numéraire market and F a contingent claim at time $\tau \in \mathcal{T}_{[0, T]}$. Suppose there exists $\vartheta \in \mathcal{U}$ such that $V_\tau(\vartheta) \geq F$ P-a.s. Then for each $\varepsilon \in (0, 1)$, each $\delta > 0$ and each positive contingent claim C at time 0, there exists $\tilde{\vartheta} \in \mathbf{b}\mathcal{U}$ satisfying*

$$V_0(\tilde{\vartheta}) = V_0(\vartheta) + \delta C \quad \text{and} \quad \mathbb{P}[V_\tau(\tilde{\vartheta}) \geq F] \geq 1 - \varepsilon. \quad (7.1)$$

If \mathcal{S} is in addition simple predictable and continuous, $\tilde{\vartheta}$ can be chosen in $\mathbf{s}\mathcal{U}$.

Proof. Throughout this proof, denote by $\|\cdot\|$ the maximum norm in \mathbb{R}^N , where N is the dimension of \mathcal{S} . Let η be a bounded (or simple predictable) numéraire strategy such that $S^{(\eta)}$ is bounded. Fix ε , δ and C as above. Since all contingent claims at time 0 coincide up to a constant (this uses that \mathcal{F}_0 is P-trivial), we may assume without loss of generality that $C = V_0(\eta)$. Choose $K > 0$ large enough that $\sup_{t \in [0, T]} \max(\|\eta_t\|, \|S_t^{(\eta)}\|) \leq K$ P-a.s.

Step 1. By an approximation argument—which is different in both cases—, we construct a bounded (or simple predictable) process $\tilde{\zeta}^{(1)}$ such that

$$\tilde{\zeta}_0^{(1)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(1)} \bullet S^{(\eta)} \geq -\delta \text{ P-a.s.}, \quad (7.2)$$

$$\mathbb{P} \left[\tilde{\zeta}_0^{(1)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(1)} \bullet S_T^{(\eta)} \geq V_T(\vartheta)(S^{(\eta)}) - \delta \right] \geq 1 - \frac{\varepsilon}{2} \quad (7.3)$$

If \mathcal{S} is a general bounded numéraire market, set $\zeta^{(n)} := \vartheta \mathbf{1}_{\{\|\vartheta\| \leq n\}}$, $n \in \mathbb{N}$. Then $\zeta^{(n)} \bullet S^{(\eta)}$ converges to $\vartheta \bullet S^{(\eta)}$ in the *semimartingale topology* (see Memin [63, Lemme V.3]) and a fortiori *uniformly in probability* on compact intervals (see Emery [20, p. 264]). Hence, there exists $M \in \mathbb{N}$ such that

$$\mathbb{P} \left[\sup_{0 \leq t \leq T} \left| \zeta^{(M)} \bullet S_t^{(\eta)} - \vartheta \bullet S_t^{(\eta)} \right| \geq \delta \right] \leq \varepsilon/2.$$

Since \mathcal{F}_0 is trivial, we may assume (after possibly enlarging M) that in addition $\zeta_0^{(M)} = \vartheta_0$. Define the stopping time

$$\sigma_1 := \inf \left\{ t \in [0, T] : \sup_{0 \leq s \leq t} \left| \zeta^{(M)} \bullet S_s^{(\eta)} - \vartheta \bullet S_s^{(\eta)} \right| \geq \delta \right\}.$$

By construction, $\sigma_1 > 0$ \mathbb{P} -a.s. and $\mathbb{P}[\sigma_1 = \infty] \geq 1 - \varepsilon/2$. Set

$$\tilde{\zeta}^{(1)} := \zeta^{(M)} \mathbf{1}_{[0, \sigma_1]}.$$

Then $\|\tilde{\zeta}^{(1)}\| \leq M$, and

$$\tilde{\zeta}_0^{(1)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(1)} \bullet S_T^{(\eta)} \geq V_T(\vartheta)(S^{(\eta)}) - \delta \text{ on } \{\sigma_1 = +\infty\}, \quad (7.4)$$

which gives (7.3). To establish (7.2), due to $\zeta_0^{(M)} = \vartheta_0$, it suffices to show that

$$\tilde{\zeta}^{(1)} \bullet S^{(\eta)} \geq -\vartheta_0 \cdot S_0^{(\eta)} - \delta.$$

Fix $0 \leq t \leq T$. Then on $\{t < \sigma_1\}$,

$$\begin{aligned} \tilde{\zeta}^{(1)} \bullet S_t^{(\eta)} &= \zeta^{(M)} \bullet S_t^{(\eta)} \geq \vartheta \bullet S_t^{(\eta)} - \sup_{0 \leq s \leq t} |\zeta^{(M)} \bullet S_s^{(\eta)} - \vartheta \bullet S_s^{(\eta)}| \\ &\geq V_t(\vartheta)(S^{(\eta)}) - \vartheta_0 \cdot S_0^{(\eta)} - \delta \geq 0 - \vartheta_0 \cdot S_0^{(\eta)} - \delta \text{ } \mathbb{P}\text{-a.s.}, \end{aligned}$$

on $\{t \geq \sigma_1\} \cap \{\|\vartheta_{\sigma_1}\| \leq M\}$,

$$\begin{aligned} \tilde{\zeta}^{(1)} \bullet S_t^{(\eta)} &= \zeta^{(M)} \bullet S_{\sigma_1}^{(\eta)} = \zeta^{(M)} \bullet S_{\sigma_1-}^{(\eta)} + \zeta_{\sigma_1}^{(M)} \cdot \Delta S_{\sigma_1}^{(\eta)} \\ &\geq \vartheta \bullet S_{\sigma_1-}^{(\eta)} + \vartheta_{\sigma_1} \cdot \Delta S_{\sigma_1}^{(\eta)} - \sup_{0 \leq s < \sigma_1} |\zeta^{(M)} \bullet S_s^{(\eta)} - \vartheta \bullet S_s^{(\eta)}| \\ &\geq V_{\sigma_1}(\vartheta)(S^{(\eta)}) - \vartheta_0 \cdot S_0^{(\eta)} - \delta \geq 0 - \vartheta_0 \cdot S_0^{(\eta)} - \delta \text{ } \mathbb{P}\text{-a.s.}, \end{aligned}$$

and on $\{t \geq \sigma_1\} \cap \{\|\vartheta_{\sigma_1}\| > M\}$,

$$\begin{aligned} \tilde{\zeta}^{(1)} \bullet S_t^{(\eta)} &= \zeta^{(M)} \bullet S_{\sigma_1}^{(\eta)} = \zeta^{(M)} \bullet S_{\sigma_1-}^{(\eta)} + 0 \cdot \Delta S_{\sigma_1}^{(\eta)} \\ &\geq \vartheta \bullet S_{\sigma_1-}^{(\eta)} - \sup_{0 \leq s < \sigma_1} |\zeta^{(M)} \bullet S_s^{(\eta)} - \vartheta \bullet S_s^{(\eta)}| \\ &\geq V_{\sigma_1-}(\vartheta)(S^{(\eta)}) - \vartheta_0 \cdot S_0^{(\eta)} - \delta \geq 0 - \vartheta_0 \cdot S_0^{(\eta)} - \delta \text{ } \mathbb{P}\text{-a.s.} \end{aligned}$$

If \mathcal{S} is simple predictable and continuous, by (a trivial extension of) Stricker [78, Proposition 2], there is a sequence $(\zeta^{(n)})_{n \in \mathbb{N}}$ in $\mathbf{s}L(S^{(\eta)})$ such that $\zeta_0^{(n)} = \zeta_0$,

$n \in \mathbb{N}$, and the sequence of stochastic integrals $(\zeta^{(n)} \bullet S^{(n)})_{n \in \mathbb{N}}$ converges to $\vartheta \bullet S^{(n)}$ in the *semimartingale topology* and a fortiori *uniformly in probability* on compact intervals. Define M , σ_1 and $\tilde{\zeta}^{(1)}$ as above—noting that now $\tilde{\zeta}^{(1)}$ is simple predictable. This gives (7.3), and (7.2) follows by a similar (but easier) argument as above using that $\Delta S_{\sigma_1}^{(\eta)} = 0$ on $\{\sigma_1 < \infty\}$ because $S^{(n)}$ has \mathbb{P} -a.s. continuous paths.

Step 2. Set

$$\tilde{\zeta}^{(2)} := \tilde{\zeta}^{(1)} + \delta\eta.$$

Then $\tilde{\zeta}^{(2)}$ is bounded (or simple predictable) and satisfies

$$\|\tilde{\zeta}^{(2)}\| \leq \|\tilde{\zeta}^{(1)}\| + \delta\|\eta\| \leq M + \delta K.$$

Moreover, by construction and the first step,

$$\tilde{\zeta}_0^{(2)} \cdot S_0^{(\eta)} = V_0(\vartheta)(S^{(n)}) + \delta V_0(\eta)(S^{(n)}), \quad (7.5)$$

$$\tilde{\zeta}^{(2)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(2)} \bullet S^{(n)} \geq 0 \text{ } \mathbb{P}\text{-a.s.}, \quad (7.6)$$

$$\tilde{\zeta}^{(2)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(2)} \bullet S^{(n)} \geq V(\vartheta)(S^{(n)}) \text{ on } \{\sigma_1 = \infty\}. \quad (7.7)$$

Step 3. Choose $\tilde{M} \in \mathbb{N}$ large enough that

$$\mathbb{P} \left[\sup_{0 \leq t \leq T} \left| \tilde{\zeta}^{(2)} \bullet S_t^{(\eta)} \right| \geq \tilde{M} \right] \leq \frac{\varepsilon}{2}$$

and define the stopping time

$$\sigma_2 := \inf \left\{ t \in [0, T] : \sup_{0 \leq s \leq t} \left| \tilde{\zeta}^{(2)} \bullet S_s^{(\eta)} \right| \geq \tilde{M} \right\}.$$

By construction, $\sigma_2 > 0$ \mathbb{P} -a.s. and $\mathbb{P}[\sigma_2 = \infty] \geq 1 - \varepsilon/2$. Moreover, on $\{\sigma_2 < \infty\}$,

$$\begin{aligned} |\Delta \tilde{\zeta}^{(2)} \bullet S_{\sigma_2}^{(\eta)}| &= |\tilde{\zeta}_{\sigma_2}^{(2)} \cdot \Delta S_{\sigma_2}^{(\eta)}| \leq N \|\tilde{\zeta}_{\sigma_2}^{(2)}\| (\|S_{\sigma_2}^{(\eta)}\| + \|S_{\sigma_2-}^{(\eta)}\|) \\ &\leq N(M + \delta K)(2K) = 2NK(M + \delta K). \end{aligned} \quad (7.8)$$

Set

$$\tilde{\zeta}^{(3)} := \tilde{\zeta}^{(2)} \mathbf{1}_{[0, \sigma_2]}.$$

Then $\tilde{\zeta}^{(3)}$ is bounded (or simple predictable) and satisfies $\|\tilde{\zeta}^{(3)}\| \leq M + \delta K$. In addition, by construction, (7.5)–(7.7), the definition of σ_2 and (7.8),

$$\tilde{\zeta}_0^{(3)} \cdot S_0^{(\eta)} = V_0(\vartheta)(S^{(n)}) + \delta V_0(\eta)(S^{(n)}), \quad (7.9)$$

$$\tilde{\zeta}^{(3)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(3)} \bullet S^{(n)} \geq 0 \text{ } \mathbb{P}\text{-a.s.}, \quad (7.10)$$

$$\tilde{\zeta}^{(3)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(3)} \bullet S^{(n)} \geq V(\vartheta)(S^{(n)}) \text{ on } \{\sigma_1 \wedge \sigma_2 = \infty\}, \quad (7.11)$$

$$\sup_{0 \leq t \leq T} \left| \tilde{\zeta}^{(3)} \bullet S_t^{(\eta)} \right| \leq \tilde{M} + 2NK(M + \delta K) \text{ } \mathbb{P}\text{-a.s.} \quad (7.12)$$

Step 4. Since η , $S^{(\eta)}$, $\tilde{\zeta}^{(3)}$ and $\tilde{\zeta}^{(3)} \bullet S^{(\eta)}$ are bounded (and η is simple predictable), Theorem 3.7 gives $\tilde{\vartheta} \in \mathbf{bL}^{\text{sf}}$ (or $\tilde{\vartheta} \in \mathbf{sL}^{\text{sf}}$) with

$$V(\tilde{\vartheta})(S^{(\eta)}) = \tilde{\zeta}^{(3)} \cdot S_0^{(\eta)} + \tilde{\zeta}^{(3)} \bullet S^{(\eta)} \text{ P-a.s.}$$

Moreover, (7.10) gives $\tilde{\vartheta} \in \mathbf{bU}$ (or $\tilde{\vartheta} \in \mathbf{sU}$), and the claim follows from (7.9) and (7.11) together with $\mathbb{P}[\sigma_1 \wedge \sigma_2 = \infty] \geq 1 - \varepsilon$. \square

Now we can prove the main result of this section.

Theorem 7.3. *Let \mathcal{S} be a bounded numéraire market and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. Then*

$$\Pi^*(F | \mathcal{U}) = \Pi^*(F | \mathbf{bU}). \quad (7.13)$$

If \mathcal{S} is in addition simple predictable and continuous, then

$$\Pi^*(F | \mathcal{U}) = \Pi^*(F | \mathbf{bU}) = \Pi^*(F | \mathbf{sU}). \quad (7.14)$$

Proof. We only establish (7.13); (7.14) follows by a similar argument.

Since $\mathbf{bU} \subset \mathcal{U}$, the inequality “ \leq ” in (7.13) follows easily from the definitions of quantile and limit quantile superreplication prices.

For the reverse inequality “ \geq ”, by the definition of limit quantile superreplication prices, it suffices to show that, for all $\varepsilon \in (0, 1)$,

$$\Pi^\varepsilon(F | \mathbf{bU}) \leq \Pi^*(F | \mathcal{U}). \quad (7.15)$$

So let $\varepsilon \in (0, 1)$ be given. By Proposition 6.14, there is $A \in \mathcal{F}_\tau$ with $\mathbb{P}[A] \leq \varepsilon/2$ such that $\Pi^*(F | \mathcal{U}) \geq \Pi(F\mathbf{1}_{A^c} | \mathcal{U})$. In order to establish (7.15), it suffices to show that

$$\Pi^\varepsilon(F | \mathbf{bU}) \leq \Pi(F\mathbf{1}_{A^c} | \mathcal{U}). \quad (7.16)$$

We may assume without loss of generality that $\Pi(F\mathbf{1}_{A^c} | \mathcal{U}) < \infty$. Let ϑ be a bounded numéraire strategy such that also $S^{(\vartheta)}$ is bounded, $\delta > 0$ and C a contingent claim at time 0. By Proposition 6.3, there exists $\vartheta \in \mathcal{U}$ such that

$$V_0(\vartheta) \leq \Pi_\tau(F\mathbf{1}_{A^c} | \mathcal{U}) + \frac{\delta}{2}C \quad \text{and} \quad V_\tau(\vartheta) \geq F\mathbf{1}_{A^c} \text{ P-a.s.} \quad (7.17)$$

Next, by Lemma 7.2, there exists $\tilde{\vartheta} \in \mathbf{bU}$ such that

$$V_0(\tilde{\vartheta}) = V_0(\vartheta) + \frac{\delta}{2}C \quad \text{and} \quad \mathbb{P}[V_\tau(\tilde{\vartheta}) \geq F\mathbf{1}_{A^c}] \geq 1 - \varepsilon. \quad (7.18)$$

Thus,

$$\Pi^\varepsilon(F | \mathbf{bU}) \leq \Pi(F\mathbf{1}_{A^c} | \mathcal{U}) + \delta C,$$

and letting $\delta \searrow 0$ yields (7.16). \square

Chapter III

Numéraire-independent no-arbitrage (NINA)

In this chapter, we study the fundamental concept of (no-)arbitrage in our numéraire-independent framework. In Section 1, we introduce a very general and *quantitative* notion of arbitrage, called *gratis events*, based on limit quantile superreplication prices. The results in Section 1 also lay the ground for Chapters IV and V. For the most relevant case of undefaultable strategies, we derive an equivalent and simpler characterisation of the absence of gratis events based on ordinary superreplication prices in Section 2. We call the corresponding notion *numéraire-independent no-arbitrage (NINA)*. In Section 3, we compare NINA to classic notions of no-arbitrage including no-arbitrage (NA), no free lunch with vanishing risk (NFLVR) and no unbounded profit with bounded risk (NUPBR) by presenting a new unifying characterisation of those concepts in terms of *maximal strategies*. The results in Section 3 are also foundational for Chapter VI. The material presented in this chapter is for the main part taken from [31] and [32].

1 Gratis events

Loosely speaking, arbitrage means “*making a profit out of nothing without risk*”. If we fix a strategy cone Γ , e.g. $\Gamma = \mathcal{U}$, and interpret this to describe the investments which we consider “without risk”, a rather intuitive translation of this catchphrase into a mathematical definition is the existence of a nonzero contingent claim F (“the profit”) which can be superreplicated for free (“out of nothing”) using strategies in Γ .

Before making the above idea precise, we introduce the *support* of a contingent claim.

Definition 1.1. Let \mathcal{S} be a market, Γ a strategy cone and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. A set $A \in \mathcal{F}_\tau$ satisfying $A = \{F > 0\}$ P-a.s. is called a *support* of F . We write $A = \text{supp } F$ P-a.s. and say that F is *supported on* A .

Clearly, the support of a contingent claim F is P-a.s. unique.

For the next definition, recall that the collection of all \mathbb{P} -null sets in \mathcal{F}_T is denoted by \mathcal{N} .

Definition 1.2. Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. A set $A \in \mathcal{F}_\tau \setminus \mathcal{N}$ is called a *gratis event of \mathcal{S} at time τ for Γ* if there exists a contingent claim F at time τ such that

$$A = \text{supp } F \text{ } \mathbb{P}\text{-a.s.} \quad \text{and} \quad \Pi^*(F | \Gamma) = 0.$$

We denote by $\mathcal{G}_\tau(\Gamma)$ the collection of all gratis events of \mathcal{S} at time τ for Γ .

Remark 1.3. One could define an analogous notion of gratis events based on *ordinary* superreplication prices. It turns out, however, that this is far less useful than the above concept based on limit quantile superreplication prices. The reason is that ordinary (as opposed to limit quantile) superreplication prices do not satisfy a monotone convergence property (cf. Lemma II.6.18). Notwithstanding, in the important case of undefaultable strategies, the *absence* of gratis events can equivalently be characterised by ordinary superreplication prices; see Section 2.

First, we collect two basic properties of the set $\mathcal{G}_\tau(\Gamma)$, which follow immediately from subadditivity and monotonicity of limit quantile superreplication prices (cf. Proposition II.6.16).

Proposition 1.4. *Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Then $\mathcal{G}_\tau(\Gamma)$ is closed under finite unions. Moreover, if $A \in \mathcal{G}_\tau(\Gamma)$ and $B \in \mathcal{F}_\tau \setminus \mathcal{N}$ with $B \subset A$ \mathbb{P} -a.s., then $B \in \mathcal{G}_\tau(\Gamma)$.*

Second, we establish that the limit quantile superreplication price of *every* contingent claim supported on a gratis event is indeed *gratis*.

Proposition 1.5. *Let \mathcal{S} be a market, Γ a strategy cone and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ with $\text{supp } F \in \mathcal{G}_\tau(\Gamma)$. Then*

$$\Pi^*(F | \Gamma) = 0.$$

Proof. Since $\text{supp } F \in \mathcal{G}_\tau(\Gamma)$, by the definition of gratis events, there exists a contingent claim G at time τ with $\text{supp } G = \text{supp } F$ \mathbb{P} -a.s. and $\Pi^*(G | \Gamma) = 0$. For $n \in \mathbb{N}$, set $F_n := F \mathbf{1}_{\{F \leq nG\}}$. Then $F_n \leq F_{n+1}$ \mathbb{P} -a.s., and $F = \lim_{n \rightarrow \infty} F_n$ \mathbb{P} -a.s. Moreover, since $F_n \leq nG$ \mathbb{P} -a.s., monotonicity and positive homogeneity of limit quantile superreplication prices (cf. Proposition II.6.16) yield

$$\Pi^*(F_n | \Gamma) \leq \Pi^*(nG | \Gamma) = n\Pi^*(G | \Gamma) = 0.$$

Now the claim follows from Lemma II.6.18. □

The last result shows that if $A \in \mathcal{G}_\tau(\Gamma)$ is a gratis event and we fix *any* representative $S \in \mathcal{S}$ and consider the *Arrow-Debreu type security* $\mathbf{1}_A$ in the currency unit determined by S , i.e., the contingent claim F at time τ from Proposition II.5.3 satisfying $F(S) = \mathbf{1}_A$, then its limit quantile superreplication price is 0.

Remark 1.6. (a) The notion of gratis events can be seen as a numéraire-independent version of the notion of *cheap thrills* introduced by Loewenstein and Willard [56, Definition 2]. Indeed, fix *any* representative $S \in \mathcal{S}$ and *any* set $A \in \mathcal{G}_\tau(\Gamma)$. Let F be the contingent claim at time τ satisfying $F(S) = \mathbf{1}_A$ \mathbb{P} -a.s. and C the positive contingent claim at time 0 satisfying $C(S) = 1$ (cf. Proposition II.5.3). Then for each $n \in \mathbb{N}$, $\Pi^*(nF | \Gamma) = n\Pi^*(F | \Gamma) = 0$, and so Corollary II.6.15 gives an undefaultable strategy ϑ_n satisfying

$$V_0(\vartheta_n)(S) \leq \frac{1}{n} \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta_n)(S) \geq n\mathbf{1}_A] \geq 1 - \frac{1}{n}.$$

Hence, on a gratis event A , in *any currency unit*, we can get *arbitrarily much* with *probability almost 1* with *arbitrarily low initial investment* using strategies having *nonnegative value processes*.

(b) If we fix $S \in \mathcal{S}$, $\tau \in \mathcal{T}_{[0,T]}$ and Γ and define the generalised coherent risk measure $\rho_\tau : \mathbf{L}_-^0(\mathcal{F}_\tau) \rightarrow [0, \infty]$ by $\rho_\tau(X) := \Pi^*(F | \Gamma)(S)$, where F is the contingent claim at time τ from Proposition II.5.3 satisfying $F(S) = -X$ (cf. Remark II.6.17), then $\mathcal{G}_\tau(\Gamma) = \emptyset$ is equivalent to saying that the generalised risk measure ρ_τ is *relevant* or *sensitive*; see [25, Definition 4.32]. This kind of connection between sensitive generalised risk measures and absence of arbitrage deserves a more careful analysis, which we postpone to future research.

The following corollary provides a characterisation of gratis events solely in terms of strategies.

Corollary 1.7. *Let \mathcal{S} be a market, Γ a strategy cone and F a nonzero contingent claim at time τ . Then $\text{supp } F \in \mathcal{G}_\tau(\Gamma)$ if and only if for each $\varepsilon \in (0, 1)$, each $\delta > 0$ and each positive contingent claim C at time 0, there exists $\vartheta \in \Gamma$ satisfying*

$$V_0(\vartheta) \leq \delta C \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta) \geq F] \geq 1 - \varepsilon. \quad (1.1)$$

Proof. First, assume that $\text{supp } F \in \mathcal{G}_\tau(\Gamma)$. Proposition 1.5 gives $\Pi^*(F | \Gamma) = 0$, and (1.1) follows from Corollary II.6.15. Conversely, assuming (1.1), letting first $\delta \searrow 0$ establishes $\Pi^\varepsilon(F | \Gamma) = 0$ for all $\varepsilon \in (0, 1)$, and letting then $\varepsilon \searrow 0$ yields $\Pi^*(F | \Gamma) = 0$, which gives $\text{supp } F \in \mathcal{G}_\tau(\Gamma)$. \square

Next, we show that limit quantile superreplication prices of (generalised) contingent claims are determined “outside of gratis events”.

Proposition 1.8. *Let \mathcal{S} be a market, Γ a strategy cone and F a generalised contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. Then for each $A \in \mathcal{G}_\tau(\Gamma)$,*

$$\Pi^*(F | \Gamma) = \Pi^*(F - \infty\mathbf{1}_A | \Gamma).$$

Proof. We may assume without loss of generality that $\mathcal{G}_\tau(\Gamma) \neq \emptyset$. Pick $A \in \mathcal{G}_\tau(\Gamma)$. By monotonicity and the definition of limit quantile superreplication prices, it suffices to show that for each $\varepsilon \in (0, 1)$, each $\delta > 0$, and each positive contingent claim C at time 0,

$$\Pi^\varepsilon(F | \Gamma) \leq \Pi^*(F - \infty\mathbf{1}_A | \Gamma) + \delta C. \quad (1.2)$$

So let ε, δ and C be as above. It suffices to consider the case $\Pi^*(F - \infty \mathbb{1}_A | \Gamma) < \infty$. By Corollary II.6.15, there exists $\vartheta^{(1)} \in \Gamma$ satisfying

$$V_0(\vartheta^{(1)}) \leq \Pi^*(F - \infty \mathbb{1}_A | \Gamma) + \frac{\delta}{2}C \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta^{(1)}) \geq F - \infty \mathbb{1}_A] \geq 1 - \frac{\varepsilon}{2}.$$

Set $G := (F \mathbb{1}_{\{F \geq 0\}} + |V_\tau(\vartheta^{(1)})|) \mathbb{1}_A$. Then G is a contingent claim at time τ with $\{G > 0\} \subset A$, and so $\Pi^*(G | \Gamma) = 0$ by Propositions 1.4 and 1.5 if $\mathbb{P}[G > 0] > 0$, and trivially if $G = 0$ \mathbb{P} -a.s. Hence, by Corollary 1.7, there exists $\vartheta^{(2)} \in \Gamma$ satisfying

$$V_0(\vartheta^{(2)}) \leq \frac{\delta}{2}C \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta^{(2)}) \geq G] \geq 1 - \frac{\varepsilon}{2}.$$

Set $\vartheta := \vartheta^{(1)} + \vartheta^{(2)} \in \Gamma$. Then $V_0(\vartheta) \leq \Pi^*(F - \infty \mathbb{1}_A | \Gamma) + \delta C$ and

$$\mathbb{P}[V_\tau(\vartheta) \geq F] \geq \mathbb{P}[G + V_\tau(\vartheta^{(2)}) \geq F] - \frac{\varepsilon}{2} = 1 - \varepsilon.$$

This gives (1.2). □

We proceed to show that the set $\mathcal{G}_\tau(\Gamma)$ is also closed under *countable* unions.

Lemma 1.9. *Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Then $\mathcal{G}_\tau(\Gamma)$ is closed under countable unions.*

Proof. We may assume without loss of generality that $\mathcal{G}_\tau(\Gamma) \neq \emptyset$. Let F be a positive contingent claim at time τ and $(A_n)_{n \in \mathbb{N}}$ a sequence in $\mathcal{G}_\tau(\Gamma)$. For $n \in \mathbb{N}$, define $B_n := \bigcup_{1 \leq k \leq n} A_k$, and set $A := \bigcup_{n \in \mathbb{N}} A_n = \bigcup_{n \in \mathbb{N}} B_n$. Then $(B_n)_{n \in \mathbb{N}}$ is a sequence in $\mathcal{G}_\tau(\Gamma)$ by Proposition 1.4 and $(F \mathbb{1}_{B_n})_{n \in \mathbb{N}}$ is a monotone increasing sequence of contingent claims satisfying

$$\lim_{n \rightarrow \infty} F \mathbb{1}_{B_n} = F \mathbb{1}_A.$$

Thus, Proposition 1.5 and Lemma II.6.18 imply that

$$\Pi^*(F \mathbb{1}_A | \Gamma) = \lim_{n \rightarrow \infty} \Pi^*(F \mathbb{1}_{B_n} | \Gamma) = 0. \quad \square$$

An important consequence of Lemma 1.9 is the existence of maximal elements in $\mathcal{G}_\tau(\Gamma)$.

Definition 1.10. Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Assume that $\mathcal{G}_\tau(\Gamma) \neq \emptyset$. Then $G \in \mathcal{G}_\tau(\Gamma)$ is called a *maximal gratis event at time τ for Γ* if $A \subset G$ \mathbb{P} -a.s. for all $A \in \mathcal{G}_\tau(\Gamma)$.

It is clear from the definition that maximal gratis events, if they exist, are \mathbb{P} -a.s. unique.

Corollary 1.11. *Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Assume that $\mathcal{G}_\tau(\Gamma) \neq \emptyset$. Then there exists a maximal gratis event at time τ for Γ .*

Proof. Since $\mathcal{G}_\tau(\Gamma)$ is closed under countable unions by Lemma 1.9, this is a standard argument. Set $\gamma := \sup_{G \in \mathcal{G}_\tau(\Gamma)} \mathbb{P}[G]$ and let $(G_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{G}_\tau(\Gamma)$ satisfying $\lim_{n \rightarrow \infty} \mathbb{P}[G_n] = \gamma$. Set $G := \bigcup_{n \in \mathbb{N}} G_n$. Then $G \in \mathcal{G}_\tau(\Gamma)$ by Lemma 1.9 and $\mathbb{P}[G] = \gamma$. Seeking a contradiction, suppose there exists $A \in \mathcal{G}_\tau(\Gamma)$ with $\mathbb{P}[A \setminus G] > 0$. Then $G \cup A \in \mathcal{G}_\tau(\Gamma)$ by Proposition 1.4 and $\mathbb{P}[G \cup A] > \gamma$, which is a contradiction. □

2 Numéraire-independent no-arbitrage (NINA)

It is clear from an economic perspective, that a “reasonable” market should not admit gratis events for undefaultable strategies. So we might say that a market satisfies *numéraire-independent no-arbitrage (NINA)* if $\mathcal{G}_\tau(\mathcal{U}) = \emptyset$ for all stopping times $\tau \in \mathcal{T}_{[0,T]}$. But even though this definition is economically compelling, it is not practical as on the one hand, it involves a condition on *all* stopping times $\tau \in \mathcal{T}_{[0,T]}$, and on the other hand, it is based on the rather complicated notion of limit quantile superreplication prices. The goal of this section is to derive a simpler equivalent characterisation, which we then take as definition of NINA.

First, we consider the sets $\mathcal{G}_\tau(\Gamma)$ as a function of $\tau \in \mathcal{T}_{[0,T]}$, and we do this—also for future reference in Chapter IV—for general strategy cones Γ . In order to get interesting results, we have to assume that \mathcal{S} is a numéraire market and that Γ is rich enough to allow at each stopping time to *switch* to some numéraire strategy.

Definition 2.1. Let \mathcal{S} be a numéraire market. A strategy cone Γ is said to *allow switching to numéraire strategies* if for all $\tau \in \mathcal{T}_{[0,T]}$, there exists a numéraire strategy η , called (a) *switching numéraire strategy at time τ* , which may depend on τ and which is such that for all $\vartheta \in \Gamma$,

$$\vartheta \mathbf{1}_{[0,\tau]} + V_\tau(\vartheta)(S^{(n)}\eta \mathbf{1}_{] \tau, T]}) \in \Gamma.$$

If \mathcal{S} is a numéraire market, L^{sf} and \mathcal{U} allow switching to numéraire strategies, and any numéraire strategy can be taken as switching numéraire strategy at all stopping times. Moreover, if \mathcal{S} is a *bounded* or *simple predictable* numéraire market (cf. Definition II.3.4), $\mathbf{b}L^{\text{sf}}$ and $\mathbf{b}\mathcal{U}$ or $\mathbf{s}L^{\text{sf}}$ and $\mathbf{s}\mathcal{U}$ allow switching to numéraire strategies, and any bounded or simple predictable numéraire strategy with a bounded numéraire representative can be taken as switching numéraire strategy at all stopping times. In particular, if \mathcal{S} is a nonnegative market, L^{sf} , \mathcal{U} , $\mathbf{b}L^{\text{sf}}$, $\mathbf{b}\mathcal{U}$, $\mathbf{s}L^{\text{sf}}$ and $\mathbf{s}\mathcal{U}$ allow switching to numéraire strategies, and the market portfolio $\eta^{\mathcal{S}} = (1, \dots, 1)$ can be taken as switching numéraire strategy at all stopping times. Another important class of strategy cones which allow switching to numéraires are *undefaultable strategies outside the gratis events of another strategy cone* Γ ; see Chapter V.3.

For strategy cones which allow switching to numéraire strategies, the notion of gratis events is *time-consistent* in the sense that gratis events propagate forward in time.

Proposition 2.2. *Let \mathcal{S} be a numéraire market, Γ a strategy cone which allows switching to numéraire strategies and $\tau_1, \tau_2 \in \mathcal{T}_{[0,T]}$ with $\tau_1 \leq \tau_2$ \mathbb{P} -a.s. Then*

$$\mathcal{G}_{\tau_1}(\Gamma) \subset \mathcal{G}_{\tau_2}(\Gamma).$$

In particular, $\mathcal{G}_\tau(\Gamma) \subset \mathcal{G}_T(\Gamma)$ for all $\tau \in \mathcal{T}_{[0,T]}$.

Proof. Let η be a switching numéraire strategy at time τ_1 . We may assume without loss of generality that there exists $A \in \mathcal{G}_{\tau_1}(\Gamma)$ with $\mathbb{P}[A] > 0$. By the

definition of gratis events, there exists a contingent claim F_1 at time τ_1 with $\text{supp } F_1 = A$ \mathbb{P} -a.s. Set $g := F_1(S^{(n)})$. Then g is \mathcal{F}_{τ_1} -measurable because F_1 is a contingent claim at time τ_1 and also \mathcal{F}_{τ_2} -measurable since $\mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2}$. Let F_2 be the contingent claim at time τ_2 from Proposition II.5.3 satisfying $F_2(S^{(n)}) = g$. Then $\text{supp } F_2 = A$ \mathbb{P} -a.s. We proceed to show that $\Pi^*(F_2 | \Gamma) = 0$. To this end, let $\varepsilon \in (0, 1)$, $\delta > 0$, and C be a positive contingent claim at time 0. By Corollary 1.7, there exists $\vartheta \in \Gamma$ such that

$$V_0(\vartheta) \leq \delta C \quad \mathbb{P}\text{-a.s.} \quad \text{and} \quad \mathbb{P}[V_{\tau_1}(\vartheta) \geq F_1] \geq 1 - \varepsilon.$$

Set $\tilde{\vartheta} := \vartheta \mathbf{1}_{[0, \tau_1]} + V_{\tau_1}(\vartheta)(S^{(n)}) \eta \mathbf{1}_{] \tau_1, T]} \in \Gamma$. Then $V_0(\tilde{\vartheta}) = V_0(\vartheta) \leq \delta C$. Moreover,

$$\mathbb{P}[V_{\tau_2}(\tilde{\vartheta}) \geq F_2] = \mathbb{P}[V_{\tau_2}(\tilde{\vartheta})(S^{(n)}) \geq g] = \mathbb{P}[V_{\tau_1}(\vartheta)(S^{(n)}) \geq g] \geq 1 - \varepsilon.$$

Now the claim follows from Corollary 1.7. \square

Remark 2.3. In Lemma IV.1.3, we prove a refinement of Proposition 2.2, showing that $\mathcal{G}_{\tau_1}(\Gamma) \cap \{\tau_1 \leq \tau_2\} \subset \mathcal{G}_{\tau_2}(\Gamma) \cup \mathcal{N}$ for *arbitrary* stopping times $\tau_1, \tau_2 \in \mathcal{T}_{[0, T]}$ provided that $\Gamma \subset \mathcal{U}$.

Thanks to Proposition 2.2, in order to check that $\mathcal{G}_{\tau}(\mathcal{U}) = \emptyset$ for all stopping times $\tau \in \mathcal{T}_{[0, T]}$, it suffices to check that $\mathcal{G}_T(\mathcal{U}) = \emptyset$. By the definition of gratis events, this is equivalent to

$$\Pi^*(F | \mathcal{U}) > 0 \tag{2.1}$$

for all nonzero contingent claims F at time T . One drawback of condition (2.1) is that it involves the rather complicated notion of limit quantile superreplication prices. The next result shows that (2.1) is equivalent to the same condition for ordinary superreplication prices.

Proposition 2.4. *Let \mathcal{S} be a market, $\Gamma \subset \mathcal{U}$ a strategy cone and $\tau \in \mathcal{T}_{[0, T]}$. Then $\Pi^*(F | \Gamma) > 0$ for all nonzero contingent claims F at time τ if and only if $\Pi(F | \Gamma) > 0$ for all nonzero contingent claims F at time τ .*

Note that Proposition 2.4 does not say that for some *fixed* nonzero contingent claim F at time τ , $\Pi^*(F | \Gamma) > 0$ if and only if $\Pi(F | \Gamma) > 0$. This is wrong as can be seen from Example II.6.12.

Proof. First, if $\Pi^*(F | \Gamma) > 0$ for all nonzero contingent claim F at time τ , then a fortiori $\Pi(F | \Gamma) > 0$ for all nonzero contingent claim F at time τ by (II.6.2). Conversely, by way of contradiction, suppose $\Pi(F | \Gamma) > 0$ for all nonzero contingent claim F at time τ but $\Pi^*(\tilde{F} | \Gamma) = 0$ for some nonzero contingent claim \tilde{F} at time τ . Set $\varepsilon := \frac{1}{2} \mathbb{P}[\text{supp } \tilde{F}]$. Then $\varepsilon \in (0, 1)$, and by Proposition II.6.14, there exists $A \in \mathcal{F}_{\tau}$ with $\mathbb{P}[A] \leq \varepsilon$ such that

$$\Pi(\tilde{F} \mathbf{1}_{A^c} | \Gamma) \leq \Pi^*(\tilde{F} | \Gamma) = 0.$$

Since $\tilde{F} \mathbf{1}_{A^c}$ is a nonzero contingent claim, we arrive at a contradiction. \square

After the above preparations, we can now state our definition of *numéraire-independent no-arbitrage (NINA)*.

Definition 2.5. A market \mathcal{S} is said to satisfy *numéraire-independent no-arbitrage (NINA)* if for all nonzero contingent claims F at time T ,

$$\Pi(F | \mathcal{U}) > 0.$$

We finish this section by providing a couple of equivalent characterisations of NINA. From an economic perspective, in particular the characterisations in (d) and (e) are interesting, as they only involve bounded or simple predictable undefaultable strategies.

Lemma 2.6. *Let \mathcal{S} be a numéraire market. Then the following are equivalent:*

- (a) \mathcal{S} satisfies NINA.
- (b) For all nonzero contingent claims F at time T , $\Pi^*(F | \mathcal{U}) > 0$.
- (c) For all stopping times $\tau \in \mathcal{T}_{[0,T]}$ and all nonzero contingent claims F at time τ , $\Pi^*(F | \mathcal{U}) > 0$.

If \mathcal{S} is a bounded numéraire market, (a) – (c) are equivalent to

- (d) For all nonzero contingent claims F at time T , $\Pi^*(F | \mathbf{b}\mathcal{U}) > 0$.

If \mathcal{S} is continuous and simple predictable, (a) – (d) are equivalent to

- (e) For all nonzero contingent claims F at time T , $\Pi^*(F | \mathbf{s}\mathcal{U}) > 0$.

Proof. The equivalence (a) \Leftrightarrow (b) follows from Proposition 2.4, and the equivalence (b) \Leftrightarrow (c) follows from Proposition 2.2 since \mathcal{U} allows switching to numéraire strategies. The equivalence (c) \Leftrightarrow (d) if \mathcal{S} is a bounded numéraire market and the equivalence (c) \Leftrightarrow (e) if \mathcal{S} is a continuous simple predictable numéraire market follow from Theorem II.7.3. \square

3 Maximal strategies

The goal of this section is to compare our no-arbitrage concept of numéraire-independent no-arbitrage (NINA) to classic no-arbitrage notions including no-arbitrage (NA), no-free lunch with vanishing risk (NFLVR) and no unbounded profit with bounded risk (NUPBR). To this end, we develop a new unifying characterisation of those concepts in terms of *maximal strategies*, which explains in particular how the classic notions of no-arbitrage depend on the choice of currency unit and numéraire.

To motivate the concept of maximal strategies, fix a strategy cone Γ of possible trading opportunities. A strategy $\vartheta \in \Gamma$ can be considered as a “reasonable investment” from that class on $\llbracket 0, \tau \rrbracket$ for some stopping time $\tau \in \mathcal{T}_{[0,T]}$ only if by using self-financing strategies in Γ , (1) one cannot create *more wealth* at time

τ with the *same* (or a lower) initial investment, and (2) one cannot create the *same wealth* at time τ with a *lower* initial investment. It is natural to call such a strategy ϑ *maximal*—and this can be made mathematically precise in different ways leading to the two notions of *weakly* and *strongly* maximal strategies; see also Remarks 3.2 and 3.12 below for more comments.

3.1 Weakly maximal strategies

First, we introduce a weak notion of maximality.

Definition 3.1. Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. A strategy $\vartheta \in \Gamma$ is called *weakly maximal at time τ for Γ* if there does not exist $\tilde{\vartheta} \in \Gamma$ with

$$V_0(\tilde{\vartheta}) \leq V_0(\vartheta) \quad \text{and} \quad V_\tau(\tilde{\vartheta}) \geq V_\tau(\vartheta) \quad \mathbb{P}\text{-a.s.}$$

such that one of the two inequalities is strict with positive probability. If $\vartheta \in \Gamma$ is weakly maximal at *each* time $\tau \in \mathcal{T}_{[0,T]}$ for Γ , it is called *weakly maximal for Γ* . We often omit the qualifier “for Γ ”.

Remark 3.2. The *terminology* “maximal strategies” goes back at least to Delbaen and Schachermayer [12]. However, the setting there is different from here so that also maximality has a different meaning. More precisely, [12] is cast in the classic setup, i.e., it considers a classic model $S = (1, X)$, uses for Γ the class \mathcal{A} of classic admissible strategies which is not numéraire-independent (cf. the discussion after Definition II.4.3), and also imposes ex ante the absence-of-arbitrage condition NFLVR. In our terminology, a maximal strategy in the sense of [12] is then weakly maximal at time T for \mathcal{A} . For this reason, some of our results on *weakly* maximal strategies are similar to results in [12].

The zero strategy plays a key role when studying weakly maximal strategies. This is easy.

Proposition 3.3. *Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Then $\vartheta \in \Gamma$ is weakly maximal at time τ only if the zero strategy 0 is weakly maximal at time τ .*

We illustrate by a counter-example that the converse of Proposition 3.3 is false.

Example 3.4. Let $W = (W_t)_{t \in [0,1]}$ be a Brownian motion on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0,1]}, \mathbb{P})$. Let \mathcal{S} be the market generated by the classic model $S = (1, X_t)_{t \in [0,1]}$, where

$$X_t := 1 + \int_0^{\tau \wedge t} \frac{1}{1-s} dW_s \quad \text{and} \quad \tau := \inf \left\{ t \in [0, 1] : \int_0^t \frac{1}{1-s} dW_s = -\frac{1}{2} \right\}.$$

Note that $S^2 = X$ is a strict local martingale with $S_1^2 = \frac{1}{2}$ \mathbb{P} -a.s. We claim that the zero strategy 0 is weakly maximal at time 1 for \mathcal{U} but that the numéraire strategy

$e_2 = (0, 1)$ is not. For the first claim, suppose by way of contradiction that there is $\vartheta \in \mathcal{U}$ with $V_0(\vartheta) = 0$ and $\mathbb{P}[V_1(\vartheta) > 0] > 0$. (We cannot have $V_0(\vartheta) < 0$ because $\vartheta \in \mathcal{U}$.) Then $V(\vartheta)(S) = \vartheta \bullet S$ is a nonnegative local martingale and a supermartingale. Hence, $\mathbb{E}[V_1(\vartheta)(S)] \leq V_0(\vartheta)(S) = 0$, and we arrive at a contradiction. For the second claim, consider the numéraire strategy $e_1 = (1, 0)$. Then $V_0(e_1)(S) = V_0(e_2)(S) = 1$ and $V_1(e_1)(S) = 1 > \frac{1}{2} = V_1(e_2)(S)$, which shows that e_2 fails to be weakly maximal for \mathcal{U} . More precisely, and anticipating ourselves, since $V(e_1)(S) \equiv 1$ is (trivially) a martingale, it follows from Theorem VI.2.1 below that e_1 is a *dominating maximal strategy* for e_2 (cf. Theorem VI.1.1 below).

The following result shows that strategies which are not weakly maximal always fail to satisfy condition (1) of a “reasonable investment” (see above) under a mild technical condition. We omit the straightforward proof.

Proposition 3.5. *Let \mathcal{S} be a numéraire market, $\tau \in \mathcal{T}_{[0,T]}$ and Γ a strategy cone containing a numéraire strategy. If $\vartheta \in \Gamma$ is not weakly maximal at time τ , there exists $\tilde{\vartheta} \in \Gamma$ such that*

$$V_0(\tilde{\vartheta}) = V_0(\vartheta), \quad V_\tau(\tilde{\vartheta}) \geq V_\tau(\vartheta) \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[V_\tau(\tilde{\vartheta}) > V_\tau(\vartheta)] > 0.$$

Our next goal is to establish that the notion of weak maximality is *time-consistent*. To this end, we have to assume that \mathcal{S} is a numéraire market and that Γ is rich enough to allow *switching to a dominated strategy* at each stopping time. This is a weak analogue of *predictable convexity*, considered in the standard framework for the set of (constrained) admissible strategies.

Definition 3.6. Let \mathcal{S} be a numéraire market. A strategy cone Γ is said to *allow switching to dominated strategies* if for all $\tau \in \mathcal{T}_{[0,T]}$, there exists a numéraire strategy η , called (a) *switching numéraire strategy at time τ* , which may depend on τ and which is such that for all $\vartheta, \tilde{\vartheta} \in \Gamma$,

$$\vartheta \mathbb{1}_{[0,\tau]} + \left(\mathbb{1}_{\{V_\tau(\vartheta) < V_\tau(\tilde{\vartheta})\}} \vartheta + \mathbb{1}_{\{V_\tau(\vartheta) \geq V_\tau(\tilde{\vartheta})\}} \left(\tilde{\vartheta} + V_\tau(\vartheta - \tilde{\vartheta})(S^{(\eta)}\eta) \right) \right) \mathbb{1}_{[\tau,T]} \in \Gamma.$$

If \mathcal{S} is a numéraire market, then \mathcal{U} allows switching to dominated strategies. Indeed, any numéraire strategy can be taken as switching numéraire strategy at all stopping times.

Remark 3.7. If \mathcal{S} is a numéraire market and $\Gamma \subset \mathcal{U}$ allows switching to dominated strategies, it allows in particular switching to numéraire strategies (cf. Definition 2.1). Indeed, take $\tilde{\vartheta} = 0$ in Definition 3.6. Therefore, it is justified to call η in both cases a “switching numéraire strategy”.

Proposition 3.8. *Let \mathcal{S} be a numéraire market, $\tau_1 \leq \tau_2 \in \mathcal{T}_{[0,T]}$ stopping times and Γ a strategy cone which allows switching to dominated strategies. If $\vartheta \in \Gamma$ is weakly maximal at time τ_2 , it is also weakly maximal at time τ_1 . As a consequence, $\vartheta \in \Gamma$ is weakly maximal if and only if it is weakly maximal at time T .*

Proof. By contraposition, suppose $\vartheta \in \Gamma$ is not weakly maximal at time τ_1 . By Proposition 3.5, there exists $\tilde{\vartheta}$ with $V_0(\tilde{\vartheta}) = V_0(\vartheta)$, $V_{\tau_1}(\tilde{\vartheta}) \geq V_{\tau_1}(\vartheta)$ \mathbb{P} -a.s. and $\mathbb{P}[V_{\tau_1}(\tilde{\vartheta}) > V_{\tau_1}(\vartheta)] > 0$. Let η be a switching numéraire strategy at time τ . Set $\hat{\vartheta} := \tilde{\vartheta}\mathbb{1}_{[0, \tau_1]} + \left(\vartheta + V_{\tau_1}(\tilde{\vartheta} - \vartheta)(S^{(\eta)})\eta\right)\mathbb{1}_{] \tau_1, T]} \in \Gamma$. Then $V_0(\hat{\vartheta}) = V_0(\tilde{\vartheta}) = V_0(\vartheta)$ and $V_{\tau_2}(\hat{\vartheta}) = V_{\tau_2}(\vartheta) + \left(V_{\tau_1}(\tilde{\vartheta} - \vartheta)(S^{(\eta)})\right)V_{\tau_2}(\eta)$ \mathbb{P} -a.s. Since $\mathbb{P}[V_{\tau_1}(\tilde{\vartheta} - \vartheta)(S^{(\eta)}) > 0] > 0$, ϑ is not weakly maximal at time τ_2 , and we arrive at a contradiction. \square

We proceed to show that if a weakly maximal strategy is dominated by another strategy at some stopping time $\tau \in \mathcal{T}_{[0, T]}$, it is also dominated at all earlier stopping times $\sigma \leq \tau$.

Proposition 3.9. *Let \mathcal{S} be a numéraire market, Γ a strategy cone which allows switching to dominated strategies and $\tau \in \mathcal{T}_{[0, T]}$. Suppose that $\vartheta \in \Gamma$ is weakly maximal at time τ and $\tilde{\vartheta} \in \Gamma$ satisfies $V_{\tau}(\tilde{\vartheta}) \geq V_{\tau}(\vartheta)$ \mathbb{P} -a.s. Then $V(\tilde{\vartheta}) \geq V(\vartheta)$ \mathbb{P} -a.s. on $[0, \tau]$.*

Proof. Seeking a contradiction, suppose there is a stopping time $\sigma \leq \tau$ with $\mathbb{P}[V_{\sigma}(\tilde{\vartheta}) < V_{\sigma}(\vartheta)] > 0$. Let η be a switching numéraire strategy at time σ , and set $\hat{\vartheta} := \vartheta\mathbb{1}_{[0, \sigma]} + \left(\mathbb{1}_{\{V_{\sigma}(\vartheta) < V_{\sigma}(\tilde{\vartheta})\}}\vartheta + \mathbb{1}_{\{V_{\sigma}(\vartheta) \geq V_{\sigma}(\tilde{\vartheta})\}}(\tilde{\vartheta} + V_{\sigma}(\vartheta - \tilde{\vartheta})(S^{(\eta)})\eta)\right)\mathbb{1}_{] \sigma, T]} \in \Gamma$. Then $V_0(\hat{\vartheta}) = V_0(\vartheta)$, and using that $V_{\tau}(\tilde{\vartheta}) \geq V_{\tau}(\vartheta)$ \mathbb{P} -a.s. gives

$$\begin{aligned} V_{\tau}(\hat{\vartheta}) &= \mathbb{1}_{\{V_{\sigma}(\vartheta) < V_{\sigma}(\tilde{\vartheta})\}}V_{\tau}(\vartheta) + \mathbb{1}_{\{V_{\sigma}(\vartheta) \geq V_{\sigma}(\tilde{\vartheta})\}}\left(V_{\tau}(\tilde{\vartheta}) + V_{\sigma}(\vartheta - \tilde{\vartheta})(S^{(\eta)})V_{\tau}(\eta)\right) \\ &\geq \mathbb{1}_{\{V_{\sigma}(\vartheta) < V_{\sigma}(\tilde{\vartheta})\}}V_{\tau}(\vartheta) + \mathbb{1}_{\{V_{\sigma}(\vartheta) \geq V_{\sigma}(\tilde{\vartheta})\}}\left(V_{\tau}(\vartheta) + V_{\sigma}(\vartheta - \tilde{\vartheta})(S^{(\eta)})V_{\tau}(\eta)\right) \\ &= V_{\tau}(\vartheta) + \mathbb{1}_{\{V_{\sigma}(\vartheta) > V_{\sigma}(\tilde{\vartheta})\}}V_{\sigma}(\vartheta - \tilde{\vartheta})(S^{(\eta)})V_{\tau}(\eta) \text{ } \mathbb{P}\text{-a.s.} \end{aligned}$$

Since $\mathbb{P}[V_{\sigma}(\vartheta - \tilde{\vartheta})(S^{(\eta)}) > 0] > 0$, ϑ fails to be weakly maximal at time τ , and we arrive at a contradiction. \square

An important consequence of the above result is that weakly maximal strategies in \mathcal{U} form a convex cone. This can be seen as a numéraire-independent version of [12, Theorem 2.12].

Corollary 3.10. *Let \mathcal{S} be a numéraire market and $\tau \in \mathcal{T}_{[0, T]}$. If $\vartheta^{(1)}, \vartheta^{(2)} \in \mathcal{U}$ are weakly maximal at time τ for \mathcal{U} , then $\vartheta^{(1)} + \vartheta^{(2)}$ is so, too. As a consequence, weakly maximal strategies in \mathcal{U} form a convex cone.*

Proof. Seeking a contradiction, suppose there are $\vartheta^{(1)}, \vartheta^{(2)} \in \mathcal{U}$ such that $\vartheta^{(1)}$ and $\vartheta^{(2)}$ but not $\vartheta^{(1)} + \vartheta^{(2)}$ are weakly maximal at time τ . By Proposition 3.5, there is $\tilde{\vartheta} \in \mathcal{U}$ satisfying $V_0(\tilde{\vartheta}) = V_0(\vartheta^{(1)} + \vartheta^{(2)})$, $V_{\tau}(\tilde{\vartheta}) \geq V_{\tau}(\vartheta^{(1)} + \vartheta^{(2)})$ \mathbb{P} -a.s. and $\mathbb{P}[V_{\tau}(\tilde{\vartheta}) > V_{\tau}(\vartheta^{(1)} + \vartheta^{(2)})] > 0$. Let $\eta \in \mathcal{U}$ be a numéraire strategy and $\vartheta := (\vartheta^{(1)} - \vartheta^{(2)})\mathbb{1}_{[0, \tau]} + (V_{\tau}(\tilde{\vartheta} - \vartheta^{(1)} + \vartheta^{(2)})(S^{(\eta)})\eta)\mathbb{1}_{] \tau, T]}$. Then ϑ is in \mathcal{U} by Proposition 3.9, $V_0(\vartheta) = V_0(\vartheta^{(1)})$, $V_{\tau}(\vartheta) \geq V_{\tau}(\vartheta^{(1)})$ \mathbb{P} -a.s. and $\mathbb{P}[V_{\tau}(\vartheta) > V_{\tau}(\vartheta^{(1)})] > 0$. Thus, we arrive at a contradiction. \square

3.2 Strongly maximal strategies

Next, we study a strong notion of maximality.

Definition 3.11. Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. A strategy $\vartheta \in \Gamma$ is called *strongly maximal at time τ for Γ* if $\Pi(V_\tau(\vartheta) | \Gamma) = V_0(\vartheta)$ and if for all nonzero contingent claims F at time τ ,

$$\Pi(V_\tau(\vartheta) + F | \Gamma) > V_0(\vartheta). \quad (3.1)$$

If $\vartheta \in \Gamma$ is strongly maximal at *each* time $\tau \in \mathcal{T}_{[0,T]}$ for Γ , it is called *strongly maximal for Γ* . We often omit the qualifier “for Γ ”.

Remark 3.12. (a) The concept of strongly maximal strategies is to the best of our knowledge entirely new and has also not been considered in any form in the standard framework before.

(b) One could in principle define two notions of strong maximality: one using ordinary superreplication prices as above and one using limit quantile superreplication prices. This has been done in [31]. However, it turns out that both concepts coincide for undefaultable strategies; see [31, Proposition 5.6]. For this reason, we only consider here the more relevant and technically easier notion based on ordinary superreplication prices.

The following result justifies our terminology.

Proposition 3.13. *Let \mathcal{S} be a market, $\tau \in \mathcal{T}_{[0,T]}$ and Γ a strategy cone. If $\vartheta \in \Gamma$ is strongly maximal at time τ , then it is also weakly maximal at time τ .*

Proof. By way of contradiction, suppose there is $\vartheta \in \Gamma$ which is strongly but not weakly maximal at time τ . Then there exists $\tilde{\vartheta} \in \Gamma$ with $V_0(\tilde{\vartheta}) \leq V_0(\vartheta)$ and $V_\tau(\tilde{\vartheta}) \geq V_\tau(\vartheta)$ \mathbb{P} -a.s., and either $V_0(\tilde{\vartheta}) < V_0(\vartheta)$ or $\mathbb{P}[V_\tau(\tilde{\vartheta}) > V_\tau(\vartheta)] > 0$. In the first case, we arrive at the contradiction $\Pi(V_\tau(\vartheta) | \Gamma) \leq V_0(\tilde{\vartheta}) < V_0(\vartheta)$, and in the second case, setting $F = V_\tau(\tilde{\vartheta}) - V_\tau(\vartheta)$, we arrive at the contradiction $\Pi(V_\tau(\vartheta) + F | \Gamma) = \Pi(V_\tau(\tilde{\vartheta}) | \Gamma) \leq V_0(\tilde{\vartheta}) = V_0(\vartheta)$. \square

The following counter-example, which is inspired by [33], shows that the converse of Proposition 3.13 does not hold.

Example 3.14. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting a random variable U , uniformly distributed on $(0, 1)$. Set $\mathcal{F}_t^0 := \sigma(\{U \leq u\} : u \leq t)$ and $\mathcal{F}_t := \sigma(\mathcal{F}_t^0, \mathcal{N})$, $t \in [0, 1]$, where \mathcal{N} denotes the \mathbb{P} -null sets in \mathcal{F}_1^0 . It is not difficult to check that $(\mathcal{F}_t)_{t \in [0,1]}$ satisfies the usual conditions, U is an $(\mathcal{F}_t)_{t \in [0,1]}$ -stopping time, and for each $(\mathcal{F}_t)_{t \in [0,1]}$ -predictable process $H = (H_t)_{t \in [0,1]}$, there is a deterministic measurable function $h : [0, 1] \rightarrow \mathbb{R}$ with $H \mathbb{1}_{[0,U]} = h \mathbb{1}_{[0,U]}$ \mathbb{P} -a.s. Let \mathcal{S} be the market generated by the classic model $S = (1, X_t)_{t \in [0,1]}$, where $X_t = A(t) \mathbb{1}_{\{t < U\}} + B(U) \mathbb{1}_{\{t \geq U\}}$ and $A, B : [0, 1] \rightarrow \mathbb{R}$ are given by

$$A(t) := \begin{cases} 1 + s - \frac{2}{3}s^2, & t \in [0, \frac{1}{2}] \\ \frac{4}{3}, & t \in (\frac{1}{2}, 1] \end{cases} \quad \text{and} \quad B(t) := \begin{cases} \frac{1}{3} + 3s - 2s^2, & t \in [0, \frac{1}{2}] \\ \frac{4}{3}, & t \in (\frac{1}{2}, 1]. \end{cases}$$

Note that $S^2 = X$ is constant after $U \wedge \frac{1}{2}$. We claim that each $\vartheta \in \mathcal{U}$ is weakly but not strongly maximal at time 1 for \mathcal{U} .

For the first claim, fix $\vartheta \in \mathcal{U}$ and let $\tilde{\vartheta} \in \mathcal{U}$ be such that $V_0(\tilde{\vartheta}) = V_0(\vartheta)$ and $V_1(\tilde{\vartheta}) \geq V_1(\vartheta)$ \mathbb{P} -a.s. Let $h : [0, 1] \rightarrow \mathbb{R}$ be a measurable function satisfying $h\mathbb{1}_{[0, U]} = (\tilde{\vartheta}^2 - \vartheta^2)\mathbb{1}_{[0, U]}$. Then

$$\begin{aligned} V_1(\tilde{\vartheta})(S) - V_1(\vartheta)(S) &= (\tilde{\vartheta}^2 - \vartheta^2) \bullet S_1^2 = (\tilde{\vartheta}^2 - \vartheta^2) \bullet S_{U \wedge \frac{1}{2}}^2 \\ &= \int_0^{U \wedge \frac{1}{2}} h(s) dA(s) + h(U)(B(U) - A(U))\mathbb{1}_{\{U < 1/2\}} \quad \mathbb{P}\text{-a.s.} \end{aligned}$$

The assumption that $V_1(\tilde{\vartheta})(S) \geq V_1(\vartheta)(S)$ \mathbb{P} -a.s. implies that

$$\int_0^t h(s) \left(1 - \frac{4}{3}s\right) ds \geq h(t) \frac{2}{3}(1-t)(1-2t) \quad \text{for a.e. } t \in (0, 1/2), \quad (3.2)$$

$$\int_0^{\frac{1}{2}} h(s) \left(1 - \frac{4}{3}s\right) ds \geq 0. \quad (3.3)$$

A version of Gronwall's inequality [74, Lemma D.2] implies that $h(t) \leq 0$ for a.e. $t \in (0, 1/2)$. (Clearly, (3.2) remains valid when replacing h by h^+ .) This together with (3.3) shows that $h(t) = 0$ for a.e. $t \in (0, 1/2)$. Hence $V_1(\tilde{\vartheta}) = V_1(\vartheta)$ \mathbb{P} -a.s., and the claim follows from Proposition 3.5.

For the second claim, by Proposition 3.15 below, it suffices to show that the zero strategy 0 is not strongly maximal at time 1 for \mathcal{U} . Let F be the contingent claim at time 1 from Proposition II.5.3 satisfying $F(S) = \mathbb{1}_{\{U \geq 1/2\}}$. We claim that $\Pi(F | \mathcal{U}) = 0$. Since F is nonzero, this establishes the claim. By Remark II.6.2 and the definition of ordinary superreplication prices, it suffices to show that for each $\delta > 0$, there is $\vartheta \in \mathcal{U}$ with $V_0(\vartheta)(S) \leq \delta$ and $V_1(\vartheta)(S) \geq F(S)$ \mathbb{P} -a.s. To this end, by Corollary II.3.8 (b), it suffices to find $\zeta \in L(X)$ with

$$\delta + \zeta \bullet X \geq 0 \quad \mathbb{P}\text{-a.s.} \quad \text{and} \quad \delta + \zeta \bullet X_1 \geq F(S) = \mathbb{1}_{\{U \geq 1/2\}} \quad \mathbb{P}\text{-a.s.} \quad (3.4)$$

So let $\delta \in (0, \frac{2}{3})$. Set $\zeta_t := \frac{3}{2} \frac{\delta}{1-2t} \mathbb{1}_{[\ell(\delta), r(\delta)]}$, $t \in [0, 1]$, where $\ell(\delta) := \frac{1}{2}(1 - \frac{3}{2}\delta)$ and $r(\delta) = \frac{1}{2}(1 - \frac{3}{2}\delta \exp(-\frac{4}{\delta}))$. Note that $0 < \ell(\delta) < r(\delta) < \frac{1}{2}$. To establish the first part of (3.4), it suffices to consider $\delta + \zeta \bullet X$ on $\{U > \ell(\delta)\} \times (\ell(\delta), r(\delta)]$. So fix $t \in (\ell(\delta), r(\delta)]$. Then on $\{t < U\}$,

$$\delta + \zeta \bullet X_t = \delta + \int_{\ell(\delta)}^t \zeta_s dA(s) = \delta + \int_{\ell(\delta)}^t \frac{3}{2} \frac{\delta}{1-2s} \left(1 - \frac{4}{3}s\right) ds \geq \delta \geq 0 \quad \mathbb{P}\text{-a.s.},$$

and on $\{t \geq U > \ell(\delta)\}$,

$$\begin{aligned} \delta + \zeta \bullet X_t &= \delta + \int_{\ell(\delta)}^U \zeta_s dA(s) + \zeta_U(B(U) - A(U)) \\ &= \delta + \int_{\ell(\delta)}^U \frac{3}{2} \frac{\delta}{1-2s} \left(1 - \frac{4}{3}s\right) ds + \frac{3}{2} \frac{\delta}{1-2U} \frac{2}{3}(U-1)(1-2U) \\ &= \delta + \delta \left(U - \ell(\delta) + \frac{1}{4} \log \frac{1-2\ell(\delta)}{1-2U} \right) - \delta(1-U) \geq 0 \quad \mathbb{P}\text{-a.s.} \end{aligned}$$

Finally, on $\{U \geq 1/2\}$,

$$\begin{aligned} \delta + \zeta \bullet X_1 &= \delta + \int_{\ell(\delta)}^{r(\delta)} \frac{3}{2} \frac{\delta}{1-2s} \left(1 - \frac{4}{3}s\right) \\ &= \delta + \delta \left(r(\delta) - \ell(\delta) + \frac{1}{4} \log \frac{1-2\ell(\delta)}{1-2r(\delta)} \right) \geq 1 = F(S) \text{ P-a.s.}, \end{aligned}$$

which gives the second part of (3.4). This ends the example.

As in the case of weakly maximal strategies (cf. Proposition 3.3), the zero strategy plays a fundamental role when studying strongly maximal strategies. We omit the easy proof.

Proposition 3.15. *Let \mathcal{S} be a market, Γ a strategy cone and $\tau \in \mathcal{T}_{[0,T]}$. Then $\vartheta \in \Gamma$ is strongly maximal at time τ only if the zero strategy 0 is strongly maximal at time τ .*

We demonstrate by a counter-example that the converse of Proposition 3.15 is false.

Example 3.16. Consider the setup of Example 3.4. We claim that the zero strategy 0 is strongly maximal at time 1 for \mathcal{U} but that the numéraire strategy $e_1 = (0, 1)$ is not. The second claim follows immediately from Proposition 3.13 and the fact that e_1 is not even weakly maximal at time 1 for \mathcal{U} (see Example 3.4). For the first claim, seeking a contradiction, suppose there exists a nonzero contingent claim F at time 1 with $\Pi(F|\mathcal{U}) = 0$. Let $0 < \delta < \mathbb{E}[F(S)]$ and C be the positive contingent claim at time 0 from Proposition II.5.3 satisfying $C(S) = 1$. Then Proposition II.6.3 gives $\vartheta \in \mathcal{U}$ with $V_0(\vartheta) \leq \delta C$ and $V_1(\vartheta) \geq F$ P-a.s. Since $V(\vartheta) = \vartheta_0 \cdot S_0 + \vartheta \bullet S$ is a nonnegative local martingale and a supermartingale, we arrive at the contradiction

$$V_0(\vartheta)(S) \geq \mathbb{E}[V_1(\vartheta)(S)] \geq \mathbb{E}[F(S)] > \delta = \delta C(S) \geq V_0(\vartheta)(S).$$

Also Propositions 3.5 and 3.8 have analogues for strongly maximal strategies; the proof of Propositions 3.17 is easy and hence omitted.

Proposition 3.17. *Let \mathcal{S} be a numéraire market, $\tau \in \mathcal{T}_{[0,T]}$ and Γ a strategy cone containing a numéraire strategy. If $\vartheta \in \Gamma$ is not strongly maximal at time τ , there exists a nonzero contingent claim F at time τ such that $\Pi(V_\tau(\vartheta) + F | \Gamma) \leq V_0(\vartheta)$.*

Proposition 3.18. *Let \mathcal{S} be a numéraire market, $\tau_1 \leq \tau_2 \in \mathcal{T}_{[0,T]}$ and Γ a strategy cone which allows switching to dominated strategies. If $\vartheta \in \Gamma$ is strongly maximal at time τ_2 , it is also strongly maximal at time τ_1 . As a consequence, $\vartheta \in \Gamma$ is strongly maximal if and only if it is strongly maximal at time T .*

Proof. The argument is somewhat similar to the proof of Proposition 2.2. By contraposition, suppose that $\vartheta \in \Gamma$ is not strongly maximal at time τ_1 . By Proposition 3.17, there exists a nonzero contingent claim F_1 at time τ_1 satisfying

$\Pi(V_{\tau_1}(\vartheta) + F_1 | \Gamma) \leq V_0(\vartheta)$. Let η be a switching numéraire strategy at time τ_1 and set $g = F(S^{(\eta)})$. Then g is \mathcal{F}_{τ_1} -measurable because F_1 is a contingent claim at time τ_1 and also \mathcal{F}_{τ_2} -measurable since $\mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2}$. Let F_2 be the contingent claim at time τ_2 from Proposition II.5.3 satisfying $F_2(S^{(\eta)}) = g$, $\delta > 0$, and C a positive contingent claim at time 0. By Proposition II.6.3, there exist $\tilde{\vartheta} \in \Gamma$ with $V_0(\tilde{\vartheta}) \leq V_0(\vartheta) + \delta C$ and $V_{\tau_1}(\tilde{\vartheta}) \geq V_{\tau_1}(\vartheta) + F_1$ \mathbb{P} -a.s. Define $\hat{\vartheta} \in \Gamma$ by $\hat{\vartheta} := \tilde{\vartheta} \mathbf{1}_{\llbracket 0, \tau_1 \rrbracket} + (\vartheta + V_{\tau_1}(\tilde{\vartheta} - \vartheta)(S^{(\eta)})\eta) \mathbf{1}_{\llbracket \tau_1, T \rrbracket}$. Then $V_0(\hat{\vartheta}) = V_0(\tilde{\vartheta}) \leq V_0(\vartheta) + \delta C$, and using that $V(\eta)(S^{(\eta)}) \equiv 1$ and $V_{\tau_1}(\tilde{\vartheta} - \vartheta) \geq F_1$ gives

$$\begin{aligned} V_{\tau_2}(\hat{\vartheta})(S^{(\eta)}) &= V_{\tau_2}(\vartheta)(S^{(\eta)}) + V_{\tau_1}(\tilde{\vartheta} - \vartheta)(S^{(\eta)}) \geq V_{\tau_2}(\vartheta)(S^{(\eta)}) + F_1(S^{(\eta)}) \\ &= V_{\tau_2}(\vartheta)(S^{(\eta)}) + g = V_{\tau_2}(\vartheta)(S^{(\eta)}) + F_2(S^{(\eta)}) \text{ } \mathbb{P}\text{-a.s.} \end{aligned}$$

This yields $\Pi(V_{\tau_2}(\vartheta) + F_2 | \Gamma) \leq V_0(\vartheta) + \delta C$. Letting $\delta \searrow 0$ shows that ϑ is not strongly maximal at time τ_2 . \square

For the important special case $\Gamma = \mathcal{U}$, we almost have a converse to Proposition 3.13.

Proposition 3.19. *Let \mathcal{S} be a numéraire market and $\tau \in \mathcal{T}_{[0, T]}$. Suppose that the zero strategy 0 is strongly maximal for \mathcal{U} at time τ . Then every strategy $\vartheta \in \mathcal{U}$ that is weakly maximal for \mathcal{U} at time τ is also strongly maximal for \mathcal{U} at time τ .*

Proof. Seeking a contradiction, suppose $\vartheta \in \mathcal{U}$ is weakly but not strongly maximal at time τ for \mathcal{U} . By Proposition 3.17, there is a nonzero contingent claim F at time τ with $\Pi(V_{\tau}(\vartheta) + F | \mathcal{U}) \leq V_0(\vartheta)$. Let $\delta > 0$ and C be a positive contingent claim at time 0. By Proposition II.6.3, there is $\tilde{\vartheta} \in \mathcal{U}$ with $V_0(\tilde{\vartheta}) \leq V_0(\vartheta) + \delta C$ and $V_{\tau}(\tilde{\vartheta}) \geq V_{\tau}(\vartheta) + F$ \mathbb{P} -a.s. Let η be a numéraire strategy and define $\bar{\vartheta} \in \Gamma$ by $\bar{\vartheta} := \tilde{\vartheta} \mathbf{1}_{\llbracket 0, \tau \rrbracket} + (\vartheta + V_{\tau}(\tilde{\vartheta} - \vartheta)(S^{(\eta)})\eta) \mathbf{1}_{\llbracket \tau, T \rrbracket}$. Proposition 3.9 gives $\hat{\vartheta} := \bar{\vartheta} - \vartheta \in \mathcal{U}$. Moreover, $V_0(\hat{\vartheta}) \leq \delta C$ and $V_{\tau}(\hat{\vartheta}) \geq F$ \mathbb{P} -a.s., and so $\Pi(F | \mathcal{U}) \leq \delta C$. Letting $\delta \searrow 0$ gives $\Pi(F | \mathcal{U}) = 0$, which implies that 0 is not strongly maximal at time τ for \mathcal{U} . Hence, we arrive at a contradiction. \square

3.3 No-arbitrage and maximal strategies

We now use our paradigm of maximal strategies to compare our concept of numéraire-independent no-arbitrage (NINA) to classic no-arbitrage conditions like NA, NFLVR or NUPBR.¹ To this end, we first introduce some compact notation.

Definition 3.20. Let \mathcal{S} be a market and $\vartheta \in \mathcal{U}$. We say that we have

- $wm(\vartheta)$ if ϑ is *weakly* maximal for \mathcal{U} .
- $sm(\vartheta)$ if ϑ is *strongly* maximal for \mathcal{U} .

The implication structure between $wm(0)$, $wm(\vartheta)$, $sm(0)$ and $sm(\vartheta)$ follows directly from the results and counter-examples in Sections 3.1 and 3.2.

¹For an excellent guided tour through the zoo of classic notions of no-arbitrage for a *continuous* classic model $S = (1, X)$, we refer to the Ph.D. thesis of Hulley [35].

Proposition 3.21. *Let \mathcal{S} be a numéraire market and $\vartheta \in \mathcal{U}$. Then:*

$$\begin{array}{ccc} (wm(\vartheta) \ \& \ sm(0)) & \Leftrightarrow & \begin{array}{ccc} sm(\vartheta) & \xRightarrow{\neq} & wm(\vartheta) \\ \Downarrow & & \Downarrow \\ sm(0) & \xRightarrow{\neq} & wm(0). \end{array} \end{array}$$

If we want to interpret $wm(\vartheta)$ or $sm(\vartheta)$ as no-arbitrage conditions for the market \mathcal{S} , the above result gives two insights. First, $wm(0)$ or $sm(0)$ are the weakest of such conditions. Second, if we have $wm(\vartheta)$ or $sm(\vartheta)$ for some fixed $\vartheta \in \mathcal{U}$, we do not have $wm(\tilde{\vartheta})$ or $sm(\tilde{\vartheta})$ for a different $\tilde{\vartheta} \in \mathcal{U}$ in general. Therefore, neither $wm(\vartheta)$ nor $sm(\vartheta)$ for $\vartheta \neq 0$ are good notions of no-arbitrage, as they crucially depend on an arbitrary choice of some $\vartheta \in \mathcal{U}$. (Anticipating ourselves a bit, we note that in the standard framework one chooses for ϑ the numéraire strategy $e_1 = (1, 0, \dots, 0)$.) So we are left with $wm(0)$ and $sm(0)$. One can show that in finite discrete time, both concepts are equivalent. In infinite discrete and continuous time, however, $wm(0)$ is too weak.² As Example 3.16 demonstrates, under $wm(0)$ one may still become arbitrarily rich with positive probability (1/2 in the example) by using an undefaultable strategy ϑ with initial cost $V(\vartheta)(S) \leq \delta$, where $\delta > 0$ can be chosen arbitrarily small. From an economic perspective, such a situation should clearly be called arbitrage. This also shows that in the standard framework, i.e., for a classic model $S = (1, X)$, working with 0-admissible strategies, i.e., the set $\mathcal{A}_0 := \{\vartheta \in L(X) : \vartheta_0 = 0 \text{ and } \vartheta \bullet X \geq 0\}$, does not lead to a good concept of no-arbitrage since under $wm(0)$, $\{\vartheta \bullet X_T : \vartheta \in \mathcal{A}_0\} = \{0\}$, and so the corresponding notions “0-NA” or “0-NFLVR” are trivially satisfied even in a situation as above. Thus $sm(0)$ is the only candidate for a numéraire-independent notion of no-arbitrage, and by Proposition 3.18 this is exactly what we have called numéraire-independent no-arbitrage (NINA) in Definition 2.5. Let us stress again that NINA or $sm(0)$ are completely general concepts. They do not assume that a numéraire strategy for \mathcal{S} exists, they do not involve any topological concept, and they are economically compelling: *Every nonzero contingent claim must have a positive (superreplication) price.*

We proceed to show how the classic no-arbitrage conditions *no-arbitrage (NA)* [9], *no free lunch with vanishing risk (NFLVR)* [9], *BK* or *no unbounded profit with bounded risk (NUPBR)* [42, 46], *no-arbitrage of the first kind (NA₁)* [48] and *no cheap thrills* [56] can be neatly characterized by our concept of maximal strategies. They are formulated for a classic model $S = (1, X)$, i.e., for the numéraire representative $S^{(e_1)}$ of $e_1 = (1, 0, \dots, 0)$; cf. Example II.3.6. To compare them to our framework, we have to generalise them to *general* numéraire strategies (as opposed to $\eta = e_1$). Let us repeat that all those classic notions depend a priori *by their very definition* on an initial choice of a numéraire strategy η . We indicate this—despite its clumsiness—also in the notation, and speak henceforth of $NA^{(\eta)}$ instead of NA, $NFLVR^{(\eta)}$ instead of NFLVR, and so forth.

²In the classic setup and for *continuous* markets, the analogue of $wm(0)$ has been studied under the names NA^+ and NA_+ by Strasser [77] and Hulley [35], respectively.

Definition 3.22. Let \mathcal{S} be a numéraire market and η a numéraire strategy. Set

$$\begin{aligned} K_0(S^{(\eta)}) &:= \{\zeta \bullet S_T^{(\eta)} : \zeta \in L^{ad}(S^{(\eta)})\}, \\ K_0^1(S^{(\eta)}) &:= \{\zeta \bullet S_T^{(\eta)} : \zeta \in L^{ad}(S^{(\eta)}, 1)\}, \\ C(S^{(\eta)}) &:= \{g \in \mathbf{L}^\infty(\mathcal{F}_T) : g \leq f \text{ for some } f \in K_0(S^{(\eta)})\} \\ &= (K_0(S^{(\eta)}) - \mathbf{L}_+^0(\mathcal{F}_T)) \cap \mathbf{L}^\infty(\mathcal{F}_T). \end{aligned}$$

Then $S^{(\eta)}$ is said to satisfy

- $NA^{(\eta)}$ if $K_0(S^{(\eta)}) \cap \mathbf{L}_+^0(\mathcal{F}_T) = \{0\}$,
- $NFLVR^{(\eta)}$ if $\overline{C(S^{(\eta)})} \cap \mathbf{L}_+^\infty(\mathcal{F}_T) = \{0\}$, where $\overline{C(S^{(\eta)})}$ denotes the closure of $C(S^{(\eta)})$ in the norm topology of $\mathbf{L}^\infty(\mathcal{F}_T)$,
- $BK^{(\eta)}$ or $NUPBR^{(\eta)}$ if $K_0^1(S^{(\eta)})$ is bounded in probability,
- $NA_1^{(\eta)}$ if there does not exist a nonzero random variable $\xi \in \mathbf{L}_+^0(\mathcal{F}_T)$ such that for all $x > 0$, there exists $\zeta \in L^{ad}(S^{(\eta)}, x)$ with $x + \zeta \bullet S_T^{(\eta)} \geq \xi$ \mathbb{P} -a.s.,
- (η) -no cheap thrills if there are no $A \in \mathcal{F}_T$ with $\mathbb{P}[A] > 0$ and strategies $\zeta^{(n)} \in L^{ad}(S^{(\eta)}, 1/n)$, $n \in \mathbb{N}$, with $\lim_{n \rightarrow \infty} \zeta^{(n)} \bullet S^{(\eta)} = \infty$ on A \mathbb{P} -a.s.

Remark 3.23. The above notions are good generalisations of the corresponding classic concepts. Indeed, if we interpret $S^{(\eta),1}, \dots, S^{(\eta),N}$ as $d = N$ discounted “risky” assets in a standard framework, we recover the classic definitions; cf. the discussion after Definition II.4.3.

It is easy to check that $NA_1^{(\eta)}$ and (η) -no cheap thrills are equivalent. It is shown in [47, Proposition 1] that the latter two are also equivalent to $BK^{(\eta)}/NUPBR^{(\eta)}$. Moreover, it is not difficult to establish that a numéraire representative $S^{(\eta)}$ satisfies $NA_1^{(\eta)}/(\eta)$ -no cheap thrills (and hence also $BK^{(\eta)}/NUPBR^{(\eta)}$) if and only if we have $sm(0)$, i.e., if and only if the market \mathcal{S} satisfies NINA. However, strictly speaking, one cannot say that $NA_1^{(\eta)}/(\eta)$ -no cheap thrills/ $BK^{(\eta)}/NUPBR^{(\eta)}$ are equivalent to NINA *since the former classic notions all depend on an initial choice of a numéraire strategy η* . In particular, if \mathcal{S} fails to be a numéraire market, it is not immediately clear how to formulate corresponding notions, whereas NINA clearly depends neither on the choice nor the existence of a numéraire strategy.

We proceed to characterize the above concepts in terms of maximal strategies in our sense.

Proposition 3.24. *Let \mathcal{S} be a numéraire market and η a fixed numéraire strategy. Then:*

- (a) $S^{(\eta)}$ satisfies $NA^{(\eta)}$ if and only if we have $wm(\eta)$.
- (b) $S^{(\eta)}$ satisfies $BK^{(\eta)}/NUPBR^{(\eta)}/NA_1^{(\eta)}/(\eta)$ -no cheap thrills if and only if we have $sm(0)$.

(c) $S^{(\eta)}$ satisfies $NFLVR^{(\eta)}$ if and only if we have $sm(\eta)$.

Proof. (a) follows easily from Proposition II.4.4, (b) follows from the discussion preceding the result, and (c) follows from (a) and (b) using the equivalence of $s(\eta)$ and $(wm(\eta) \& sm(0))$ (Proposition 3.21) and the equivalence of $NFLVR^{(\eta)}$ and $(NA^{(\eta)} \& BK^{(\eta)})$ [42, Lemma 2.2]. \square

Remark 3.25. (a) The above result shows in particular that the most familiar no-arbitrage conditions $NA = NA^{(e_1)}$ and $NFLVR = NFLVR^{(e_1)}$ for a classic model $S = (1, X)$ (cf. Example II.3.6) implicitly assume that $e_1 = (1, 0, \dots, 0)$, the buy-and-hold strategy of the “bank account”, is weakly or strongly maximal. From an economic perspective, this means that investing in the bank account is a “reasonable investment” (cf. the introduction to Section 3.1). This assumption is at least debatable, and so we consider NINA a more natural no-arbitrage condition.

(b) For a market \mathcal{S} generated by a classic model $S = (1, X)$, our notion of NINA for the market \mathcal{S} is equivalent to the notion of $NUPBR = NUPBR^{(e_1)}$ for the numéraire representative $S^{(e_1)} = S$. So why do we not just stick to $NUPBR$ and S ? First, \mathcal{S} might not be generated by a classic model, so this identification is not always possible. Second, even if \mathcal{S} is generated by a classic model, it is not clear from the definition that $NUPBR = NUPBR^{(e_1)}$ is a numéraire-independent property.³ Third, and this the crucial point, for a dual characterisation of NINA for a market \mathcal{S} (generated by a classic model or otherwise), sticking to *one* representative S will not do; see Chapter VI.1.

³We are not aware of any reference in the literature, where this statement is *rigorously* proved. It is stated in [73, before Proposition 2.7, italics added] that $NUPBR$ is a “numéraire-free property in a *certain sense*”, but this is not made precise. Moreover, in [73, Proposition 2.7] there is a change of dimension, and so this result is *not* numéraire-independent in our sense.

Chapter IV

Separating stopping times for markets failing NINA

In this chapter, we study numéraire markets failing numéraire-independent no-arbitrage (NINA); see Definition III.2.5. We seek to find a stopping time σ such that (1) strictly before σ , we can never “make a profit out of nothing without risk”, and (2) immediately after σ , we always can. It is natural to call σ a *separating stopping time* for the market \mathcal{S} .

After making the above concept mathematically precise and establishing a preliminary lemma in Section 1, we show the existence of a *smallest* separating stopping time in Section 2; the existence of a *largest* separating stopping time is proved in Section 3. We conclude the chapter by illustrating our results with several examples in Section 4.

1 Conceptual prélude

We first introduce some additional notation. For the next definition, recall the notion of gratis events from Definition III.1.2.

Definition 1.1. Let \mathcal{S} be a market and Γ a strategy cone. A stopping time $\tau \in \mathcal{T}_{[0,T]}$ is called a *gratis event time of \mathcal{S} for Γ* if $\mathcal{G}_\tau(\Gamma) \neq \emptyset$ and a *no gratis event time of \mathcal{S} for Γ* if $\mathcal{G}_\tau(\Gamma) = \emptyset$. We write $\mathcal{T}_{\text{ge}}(\Gamma)$ and $\mathcal{T}_{\text{nge}}(\Gamma)$ for the respective subsets of $\mathcal{T}_{[0,T]}$.

Clearly, $\mathcal{T}_{[0,T]}$ is the disjoint union of $\mathcal{T}_{\text{ge}}(\Gamma)$ and $\mathcal{T}_{\text{nge}}(\Gamma)$. As explained above, our goal is to find a stopping time σ which “separates” the two sets $\mathcal{T}_{\text{ge}}(\Gamma)$ and $\mathcal{T}_{\text{nge}}(\Gamma)$. Since the set $\mathcal{T}_{[0,T]}$ is partially but not totally ordered, this is a delicate issue. For the next definition, recall that \mathcal{N} denotes the collection of all \mathbb{P} -null sets in \mathcal{F}_T .

Definition 1.2. Let \mathcal{S} be a market and Γ a strategy cone. A stopping time $\sigma \in \mathcal{T}_{[0,T]}$ is called a *separating stopping time of \mathcal{S} for Γ* if

$$(1) \quad \mathcal{G}_\tau(\Gamma) \cap \{\tau < \sigma\} \subset \mathcal{N}, \quad \text{for all } \tau \in \mathcal{T}_{[0,T]} \text{ with } \mathbb{P}[\tau < \sigma] > 0, \quad (1.1)$$

$$(2) \quad \mathcal{G}_\tau(\Gamma) \cap \{\tau > \sigma\} \not\subset \mathcal{N}, \quad \text{for all } \tau \in \mathcal{T}_{[0,T]} \text{ with } \mathbb{P}[\tau > \sigma] > 0. \quad (1.2)$$

Let us briefly comment on the above definition. Property (1.1) encodes the idea that strictly before σ , we cannot “make a profit out of nothing without risk”. This implies in particular that $\tau \in \mathcal{T}_{\text{nge}}$ if $\tau < \sigma$ \mathbb{P} -a.s. The interpretation of (1.2) is that immediately after time σ , we can “make a profit out of nothing without risk” in the sense that there exists a gratis event $A \subset \{\tau > \sigma\}$ at time τ for Γ —without specifying, however, how large A is. Anticipating ourselves a bit, we note that the extreme case that $A = \{\tau > \sigma\}$ corresponds to the largest separating stopping time; see Section 3. Moreover, (1.2) implies that $\sigma \geq \tau$ for all $\tau \in \mathcal{T}_{\text{nge}}$. This is the starting point for the definition of the smallest separating stopping time; see Section 2. Finally, note that nothing is said about σ itself—we may have $\sigma \in \mathcal{T}_{\text{nge}}$ or $\sigma \in \mathcal{T}_{\text{ge}}$.

The cornerstone of the subsequent results is a refinement of Proposition III.2.2. To this end, recall from Definition III.2.1 that a strategy cone Γ is said to allow switching to numéraire strategies if for all $\tau \in \mathcal{T}_{[0,T]}$, there exists a (switching) numéraire strategy η such that $\vartheta \mathbf{1}_{[0,\tau]} + V_\tau(\vartheta)(S^{(n)})\eta \mathbf{1}_{] \tau, T]} \in \Gamma$ for all $\vartheta \in \Gamma$. Also recall that the key example of such a strategy cone is $\Gamma = \mathcal{U}$, the collection of all undefaultable strategies (cf. Definition II.4.2).

Lemma 1.3. *Let \mathcal{S} be a numéraire market, $\Gamma \subset \mathcal{U}$ a strategy cone which allows switching to numéraire strategies and $\tau_1, \tau_2 \in \mathcal{T}_{[0,T]}$. Then*

- (a) $\mathcal{G}_{\tau_1}(\Gamma) \cap \{\tau_1 \leq \tau_2\} \subset \mathcal{G}_{\tau_2}(\Gamma) \cup \mathcal{N}$,
- (b) $\mathcal{G}_{\tau_1}(\Gamma) \cap \{\tau_1 < \tau_2\} \subset \mathcal{G}_{\tau_2}(\Gamma) \cup \mathcal{N}$,
- (c) $\mathcal{G}_{\tau_1}(\Gamma) \cap \{\tau_1 = \tau_2\} \subset \mathcal{G}_{\tau_2}(\Gamma) \cup \mathcal{N}$.

Moreover, if $\tau_1 \leq \tau_2$ \mathbb{P} -a.s., $\mathcal{G}_{\tau_1}(\Gamma) \subset \mathcal{G}_{\tau_2}(\Gamma)$.

Note that the assumption $\Gamma \subset \mathcal{U}$ is crucial.

Proof. We only establish (a); the proofs of (b) and (c) are analogous.

Let η be a switching numéraire strategy at time τ_1 . We may assume without loss of generality that there exists $A \in \mathcal{G}_{\tau_1}(\Gamma) \cap \{\tau_1 \leq \tau_2\}$ with $\mathbb{P}[A] > 0$. By the definition of gratis events, there exists a contingent claim F_1 at time τ_1 with $\text{supp } F_1 = A$ \mathbb{P} -a.s. Set $g := F_1(S^{(n)})$. We claim that g is \mathcal{F}_{τ_2} -measurable. Indeed, let $c > 0$. Then $\{g > c\} \in \mathcal{F}_{\tau_1}$ because F_1 is a contingent claim at time τ_1 . \mathbb{P} -completeness of \mathcal{F}_{τ_2} and the fact that $\{g > 0\} = A \subset \{\tau_1 \leq \tau_2\}$ \mathbb{P} -a.s. give $\{g > c\} = \{g > c\} \cap \{g > 0\} \in \mathcal{F}_{\tau_2}$, and the claim follows from nonnegativity of g . Let F_2 be the contingent claim at time τ_2 from Proposition II.5.3 satisfying $F_2(S^{(n)}) = g$. Then $\text{supp } F_2 = A$ \mathbb{P} -a.s. We proceed to show that $\Pi^*(F_2 | \Gamma) = 0$. To this end, let $\varepsilon \in (0, 1)$, $\delta > 0$ and C be a positive contingent claim at time 0. By Corollary III.1.7, there exists $\vartheta \in \Gamma$ such that

$$V_0(\vartheta) \leq \delta C \quad \mathbb{P}\text{-a.s.} \quad \text{and} \quad \mathbb{P}[V_{\tau_1}(\vartheta) \geq F_1] \geq 1 - \varepsilon.$$

Set $\tilde{\vartheta} := \vartheta \mathbf{1}_{[0,\tau_1]} + V_{\tau_1}(\vartheta)(S^{(n)})\eta \mathbf{1}_{] \tau_1, T]} \in \Gamma$. Then $V_0(\tilde{\vartheta}) = V_0(\vartheta) \leq \delta C$ and $V_{\tau_2}(\tilde{\vartheta})(S^{(n)}) \geq 0 = g$ on $\{g = 0\}$ because $\Gamma \subset \mathcal{U}$. This together with the fact that

$\{g > 0\} \subset \{\tau_1 \leq \tau_2\}$ \mathbb{P} -a.s. gives

$$\begin{aligned} \mathbb{P}[V_{\tau_2}(\tilde{\vartheta}) < F_2] &= \mathbb{P}[V_{\tau_2}(\tilde{\vartheta})(S^{(n)}) < g] = \mathbb{P}[\{V_{\tau_2}(\tilde{\vartheta})(S^{(n)}) < g\} \cap \{g > 0\}] \\ &\leq \mathbb{P}[\{V_{\tau_2}(\tilde{\vartheta})(S^{(n)}) < g\} \cap \{\tau_1 \leq \tau_2\}] \\ &= \mathbb{P}[\{V_{\tau_1}(\vartheta)(S^{(n)}) < g\} \cap \{\tau_1 \leq \tau_2\}] \\ &\leq \mathbb{P}[V_{\tau_1}(\vartheta)(S^{(n)}) < g] = \mathbb{P}[V_{\tau_1}(\vartheta) < F_1] \leq \varepsilon. \end{aligned}$$

Now the claim follows from Corollary III.1.7. \square

2 The lower separating stopping time

In this section, we show the existence of a smallest separating stopping time under a mild technical condition. As already noted in the discussion after Definition 1.2, if a separating stopping time σ for a strategy cone Γ exists, then necessarily $\sigma \geq \tau$ for all $\tau \in \mathcal{T}_{\text{nge}}$. So a natural candidate for a separating stopping time σ is the smallest stopping time that dominates all stopping times in \mathcal{T}_{nge} .

Theorem 2.1. *Let \mathcal{S} be a numéraire market and $\Gamma \subset \mathcal{U}$ a strategy cone which allows switching to numéraire strategies. Then there exists a smallest separating stopping time for Γ , which we denote by $\underline{\sigma}_\Gamma$ and call the lower separating stopping time of \mathcal{S} for Γ . It is explicitly given by*

$$\underline{\sigma}_\Gamma := \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{\text{nge}}(\Gamma)} \tau. \quad (2.1)$$

Note that for $\mathcal{T}_{\text{ge}}(\Gamma) = \emptyset$, trivially $\sigma_\Gamma = T$.

Proof. First, we show that the set $\mathcal{T}_{\text{nge}}(\Gamma)$ is closed under taking maxima. Let $\tau_1, \tau_2 \in \mathcal{T}_{\text{nge}}(\Gamma)$ and set $\tilde{\tau} := \tau_1 \vee \tau_2$. Then by Lemma 1.3 (c) and the fact that $\mathcal{G}_{\tau_i}(\Gamma) = \emptyset$, $i = 1, 2$,

$$\mathcal{G}_{\tilde{\tau}}(\Gamma) = (\mathcal{G}_{\tilde{\tau}}(\Gamma) \cap \{\tilde{\tau} = \tau_1\}) \cup (\mathcal{G}_{\tilde{\tau}}(\Gamma) \cap \{\tilde{\tau} = \tau_2\}) \subset \mathcal{N},$$

and so $\tilde{\tau} \in \mathcal{T}_{\text{nge}}(\Gamma)$.

Next, define $\underline{\sigma}_\Gamma$ by (2.1). We proceed to establish that this is indeed a separating stopping time for Γ . That it is a stopping time follows from the first step because closedness of $\mathcal{T}_{\text{nge}}(\Gamma)$ under taking maxima implies that there exists a nondecreasing sequence $(\tau_n)_{n \in \mathbb{N}}$ in $\mathcal{T}_{\text{nge}}(\Gamma)$ such that

$$\underline{\sigma}_\Gamma = \lim_{n \rightarrow \infty} \tau_n \quad \mathbb{P}\text{-a.s.} \quad (2.2)$$

To establish (1.1) for $\underline{\sigma}_\Gamma$, let $\tau \in \mathcal{T}_{[0, T]}$ with $\mathbb{P}[\tau < \sigma] > 0$ and $(\tau_n)_{n \in \mathbb{N}}$ be a nondecreasing sequence in $\mathcal{T}_{\text{nge}}(\Gamma)$ satisfying (2.2). Then by Lemma 1.3 (b) and the fact that each $\mathcal{G}_{\tau_n}(\Gamma) = \emptyset$,

$$\mathcal{G}_\tau(\Gamma) \cap \{\tau < \underline{\sigma}_\Gamma\} \subset \bigcup_{n \in \mathbb{N}} \mathcal{G}_\tau(\Gamma) \cap \{\tau < \tau_n\} \cup \mathcal{N} = \mathcal{N}.$$

Property (1.2) for $\underline{\sigma}_\Gamma$ follows immediately from the definition of $\underline{\sigma}_\Gamma$.

Finally, it follows from the discussion preceding this theorem and the definition of the essential supremum, that $\underline{\sigma}_\Gamma$ defined by (2.1) is the smallest separating stopping time for Γ . \square

Note that we either have $\underline{\sigma}_\Gamma \in \mathcal{T}_{\text{nge}}(\Gamma)$ or $\underline{\sigma}_\Gamma \in \mathcal{T}_{\text{ge}}(\Gamma)$; see the examples in Section 4. In the first case, $\underline{\sigma}_\Gamma$ can be interpreted as the last time, when there are no gratis events in the market. In the second case, $\underline{\sigma}_\Gamma$ can be interpreted as the first time, when there are gratis events in the market. Moreover, there exists a nondecreasing sequence $(\tau_n)_{n \in \mathbb{N}}$ in $\mathcal{T}_{\text{nge}}(\Gamma)$ which foretells when gratis events “enter the market”. More precisely, it is easy to check that $\underline{\sigma}_\Gamma$ is predictable on $G_{\underline{\sigma}_\Gamma}$, where $G_{\underline{\sigma}_\Gamma} \in \mathcal{G}_{\underline{\sigma}_\Gamma}(\Gamma)$ is a maximal gratis event at time $\underline{\sigma}_\Gamma$ for Γ (cf. Definition III.1.10 and Corollary III.1.11).

3 The upper separating stopping time

In this section, we show the existence of a largest separating stopping time. As already noted in the discussion after Definition 1.2, property (1.2) of a separating stopping time σ for Γ gives for each $\tau \in \mathcal{T}_{[0,T]}$ with $\mathbb{P}[\tau > \sigma] > 0$ a gratis event $A_\tau \subset \{\tau > \sigma\}$ at time τ for Γ . The extreme case is of course that $\{\tau > \sigma\}$ itself is a gratis event at time τ for Γ . From an economic perspective, this means that strictly after time σ , in *any currency unit*, one can get *arbitrarily much* with *probability almost 1* with an *arbitrarily low initial investment* using strategies having *nonnegative value processes* (cf. Proposition III.1.5 and Remark III.1.6 (a)).

Theorem 3.1. *Let \mathcal{S} be a numéraire market and $\Gamma \subset \mathcal{U}$ a strategy cone which allows switching to numéraire strategies. Then there exists a largest separating stopping time for Γ , which we denote by $\bar{\sigma}_\Gamma$ and call the upper separating stopping time of \mathcal{S} for Γ . It satisfies*

$$(2') \quad \{\tau > \bar{\sigma}_\Gamma\} \subset \mathcal{G}_\tau(\Gamma), \quad \text{for all } \tau \in \mathcal{T}_{[0,T]} \text{ with } \mathbb{P}[\tau > \bar{\sigma}_\Gamma] > 0. \quad (3.1)$$

Note that for $\mathcal{T}_{\text{ge}}(\Gamma) = \emptyset$, trivially $\bar{\sigma}_\Gamma = T$.

Proof. First, we show that if a separating stopping time σ for Γ satisfies (3.1), then it is the largest one. Seeking a contradiction, suppose there exists another separating stopping time $\tilde{\sigma}$ for Γ with $\mathbb{P}[\tilde{\sigma} > \sigma] > 0$. Then there is $r \in (0, T)$ with $\mathbb{P}[\sigma < r < \tilde{\sigma}] > 0$. By property (3.1) for σ , $\{\sigma < r\} \in \mathcal{G}_r(\Gamma)$, and so $\{\sigma < r < \tilde{\sigma}\} \in \mathcal{G}_r(\Gamma) \cap \{r < \tilde{\sigma}\}$ by Proposition III.1.4. On the other hand, $\mathcal{G}_r(\Gamma) \cap \{r < \tilde{\sigma}\} \subset \mathcal{N}$ by (1.1) for $\tilde{\sigma}$, and we arrive at a contradiction.

Next, we establish the existence of a separating stopping time for Γ satisfying (3.1). To this end, for $\tau \in \mathcal{T}_{\text{nge}}(\Gamma)$, set $G_\tau := \emptyset$, and for $\tau \in \mathcal{T}_{\text{ge}}(\Gamma)$, let $G_\tau \in \mathcal{G}_\tau(\Gamma)$ be a maximal gratis event (cf. Definition III.1.10 and Corollary III.1.11). Define the process $Z = (Z_t)_{t \in [0,T]}$ by $Z_t := \mathbb{1}_{G_t}$. Intuitively, we would like to define $\bar{\sigma}_\Gamma$ as the first time that Z takes the value 1. For technical reasons, however, things

are not so easy because Z might not have càdlàg paths. Therefore, we have to argue more carefully and work with suitable modifications of the process Z . To this end, set $L_0 := \emptyset$, and for $0 < t \leq T$, define

$$L_t := \bigcup_{\substack{r < t \\ r \text{ rational}}} G_r.$$

Moreover, set $U_T := \Omega$, and for $0 \leq t < T$ define

$$U_t := \bigcap_{\substack{r > t \\ r \text{ rational}}} L_r.$$

Then for all $0 \leq s < t \leq T$,

$$L_s \subset U_s \subset L_t \subset U_t.$$

In addition, for any increasing sequence $(t_n)_{n \in \mathbb{N}}$ in $[0, T]$ with $\lim_{n \rightarrow \infty} t_n = t$,

$$L_t = \bigcup_{\substack{r < t \\ r \text{ rational}}} G_r \subset \bigcup_{\substack{r < t \\ r \text{ rational}}} L_{r + \frac{t-r}{2}} \subset \bigcup_{n \in \mathbb{N}} L_{t_n} \subset \bigcup_{n \in \mathbb{N}} U_{t_n} \subset U_t. \quad (3.2)$$

Likewise, for any decreasing sequence $(t_n)_{n \in \mathbb{N}}$ in $[0, T]$ with $\lim_{n \rightarrow \infty} t_n = t$,

$$U_t \subset \bigcap_{n \in \mathbb{N}} U_{t_n} \subset \bigcap_{n \in \mathbb{N}} L_{t_n} \subset \bigcap_{\substack{r > t \\ r \text{ rational}}} L_r = U_t. \quad (3.3)$$

Define the processes $X = (X_t)_{t \in [0, T]}$ and $Y = (Y_t)_{t \in [0, T]}$ by $X_t := \mathbf{1}_{L_t}$ and $Y_t := \mathbf{1}_{U_t}$. Then X and Y are $\{0, 1\}$ -valued increasing adapted processes; (3.2) implies that for *all* ω and all $0 < t \leq T$,

$$X_t = \lim_{s \uparrow t} X_s = \lim_{s \uparrow t} Y_s,$$

and (3.3) implies that for *all* ω and for all $0 \leq t < T$,

$$Y_t = \lim_{s \downarrow t} Y_s = \lim_{s \downarrow t} X_s. \quad (3.4)$$

We proceed to show that for all $\tau \in \mathcal{T}_{[0, T]}$,

$$X_\tau \leq \mathbf{1}_{G_\tau} \leq Y_\tau \text{ } \mathbb{P}\text{-a.s.} \quad (3.5)$$

For deterministic $t \in [0, T]$, by construction of G_t and U_t and Lemma 1.3,

$$L_t \subset G_t \subset U_t \text{ } \mathbb{P}\text{-a.s.},$$

and so (3.5) holds for all deterministic $t \in [0, T]$. For a general stopping time $\tau \in \mathcal{T}_{[0, T]}$, we establish (3.5) by an approximation argument. To this end, for $n \in \mathbb{N}$, define the processes $X^{(n)} = (X_t^{(n)})_{t \in [0, T]}$ and $Y^{(n)} = (Y_t^{(n)})_{t \in [0, T]}$ by

$$\begin{aligned} X_t^{(n)} &= \sum_{k=1}^{2^n} X_{\frac{k-1}{2^n}T} \mathbf{1}_{\{\frac{k-1}{2^n}T \leq t < \frac{k}{2^n}T\}} + X_T \mathbf{1}_{\{t=T\}}, \\ Y_t^{(n)} &= Y_0 \mathbf{1}_{\{t=0\}} + \sum_{k=1}^{2^n} Y_{\frac{k}{2^n}T} \mathbf{1}_{\{\frac{k-1}{2^n}T < t \leq \frac{k}{2^n}T\}}. \end{aligned}$$

Left-continuity of X and right-continuity of Y give for *all* ω ,

$$\lim_{n \rightarrow \infty} X^{(n)} = X \quad \text{and} \quad \lim_{n \rightarrow \infty} Y^{(n)} = Y.$$

Thus, to establish (3.5), it suffices to show that for all $n \in \mathbb{N}$ and all $\tau \in T_{[0, T]}$,

$$X_\tau^{(n)} \leq \mathbf{1}_{G_\tau} \leq Y_\tau^{(n)} \quad \mathbb{P}\text{-a.s.}$$

Let $\tau \in \mathcal{T}_{[0, T]}$ and $n \in \mathbb{N}$. By Lemma 1.3 and (3.5) for deterministic $t \in [0, T]$,

$$\begin{aligned} X_\tau^{(n)} &= \sum_{k=1}^{2^n} \mathbf{1}_{L_{\frac{k-1}{2^n}T}} \mathbf{1}_{\{\frac{k-1}{2^n}T \leq \tau < \frac{k}{2^n}T\}} + \mathbf{1}_{L_T} \mathbf{1}_{\{\tau=T\}} \\ &\leq \sum_{k=1}^{2^n} \mathbf{1}_{G_{\frac{k-1}{2^n}T}} \mathbf{1}_{\{\frac{k-1}{2^n}T \leq \tau < \frac{k}{2^n}T\}} + \mathbf{1}_{G_T} \mathbf{1}_{\{\tau=T\}} \\ &\leq \sum_{k=1}^{2^n} \mathbf{1}_{G_\tau} \mathbf{1}_{\{\frac{k-1}{2^n}T \leq \tau < \frac{k}{2^n}T\}} + \mathbf{1}_{G_\tau} \mathbf{1}_{\{\tau=T\}} \\ &= \mathbf{1}_{G_\tau} \\ &= \mathbf{1}_{G_\tau} \mathbf{1}_{\{\tau=0\}} + \sum_{k=1}^{2^n} \mathbf{1}_{G_\tau} \mathbf{1}_{\{\frac{k-1}{2^n}T < \tau \leq \frac{k}{2^n}T\}} \\ &\leq \mathbf{1}_{G_0} \mathbf{1}_{\{\tau=0\}} + \sum_{k=1}^{2^n} \mathbf{1}_{G_{\frac{k}{2^n}T}} \mathbf{1}_{\{\frac{k-1}{2^n}T < \tau \leq \frac{k}{2^n}T\}} \\ &\leq \mathbf{1}_{U_0} \mathbf{1}_{\{\tau=0\}} + \sum_{k=1}^{2^n} \mathbf{1}_{U_{\frac{k}{2^n}T}} \mathbf{1}_{\{\frac{k-1}{2^n}T < \tau \leq \frac{k}{2^n}T\}} \\ &= Y_\tau^{(n)} \quad \mathbb{P}\text{-a.s.} \end{aligned}$$

Finally, define the stopping time $\bar{\sigma}_\Gamma \in \mathcal{T}_{[0, T]}$ by

$$\bar{\sigma}_\Gamma := \inf\{t \geq 0 : X_t = 1\} = \inf\{t \geq 0 : Y_t = 1\}, \quad (3.6)$$

where the equality follows from (3.4). We proceed to show that $\bar{\sigma}_\Gamma$ satisfies (1.1) and (3.1). To establish (1.1), let $\tau \in \mathcal{T}_{[0, T]}$ with $\mathbb{P}[\tau < \bar{\sigma}_\Gamma] > 0$. We may assume without loss of generality that $\mathcal{G}_\tau(\Gamma) \neq \emptyset$. So let $A \in \mathcal{G}_\tau(\Gamma)$. Then by maximality of G_τ in $\mathcal{G}_\tau(\Gamma)$, (3.5) and (3.6),

$$\mathbf{1}_{A \cap \{\tau < \bar{\sigma}_\Gamma\}} = \mathbf{1}_A \mathbf{1}_{\{\tau < \bar{\sigma}_\Gamma\}} \leq \mathbf{1}_{G_\tau} \mathbf{1}_{\{\tau < \bar{\sigma}_\Gamma\}} \leq Y_\tau \mathbf{1}_{\{\tau < \bar{\sigma}_\Gamma\}} = 0 \quad \mathbb{P}\text{-a.s.}$$

To establish (3.1), let $\tau \in \mathcal{T}_{[0, T]}$ with $\mathbb{P}[\tau > \bar{\sigma}_\Gamma] > 0$. Then (3.6) and (3.5) yield

$$\mathbf{1}_{\{\tau > \bar{\sigma}_\Gamma\}} = X_\tau \mathbf{1}_{\{\tau > \bar{\sigma}_\Gamma\}} \leq X_\tau \leq \mathbf{1}_{G_\tau} \quad \mathbb{P}\text{-a.s.}$$

Thus, $\mathbb{P}[G_\tau] > 0$, $\mathcal{G}_\tau(\Gamma) \neq \emptyset$, and the claim follows from Proposition III.1.4. \square

As in the case of the lower separating stopping time for Γ , we either have $\bar{\sigma}_\Gamma \in \mathcal{T}_{\text{nge}}(\Gamma)$ or $\bar{\sigma}_\Gamma \in \mathcal{T}_{\text{ge}}(\Gamma)$; see the examples in Section 4. In the first case, the lower and upper stopping time for Γ coincide.

Corollary 3.2. *Let \mathcal{S} be a numéraire market and $\Gamma \subset \mathcal{U}$ a strategy cone which allows switching to numéraire strategies. If $\bar{\sigma}_\Gamma \in \mathcal{T}_{\text{nge}}(\Gamma)$, then $\underline{\sigma}_\Gamma = \bar{\sigma}_\Gamma$ and $\underline{\sigma}_\Gamma \in \mathcal{T}_{\text{nge}}(\Gamma)$.*

Proof. The inequality $\underline{\sigma}_\Gamma \leq \bar{\sigma}_\Gamma$ is trivial, and the other inequality follows immediately from the fact that $\bar{\sigma}_\Gamma \in \mathcal{T}_{\text{nge}}(\Gamma)$ and the representation (2.1) of $\underline{\sigma}_\Gamma$. \square

4 Examples

We illustrate the above results by several examples. They show that the lower and the upper separating stopping time may coincide or be different. In the first case, $\underline{\sigma}_\Gamma = \bar{\sigma}_\Gamma$ may be either in $\mathcal{T}_{\text{nge}}(\Gamma)$ or in $\mathcal{T}_{\text{ge}}(\Gamma)$. In the second case, by Corollary 3.2, we necessarily have $\bar{\sigma}_\Gamma \in \mathcal{T}_{\text{ge}}(\Gamma)$, but $\underline{\sigma}_\Gamma$ may be either in $\mathcal{T}_{\text{nge}}(\Gamma)$ or in $\mathcal{T}_{\text{ge}}(\Gamma)$.

Example 4.1. Let $(N_t)_{t \in [0,1]}$ be a standard Poisson process with parameter $\lambda > 0$ on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0,1]}, \mathbb{P})$ satisfying the usual conditions with \mathcal{F}_0 being \mathbb{P} -a.s. trivial. Let \mathcal{S} be the market generated by the classic model $S = (1, X_t)_{t \in [0,1]}$, where

- (a) $X_t := 1 + \mathbb{1}_{\{t \geq 1/2\}}$,
- (b) $X_t := 1 + (t - 1/2)\mathbb{1}_{\{t \geq 1/2\}}$,
- (c) $X_t := 1 + N_t$,
- (d) $X_t := 1 + N_t \mathbb{1}_{\{t \geq 1/2\}}$.

Then, with σ denoting the first jump time of N , i.e., $\sigma = \inf\{t \geq 0 : N_t = 1\}$,

- (a) $\underline{\sigma}_\mathcal{U} = \bar{\sigma}_\mathcal{U} = 1/2 \in \mathcal{T}_{\text{ge}}(\mathcal{U})$.
- (b) $\underline{\sigma}_\mathcal{U} = \bar{\sigma}_\mathcal{U} = 1/2 \in \mathcal{T}_{\text{nge}}(\mathcal{U})$.
- (c) $\underline{\sigma}_\mathcal{U} = 0 \in \mathcal{T}_{\text{nge}}(\mathcal{U})$ and $\bar{\sigma}_\mathcal{U} = \sigma \wedge 1 \in \mathcal{T}_{\text{ge}}(\mathcal{U})$.
- (d) $\underline{\sigma}_\mathcal{U} = 1/2 \in \mathcal{T}_{\text{ge}}(\mathcal{U})$ and $\bar{\sigma}_\mathcal{U} = (\sigma \vee 1/2) \wedge 1 \in \mathcal{T}_{\text{ge}}(\mathcal{U})$.

We only establish the most difficult case (d); the arguments for (a) and (b) are easy, and the argument for (c) is similar to the one for (d).

To establish that $\underline{\sigma}_\mathcal{U} = 1/2$, by the explicit representation (2.1), it suffices to show that there exists a nondecreasing sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times in $\mathcal{T}_{\text{nge}}(\mathcal{U})$ with $\lim_{n \rightarrow \infty} \tau_n = 1/2$ and that all $\tau \in \mathcal{T}_{[0,T]}$ with $\mathbb{P}[\tau > 1/2] > 0$ are in $\mathcal{T}_{\text{ge}}(\mathcal{U})$. The first assertion follows from considering $\tau_n := 1/2 - \frac{1}{1+n}$, $n \in \mathbb{N}$. For the second assertion, take $\tau \in \mathcal{T}_{[0,T]}$ with $\mathbb{P}[\tau > 1/2] > 0$. Set $\vartheta^{(1)} := (-1, 1)$ and note that $\vartheta^{(1)} \in \mathcal{U}$ with $V_0(\vartheta^{(1)}) = 0$. We proceed to show that $\mathbb{E}[V_\tau(\vartheta^{(1)})(S)] > 0$, which immediately gives $\tau \in \mathcal{T}_{\text{ge}}(\mathcal{U})$. To this end, denote by $(\tilde{N}_t)_{t \in [0,1]}$ the compensated Poisson process, i.e., $\tilde{N}_t = N_t - \lambda t$, $t \in [0, 1]$, and

define the process $(M_t)_{t \in [0,1]}$ by $M_t := (\tilde{N}_t - \tilde{N}_{1/2})\mathbf{1}_{\{t \geq 1/2\}}$. Then M is a uniformly integrable martingale since \tilde{N} is so. Monotonicity of N and the optional stopping theorem for M give

$$\begin{aligned} \mathbb{E}[V_\tau(\vartheta^{(1)})(S)] &= \mathbb{E}[N_\tau \mathbf{1}_{\{\tau \geq 1/2\}}] \geq \mathbb{E}[(N_\tau - N_{1/2})\mathbf{1}_{\{\tau \geq 1/2\}}] \\ &= \mathbb{E}[M_\tau + \lambda(\tau - 1/2)\mathbf{1}_{\{\tau \geq 1/2\}}] = 0 + \lambda \mathbb{E}[(\tau - 1/2)\mathbf{1}_{\{\tau \geq 1/2\}}] \\ &= \lambda \mathbb{E}[(\tau - 1/2)\mathbf{1}_{\{\tau > 1/2\}}] > 0, \end{aligned}$$

where the last inequality follows from the fact that $\mathbb{P}[\tau > 1/2] > 0$.

To establish that $\bar{\sigma}_{\mathcal{U}} = (\sigma \vee 1/2) \wedge 1$, we have to check (1.1) and (3.1). For, (1.1), take $\tau \in \mathcal{T}_{[0,T]}$ with $\mathbb{P}[\tau < (\sigma \vee 1/2) \wedge 1] > 0$. Seeking a contradiction, we may assume without loss of generality that $\mathcal{G}_\tau(\mathcal{U}) \neq \emptyset$ and that there is $A \in \mathcal{G}_\tau(\mathcal{U})$ with $A \subset \{\tau < (\sigma \vee 1/2) \wedge 1\}$ P-a.s. By the definition of gratis events, there exist a nonzero contingent claim F at time τ with $\{F > 0\} = A$ P-a.s. Fix $S \in \mathcal{S}$ and let C be the contingent claim at time 0 from Proposition II.5.3 satisfying $C(S) = 1$. Choose $\varepsilon > 0$ small enough that $\mathbb{P}[F(S) \geq \varepsilon] \geq \varepsilon$, and note that $\mathbb{P}[A] \geq \varepsilon$. By Corollary III.1.7, there is $\vartheta \in \mathcal{U}$ such that

$$V_0(\vartheta) \leq \frac{\varepsilon}{2}C \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta) \geq F] \geq 1 - \frac{\varepsilon}{2}$$

This together with the choice of ε implies that

$$\mathbb{P}[\{V_\tau(\vartheta)(S) \geq \varepsilon\} \cap A] \geq \mathbb{P}[A] - \frac{\varepsilon}{2} > 0.$$

By the definition of X in (d), we arrive at the contradiction

$$V_\tau(\vartheta)(S) = V_0(\vartheta)(S) \leq \frac{\varepsilon}{2} < \varepsilon \quad \text{on } A \subset \{\tau < (\sigma \vee 1/2) \wedge 1\} \text{ P-a.s.}$$

For (3.1), take $\tau \in \mathcal{T}_{[0,T]}$ with $\mathbb{P}[\tau > (\sigma \vee 1/2) \wedge 1] > 0$. Set $\vartheta^{(1)} := (-1, 1) \in \mathcal{U}$ and $\vartheta^{(2)} := (1, 0) \in \mathcal{U}$. Then

$$V_0(\vartheta^{(1)}) = 0 \quad \text{and} \quad V_\tau(\vartheta^{(1)}) \geq V_\tau(\vartheta^{(2)})\mathbf{1}_{\{\tau > (\sigma \vee 1/2) \wedge 1\}},$$

and so $\{(\sigma \vee 1/2) \wedge 1 < \tau\} \in \mathcal{G}_\tau(\mathcal{U})$ since

$$\text{supp } V_\tau(\vartheta^{(2)})\mathbf{1}_{\{\tau > (\sigma \vee 1/2) \wedge 1\}} = \{(\sigma \vee 1/2) \wedge 1 < \tau\} \text{ P-a.s.}$$

Finally, $\bar{\sigma}_{\mathcal{U}} \in \mathcal{T}_{\text{ge}}(\Gamma)$ by Corollary 3.2 and the fact that $\mathbb{P}[\underline{\sigma}_{\mathcal{U}} < \bar{\sigma}_{\mathcal{U}}] > 0$. To show that $\underline{\sigma}_{\mathcal{U}} \in \mathcal{T}_{\text{ge}}(\Gamma)$, let $\vartheta^{(1)} = (-1, 1)$ and $\vartheta^{(2)} = (1, 0)$ be as above. Then $F := \mathbf{1}_{\{\sigma \leq 1/2\}}V_{1/2}(\vartheta^{(2)})$ is a nonzero contingent claim at time 1/2 by the fact that σ is exponentially distributed with parameter λ and $V_{1/2}(\vartheta^{(2)})$ is a positive contingent claim. As $V_0(\vartheta^{(1)}) = 0$ and $V_{1/2}(\vartheta^{(1)}) = F$, this implies that $\mathcal{G}_{1/2}(\mathcal{U}) \neq \emptyset$. This ends the example.

Chapter V

Absolutely continuous measures for markets failing NINA

In this chapter, we study markets which fail numéraire-independent no-arbitrage (NINA) (see Definition III.2.5) under the physical measure \mathbb{P} . We seek to answer the question whether it is then possible to find an *absolutely continuous* probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T such that the market satisfies NINA *under* \mathbb{Q} .

The basic idea how to tackle this problem is rather simple: Pick a maximal gratis event $G_T \in \mathcal{G}_T(\mathcal{U}(S))$ and assume that $\mathbb{P}[G_T] < 1$; note that if $\mathbb{P}[G_T] = 1$, there is no hope to construct \mathbb{Q} (cf. Corollaries 2.4 and 3.13). Define the probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T by $d\mathbb{Q} := \mathbb{1}_{G_T^c} \frac{1}{\mathbb{P}[G_T^c]} d\mathbb{P}$. Then \mathbb{Q} is supported outside of $\mathcal{G}_T(\mathcal{U}(S))$, and one might hope that the market satisfies NINA under \mathbb{Q} .

After addressing some technical issues related to an absolutely continuous change of measure in Section 1, we show that the above procedure works perfectly well for *continuous* markets. For *general* markets, however, things are not so easy. The crucial problem is that a strategy which is undefaultable under \mathbb{Q} might not be so under \mathbb{P} , and the obvious quick fix of stopping a strategy when it becomes negative under \mathbb{P} only works for continuous markets. This problem can even arise in finite discrete time with countable Ω ; see Example 3.1. Therefore, we have to refine the argument for general markets. This is carried out in Section 3.

A problem somewhat similar to the one considered in this chapter has been studied by Levental and Skorokhod [54], Delbaen and Schachermayer [10] and Strasser [77]. They consider a classic model $S = (1, X)$ (II.1.3) which is assumed to satisfy NA but might fail NFLVR under \mathbb{P} . They show that if X has continuous paths, there exists an absolutely continuous measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T such that X satisfies NFLVR under \mathbb{Q} .

1 Technical preliminaries

When passing from \mathbb{P} to a probability measure \mathbb{Q} which is absolutely continuous but not equivalent to \mathbb{P} on \mathcal{F}_T , several (mainly technical) issues have to be addressed: If we start with a market \mathcal{S} under \mathbb{P} , what is the corresponding market under \mathbb{Q} ? If ϑ is a self-financing strategy under \mathbb{Q} , does there exist a self-financing

strategy $\tilde{\vartheta}$ under \mathbb{P} such that ϑ and $\tilde{\vartheta}$ are indistinguishable under \mathbb{Q} ? How are superreplication prices under \mathbb{P} related to superreplication prices under \mathbb{Q} ?

The goal of this section is to answer the above and other questions and to lay the ground for the following sections. Throughout this section, let $\mathbb{Q} \ll \mathbb{P}$ be a probability measure that is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_T .

1.1 Semimartingales under \mathbb{P} and \mathbb{Q}

We write $({}^{\mathbb{Q}}\mathcal{F}_t)_{0 \leq t \leq T}$ for the completion of the filtration $(\mathcal{F}_t)_{0 \leq t \leq T}$ with respect to \mathbb{Q} . Then $\mathcal{F}_t \subset {}^{\mathbb{Q}}\mathcal{F}_t$ for all $0 \leq t \leq T$ and $(\Omega, {}^{\mathbb{Q}}\mathcal{F}, ({}^{\mathbb{Q}}\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{Q})$ satisfies the usual conditions with ${}^{\mathbb{Q}}\mathcal{F}_0$ being \mathbb{Q} -trivial. We write $\mathbb{E}_{\mathbb{Q}}[\cdot]$ for the expectation under \mathbb{Q} , ${}^{\mathbb{Q}}\mathcal{N}$ for the set of all \mathbb{Q} -null sets in ${}^{\mathbb{Q}}\mathcal{F}_T$, ${}^{\mathbb{Q}}\mathcal{T}_{[0,T]}$ for the set of all \mathbb{Q} -stopping times taking values in $[0, T]$, and so forth. For an \mathbb{R}^N -valued \mathbb{Q} -semimartingale $X = (X^1, \dots, X^N)$, we denote by ${}^{\mathbb{Q}}L(X)$ the set of all \mathbb{R}^N -valued \mathbb{Q} -predictable processes that are \mathbb{Q} -integrable with respect to X , and for $\zeta \in {}^{\mathbb{Q}}L(X)$, we denote by $\zeta \bullet^{\mathbb{Q}} X$ the stochastic integral $(\int_{(0,t]} \zeta_u dX_u)_{0 \leq t \leq T}$ computed under \mathbb{Q} .

The following technical result recalls the relationship between \mathbb{P} - and \mathbb{Q} -semimartingales and \mathbb{P} - and \mathbb{Q} -predictable processes.

Proposition 1.1. *Let X be an \mathbb{R}^N -valued \mathbb{P} -semimartingale. Then X is a \mathbb{Q} -semimartingale, $L(X) \subset {}^{\mathbb{Q}}L(X)$, and for each $\zeta \in L(X) \subset {}^{\mathbb{Q}}L(X)$, $\zeta \bullet^{\mathbb{Q}} X$ is \mathbb{Q} -indistinguishable from $\zeta \bullet X$. Moreover, for all bounded \mathbb{Q} -predictable processes ζ , there is a bounded \mathbb{P} -predictable process $\tilde{\zeta}$ which is \mathbb{Q} -indistinguishable from ζ .*

Proof. The first three claims follow from [38, Théorème 7.24 (c)] and [63, Lemme V.2]. For the final claim, let ζ be bounded \mathbb{Q} -predictable. Choose $K > 0$ large enough that $\sup_{t \in [0, T]} \|\zeta_t\| \leq K$ \mathbb{Q} -a.s., where $\|\cdot\|$ denotes any norm in \mathbb{R}^N . By [38, Proposition 1.1 (b)], there exists a \mathbb{P} -predictable process $\hat{\zeta}$ which is \mathbb{Q} -indistinguishable from ζ . Set $\tilde{\zeta} := \hat{\zeta} \mathbf{1}_{\{\|\hat{\zeta}\| \leq K\}}$. Then $\tilde{\zeta}$ is bounded \mathbb{P} -predictable and \mathbb{Q} -indistinguishable from ζ since

$$\left\{ \sup_{t \in [0, t]} \|\tilde{\zeta}_t\| > K \right\} \subset \left\{ \hat{\zeta} \neq \zeta \right\} \cup \left\{ \sup_{t \in [0, T]} \|\zeta_t\| > K \right\}. \quad \square$$

1.2 Markets, contingent claims, etc. under \mathbb{P} and \mathbb{Q}

Any \mathbb{P} -exchange rate process is a fortiori a \mathbb{Q} -exchange rate process, but in general, there are more \mathbb{Q} -exchange rate processes than \mathbb{P} -exchange rate processes. Therefore, if we want to consider a \mathbb{P} -market \mathcal{S} under \mathbb{Q} , we have to extend \mathcal{S} in an appropriate way.

Definition 1.2. Let \mathcal{S} be a \mathbb{P} -market. The \mathbb{Q} -market corresponding to \mathcal{S} , denoted by ${}^{\mathbb{Q}}\mathcal{S}$, is the \mathbb{Q} -market generated by some (and hence every) $S \in \mathcal{S}$.

Clearly, $\mathcal{S} \subset {}^{\mathbb{Q}}\mathcal{S}$. We denote the set of all \mathbb{Q} -self-financing strategies for ${}^{\mathbb{Q}}\mathcal{S}$ by $L^{\text{sf}}({}^{\mathbb{Q}}\mathcal{S})$, and the set of all \mathbb{Q} -undefaultable strategies for ${}^{\mathbb{Q}}\mathcal{S}$ by $\mathcal{U}({}^{\mathbb{Q}}\mathcal{S})$. Note that if $\vartheta \in L^{\text{sf}}({}^{\mathbb{Q}}\mathcal{S})$ is \mathbb{P} -predictable, it is not necessarily in $L^{\text{sf}}(\mathcal{S})$, and likewise,

if $\vartheta \in \mathcal{U}(\mathbb{Q}\mathcal{S})$ is \mathbb{P} -predictable, it is not necessarily in $\mathcal{U}(\mathcal{S})$ (or in $L^{\text{sf}}(\mathcal{S})$). This is because \mathbb{Q} is in general not equivalent to \mathbb{P} on \mathcal{F}_T , and so there are \mathbb{Q} -nullsets which have positive probability under \mathbb{P} .

We proceed to link properties of \mathbb{P} -markets to \mathbb{Q} -markets. For convenience, we simply write market instead of \mathbb{P} -market in the sequel.

First, we show that \mathbb{P} -self-financing strategies are \mathbb{Q} -self-financing, too.

Proposition 1.3. *Let \mathcal{S} be a market. Then $L^{\text{sf}}(\mathcal{S}) \subset L^{\text{sf}}(\mathbb{Q}\mathcal{S})$.*

Proof. Let $\vartheta \in L^{\text{sf}}(\mathcal{S})$. Pick $S \in \mathcal{S} \subset \mathbb{Q}\mathcal{S}$. Then by Proposition 1.1,

$$\vartheta \cdot S = \vartheta_0 \cdot S_0 + \vartheta \bullet S = \vartheta_0 \cdot S_0 + \vartheta \bullet^{\mathbb{Q}} S \quad \mathbb{Q}\text{-a.s.},$$

and the claim follows from Lemma II.2.5. \square

Corollary 1.4. *Let \mathcal{S} be a market. Then a strategy cone for \mathcal{S} is also a strategy cone for $\mathbb{Q}\mathcal{S}$. Moreover, $\mathcal{U}(\mathcal{S}) \subset \mathcal{U}(\mathbb{Q}\mathcal{S})$ and $\mathbf{b}\mathcal{U}(\mathcal{S}) \subset \mathbf{b}\mathcal{U}(\mathbb{Q}\mathcal{S})$.*

Next, we establish the technically important result, that the value process of a bounded \mathbb{Q} -self-financing strategy is \mathbb{Q} -indistinguishable from the value process of a bounded \mathbb{P} -self-financing strategy under a mild technical condition. To this end, recall from Definition II.3.4 that a market \mathcal{S} is called a bounded numéraire market if there exist a bounded numéraire strategy η such that also $S^{(\eta)}$ is bounded.

Lemma 1.5. *Let \mathcal{S} be a bounded numéraire market. Then for each $\vartheta \in \mathbf{b}L^{\text{sf}}(\mathbb{Q}\mathcal{S})$, there exists $\tilde{\vartheta} \in \mathbf{b}L^{\text{sf}}(\mathcal{S}) \subset \mathbf{b}L^{\text{sf}}(\mathbb{Q}\mathcal{S})$ such that $V(\tilde{\vartheta})(S)$ is \mathbb{Q} -indistinguishable from $V(\vartheta)(S)$ for some (and hence every) $S \in \mathbb{Q}\mathcal{S}$.*

Note that Lemma 1.5 does not say that $\tilde{\vartheta}$ itself is \mathbb{Q} -indistinguishable from ϑ ; the assertion only holds for the value processes $V(\tilde{\vartheta})$ and $V(\vartheta)$.

Proof. Let $\vartheta \in \mathbf{b}L^{\text{sf}}(\mathbb{Q}\mathcal{S})$ and η be a bounded numéraire strategy for \mathcal{S} such that also $S^{(\eta)} \in \mathcal{S} \subset \mathbb{Q}\mathcal{S}$ is bounded. By Proposition 1.1, there exists a bounded \mathbb{P} -predictable process $\tilde{\zeta}$ which is \mathbb{Q} -indistinguishable from ϑ , and by Theorem II.3.7, there exist $\tilde{\vartheta} \in \mathbf{b}L^{\text{sf}}(\mathcal{S})$ such that

$$V(\tilde{\vartheta})(S^{(\eta)}) = \tilde{\vartheta}_0 \cdot S_0^{(\eta)} + \tilde{\vartheta} \bullet S^{(\eta)} = \tilde{\zeta}_0 \cdot S_0^{(\eta)} + \tilde{\zeta} \bullet S^{(\eta)} \quad \mathbb{P}\text{-a.s.}$$

Thus, by Proposition 1.1 and the fact that $\tilde{\zeta}$ is \mathbb{Q} -indistinguishable from ϑ ,

$$\begin{aligned} V(\tilde{\vartheta})(S^{(\eta)}) &= \tilde{\vartheta}_0 \cdot S_0^{(\eta)} + \tilde{\vartheta} \bullet^{\mathbb{Q}} S^{(\eta)} = \tilde{\zeta}_0 \cdot S_0^{(\eta)} + \tilde{\zeta} \bullet^{\mathbb{Q}} S^{(\eta)} \\ &= \vartheta_0 \cdot S_0^{(\eta)} + \vartheta \bullet^{\mathbb{Q}} S^{(\eta)} = V(\vartheta)(S^{(\eta)}) \quad \mathbb{Q}\text{-a.s.} \end{aligned} \quad \square$$

We now consider the relationship between (generalised) \mathbb{P} -contingent claims for the market \mathcal{S} and (generalised) \mathbb{Q} -contingent claims for the \mathbb{Q} -market $\mathbb{Q}\mathcal{S}$.

Proposition 1.6. *Let \mathcal{S} be a market and F a generalised \mathbb{P} -contingent claim at time $\tau \in \mathcal{T}_{[0,T]} \subset \mathbb{Q}\mathcal{T}_{[0,T]}$ for \mathcal{S} . Then there exists a \mathbb{Q} -a.s. unique generalised \mathbb{Q} -contingent claim ${}^{\mathbb{Q}}F$ at time τ for $\mathbb{Q}\mathcal{S}$ satisfying ${}^{\mathbb{Q}}F(S) = F(S)$ for all $S \in \mathcal{S}$. Moreover, if F is improper, defaultable, a contingent claim or a positive contingent claim, then ${}^{\mathbb{Q}}F$ is so, too.*

Proof. Fix $S \in \mathcal{S} \subset {}^{\mathbb{Q}}\mathcal{S}$ and set $g := F(S)$. By Proposition II.5.3, there exists a \mathbb{Q} -a.s. unique \mathbb{Q} -contingent claim ${}^{\mathbb{Q}}F$ at time τ such that ${}^{\mathbb{Q}}F(S) = g = F(S)$. This gives uniqueness of ${}^{\mathbb{Q}}F$. Existence follows from the exchange rate consistency (II.5.1) of F and ${}^{\mathbb{Q}}F$. The remaining claims are straightforward. \square

The next result gives a converse of Proposition 1.6 by showing that every (generalised) \mathbb{Q} -contingent claim for ${}^{\mathbb{Q}}\mathcal{S}$ can be identified with a (generalised) \mathbb{P} -contingent claim for \mathcal{S} , provided of course that the stopping time $\tau \in {}^{\mathbb{Q}}\mathcal{S}$ is also a \mathbb{P} -stopping time, e.g. a deterministic time.

Proposition 1.7. *Let \mathcal{S} be a market and F a generalised \mathbb{Q} -contingent claim at time $\tau \in \mathcal{T}_{[0,T]} \subset {}^{\mathbb{Q}}\mathcal{T}_{[0,T]}$ for ${}^{\mathbb{Q}}\mathcal{S}$. Then there exists a generalised \mathbb{P} -contingent claim \tilde{F} at time τ for \mathcal{S} such that $F = {}^{\mathbb{Q}}\tilde{F}$. Moreover, if F is improper, defaultable, a contingent claim or a positive contingent claim, \tilde{F} can be chosen so, too.*

Proof. Fix $S \in \mathcal{S} \subset {}^{\mathbb{Q}}\mathcal{S}$. Then $g := F(S)$ is ${}^{\mathbb{Q}}\mathcal{F}_\tau = \sigma(\mathcal{F}_\tau, {}^{\mathbb{Q}}\mathcal{N})$ -measurable, and hence there exists an \mathcal{F}_τ -measurable function h satisfying $g = h$ \mathbb{Q} -a.s. (see [44, Lemma 1.25]). Let Z_τ be a version of the Radon-Nikodým density of \mathbb{Q} with respect to \mathbb{P} on $\mathcal{F}_\tau \subset {}^{\mathbb{Q}}\mathcal{F}_\tau$. Set $\tilde{h} := h\mathbb{1}_{\{Z_\tau > 0\}} + \mathbb{1}_{\{Z_\tau = 0\}}$ and let \tilde{F} be the \mathbb{P} -a.s. unique generalised \mathbb{P} -contingent claim at time τ from Proposition II.5.3 satisfying $\tilde{F}(S) = \tilde{h}$ \mathbb{P} -a.s. Then $\tilde{F}(S) = \tilde{h} = h = g = F(S)$ \mathbb{Q} -a.s. by construction. Thus $F = {}^{\mathbb{Q}}\tilde{F}$ by the uniqueness statement in Proposition 1.6. The remaining claims follow from the construction of \tilde{F} . \square

Finally, we compare (ordinary and limit quantile) superreplication prices and gratis events under \mathbb{P} and \mathbb{Q} . To this end, for a \mathbb{Q} -contingent claim F at time $\tau \in {}^{\mathbb{Q}}\mathcal{T}_{[0,T]}$ and a strategy cone Γ for ${}^{\mathbb{Q}}\mathcal{S}$, we write ${}^{\mathbb{Q}}\Pi(F|\Gamma)$ and ${}^{\mathbb{Q}}\Pi^*(F|\Gamma)$ for the ordinary and limit quantile \mathbb{Q} -superreplication price of F for Γ . Moreover, we write ${}^{\mathbb{Q}}\mathcal{G}_\tau(\Gamma)$ for the collection of all \mathbb{Q} -gratis events for ${}^{\mathbb{Q}}\mathcal{S}$ at time τ for Γ .

Proposition 1.8. *Let \mathcal{S} be a market, Γ a strategy cone for \mathcal{S} (and for ${}^{\mathbb{Q}}\mathcal{S}$) and F a generalised \mathbb{P} -contingent claim at time $\tau \in \mathcal{T}_{[0,T]} \subset {}^{\mathbb{Q}}\mathcal{T}_{[0,T]}$ for \mathcal{S} . Then*

- (a) ${}^{\mathbb{Q}}\Pi({}^{\mathbb{Q}}F|\Gamma) \leq \mathbb{Q}(\Pi(F|\Gamma))$,
- (b) ${}^{\mathbb{Q}}\Pi^*({}^{\mathbb{Q}}F|\Gamma) \leq \mathbb{Q}(\Pi^*(F|\Gamma))$.

Note that we cannot compare the generalised \mathbb{Q} -contingent claims ${}^{\mathbb{Q}}\Pi({}^{\mathbb{Q}}F|\Gamma)$ or ${}^{\mathbb{Q}}\Pi^*({}^{\mathbb{Q}}F|\Gamma)$ to the generalised \mathbb{P} -contingent claims $\Pi(F|\Gamma)$ or $\Pi^*(F|\Gamma)$ directly.

Proof. By the exchange rate consistency (II.5.1), it suffices to prove both statements for fixed $S \in \mathcal{S} \subset {}^{\mathbb{Q}}\mathcal{S}$. Then, part (a) is immediate, and part (b) follows from a similar argument as the one used in the proof of Proposition II.6.20. \square

Corollary 1.9. *Let \mathcal{S} be a market, Γ a strategy cone for \mathcal{S} (and for ${}^{\mathbb{Q}}\mathcal{S}$) and $\tau \in \mathcal{T}_{[0,T]} \subset {}^{\mathbb{Q}}\mathcal{T}_{[0,T]}$. Then*

$$\mathcal{G}_\tau(\Gamma) \subset {}^{\mathbb{Q}}\mathcal{G}_\tau(\Gamma).$$

2 Continuous markets

After these technical preliminaries, we now address the problem posed at the beginning of the chapter for *continuous* markets. Recall from Definition II.1.3 that a market is called continuous if there exists a representative $S \in \mathcal{S}$ which has \mathbb{P} -a.s. continuous trajectories.

Before proving the main result of this section, we introduce the important concept of a *default time* of a self-financing strategy.

Definition 2.1. Let \mathcal{S} be a market and $\vartheta \in L^{\text{sf}}(\mathcal{S})$. A *default time* of ϑ is a stopping time δ_ϑ satisfying

$$\delta_\vartheta = \inf\{t \in [0, T] : V_t(\vartheta) < 0\} \quad \mathbb{P}\text{-a.s.} \quad (2.1)$$

It is straightforward to check that a default time δ_ϑ of ϑ exists and is \mathbb{P} -a.s. unique.¹

One distinguishing property of continuous markets is that the value process of any self-financing strategy ϑ vanishes at any default time δ_ϑ (if the latter is finite).

Proposition 2.2. Let \mathcal{S} be a continuous market and $\vartheta \in L^{\text{sf}}(\mathcal{S})$. Then, for any default time δ_ϑ of \mathcal{S} ,

$$V_{\delta_\vartheta}(\vartheta) = 0 \quad \mathbb{P}\text{-a.s. on } \{\delta_\vartheta < \infty\} \quad \text{and} \quad \vartheta \mathbf{1}_{[0, \delta_\vartheta]} \in \mathcal{U}(\mathcal{S}).$$

Proof. By the exchange rate consistency (II.2.1) of $V_{\delta_\vartheta}(\vartheta)$, it suffices to show the first claim for some representative $S \in \mathcal{S}$. Fix a continuous representative $S \in \mathcal{S}$. Then $V(\vartheta)(S)$ has \mathbb{P} -a.s. continuous trajectories, and by definition of δ_ϑ ,

$$V_{\delta_\vartheta}(\vartheta)(S) = 0 \quad \mathbb{P}\text{-a.s.}$$

The second claim follows immediately from the first one. \square

We proceed to prove the main result of this section.

Theorem 2.3. Let \mathcal{S} be a continuous bounded numéraire market which fails NINA. Let $G_T \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ be a maximal gratis event at time T and assume that $\mathbb{P}[G_T] < 1$. Define the probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T by $d\mathbb{Q} := \mathbf{1}_{G_T^c} \frac{1}{\mathbb{P}[G_T^c]} d\mathbb{P}$. Then ${}^{\mathbb{Q}}\mathcal{S}$ satisfies NINA.

Proof. By Lemma III.2.6, it suffices to show that ${}^{\mathbb{Q}}\Pi^*(F | \mathbf{b}\mathcal{U}({}^{\mathbb{Q}}\mathcal{S})) > 0$ for all nonzero \mathbb{Q} -contingent claims F at time T . Seeking a contradiction, suppose that there exists a nonzero \mathbb{Q} -contingent claim F at time T satisfying

$${}^{\mathbb{Q}}\Pi^*(F | \mathbf{b}\mathcal{U}({}^{\mathbb{Q}}\mathcal{S})) = 0. \quad (2.2)$$

¹Fix $S \in \mathcal{S}$ and set $\delta_\vartheta := \inf\{t \in [0, T] : V_t(\vartheta)(S) < 0\}$. Then δ_ϑ satisfies (2.1) by right-continuity and the exchange rate consistency (II.2.1) of value processes. The latter also gives \mathbb{P} -a.s.-uniqueness of δ_ϑ .

By Proposition 1.7, $F = {}^{\mathbb{Q}}\tilde{F}$ for some \mathbb{P} -contingent claim \tilde{F} at time T . Note that \tilde{F} is also nonzero. We proceed to show that

$$\Pi^*(\tilde{F} | \mathbf{b}\mathcal{U}(\mathcal{S})) = 0. \quad (2.3)$$

To this end, by Proposition III.1.8 and Corollary III.1.7, it suffices to show that for each $\varepsilon \in (0, 1)$, each $\delta > 0$ and each positive \mathbb{P} -contingent claim C at time 0,

$$\Pi^\varepsilon(\tilde{F} - \infty\mathbf{1}_{G_T} | \mathbf{b}\mathcal{U}(\mathcal{S})) \leq \delta C. \quad (2.4)$$

Let ε, δ, C be as above and pick $S \in \mathcal{S} \subset {}^{\mathbb{Q}}\mathcal{S}$. Then by (2.2), the fact that $F = {}^{\mathbb{Q}}\tilde{F}$ and Corollary III.1.7, there exists $\vartheta \in \mathbf{b}\mathcal{U}({}^{\mathbb{Q}}\mathcal{S})$ such that

$$V_0(\vartheta)(S) \leq \delta C(S) \quad \text{and} \quad \mathbb{Q}[V_T(\vartheta)(S) \geq \tilde{F}(S)] \geq 1 - \varepsilon \quad (2.5)$$

By Lemma 1.5, we may assume without loss of generality that $\vartheta \in \mathbf{b}L^{\text{sf}}(\mathcal{S})$. Set $\tilde{\vartheta} := \vartheta\mathbf{1}_{[0, \delta_\vartheta]}$, where δ_ϑ is a default time of ϑ . Then $\tilde{\vartheta} \in \mathbf{b}\mathcal{U}(\mathcal{S})$ by Proposition 2.2 and $\tilde{\vartheta} = \vartheta$ \mathbb{Q} -a.s. because $\delta_\vartheta = +\infty$ \mathbb{Q} -a.s. by the fact that $\vartheta \in \mathbf{b}\mathcal{U}({}^{\mathbb{Q}}\mathcal{S})$. Now (2.5) and the definition of \mathbb{Q} give $V_0(\tilde{\vartheta})(S) \leq \delta C(S)$ and

$$\begin{aligned} \mathbb{P}[V_T(\tilde{\vartheta})(S) \geq \tilde{F}(S) - \infty\mathbf{1}_{G_T}] &= \mathbb{P}[G_T] + \mathbb{P}[(G_T)^c]\mathbb{Q}[V_T(\tilde{\vartheta})(S) \geq \tilde{F}(S)] \\ &= \mathbb{P}[G_T] + \mathbb{P}[(G_T)^c]\mathbb{Q}[V_T(\vartheta)(S) \geq \tilde{F}(S)] \\ &\geq \mathbb{P}[G_T] + \mathbb{P}[(G_T)^c](1 - \varepsilon) \\ &\geq 1 - \varepsilon. \end{aligned}$$

Thus, $\{\tilde{F}(S) > 0\} \in \mathcal{G}_T(\mathbf{b}\mathcal{U}(\mathcal{S})) \subset \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$, and so $\{\tilde{F}(S) > 0\} \subset G_T$ \mathbb{P} -a.s. By the fact that F is a nonzero \mathbb{Q} -contingent claim,

$$\mathbb{P}[\{F(S) > 0\} \setminus G_T] = \mathbb{P}[(G_T)^c]\mathbb{Q}[\tilde{F}(S) > 0] = \mathbb{P}[(G_T)^c]\mathbb{Q}[F(S) > 0] > 0,$$

and we arrive at a contradiction. \square

The following corollary shows that the measure \mathbb{Q} in Theorem 2.3 has the largest support among all absolutely continuous measures $\tilde{\mathbb{Q}} \ll \mathbb{P}$ on \mathcal{F}_T for which ${}^{\tilde{\mathbb{Q}}}\mathcal{S}$ satisfies NINA. In particular, if $\Omega \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$, an absolutely continuous probability measure $\tilde{\mathbb{Q}} \ll \mathbb{P}$ on \mathcal{F}_T such that ${}^{\tilde{\mathbb{Q}}}\mathcal{S}$ satisfies NINA does not exist.

Corollary 2.4. *Let \mathcal{S} be a continuous bounded numéraire market which fails NINA and $G_T \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ a maximal gratis event at time T . Suppose there exists a probability measure $\tilde{\mathbb{Q}} \ll \mathbb{P}$ on \mathcal{F}_T such that ${}^{\tilde{\mathbb{Q}}}\mathcal{S}$ satisfies NINA. Then $\mathbb{P}[G_T] < 1$ and $d\tilde{\mathbb{Q}} \ll d\mathbb{Q}$ on \mathcal{F}_T , where \mathbb{Q} is as in Theorem 2.3*

Proof. It suffices to show that $\tilde{\mathbb{Q}}[A] = 0$ for all $A \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$, since this together with $\tilde{\mathbb{Q}} \ll \mathbb{P}$ on \mathcal{F}_T first implies that $\mathbb{P}[G_T] < 1$ and then yields $\tilde{\mathbb{Q}} \ll \mathbb{Q}$ by the definition of \mathbb{Q} .

Seeking a contradiction, suppose there exists $A \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ with $\tilde{\mathbb{Q}}[A] > 0$. Let F be a contingent claim at time T with $\{F > 0\} = A$ \mathbb{P} -a.s., and ${}^{\tilde{\mathbb{Q}}}\tilde{F}$ the corresponding \mathbb{Q} -contingent claim from Proposition 1.6. Then $\tilde{\mathbb{Q}} \ll \mathbb{P}$ gives

$\{\tilde{Q}F > 0\} = A$ \tilde{Q} -a.s., and so $\tilde{Q}F$ is a nonzero \tilde{Q} -contingent claim. Now Corollary 1.4 and Proposition 1.8 yield


$$\tilde{Q}\Pi^*(\tilde{Q}F | \mathcal{U}(\tilde{Q}\mathcal{S})) \leq \tilde{Q}\Pi^*(\tilde{Q}F | \mathcal{U}(\mathcal{S})) \leq \tilde{Q}(\Pi^*(F | \mathcal{U}(\mathcal{S}))) = \tilde{Q}0 = 0.$$

But this is a contradiction to the hypothesis that $\tilde{Q}\mathcal{S}$ satisfies NINA. □

3 General markets

The key step in the proof of Theorem 2.3 is to pass from $\vartheta \in \mathbf{bL}^{\text{sf}}(\mathcal{S})$ to $\tilde{\vartheta} = \vartheta \mathbf{1}_{\llbracket 0, \delta_\vartheta \rrbracket} \in \mathbf{bU}(\mathcal{S})$. For this, continuity of \mathcal{S} is crucial. The following example shows that without the continuity assumption on \mathcal{S} , even in the very simple setup of a three-period model with finite Ω , Theorem 2.3 is wrong.

Example 3.1. Let \mathcal{S} be the market generated by $S = (1, X_k)_{k \in \{0, \dots, 3\}}$, where the movements of X are described by the following event tree, where the numbers beside the branches describe transition probabilities.



We assume that the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is minimal in the sense that $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ with $\{X_3 = 1.5\} = \{\omega_1\}$, $\{X_3 = 0.9\} = \{\omega_2\}$, etc., and that $(\mathcal{F}_k)_{k=0, \dots, 3}$ is the natural filtration of X . It is straightforward to check that $\mathcal{G}_3(\mathcal{U}(\mathcal{S})) = \{\{\omega_4\}\}$, and so $\{\omega_4\}$ is a maximal gratis event for $\mathcal{G}_3(\mathcal{U}(\mathcal{S}))$. If we define the probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_3 by $d\mathbb{Q} := \mathbf{1}_{\{\omega_1, \omega_2, \omega_3\}} \frac{1}{\mathbb{P}[\{\omega_1, \omega_2, \omega_3\}]} d\mathbb{P}$, then ${}^{\mathbb{Q}}\mathcal{G}_3({}^{\mathbb{Q}}\mathcal{U}(\mathcal{S})) = \{\{\omega_3\}, \{\omega_4\}, \{\omega_3, \omega_4\}\}$, and so ${}^{\mathbb{Q}}\mathcal{S}$ fails to satisfy NINA.

Let us briefly comment on what goes wrong in Example 3.1. Due to the simple setup, \mathbb{Q} -undefaultable strategies can be identified with \mathbb{P} -self-financing strategies which (possibly) default on a gratis event for \mathbb{P} -undefaultable strategies. The analogue of Theorem 2.3 fails because there are more gratis events for those strategies than for \mathbb{P} -undefaultable ones.

This insight is the motivation for the following definition.

Definition 3.2. Let \mathcal{S} be a market and Γ a strategy cone. An *undefaultable strategy outside the gratis events of Γ* is a self-financing strategy $\vartheta \in L^{\text{sf}}(\mathcal{S})$ satisfying $\{\delta_\vartheta \leq T\} \in \mathcal{G}_T(\Gamma) \cup \mathcal{N}$, where δ_ϑ is any default time of ϑ . The collection of all such strategies is denoted by $\mathcal{U}(\mathcal{S} | \Gamma)$.

If \mathcal{S} is a numéraire market and Γ a strategy cones which allows switching to numéraire strategies, then $\mathcal{G}_\tau(\Gamma) \subset \mathcal{G}_T(\Gamma)$ for all stopping times $\tau \in \mathcal{T}_{[0,T]}$ by Proposition III.2.2. This justifies our terminology—at least for strategy cones which allows switching to numéraire strategies.

The following straightforward result collects some properties of $\mathcal{U}(\mathcal{S} | \Gamma)$.

Proposition 3.3. *Let \mathcal{S} be a market and $\Gamma_1 \subset \Gamma_2$ strategy cones. Then $\mathcal{U}(\mathcal{S} | \Gamma_1)$ and $\mathcal{U}(\mathcal{S} | \Gamma_2)$ are again strategy cones and satisfy*

$$\mathcal{U}(\mathcal{S}) \subset \mathcal{U}(\mathcal{S} | \Gamma_1) \subset \mathcal{U}(\mathcal{S} | \Gamma_2) \subset L^{\text{sf}}(\mathcal{S}).$$

Moreover, if \mathcal{S} is a numéraire market, $\mathcal{U}(\mathcal{S} | \Gamma_1)$ and $\mathcal{U}(\mathcal{S} | \Gamma_2)$ allow switching to numéraire strategies.

We illustrate Definition 3.2 by elaborating on Example 3.1.

Example 3.4. Consider the market in Example 3.1. It is easy to check that

$$\begin{aligned} \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})) &= \left\{ \vartheta \in L^{\text{sf}}(\mathcal{S}) : V(\vartheta) \geq 0 \text{ on } \{\omega_1, \omega_2, \omega_3\} \right\}, \\ \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) &= \left\{ \vartheta \in L^{\text{sf}}(\mathcal{S}) : V(\vartheta) \geq 0 \text{ on } \{\omega_1, \omega_2\} \right\}, \\ \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})))) &= \left\{ \vartheta \in L^{\text{sf}}(\mathcal{S}) : V(\vartheta) \geq 0 \text{ on } \{\omega_1, \omega_2\} \right\}, \\ \mathcal{U}(\mathcal{S} | L^{\text{sf}}(\mathcal{S})) &= \left\{ \vartheta \in L^{\text{sf}}(\mathcal{S}) : V(\vartheta) \geq 0 \text{ on } \{\omega_1, \omega_2\} \right\}, \\ \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | L^{\text{sf}}(\mathcal{S}))) &= \left\{ \vartheta \in L^{\text{sf}}(\mathcal{S}) : V(\vartheta) \geq 0 \text{ on } \{\omega_1, \omega_2\} \right\}. \end{aligned}$$

In Example 3.1/3.4, $\Gamma \subsetneq \mathcal{U}(\mathcal{S} | \Gamma)$ for $\Gamma = \mathcal{U}(\mathcal{S}), \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))$, $\Gamma \supsetneq \mathcal{U}(\mathcal{S} | \Gamma)$ for $\Gamma = L^{\text{sf}}(\mathcal{S})$, and $\Gamma = \mathcal{U}(\mathcal{S} | \Gamma)$ for $\Gamma = \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) = \mathcal{U}(\mathcal{S} | L^{\text{sf}}(\mathcal{S}))$. These different possible behaviours of a strategy cone Γ are the motivation for the next definition.

Definition 3.5. Let \mathcal{S} be a market. A strategy cone Γ is called

- *sub-balanced* if $\Gamma \subset \mathcal{U}(\mathcal{S} | \Gamma)$,
- *super-balanced* if $\Gamma \supset \mathcal{U}(\mathcal{S} | \Gamma)$,
- *balanced* if $\Gamma = \mathcal{U}(\mathcal{S} | \Gamma)$.

We denote the collection of all sub-balanced, super-balanced and balanced strategy cones for \mathcal{S} by $\underline{\mathfrak{B}}(\mathcal{S})$, $\overline{\mathfrak{B}}(\mathcal{S})$ and $\mathfrak{B}(\mathcal{S})$, respectively.

It follows directly from Definition 3.2 that $\mathcal{U}(\mathcal{S}) \in \underline{\mathfrak{B}}(\mathcal{S})$ and $L^{\text{sf}}(\mathcal{S}) \in \overline{\mathfrak{B}}(\mathcal{S})$ for any market \mathcal{S} . However, it is not obvious that $\mathfrak{B}(\mathcal{S}) = \underline{\mathfrak{B}}(\mathcal{S}) \cap \overline{\mathfrak{B}}(\mathcal{S}) \neq \emptyset$ for all markets \mathcal{S} .

Using the above terminology, let us comment again on Example 3.1/3.4. The analogue of Theorem 2.3 fails because $\mathcal{U}(\mathcal{S})$ is sub-balanced but not balanced. So, a natural idea how to extend Theorem 2.3 to general markets is to take a

maximal gratis event for the smallest *balanced* strategy cone of \mathcal{S} (as opposed to $\mathcal{U}(\mathcal{S})$) in the definition of the absolutely continuous measure \mathbb{Q} . Before showing that this procedure works in full generality, let us check it in the simple setup of Example 3.1.

Example 3.6. Consider the market in Example 3.1. Then $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})))$ is balanced by Example 3.4. To see that it is the smallest balanced strategy cone, let Γ be another balanced strategy cone. Then $\mathcal{U}(\mathcal{S}) \subset \mathcal{U}(\mathcal{S} | \Gamma) = \Gamma$ by Proposition 3.3, and iterating the argument yields $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})) \subset \Gamma$ and $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \subset \Gamma$. Clearly, $\{\omega_3, \omega_4\}$ is a maximal gratis event for $\mathcal{G}_3(\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))))$. Define $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_3 by $d\mathbb{Q} := \mathbb{1}_{\{\omega_1, \omega_2\}} \frac{1}{\mathbb{P}[\{\omega_1, \omega_2\}]} d\mathbb{P}$. Then it is not difficult to check that ${}^{\mathbb{Q}}\mathcal{G}_T({}^{\mathbb{Q}}\mathcal{U}(\mathcal{S})) = \emptyset$, and so ${}^{\mathbb{Q}}\mathcal{S}$ satisfies NINA.

The next result collects some basic but important “closedness” properties of the sets $\underline{\mathfrak{B}}(\mathcal{S})$, $\overline{\mathfrak{B}}(\mathcal{S})$ and $\mathfrak{B}(\mathcal{S})$.

Proposition 3.7. *Let \mathcal{S} be a market and Γ a strategy cone in $\underline{\mathfrak{B}}(\mathcal{S})$, $\overline{\mathfrak{B}}(\mathcal{S})$, or $\mathfrak{B}(\mathcal{S})$. Then $\mathcal{U}(\mathcal{S} | \Gamma)$ is in $\underline{\mathfrak{B}}(\mathcal{S})$, $\overline{\mathfrak{B}}(\mathcal{S})$, or $\mathfrak{B}(\mathcal{S})$, too.*

Proof. The first claim follows from Proposition 3.3 with $\Gamma_1 = \Gamma$ and $\Gamma_2 = \mathcal{U}(\mathcal{S} | \Gamma)$, the second one follows from Proposition 3.3 with $\Gamma_1 = \mathcal{U}(\mathcal{S} | \Gamma)$ and $\Gamma_2 = \Gamma$, and the third one is trivial. \square

The sets $\underline{\mathfrak{B}}(\mathcal{S})$ and $\overline{\mathfrak{B}}(\mathcal{S})$ are in addition closed under arbitrary unions and intersections, respectively. More precisely, since the union (as opposed to the intersection) of cones is in general no longer a cone, we need to consider the cone *generated* by the union of the strategy cones. If $(\Gamma_i)_{i \in I}$ is a family of cones, we write $\text{cone}(\Gamma_i, i \in I)$ for the cone generated by the Γ_i .

Proposition 3.8. *Let \mathcal{S} be a market and $(\Gamma_i)_{i \in I}$ a family of strategy cones.*

- (a) *If $\Gamma_i \in \underline{\mathfrak{B}}(\mathcal{S})$ for all $i \in I$, then $\text{cone}(\Gamma_i, i \in I) \in \underline{\mathfrak{B}}(\mathcal{S})$.*
- (b) *If $\Gamma_i \in \overline{\mathfrak{B}}(\mathcal{S})$ for all $i \in I$, then $\bigcap_{i \in I} \Gamma_i \in \overline{\mathfrak{B}}(\mathcal{S})$.*

Proof. We only establish (a), part (b) follows by a similar but slightly easier argument. Let $\vartheta \in \text{cone}(\Gamma_i, i \in I)$. Then there exist $n \in \mathbb{N}$, $i_1, \dots, i_n \in I$ and $\lambda_1, \dots, \lambda_n > 0$ such that $\vartheta = \sum_{k=1}^n \lambda_k \vartheta^{(k)}$, where $\vartheta^{(k)} \in \Gamma_{i_k}$. Now the claim follows from Proposition III.1.4 because

$$\{\delta_{\vartheta^{(k)}} \leq T\} \in \mathcal{G}(\Gamma_{i_k}) \cup \mathcal{N} \subset \mathcal{G}(\text{cone}(\Gamma_i, i \in I)) \cup \mathcal{N},$$

for each $k \in \{1, \dots, n\}$, since $\Gamma_k \in \underline{\mathfrak{B}}(\mathcal{S})$ and $\Gamma_k \subset \text{cone}(\Gamma_i, i \in I)$, and so

$$\{\delta_{\vartheta} \leq T\} \subset \bigcup_{k=1}^n \{\delta_{\vartheta^{(k)}} \leq T\} \in \mathcal{G}(\text{cone}(\Gamma_i, i \in I)) \cup \mathcal{N}. \quad \square$$

Using the above properties, we can now show the existence of a smallest balanced strategy cone.

Lemma 3.9. *Let \mathcal{S} be a market. Then $\mathfrak{B}(\mathcal{S}) \neq \emptyset$. Moreover,*

$$\bar{\mathcal{U}}(\mathcal{S}) := \bigcap_{\Gamma \in \bar{\mathfrak{B}}(\mathcal{S})} \Gamma,$$

is the smallest balanced strategy cone for \mathcal{S} . It satisfies $\mathcal{U}(\mathcal{S}) \subset \bar{\mathcal{U}}(\mathcal{S})$.

We call $\vartheta \in \bar{\mathcal{U}}(\mathcal{S})$ a *generalised undefaultable strategy*. This terminology is justified by the discussion preceding Example 3.6.

Proof. Proposition 3.8 implies that $\bar{\mathcal{U}}(\mathcal{S})$ is the smallest element in $\bar{\mathfrak{B}}(\mathcal{S})$, and Proposition 3.7 shows that $\bar{\mathcal{U}}(\mathcal{S}) \supset \mathcal{U}(\mathcal{S} | \bar{\mathcal{U}}(\mathcal{S})) \in \bar{\mathfrak{B}}(\mathcal{S})$. The latter together with minimality $\bar{\mathcal{U}}(\mathcal{S})$ in $\bar{\mathfrak{B}}(\mathcal{S})$ establishes the main claim. The final claim follows from Proposition 3.3. \square

It is straightforward to check that $\bar{\mathcal{U}}(\mathcal{S}) = \mathcal{U}(\mathcal{S})$ if \mathcal{S} satisfies NINA. For continuous market, we have a general simple representation of $\bar{\mathcal{U}}(\mathcal{S})$.

Proposition 3.10. *Let \mathcal{S} be a continuous market. Then $\bar{\mathcal{U}}(\mathcal{S}) = \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))$ and $\mathcal{G}_T(\bar{\mathcal{U}}(\mathcal{S})) = \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$.*

Proof. It suffices to show that

$$\mathcal{G}_T(\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \subset \mathcal{G}_T(\mathcal{U}(\mathcal{S})). \quad (3.1)$$

Indeed, (3.1) gives $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \subset \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))$, and so $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})) \in \bar{\mathfrak{B}}(\mathcal{S})$. Since $\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S})) \subset \mathcal{U}(\mathcal{S} | \bar{\mathcal{U}}(\mathcal{S})) = \bar{\mathcal{U}}(\mathcal{S})$ by Proposition 3.3 and Lemma 3.9, the first claim follows from minimality of $\bar{\mathcal{U}}(\mathcal{S})$ in $\bar{\mathfrak{B}}(\mathcal{S})$. The second claim is an immediate consequence of (3.1) and the first claim because the inclusion $\mathcal{G}_T(\bar{\mathcal{U}}(\mathcal{S})) \supset \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ is trivial.

For (3.1), we may assume without loss of generality that $\mathcal{G}_T(\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \neq \emptyset$. Suppose by way of contradiction there exists $A \in \mathcal{G}_T(\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \setminus \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ with $\mathbb{P}[A] > 0$. By the definition of gratis events, there exist a contingent claim F at time T such that $\text{supp } F = A$ \mathbb{P} -a.s. Let $\varepsilon \in (0, 1)$, $\delta > 0$ and C be a positive contingent claim at time 0. By Corollary III.1.7, there exists $\tilde{\vartheta} \in \mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))$ such that

$$V_0(\tilde{\vartheta}) \leq \delta C \quad \text{and} \quad \mathbb{P}[V_\tau(\tilde{\vartheta}) \geq F] \geq 1 - \varepsilon.$$

Set $\vartheta := \tilde{\vartheta} \mathbf{1}_{[0, \delta_{\tilde{\vartheta}}]}$, where $\delta_{\tilde{\vartheta}}$ is a default time of $\tilde{\vartheta}$. Then $\vartheta \in \mathcal{U}(\mathcal{S})$ by Proposition 2.2 and $\{\delta_{\tilde{\vartheta}} \leq T\} \cap \{F > 0\} \in \mathcal{N}$ since $\{\delta_{\tilde{\vartheta}} \leq T\} \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$ and $\{F > 0\} \in \mathcal{G}_T(\mathcal{U}(\mathcal{S} | \mathcal{U}(\mathcal{S}))) \setminus \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$. Thus,

$$V_0(\vartheta) \leq \delta C \quad \text{and} \quad \mathbb{P}[V_\tau(\vartheta) \geq F] \geq 1 - \varepsilon.$$

By Corollary III.1.7 we arrive at the contradiction $\{F > 0\} \in \mathcal{G}_T(\mathcal{U}(\mathcal{S}))$. \square

Before stating and proving the main result of this section, we have to study generalised undefaultable strategies under an absolutely continuous change of measure (cf. Section 1).

Proposition 3.11. *Let \mathcal{S} be a market and $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T . Then*

$$\bar{\mathcal{U}}(\mathcal{S}) \subset \bar{\mathcal{U}}(\mathbb{Q}\mathcal{S}).$$

Proof. Denote by \mathfrak{U} all sub-balanced \mathbb{P} -strategy cones that are contained in $\bar{\mathcal{U}}(\mathbb{Q}\mathcal{S})$. Then $\mathfrak{U} \neq \emptyset$ because $\mathcal{U}(\mathcal{S}) \subset \mathcal{U}(\mathbb{Q}\mathcal{S}) \subset \bar{\mathcal{U}}(\mathbb{Q}\mathcal{S})$ by Corollary 1.4 and Lemma 3.9. Set $\Gamma_{\mathfrak{U}} := \text{cone}(\Gamma : \Gamma \in \mathfrak{U})$. Proposition 3.8 shows that $\Gamma_{\mathfrak{U}}$ is the largest element in $\mathfrak{U} \subset \mathfrak{B}(\mathcal{S})$. Using this, the fact that $\mathcal{U}(\mathcal{S} | \Gamma) \subset \mathcal{U}(\mathbb{Q}\mathcal{S} | \Gamma)$, Proposition 3.3 and the fact that $\bar{\mathcal{U}}(\mathbb{Q}\mathcal{S}) \in \mathfrak{B}(\mathbb{Q}\mathcal{S})$ yields

$$\Gamma_{\mathfrak{U}} \subset \mathcal{U}(\mathcal{S} | \Gamma_{\mathfrak{U}}) \subset \mathcal{U}(\mathbb{Q}\mathcal{S} | \Gamma_{\mathfrak{U}}) \subset \mathcal{U}(\mathbb{Q}\mathcal{S} | \bar{\mathcal{U}}(\mathbb{Q}\mathcal{S})) = \bar{\mathcal{U}}(\mathbb{Q}\mathcal{S}).$$

This together with Proposition 3.7 gives $\mathcal{U}(\mathcal{S} | \Gamma_{\mathfrak{U}}) \in \mathfrak{U}$. Now, the maximality of $\Gamma_{\mathfrak{U}}$ in \mathfrak{U} implies that $\Gamma_{\mathfrak{U}} \in \mathfrak{B}(\mathcal{S})$. Finally, since $\bar{\mathcal{U}}(\mathcal{S})$ is the smallest element in $\mathfrak{B}(\mathcal{S})$, we may conclude that

$$\bar{\mathcal{U}}(\mathcal{S}) \subset \Gamma_{\mathfrak{U}} \subset \bar{\mathcal{U}}(\mathbb{Q}\mathcal{S}). \quad \square$$

We proceed to prove the main result of this chapter.

Theorem 3.12. *Let \mathcal{S} be a bounded numéraire market which fails NINA. Let $G_T \in \mathcal{G}_T(\bar{\mathcal{U}}(\mathcal{S}))$ be a maximal gratis event and suppose that $\mathbb{P}[G_T] < 1$. Define the probability measure $\mathbb{Q} \ll \mathbb{P}$ on \mathcal{F}_T by $d\mathbb{Q} := \mathbf{1}_{G_T^c} \frac{1}{\mathbb{P}[G_T^c]} d\mathbb{P}$. Then $\mathbb{Q}\mathcal{S}$ satisfies NINA.*

Theorem 2.3 is a direct consequence of Theorem 3.12. Indeed, for continuous markets, it follows from Proposition 3.10 that $G_T \in \mathcal{F}_T$ is a maximal gratis event for $\mathcal{G}_T(\bar{\mathcal{U}}(\mathcal{S}))$ if and only if it is a maximal gratis event for $\mathcal{G}_T(\mathcal{U}(\mathcal{S}))$.

Proof. The proof is similar to the proof of Theorem 2.3; the main difference is that we work with $\mathbf{b}\bar{\mathcal{U}}(\mathcal{S})$ instead of $\mathbf{b}\mathcal{U}(\mathcal{S})$ and use that $\bar{\mathcal{U}}(\mathcal{S})$ is balanced.

By Lemma III.2.6, it suffices to show that ${}^{\mathbb{Q}}\Pi_T^*(F | \mathbf{b}\mathcal{U}(\mathbb{Q}\mathcal{S})) = 0$ for all nonzero \mathbb{Q} -contingent claims F at time T . Seeking a contradiction, suppose there exists a nonzero \mathbb{Q} -contingent claim F at time T satisfying

$${}^{\mathbb{Q}}\Pi_T^*(F | \mathbf{b}\mathcal{U}(\mathbb{Q}\mathcal{S})) = 0. \quad (3.2)$$

By Proposition 1.7, $F = {}^{\mathbb{Q}}\tilde{F}$, for some \mathbb{P} -contingent claim \tilde{F} at time T . Note that \tilde{F} is also nonzero. We proceed to show that

$$\Pi^*(\tilde{F} | \mathbf{b}\bar{\mathcal{U}}(\mathcal{S})) = 0. \quad (3.3)$$

To this end, by Proposition III.1.8 and Corollary III.1.7, it suffices to show that for each $\varepsilon \in (0, 1)$, each $\delta > 0$ and each positive contingent claim C at time 0,

$$\Pi^\varepsilon(\tilde{F} - \infty \mathbf{1}_{G_T} | \mathbf{b}\bar{\mathcal{U}}(\mathcal{S})) \leq \delta C. \quad (3.4)$$

So let ε, δ, C be as above and pick $S \in \mathcal{S} \subset \mathbb{Q}\mathcal{S}$. Then by (3.2), the fact that $F = {}^{\mathbb{Q}}\tilde{F}$ and Corollary III.1.7, there exists $\vartheta \in \mathbf{b}\mathcal{U}(\mathbb{Q}\mathcal{S})$ such that

$$V_0(\vartheta)(S) \leq \delta C(S) \quad \text{and} \quad \mathbb{Q}[V_T(\vartheta)(S) \geq F(S)] \geq 1 - \varepsilon \quad (3.5)$$

By Lemma 1.5, we may assume without loss of generality that $\vartheta \in \mathbf{b}L^{\text{sf}}(\mathcal{S})$. Let δ_ϑ be a default time of ϑ . As ϑ is undefaultable under \mathbb{Q} , $\{\delta_\vartheta \leq T\} \in \mathcal{G}_T(\overline{\mathcal{U}}(\mathcal{S}))$. Moreover, since $\overline{\mathcal{U}}(\mathcal{S})$ is balanced, $\vartheta \in \overline{\mathcal{U}}(\mathcal{S})$, and since ϑ is bounded, $\vartheta \in \mathbf{b}\overline{\mathcal{U}}(\mathcal{S})$. Now (3.5) and the definition of \mathbb{Q} give $V_0(\vartheta)(S) \leq \delta C(S)$ and

$$\begin{aligned} \mathbb{P}[V_T(\vartheta)(S) \geq \tilde{F}(S) - \infty \mathbb{1}_{G_T}] &= \mathbb{P}[G_T] + \mathbb{P}[(G_T)^c] \mathbb{Q}[V_T(\vartheta)(S) \geq \tilde{F}(S)] \\ &\geq \mathbb{P}[G_T] + \mathbb{P}[(G_T)^c](1 - \varepsilon) \\ &\geq 1 - \varepsilon. \end{aligned}$$

Thus, $\{\tilde{F}(S) > 0\} \in \mathcal{G}_T(\mathbf{b}\overline{\mathcal{U}}(\mathcal{S})) \subset \mathcal{G}_T(\overline{\mathcal{U}}(\mathcal{S}))$ and so $\{\tilde{F}(S) > 0\} \subset G_T$ \mathbb{P} -a.s. By the fact that F is a nonzero \mathbb{Q} -contingent claim,

$$\mathbb{P}[\{F(S) > 0\} \setminus G_T] = \mathbb{P}[(G_T)^c] \mathbb{Q}[\tilde{F}(S) > 0] = \mathbb{P}[(G_T)^c] \mathbb{Q}[F(S) > 0] > 0,$$

and we arrive at a contradiction. \square

The following corollary is the general version of Corollary 2.4.

Corollary 3.13. *Let \mathcal{S} be a bounded numéraire market which fails NINA and $G_T \in \mathcal{G}_T(\overline{\mathcal{U}}(\mathcal{S}))$ a maximal gratis event at time T . Suppose there exists a probability measure $\tilde{\mathbb{Q}} \ll \mathbb{P}$ on \mathcal{F}_T such that $\tilde{\mathbb{Q}}\mathcal{S}$ satisfies NINA. Then $\mathbb{P}[G_T] < 1$ and $d\tilde{\mathbb{Q}} \ll d\mathbb{Q}$ on \mathcal{F}_T , where \mathbb{Q} is as in Theorem 3.12.*

Corollary 2.4 is a direct consequence of Corollary 3.13. This follows as above from Proposition 3.10.

Proof. The proof is very similar to the proof of Corollary 2.4; the main difference is that we work with $\overline{\mathcal{U}}(\mathcal{S})$ instead of $\mathcal{U}(\mathcal{S})$.

It suffices to show that $\tilde{\mathbb{Q}}[A] = 0$ for all $A \in \mathcal{G}_T(\overline{\mathcal{U}}(\mathcal{S}))$. Seeking a contradiction, suppose there exists $A \in \mathcal{G}_T(\overline{\mathcal{U}}(\mathcal{S}))$ with $\tilde{\mathbb{Q}}[A] > 0$. Let F be a contingent claim at time T with $\{F > 0\} = A$ \mathbb{P} -a.s., and $\tilde{\mathbb{Q}}F$ the corresponding $\tilde{\mathbb{Q}}$ -contingent claim from Proposition 1.6. Then $\tilde{\mathbb{Q}} \ll \mathbb{P}$ gives $\{\tilde{\mathbb{Q}}F > 0\} = A$ $\tilde{\mathbb{Q}}$ -a.s. Moreover, by the fact that $\tilde{\mathbb{Q}}\mathcal{S}$ satisfies NINA, Propositions 3.11 and 1.8,

$$\begin{aligned} \tilde{\mathbb{Q}}\Pi^*(\tilde{\mathbb{Q}}F | \mathcal{U}(\tilde{\mathbb{Q}}\mathcal{S})) &= \tilde{\mathbb{Q}}\Pi^*(\tilde{\mathbb{Q}}F | \overline{\mathcal{U}}(\tilde{\mathbb{Q}}\mathcal{S})) \leq \tilde{\mathbb{Q}}\Pi^*(\tilde{\mathbb{Q}}F | \overline{\mathcal{U}}(\mathcal{S})) \\ &\leq \tilde{\mathbb{Q}}(\Pi^*(F | \overline{\mathcal{U}}(\mathcal{S}))) = \tilde{\mathbb{Q}}0 = 0. \end{aligned}$$

But this is a contradiction to the hypothesis that $\tilde{\mathbb{Q}}\mathcal{S}$ satisfies NINA. \square

Chapter VI

Dual characterisation of markets satisfying NINA

In this chapter, we study numéraire markets satisfying numéraire-independent no-arbitrage (NINA); see Definition III.2.5. After proving the existence of nonzero strongly maximal (numéraire) strategies, we derive a numéraire-independent version of the fundamental theorem of asset pricing (FTAP) in Section 1. In Section 2, we provide a dual characterisation of (weakly and strongly) maximal strategies and give conditions for the existence of (*true*) martingale representatives. In Section 3, we derive a numéraire-independent dual characterisation of superreplication prices and discuss the notion of (strongly) maximal and (strongly) attainable contingent claims. The material for this chapter is taken from [32].

1 Dominating maximal strategies and numéraire-independent FTAP

In Chapter III.3, we have argued that a strategy $\vartheta \in \mathcal{U}$ is a “reasonable investment” only if it is (weakly or strongly) maximal for \mathcal{U} . Suppose now that ϑ is not strongly maximal for \mathcal{U} . Is it then possible to replace ϑ by another “dominating” strategy $\vartheta^* \in \mathcal{U}$ which is strongly maximal, requires the same initial investment and yields the same or more wealth at the end? Clearly, by Proposition III.3.15, we can expect a positive answer to this question only if the zero strategy 0 is strongly maximal for \mathcal{U} —or equivalently if the market satisfies NINA (cf. the discussion after Proposition III.3.21). Perhaps surprisingly, this necessary condition is also sufficient for numéraire market.

Theorem 1.1. *Let \mathcal{S} be a numéraire market satisfying NINA and $\vartheta \in \mathcal{U}$. Then there exists a (non-unique) strongly maximal strategy $\vartheta^* \in \mathcal{U}$ such that*

$$V_0(\vartheta^*) = V_0(\vartheta) \quad \text{and} \quad V_T(\vartheta^*) \geq V_T(\vartheta) \quad \mathbb{P}\text{-a.s.} \quad (1.1)$$

We call ϑ^ a dominating maximal strategy for ϑ .*

By Proposition III.3.19, it suffices to find a *weakly* maximal strategy $\vartheta^* \in \mathcal{U}$ satisfying (1.1). Using this insight, we reformulate Theorem 1.1 in the language of

admissible investment processes (cf. Definition II.4.3) in order to use arguments from Delbaen and Schachermayer [9] and Kabanov [42]. For the next result recall from Definition III.3.22 the notion of $BK^{(\eta)}/NUPBR^{(\eta)}$.

Lemma 1.2. *Let \mathcal{S} be a numéraire market and η a numéraire strategy. Suppose that $S^{(\eta)}$ satisfies $BK^{(\eta)}/NUPBR^{(\eta)}$. Then for each $\zeta \in L^{ad}(S^{(\eta)}, 1)$, there is $\zeta^* \in L^{ad}(S^{(\eta)}, 1)$ such that $\zeta^* \bullet S_T \geq \zeta \bullet S_T$ \mathbb{P} -a.s. and there is no $\tilde{\zeta} \in L^{ad}(S^{(\eta)}, 1)$ with*

$$\tilde{\zeta} \bullet S_T \geq \zeta^* \bullet S_T \text{ } \mathbb{P}\text{-a.s.} \quad \text{and} \quad \mathbb{P}[\tilde{\zeta} \bullet S_T > \zeta^* \bullet S_T] > 0.$$

To see that Theorem 1.1 and Lemma 1.2 are indeed equivalent, note that by the discussion after Proposition III.3.21 and Proposition III.3.24 (b), \mathcal{S} satisfies NINA (i.e., we have $sm(0)$) if and only if $S^{(\eta)}$ satisfies $BK^{(\eta)}/NUPBR^{(\eta)}$, and by Proposition II.4.4, we can identify ζ and ζ^* in Lemma 1.2 with $\vartheta - \eta$ and $\vartheta^* - \eta$, where ϑ and ϑ^* are as in Theorem 1.1.

The quite technical proof of Lemma 1.2 is given below in Section 4. Here we just sketch the main idea. First, one constructs a sequence $(\zeta_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ such that for each n , ζ_n dominates ζ , and as n grows, ζ_n becomes more and more “close” to a weakly maximal strategy. Then, one modifies that sequence such that the modified sequence $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ converges in the semimartingale topology (see [63]) to some weakly maximal strategy $\zeta^* \in L^{ad}(S^{(\eta)}, 1)$ dominating ζ .

Our next goal is to establish the existence of strongly maximal *numéraire* strategies for numéraire markets satisfying NINA. To this end, we have to explore what happens when the value process of an undefaultable strategy becomes 0.

Definition 1.3. Let \mathcal{S} be a market and $\vartheta \in \mathcal{U}$. An *absorption time* of ϑ is a stopping time a_ϑ satisfying

$$a_\vartheta = \inf\{t \in [0, T] : V_t(\vartheta) = 0 \text{ or } V_{t-}(\vartheta) = 0\} \text{ } \mathbb{P}\text{-a.s.} \quad (1.2)$$

It is straightforward to check that an absorption time for ϑ exists and is \mathbb{P} -a.s. unique.¹ The next result explains our terminology.

Proposition 1.4. *Let \mathcal{S} be a numéraire market satisfying NINA. Then for all $\vartheta \in \mathcal{U}$ and any absorption time a_ϑ of ϑ ,*

$$V(\vartheta) \equiv 0 \text{ on } \llbracket a_\vartheta, T \rrbracket \text{ } \mathbb{P}\text{-a.s.}$$

As a consequence, $\eta \in \mathcal{U}$ is a numéraire strategy if and only if $V_T(\eta) > 0$ \mathbb{P} -a.s.

Proof. Let η be a numéraire strategy and $\vartheta \in \mathcal{U}$. For $c > 0$, let τ_c be the (\mathbb{P} -a.s. unique) stopping time with $\tau_c = \inf\{t > a_\vartheta : V_t(\vartheta) \geq cV_t(\eta)\}$ \mathbb{P} -a.s. It suffices to show that $\tau_c = +\infty$ \mathbb{P} -a.s. for all $c > 0$. Seeking a contradiction, suppose there is $c > 0$ with $\mathbb{P}[\tau_c < \infty] > 0$. Then $F := cV_{\tau_c \wedge T}(\eta)\mathbb{1}_{\{\tau_c < \infty\}}$ is a nonzero contingent

¹Fix $S \in \mathcal{S}$ and set $a_\vartheta := \inf\{t \in [0, T] : V_t(\vartheta)(S) = 0 \text{ or } V_{t-}(\vartheta)(S) = 0\}$. Then a_ϑ satisfies (1.2) by right-continuity and the exchange rate consistency (II.2.1) of value processes. The latter also gives \mathbb{P} -a.s. uniqueness of a_ϑ . The same kind of reasoning applies to the stopping time τ_c and σ_N defined in the proof of Proposition 1.4 below.

at time $\tau_c \wedge T$. For $n \in \mathbb{N}$, let σ_n be the (\mathbb{P} -a.s. unique) stopping time with $\sigma_n = \inf\{t > 0 : V_t(\vartheta) \leq \frac{1}{n}V_t(\eta)\} \wedge T$ \mathbb{P} -a.s. Then $\sigma_n \leq a_\vartheta \leq \tau_c$ \mathbb{P} -a.s. for each n . Let $\delta > 0$ and choose N large enough that $\frac{1}{N} < \min(\delta, c)$. Right-continuity of each $S \in \mathcal{S}$ gives $\sigma_N < \tau_c$ \mathbb{P} -a.s. on $\{\tau_c < \infty\}$. Set

$$\tilde{\vartheta} := \frac{1}{N}\eta\mathbf{1}_{\llbracket 0, \sigma_N \rrbracket} + \left(\vartheta + V_{\sigma_N} \left(\frac{1}{N}\eta - \vartheta \right) (S^{(\eta)})\eta \right) \mathbf{1}_{\llbracket \sigma_N, T \rrbracket} \in \mathcal{U}. \quad (1.3)$$

Then $V_0(\tilde{\vartheta}) = \frac{1}{N}V_0(\eta) \leq \delta V_0(\eta)$ and

$$\begin{aligned} V_{\tau_c \wedge T}(\tilde{\vartheta}) &\geq V_{\tau_c \wedge T}(\tilde{\vartheta})\mathbf{1}_{\{\tau_c < \infty\}} = \left(V_{\tau_c \wedge T}(\vartheta) + V_{\sigma_N} \left(\frac{1}{N}\eta - \vartheta \right) (S^{(\eta)})V_{\tau_c \wedge T}(\eta) \right) \mathbf{1}_{\{\tau_c < \infty\}} \\ &\geq cV_{\tau_c \wedge T}(\eta)\mathbf{1}_{\{\tau_c < \infty\}} = F \text{ } \mathbb{P}\text{-a.s.} \end{aligned}$$

Since $\delta > 0$ was arbitrary, $\Pi_{\tau_c \wedge T}(F | \mathcal{U}) = 0$, and so \mathcal{S} fails NINA, in contradiction to the hypothesis. \square

Remark 1.5. Once Theorem 1.10 below has been established, Proposition 1.4 follows from the *minimum principle* for nonnegative supermartingales [44, Theorem 7.32]. But since Proposition 1.4 is needed for the proof of Theorem 1.10, the above direct argument is really necessary.

An immediate but important consequence of Theorem 1.1 and Proposition 1.4 is that in numéraire markets satisfying NINA, there are “enough” strongly maximal numéraire strategies.

Theorem 1.6. *Let \mathcal{S} be a numéraire market satisfying NINA. Then for each numéraire strategy η , there exists a dominating maximal numéraire strategy η^* .*

To derive a dual characterisation of numéraire markets satisfying NINA, we have to find the analogue to the concept of an equivalent σ -martingale measure ($E\sigma$ MM) in our framework. Since we do not work with a *fixed* representative, like $S = (1, X)$ in the standard framework, the notion of an $E\sigma$ MM alone does not make sense. Dual objects are rather *pairs* (S, \mathbb{Q}) consisting of a representative S and an equivalent measure $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T under which S is a \mathbb{Q} - σ -martingale.

Definition 1.7. Let \mathcal{S} be an (N -dimensional) market and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T an equivalent probability measure. $S = (S^1, \dots, S^N) \in \mathcal{S}$ is called a \mathbb{Q} - σ -martingale representative, if S^1, \dots, S^N are σ -martingales under \mathbb{Q} . We often omit the qualifier “ \mathbb{Q} ” and call S an equivalent σ -martingale representative. We denote the set of all \mathbb{Q} - σ -martingale representatives of \mathcal{S} by $\mathcal{M}_{\mathbb{Q}}$.

Remark 1.8. (a) For $S \in \mathcal{M}_{\mathbb{Q}}$ and $\vartheta \in \mathcal{U}$, $V(\vartheta)(S)$ is a nonnegative local \mathbb{Q} -martingale and \mathbb{Q} -supermartingale by Ansel and Stricker [1, Corollaire 3.5] and Fatou’s lemma.

(b) By Bayes’ theorem for (true) martingales/ σ -martingales [45, Proposition 5.1], the existence of equivalent (true) martingale/ σ -martingale representatives depends on the measure \mathbb{P} only through its nullsets.

Before formulating the numéraire-independent version of the FTAP, we recall the classic FTAP by Delbaen and Schachermayer [13, Theorem 1.1] using our terminology.

Theorem 1.9 (Classic FTAP). *Let \mathcal{S} be a numéraire market and assume that $e_1 = (1, 0, \dots, 0)$ is a numéraire strategy. Then the following are equivalent:*

- (a) $S^{(e_1)}$ satisfies NFLVR $^{(e_1)}$, i.e., the numéraire strategy e_1 is strongly maximal for \mathcal{U} .
- (b) There exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(e_1)} \in \mathcal{M}_{\mathbb{Q}}$.

The classic FTAP *fixes* the numéraire strategy e_1 and works in units corresponding to $S^{(e_1)} = (1, X)$, cf. Example II.3.6. It says that e_1 , the buy-and-hold strategy of the “bank account”, is “good” (strongly maximal for \mathcal{U}), if and only if there is an E σ MM $\mathbb{Q} \approx \mathbb{P}$ for $S^{(e_1)}$. It is clear from the discussion after Definition II.4.3, that we can replace e_1 in Theorem 1.9 by any other numéraire strategy η . This “generalised” version of the FTAP then says that η is strongly maximal for \mathcal{U} , if and only if there is an E σ MM $\mathbb{Q} \approx \mathbb{P}$ for $S^{(\eta)}$. But this still assumes that the numéraire strategy η has *ex ante* been chosen. For a numéraire-independent FTAP, strong maximality of η cannot be an ex-ante assumption; it must be part of the dual characterisation.

Theorem 1.10 (Numéraire-independent FTAP). *Let \mathcal{S} be a numéraire market. Then the following are equivalent:*

- (a) \mathcal{S} satisfies NINA, i.e., the zero strategy 0 is strongly maximal for \mathcal{U} .
- (b) There exists a pair (η, \mathbb{Q}) , where η is a numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$. For each such pair, η is strongly maximal for \mathcal{U} .
- (c) There exists a numéraire strategy η which is strongly maximal for \mathcal{U} . For each such η , there exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$.
- (d) There exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $\mathcal{M}_{\mathbb{Q}} \neq \emptyset$.
- (e) For each $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , $\mathcal{M}_{\mathbb{Q}} \neq \emptyset$.

Proof. “(a) \Rightarrow (c)”. The first claim follows from Theorem 1.6 and the second one from (the discussion subsequent to) Theorem 1.9.

“(c) \Rightarrow (b)”. The first claim is trivial, and the second one follows from (the discussion subsequent to) Theorem 1.9.

“(b) \Rightarrow (d)”. This is trivial.

“(d) \Rightarrow (e)”. This follows from Remark 1.8 (b).

“(e) \Rightarrow (a)”. Seeking a contradiction, suppose there exists a nonzero contingent claim F at time T with $\Pi(F | \mathcal{U}) = 0$. Let $S \in \mathcal{M}_{\mathbb{P}}$, C be the contingent claim at time 0 from Proposition II.5.3 satisfying $C(S) = 1$ and $\delta := \mathbb{E}[F(S)]/2 > 0$. By Proposition II.6.3, there exists $\vartheta \in \mathcal{U}$ with $V_0(\vartheta) \leq \delta C$ and $V_T(\vartheta) \geq F$ \mathbb{P} -a.s. By the (\mathbb{P} -)supermartingale property of $V(\vartheta)(S)$, we arrive at the contradiction

$$\delta C(S) \geq V_0(\vartheta)(S) \geq \mathbb{E}[V_T(\vartheta)(S)] \geq \mathbb{E}[F(S)] = 2\delta > \delta C(S). \quad \square$$

Remark 1.11. Independently from and partly parallel to our work, the equivalence “(a) \Leftrightarrow (d)” (for $\mathbb{Q} = \mathbb{P}$) has been derived by Kardaras [48] (for $N = 2$) and by Schweizer and Takaoka [73] (for general N), where *in essence* “(a) \Leftrightarrow (b)” is shown. However, neither [48] nor [73] have recognised the numéraire-independent structure of the dual objects (η, \mathbb{Q}) in any way. We conjecture that the equivalence “(a) \Leftrightarrow (d)” remains valid even if \mathcal{S} fails to be a numéraire market; however this would require a completely different proof technique.

2 Martingale properties

In numéraire markets satisfying NINA, there is no difference between the notions of weakly and strongly maximal strategies for \mathcal{U} , this follows from Proposition III.3.19. Therefore, we call such strategies simply *maximal strategies* for \mathcal{U} in the sequel; we also often omit the qualifier “for \mathcal{U} ”.

The following result provides a dual characterisation of maximal strategies in terms of the (true) martingale property of their value processes. It can be seen as a numéraire-independent version of [11, Theorem 13], where it is shown that the value process of a maximal admissible strategy *in the sense of Delbaen and Schachermayer* is a (true) \mathbb{Q} -martingale for some $\mathbb{E}\sigma$ MM \mathbb{Q} .

Theorem 2.1. *Let \mathcal{S} be a numéraire market satisfying NINA and $\vartheta \in \mathcal{U}$. Then the following are equivalent:*

- (a) ϑ is maximal for \mathcal{U} .
- (b) There exists a pair (η, \mathbb{Q}) , where η is a numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$ and $V(\vartheta)(S^{(\eta)})$ is a \mathbb{Q} -martingale uniformly bounded by 1.
- (c) For each maximal numéraire strategy η , there exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$ and $V(\vartheta)(S^{(\eta)})$ is a \mathbb{Q} -martingale.
- (d) There exist $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{M}_{\mathbb{Q}}$ such that $V(\vartheta)(S)$ is a \mathbb{Q} -martingale.
- (e) For each $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , there exists $S \in \mathcal{M}_{\mathbb{Q}}$ such that $V(\vartheta)(S)$ is a \mathbb{Q} -martingale.

Note that by Theorem 1.10 (b), the numéraire strategy η in (b) is automatically (strongly) maximal for \mathcal{U} .

Proof. “(a) \Rightarrow (b)”. Let $\vartheta \in \mathcal{U}$ be a maximal strategy. By Theorem 1.10 (b), there exists a maximal numéraire strategy $\hat{\eta} \in \mathcal{U}$. Set $\eta := \hat{\eta} + \vartheta$. Then η is a numéraire strategy and (weakly) maximal by Corollary III.3.10. Theorem 1.10 (c) gives $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$. Hence $V(\vartheta)(S^{(\eta)})$ is a nonnegative local \mathbb{Q} -martingale. It is even a (true) \mathbb{Q} -martingale since

$$V(\vartheta)(S^{(\eta)}) \leq V(\hat{\eta} + \vartheta)(S^{(\eta)}) = V(\eta)(S^{(\eta)}) \equiv 1 \quad \mathbb{Q}\text{-a.s.}$$

“(a) \Rightarrow (c)”. Let $\vartheta \in \mathcal{U}$ be a maximal strategy and η a maximal numéraire strategy. Set $\tilde{\eta} := \eta + \vartheta$. Arguing as in “(a) \Rightarrow (b)” shows that $\tilde{\eta}$ is a maximal numéraire strategy and that there is $\tilde{\mathbb{Q}} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\tilde{\eta})} \in \mathcal{M}_{\tilde{\mathbb{Q}}}$ and $V(\vartheta)(S^{(\tilde{\eta})})$ and $V(\eta)(S^{(\tilde{\eta})})$ are $\tilde{\mathbb{Q}}$ -martingales (uniformly bounded by 1). Since $V(\eta)(S^{(\tilde{\eta})})$ is positive, we can define $\mathbb{Q} \approx \tilde{\mathbb{Q}}$ on \mathcal{F}_T by $d\mathbb{Q} := \frac{Z_T}{Z_0} d\tilde{\mathbb{Q}}$, where $Z := V(\eta)(S^{(\tilde{\eta})})$. Then $S^{(\eta)} = S^{(\tilde{\eta})}/Z$ \mathbb{P} -a.s. by (II.3.1), and the exchange rate consistency of value processes (II.2.1) (for the exchange rate process $D = \frac{1}{Z}$) gives $V(\vartheta)(S^{(\eta)}) = V(\vartheta)(S^{(\tilde{\eta})})/Z$ \mathbb{P} -a.s. Since $V(\vartheta)(S^{(\tilde{\eta})})$ and Z are $\tilde{\mathbb{Q}}$ -martingales and $S^{(\tilde{\eta})}$ is a $\tilde{\mathbb{Q}}$ - σ -martingale, $S^{(\eta)}$ is a \mathbb{Q} - σ -martingale and $V(\vartheta)(S^{(\eta)})$ is a \mathbb{Q} -martingale by Bayes’ theorem for (σ -)martingales; see [45, Proposition 5.1].

“(b), (c) \Rightarrow (d)”. This is trivial—noting that a maximal numéraire strategy exists by Theorem 1.6.

“(d) \Rightarrow (e)”. This follows from Bayes’ theorem by arguing as in the final part of “(a) \Rightarrow (c)”.

“(e) \Rightarrow (a)”. Let $S \in \mathcal{M}_{\mathbb{P}}$ be such that $V(\vartheta)(S)$ is a (\mathbb{P} -)martingale. Suppose by way of contradiction that ϑ fails to be (weakly) maximal for \mathcal{U} . By Proposition III.3.8, it then fails to be weakly maximal at time T for \mathcal{U} . Hence, by Proposition III.3.5, there is $\tilde{\vartheta} \in \mathcal{U}$ with

$$V_0(\tilde{\vartheta}) = V_0(\vartheta), \quad V_T(\tilde{\vartheta}) \geq V_T(\vartheta) \text{ } \mathbb{P}\text{-a.s.} \quad \text{and} \quad \mathbb{P}[V_T(\tilde{\vartheta}) > V_T(\vartheta)] > 0.$$

By the supermartingale property of $V(\tilde{\vartheta})(S)$ and the martingale property of $V(\vartheta)(S)$, we arrive at the contradiction

$$V_0(\tilde{\vartheta})(S) \geq \mathbb{E}[V_T(\tilde{\vartheta})(S)] > \mathbb{E}[V_T(\vartheta)(S)] = V_0(\vartheta)(S). \quad \square$$

Theorem 1.10 shows that in a numéraire market satisfying NINA, the set of all \mathbb{Q} - σ -martingale representatives $\mathcal{M}_{\mathbb{Q}}$ is nonempty for some (and hence each) $\mathbb{Q} \approx \mathbb{P}$. We now address the question under which conditions $\mathcal{M}_{\mathbb{Q}}$ contains a (true) \mathbb{Q} -martingale representative.

For the following result, recall from Definition II.4.1 that $\mathbf{b}\Gamma$ denotes the collection of all bounded strategies in a given strategy cone Γ .

Lemma 2.2. *Let \mathcal{S} be a numéraire market. Suppose there exist $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{M}_{\mathbb{Q}}$ such that S is a (true) \mathbb{Q} -martingale representative. Then:*

- (a) *For each $\vartheta \in \mathbf{b}L^{sf}$, the value process $V(\vartheta)(S)$ is a \mathbb{Q} -martingale.*
- (b) *Each $\vartheta \in \mathbf{b}\mathcal{U}$ is maximal for \mathcal{U} .*
- (c) *$S^{(\eta)}$ is an equivalent (true) martingale representative for every bounded numéraire strategy η .*

Proof. Let N be the dimension of \mathcal{S} , and denote by $\|\cdot\|$ the maximum norm in \mathbb{R}^N . By Remark 1.8 (b), we may assume without loss of generality that $\mathbb{Q} = \mathbb{P}$.

(a) Fix $\vartheta \in \mathbf{b}L^{sf}$. Then $V(\vartheta)(S)$ is a local (\mathbb{P} -)martingale, and it is even a (true) martingale since it is of class (D) , i.e., the family $(V_\tau(\vartheta)(S))_{\tau \in T_{[0,T]}}$ is

uniformly integrable. Indeed,

$$|V_\tau(\vartheta)(S)| = |\vartheta_\tau \cdot S_\tau| \leq \|\vartheta\| \sum_{i=1}^N |S_\tau^i|, \quad \tau \in T_{[0,T]},$$

and for each i , the family $(|S_\tau^i|)_{\tau \in T_{[0,T]}}$ is uniformly integrable because S^i is a (true) martingale.

(b) This follows immediately from (a) and Theorem 2.1 (d).

(c) Let η be a bounded numéraire strategy. Then by part (a), $V(\eta)(S)$ is a positive (P-)martingale. Define $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T by $d\mathbb{Q} = \frac{V_T(\eta)(S)}{V_0(\eta)(S)} d\mathbb{P}$. Then $S^{(\eta)}$ is a (true) \mathbb{Q} -martingale representative by arguing as in the proof of “(a) \Rightarrow (c)” in Theorem 2.1. \square

Remark 2.3. It follows immediately from the proof that part (a) of Lemma 2.2 does not need the existence of a numéraire strategy. It can be reformulated in an abstract setting as follows: *Let X be an N -dimensional martingale and H a bounded predictable process satisfying*

$$H \bullet X = H \cdot X - H_0 \cdot X_0. \quad (2.1)$$

Then the stochastic integral $H \bullet X$ is again a martingale. Without assumption (2.1), this is false; see Herdegen and Herrmann [33, Section 5.2] for a counterexample.

Under a mild additional hypothesis, Lemma 2.2 can be strengthened. Recall from Definition II.3.4 that a numéraire market is called *bounded* if there exists a bounded numéraire strategy η such that $S^{(\eta)}$ is bounded.

Corollary 2.4. *Let \mathcal{S} be a bounded numéraire market. Then the following are equivalent:*

- (a) *There exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{M}_{\mathbb{Q}}$ such that S is a \mathbb{Q} -martingale representative.*
- (b) *Each $\vartheta \in \mathbf{b}\mathcal{U}$ is strongly maximal for \mathcal{U} .*
- (c) *$S^{(\eta)}$ is an equivalent martingale representative for each bounded numéraire strategy η .*
- (d) *There exists a strongly maximal numéraire η such that both η and $S^{(\eta)}$ are bounded.*

Moreover, if one of the above assertions holds, \mathcal{S} satisfies NINA.

Proof. The final assertion follows from the fact that (a) implies part (d) in Theorem 1.10.

“(a) \Rightarrow (c) \Rightarrow (b) \Rightarrow (d)”. This follows from Lemma 2.2 and the hypothesis that \mathcal{S} is a bounded numéraire market.

“(d) \Rightarrow (a)”. By Theorem 1.10 (c), $S^{(\eta)}$ is a bounded equivalent σ -martingale representative and so an equivalent (true) martingale representative. \square

In the case of *nonnegative* markets, the market portfolio $\eta^S = (1, \dots, 1)$ is a bounded numéraire strategy and $S^{(\eta^S)}$ is bounded. Thus, we get a further criterion for the existence of (true) equivalent martingale representatives.

Corollary 2.5. *Let \mathcal{S} be a nonnegative (N -dimensional) market. Then the following are equivalent:*

- (a) *There exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{M}_{\mathbb{Q}}$ such that S is a \mathbb{Q} -martingale representative.*
- (b) *For each $i = 1, \dots, N$, $e_i = (0, \dots, 0, 1, 0, \dots, 0)$, the buy-and-hold strategy of the i -th asset, where the 1 is at position i , is strongly maximal for \mathcal{U} .*
- (c) *The market portfolio $\eta^S = (1, \dots, 1)$ is strongly maximal for \mathcal{U} .*

Proof. “(a) \Rightarrow (b)”. This follows from Corollary 2.4.

“(b) \Rightarrow (c)”. By Proposition III.3.13, each e_i is weakly maximal for \mathcal{U} , and so is $\eta^S = \sum_{i=1}^N e_i$ by Corollary III.3.10. Now the claim follows from Propositions III.3.15 and III.3.19.

“(c) \Rightarrow (a)”. This follows from Corollary 2.4 because η^S and $S^{(\eta^S)}$ are bounded. \square

Remark 2.6. *Weak* maximality (for \mathcal{U}) of the buy-and-hold strategies of the primary assets has been called *no-dominance* (ND) in the literature (see e.g. Merton [62] or Jarrow and Larsson [39]), and has recently been studied in the context of financial *bubbles* as a possible requirement in addition to NFLVR (cf. Jarrow et al. [41]). In a classic model $S = (1, X)$, NFLVR and ND together have been shown to be equivalent to the existence of a (true) equivalent martingale measure for X ; see [39, Theorem 3.2]. Our Corollary 2.5 can be seen as a numéraire-independent version of that result because by Proposition III.3.24 (c), *strong* maximality of e_1 implies NFLVR^(e_1) (which is just a more precise terminology for NFLVR in the classic sense; see Definition III.3.22). For a detailed discussion of the exact relationship between maximal strategies, bubbles and *market efficiency* (see [39]) in a numéraire-independent setup, we refer to Chapter VIII.

3 Dual characterisation of superreplication prices and strong attainability of contingent claims

In this section, we derive a numéraire-independent dual characterisation of superreplication prices and briefly address the topic of *attainability* of contingent claims.

First, we recall—using our terminology—the classic result by Kramkov [52] and Föllmer and Kabanov [24] on the dual representation of superreplication prices. To this end, for a strongly maximal numéraire strategy η , we denote by $\mathcal{Q}(S^{(\eta)})$ the set of all probability measures $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)}$ is a \mathbb{Q} - σ -martingale; note that $\mathcal{Q}(S^{(\eta)}) \neq \emptyset$ by Theorem 1.10 (c).

Proposition 3.1. *Let \mathcal{S} be a numéraire market satisfying NINA, η a maximal numéraire strategy and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. Suppose that $\sup_{\mathbb{Q} \in \mathcal{Q}(S^{(n)})} \mathbb{E}_{\mathbb{Q}}[F(S^{(n)})] < \infty$. Then there is $\vartheta \in \mathcal{U}$ with $V_{\tau}(\vartheta) \geq F$ \mathbb{P} -a.s. and*

$$\Pi(F(S^{(n)} | \mathcal{U})(S^{(n)}) = V_0(\vartheta)(S^{(n)}) = \sup_{\mathbb{Q} \in \mathcal{Q}(S^{(n)})} \mathbb{E}_{\mathbb{Q}}[F(S^{(n)})].$$

To derive a numéraire-independent version of the above result, we have to replace the (η -dependent) set $\mathcal{Q}(S^{(n)})$ by a numéraire-independent object. For $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $s \in \mathbb{R}^d \setminus \{0\}$, set $\mathcal{M}_{\mathbb{Q}}(s) := \{S \in \mathcal{M}_{\mathbb{Q}} : S_0 = s\}$. If η is a strongly maximal numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , then each $\tilde{\mathbb{Q}} \in \mathcal{Q}(S^{(n)})$ can be identified with $ZS^{(n)} \in \mathcal{M}_{\mathbb{Q}}(S_0^{(n)})$, where Z is the density process of $\tilde{\mathbb{Q}}$ with respect to \mathbb{Q} . Moreover, if F is a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$, then

$$\mathbb{E}_{\tilde{\mathbb{Q}}}[F(S^{(n)})] = \mathbb{E}_{\mathbb{Q}}[Z_{\tau}F(S^{(n)})] = \mathbb{E}_{\mathbb{Q}}[F(ZS^{(n)})]. \quad (3.1)$$

Note, however, that in general there are $S \in \mathcal{M}_{\mathbb{Q}}(S_0^{(n)})$, for which the unique exchange rate process D from Remark II.1.4 satisfying $S^{(n)} = DS$ \mathbb{P} -a.s. is not a density process of some $\tilde{\mathbb{Q}} \in \mathcal{Q}(S^{(n)})$ with respect to \mathbb{Q} .

After these preparations, we can prove the main result of this section.

Theorem 3.2. *Let \mathcal{S} be a numéraire market satisfying NINA and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$. Then for each $S \in \mathcal{S}$ and each $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T ,*

$$\Pi(F | \mathcal{U})(S) = \sup_{\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})]. \quad (3.2)$$

Moreover, if $\sup_{\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] < \infty$ for some $S \in \mathcal{S}$ and some $\mathbb{Q} \approx \mathbb{P}$, there exists a maximal strategy $\vartheta^* \in \mathcal{U}$ such that

$$V_0(\vartheta^*) = \Pi(F | \mathcal{U}) \quad \text{and} \quad V_{\tau}(\vartheta^*) \geq F \quad \mathbb{P}\text{-a.s.} \quad (3.3)$$

It follows from (3.2) and Remark II.6.2 that $\sup_{\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] < \infty$ holds for some $S \in \mathcal{S}$ and some $\mathbb{Q} \approx \mathbb{P}$ if and only if it holds for every $S \in \mathcal{S}$ and every $\mathbb{Q} \approx \mathbb{P}$.

Proof. First, we establish the easy inequality “ \geq ” in (3.2) in the nontrivial case $\Pi(F | \mathcal{U}) < \infty$. Fix $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{S}$ and let $\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0)$ be arbitrary; note that $\mathcal{M}_{\mathbb{Q}} \neq \emptyset$ by Theorem 1.10. Since $\Pi(F | \mathcal{U})(S) = \Pi(F | \mathcal{U})(\tilde{S})$ by Remark II.6.2, it suffices to show that $\Pi(F | \mathcal{U})(\tilde{S}) \geq \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})]$. Let $\delta > 0$ and C be a positive contingent claim at time 0. By Proposition II.6.3, there is $\vartheta \in \mathcal{U}$ with $V_0(\vartheta) \leq \Pi(F | \mathcal{U}) + \delta C$ and $V_{\tau}(\vartheta) \geq F$ \mathbb{P} -a.s. By the \mathbb{Q} -supermartingale property of $V(\vartheta)(\tilde{S})$,

$$\Pi(F | \mathcal{U})(\tilde{S}) + \delta C(\tilde{S}) \geq V_0(\vartheta)(\tilde{S}) \geq \mathbb{E}_{\mathbb{Q}}[V_{\tau}(\vartheta)(\tilde{S})] \geq \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})].$$

Letting $\delta \searrow 0$ establishes the claim.

Next, if $\sup_{\tilde{S} \in \mathcal{M}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] = \infty$ for each $S \in \mathcal{S}$ and each $\mathbb{Q} \approx \mathbb{P}$, then the claim follows from the first step. Otherwise, fix $S \in \mathcal{S}$ and $\mathbb{Q} \approx \mathbb{P}$ such that $\sup_{\tilde{S} \in \mathcal{M}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] < \infty$. Let η be a maximal numéraire strategy (which exists by Theorem 1.6) and D the unique exchange rate process from Remark II.1.4 satisfying $S^{(\eta)} = DS$ \mathbb{P} -a.s. Then the first part of the proof, the fact that $\mathcal{M}_{\mathbb{Q}}(S_0) = \frac{1}{D_0} \mathcal{M}_{\mathbb{Q}}(D_0 S_0) = \frac{1}{D_0} \mathcal{M}_{\mathbb{Q}}(S_0^{(\eta)})$, (3.1), Proposition 3.1 and Remark II.6.2 yield

$$\begin{aligned} \Pi(F | \mathcal{U})(S) &\geq \sup_{\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0)} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] = \frac{1}{D_0} \sup_{\tilde{S} \in \mathcal{M}_{\mathbb{Q}}(S_0^{(\eta)})} \mathbb{E}_{\mathbb{Q}}[F(\tilde{S})] \\ &\geq \frac{1}{D_0} \sup_{\tilde{\mathbb{Q}} \in \mathcal{Q}(S^{(\eta)})} \mathbb{E}_{\tilde{\mathbb{Q}}}[F(S^{(\eta)})] = \frac{1}{D_0} \Pi(F | \mathcal{U})(S^{(\eta)}) = \Pi(F | \mathcal{U})(S), \end{aligned}$$

which establishes (3.2). Proposition 3.1 yields $\vartheta \in \mathcal{U}$ with $V_0(\vartheta) = \Pi(F | \mathcal{U})$ and $V_T(\vartheta) \geq F$ \mathbb{P} -a.s., and Theorem 1.1 gives a (dominating) maximal strategy $\vartheta^* \in \mathcal{U}$ for ϑ satisfying (3.3). \square

The supremum in (3.2) is in general not attained. We proceed to show that it is so if and only if the inequality in (3.3) is an equality. To this end, we briefly discuss the topic of *attainability* of contingent claims. In the standard framework, it is well known that in continuous time, the existence of a replicating strategy for a contingent claim alone is not enough to guarantee that the value process of the replicating portfolio is a (true) martingale under an $E\sigma$ MM—assuming of course that the market satisfies NFLVR. We show that in our numéraire-independent framework, the good notion of attainability is the existence of a *strongly maximal* replicating strategy; for this reason we also speak of *strong attainability*. Moreover, we introduce the closely related notion of *strongly maximal contingent claims*. In the standard framework, similar concepts have been discussed by Delbaen and Schachermayer [12].

Definition 3.3. Let \mathcal{S} be a market and Γ a strategy cone. A contingent claim F at time $\tau \in \mathcal{T}_{[0,T]}$ with $\Pi(F | \Gamma) < \infty$ is called *strongly maximal for Γ* if for all nonzero contingent claims G at time τ ,

$$\Pi(F + G | \Gamma) > \Pi(F | \Gamma).$$

It is called *strongly attainable for Γ* if there exists a strategy $\vartheta \in \Gamma$ which is strongly maximal at time τ for Γ and satisfies

$$V_{\tau}(\vartheta) = F \text{ } \mathbb{P}\text{-a.s.}$$

Remark 3.4. One could of course also introduce the notion of *weakly attainable* contingent claims but this turns out to be of little use because the existence of a contingent claim which is strongly maximal for \mathcal{U} implies that the market satisfies NINA, and in this case, the notions of weakly and strongly attainable contingent claims for \mathcal{U} coincide.

It follows from the definition of strongly maximal strategies that a strongly attainable contingent claim is a fortiori strongly maximal. The following result shows that in numéraire markets satisfying NINA, the converse is also true for $\Gamma = \mathcal{U}$. This kind of result is well known in the standard framework; see [13, Theorem 5.16].

Lemma 3.5. *Let \mathcal{S} be a numéraire market satisfying NINA and F a contingent claim at time $\tau \in \mathcal{T}_{[0,T]}$ with $\Pi(F | \mathcal{U}) < \infty$. Then the following are equivalent:*

- (a) F is strongly maximal for \mathcal{U} .
- (b) F is strongly attainable for \mathcal{U} .
- (c) There exists a pair (η, \mathbb{Q}) , where η is a numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$ and $\Pi(F | \mathcal{U})(S^{(\eta)}) = \mathbb{E}_{\mathbb{Q}}[F(S^{(\eta)})]$.
- (d) For each maximal numéraire strategy η , there exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$ and $\Pi(F | \mathcal{U})(S^{(\eta)}) = \mathbb{E}_{\mathbb{Q}}[F(S^{(\eta)})]$.
- (e) There exist $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and $S \in \mathcal{M}_{\mathbb{Q}}$ such that $\Pi(F | \mathcal{U})(S) = \mathbb{E}_{\mathbb{Q}}[F(S)]$.
- (f) For each $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , there is $S \in \mathcal{M}_{\mathbb{Q}}$ with $\Pi(F | \mathcal{U})(S) = \mathbb{E}_{\mathbb{Q}}[F(S)]$.

Note that by Theorem 1.10 (b), the numéraire strategy η in (c) is automatically strongly maximal for \mathcal{U} .

Proof. “(a) \Rightarrow (b)”. By Theorem 3.2, there is a strongly maximal strategy $\vartheta \in \mathcal{U}$ such that $V_0(\vartheta) = \Pi(F | \mathcal{U})$ and $V_\tau(\vartheta) \geq F$ \mathbb{P} -a.s. Thus, $\Pi(V_\tau(\vartheta) | \mathcal{U}) \leq V_0(\vartheta)$ by the definition of superreplication prices, which together with strong maximality of F establishes the claim.

“(b) \Rightarrow (d)”. Let $\vartheta \in \mathcal{U}$ be a maximal strategy satisfying $V_\tau(\vartheta) = F$ \mathbb{P} -a.s. and η a maximal numéraire strategy. Theorem 2.1 gives $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)} \in \mathcal{M}_{\mathbb{Q}}$ and $V(\vartheta)(S^{(\eta)})$ is a true \mathbb{Q} -martingale. Theorem 3.2, the \mathbb{Q} -martingale property of $V(\vartheta)(S^{(\eta)})$ and the definition of superreplication prices yield

$$\begin{aligned} \Pi(F | \mathcal{U})(S^{(\eta)}) &\geq \mathbb{E}_{\mathbb{Q}}[F(S^{(\eta)})] = \mathbb{E}_{\mathbb{Q}}[V_\tau(\vartheta)(S^{(\eta)})] = V_0(\vartheta)(S^{(\eta)}) \\ &\geq \Pi(F | \mathcal{U})(S^{(\eta)}). \end{aligned}$$

“(d) \Rightarrow (c)”. This is easy—a strongly maximal numéraire strategy η exists by Theorem 1.6.

“(c) \Rightarrow (e)”. This is trivial setting $S := S^{(\eta)}$.

“(e) \Rightarrow (f)”. Fix $\mathbb{Q}_2 \approx \mathbb{P}$ on \mathcal{F}_T . By assumption, there exists $\mathbb{Q}_1 \approx \mathbb{P}$ on \mathcal{F}_T and $S^{(1)} \in \mathcal{M}_{\mathbb{Q}_1}$ such that

$$\Pi(F | \mathcal{U})(S^{(1)}) = \mathbb{E}_{\mathbb{Q}_1}[F(S^{(1)})]. \quad (3.4)$$

Let Z be the density process of \mathbb{Q}_2 with respect to \mathbb{Q}_1 . Set $S^{(2)} := \frac{S^{(1)}}{Z}$. Then $S_0^{(2)} = S_0^{(1)}$ and $S^{(2)} \in \mathcal{M}_{\mathbb{Q}_2}$ by Bayes' theorem for σ -martingales; see [45, Proposition 5.1]. Thus Remark II.6.2, the exchange rate consistency of contingent claims (II.5.1) and (3.4) give

$$\begin{aligned} \Pi(F | \mathcal{U})(S^{(2)}) &= \Pi(F | \mathcal{U})(S^{(1)}) = \mathbb{E}_{\mathbb{Q}_1}[F(S^{(1)})] = \mathbb{E}_{\mathbb{Q}_1} \left[Z_\tau F \left(\frac{S^{(1)}}{Z} \right) \right] \\ &= \mathbb{E}_{\mathbb{Q}_2}[F(S^{(2)})]. \end{aligned}$$

“(f) \Rightarrow (a)”. Let G be a nonzero contingent claim at time τ and $S \in \mathcal{M}_{\mathbb{P}}$ such that $\Pi(F | \mathcal{U})(S) = \mathbb{E}[F(S)]$. Theorem 3.2 gives

$$\Pi(F + G | \mathcal{U})(S) \geq \mathbb{E}[F(S) + G(S)] > \mathbb{E}[F(S)] = \Pi(F | \mathcal{U})(S). \quad \square$$

4 Appendix

The goal of this appendix is to prove Lemma 1.2. To this end, we need one further definition and one further result.

Definition 4.1. Let \mathcal{S} be a numéraire market, η a numéraire strategy, and $\zeta \in L^{ad}(S^{(\eta)}, 1)$. A sequence $(\zeta_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ is called a *maximalising sequence* for ζ if for all $n \in \mathbb{N}$,

$$\zeta_{n+1} \bullet S_T^{(\eta)} \geq \zeta_n \bullet S_T^{(\eta)} \geq \zeta \bullet S_T^{(\eta)} \quad \mathbb{P}\text{-a.s.},$$

and there is no $\tilde{\zeta}_n \in L^{ad}(S^{(\eta)}, 1)$ satisfying

$$\tilde{\zeta}_n \bullet S_T^{(\eta)} \geq \zeta_n \bullet S_T^{(\eta)} \quad \mathbb{P}\text{-a.s.} \quad \text{and} \quad \mathbb{P}[\tilde{\zeta}_n \bullet S_T^{(\eta)} \geq \zeta_n \bullet S_T^{(\eta)} + 2^{-n}] \geq 2^{-n}.$$

Proposition 4.2. Let \mathcal{S} be a numéraire market η a numéraire strategy and $\zeta_0 \in L^{ad}(S^{(\eta)}, 1)$. If $S^{(\eta)}$ satisfies $BK^{(\eta)}/NUPBR^{(\eta)}$, there exists a maximalising sequence for ζ_0 . Moreover, if $(\zeta_n)_{n \in \mathbb{N}}$ is a maximalising sequence for ζ_0 and $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ any other sequence with $\hat{\zeta}_n \in \text{conv}\{\zeta_n, \zeta_{n+1}, \dots\}$ for all $n \in \mathbb{N}$, then $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ has a subsequence which is a maximalising sequence for ζ_0 .

Proof. For $\zeta \in L^{ad}(S^{(\eta)}, 1)$ and $n \in \mathbb{N}$, set

$$\begin{aligned} B^{(n)}(\zeta) &:= \left\{ \tilde{\zeta} \in L^{ad}(S^{(\eta)}, 1) : \tilde{\zeta} \bullet S_T^{(\eta)} \geq \zeta \bullet S_T^{(\eta)} \quad \mathbb{P}\text{-a.s.} \right\}, \\ B_n^{(n)}(\zeta) &:= \left\{ \tilde{\zeta} \in B^{(n)} : \mathbb{P}[\tilde{\zeta} \bullet S_T^{(\eta)} \geq \zeta \bullet S_T^{(\eta)} + 2^{-n}] \geq 2^{-n} \right\}, \\ A_n^{(n)}(\zeta) &:= \left\{ \tilde{\zeta} \in B^{(n)}(\zeta) : B_n^{(n)}(\tilde{\zeta}) = \emptyset \right\}. \end{aligned}$$

Clearly, $(\zeta_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ is a maximalising sequence for ζ_0 if and only if for all $n \in \mathbb{N}$,

$$\zeta_n \in B^{(n)}(\zeta_{n-1}) \quad \text{and} \quad B_n^{(n)}(\zeta_n) = \emptyset. \quad (4.1)$$

For the first claim, it suffices to show that $A_n^{(\eta)}(\zeta) \neq \emptyset$ for all $\zeta \in L^{ad}(S^{(\eta)}, 1)$, since we can then recursively construct a sequence $(\zeta_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ such that $\zeta_n \in A_n^{(\eta)}(\zeta_{n-1})$ for all $n \in \mathbb{N}$, yielding (4.1) for $(\zeta_n)_{n \in \mathbb{N}}$. Seeking a contradiction, suppose there are $\tilde{\zeta}_0 \in L^{ad}(S^{(\eta)}, 1)$ and $N \in \mathbb{N}$ with $A_N^{(\eta)}(\tilde{\zeta}_0) = \emptyset$. By induction, we construct a sequence $(\tilde{\zeta}_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ satisfying

$$A_N^{(\eta)}(\tilde{\zeta}_n) = \emptyset \quad \text{and} \quad \tilde{\zeta}_n \in B_N^{(\eta)}(\tilde{\zeta}_{n-1}), \quad n \in \mathbb{N}. \quad (4.2)$$

First, $A_N^{(\eta)}(\tilde{\zeta}_0) = \emptyset$ and $\tilde{\zeta}_0 \in B^{(\eta)}(\tilde{\zeta}_0)$ give $\tilde{\zeta}_1 \in B_N^{(\eta)}(\tilde{\zeta}_0) \subset B^{(\eta)}(\tilde{\zeta}_0)$, yielding $B^{(\eta)}(\tilde{\zeta}_1) \subset B^{(\eta)}(\tilde{\zeta}_0)$ and implying $A_N^{(\eta)}(\tilde{\zeta}_1) \subset A_N^{(\eta)}(\tilde{\zeta}_0) = \emptyset$. Next, if $A_N^{(\eta)}(\tilde{\zeta}_n) = \emptyset$ for $n \in \mathbb{N}$, $\tilde{\zeta}_n \in B^{(\eta)}(\tilde{\zeta}_n)$ yields as above $\tilde{\zeta}_{n+1} \in B_N^{(\eta)}(\tilde{\zeta}_n)$ with $A_N^{(\eta)}(\tilde{\zeta}_{n+1}) = \emptyset$. Now, set $a_n := \tilde{\zeta}_n \bullet S_T^{(\eta)} - \tilde{\zeta}_{n-1} \bullet S_T^{(\eta)}$, $n \in \mathbb{N}$, and $A := \limsup_{n \rightarrow \infty} \{a_n \geq 2^{-N}\}$. Then $\mathbb{P}[A] \geq 2^{-N}$ since $\mathbb{P}[a_n \geq 2^{-N}] \geq 2^{-N}$ for each n because $\tilde{\zeta}_n \in B_N^{(\eta)}(\tilde{\zeta}_{n-1})$. Thus, $\lim_{n \rightarrow \infty} \tilde{\zeta}_n \bullet S_T^{(\eta)} = \tilde{\zeta}_0 \bullet S_T^{(\eta)} + \sum_{k=1}^{\infty} a_k \geq -1 + \infty = \infty$ P-a.s. on A , in contradiction to the hypothesis that $S^{(\eta)}$ satisfies $\text{BK}^{(\eta)}/\text{NUPBR}^{(\eta)}$.

For the second claim, define $h(n) := \min \{k \geq n : \hat{\zeta}_n \in \text{conv}\{\zeta_n, \dots, \zeta_k\}\}$, $n \in \mathbb{N}$. Moreover, set $n_1 := 1$ and define recursively $n_{k+1} := h(n_k) + 1$. Since $(\zeta_n \bullet S_T^{(\eta)})_{n \in \mathbb{N}}$ is P-a.s. increasing and $\hat{\zeta}_n \in \text{conv}\{\zeta_n, \zeta_{n+1}, \dots\}$, for all $n \in \mathbb{N}$,

$$\hat{\zeta}_n \bullet S_T^{(\eta)} \geq \zeta_n \bullet S_T^{(\eta)} \geq \zeta \bullet T \quad \text{P-a.s.}, \quad (4.3)$$

and by the definition of h , for all $k \in \mathbb{N}$,

$$\hat{\zeta}_{n_{k+1}} \bullet S_T^{(\eta)} \geq \zeta_{h(n_k)+1} \bullet S_T^{(\eta)} \geq \hat{\zeta}_{n_k} \bullet S_T^{(\eta)} \quad \text{P-a.s.} \quad (4.4)$$

Since $(\zeta_{n_k})_{k \in \mathbb{N}}$ is a fortiori a maximalising sequence for ζ , (4.3) and (4.4) imply that $(\hat{\zeta}_{n_k})_{k \in \mathbb{N}}$ is a maximalising sequence for ζ . \square

Proof of Lemma 1.2. First, by Proposition 4.2 (a), there exists a maximalising sequence $(\zeta_n)_{n \in \mathbb{N}}$ in $L^{ad}(S^{(\eta)}, 1)$ for ζ . Second, arguing as in [9, Lemma 4.5] (with the exception that we do not need to pass to convex combinations due to the monotonicity structure of a maximalising sequence) yields a subsequence, called again $(\zeta_n)_{n \in \mathbb{N}}$, such that the sequence of integral processes $(\zeta_n \bullet S^{(\eta)})_{n \in \mathbb{N}}$ converges P-a.s. in the uniform topology to some càdlàg adapted process Ξ . Third, [42]—or *in essence* [9]—gives a sequence $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ with $\hat{\zeta}_n \in \text{conv}\{\zeta_n, \zeta_{n+1}, \dots\}$ such that the sequence of integral processes $(\hat{\zeta}_n \bullet S^{(\eta)})_{n \in \mathbb{N}}$ converges to Ξ in the semimartingale topology (cf. [63]). Next, by Proposition 4.2 (b), $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ has a subsequence, called again $(\hat{\zeta}_n)_{n \in \mathbb{N}}$, which is a maximalising sequence for ζ . By Mémin's theorem (see [63], Theorem V.4), there exists $\zeta^* \in L(S^{(\eta)})$ with $\zeta_0^* = 0$ such that $\Xi = \zeta^* \bullet S^{(\eta)}$ P-a.s. Since all the $\hat{\zeta}_n$ are in $L^{ad}(S^{(\eta)}, 1)$ and convergence in the semimartingale topology implies uniform convergence in probability (on a finite time horizon), $\Xi \geq -1$ P-a.s. and $\zeta^* \in L^{ad}(S^{(\eta)}, 1)$. Moreover, since the sequence $(\hat{\zeta}_n \bullet S_T^{(\eta)})_{n \in \mathbb{N}}$ is increasing and converges in probability, we have $\zeta^* \bullet S_T^{(\eta)} \geq \hat{\zeta}_n \bullet S_T^{(\eta)} \geq \hat{\zeta} \bullet S_T$ P-a.s. for each n . Using this and the fact that $(\hat{\zeta}_n)_{n \in \mathbb{N}}$ is maximalising for ζ , it is straightforward to check that ζ^* is as desired. \square

Chapter VII

Comparison to other modelling frameworks

In this short chapter, we compare our numéraire-independent approach to modelling financial markets and studying no-arbitrage to the standard and other recent approaches to these issues. The material for this chapter is taken from [32].

1 The standard modelling framework

Let us first summarise how our approach relates to the standard framework and highlight some limitations and (hidden) assumptions of classic modelling. Recall from Example II.3.6 that a classic model $S = (1, X)$ with d “risky” assets X^1, \dots, X^d and one “riskless” asset/“bank account” 1 corresponds to the N -dimensional market \mathcal{S} generated by S , $N = d + 1$, where one works with the numéraire representative $S^{(e_1)}$ of $e_1 = (1, 0, \dots, 0)$, the buy-and-hold strategy of the bank account. This assumes that the bank account cannot default, which is a restriction. For admissible trading, one considers self-financing strategies of the form $\vartheta = \tilde{\vartheta} - ae_1$, where $\tilde{\vartheta} \in \mathcal{U}$ and $a \geq 0$; cf. Corollary II.3.8 (b) and the subsequent discussion.

The no-arbitrage notions NA or NFLVR for $S = (1, X)$ (or just for X) are better expressed in our terminology by saying that $S^{(e_1)}$ satisfies $\text{NA}^{(e_1)}$ or $\text{NFLVR}^{(e_1)}$ and are equivalent to weak or strong maximality of e_1 for \mathcal{U} ; see Chapter III.3.3 and in particular, Proposition III.3.24. They tacitly include the ex-ante assumption that investing in the bank account is a “good” investment. However, this is unduly restrictive: Even if $S^{(e_1)}$ fails $\text{NA}^{(e_1)}$ or $\text{NFLVR}^{(e_1)}$, one can make “a profit out of nothing without risk” via undefaultable (as opposed to $S^{(e_1)}$ -admissible) strategies *only if* \mathcal{S} also fails NINA; otherwise the only thing one can (and definitely should) do is to invest one’s money better than in the bank account.

By contrast, the numéraire-independent concept NINA really captures the idea of forbidding to “make a profit out of nothing without risk”: If \mathcal{S} fails NINA, with some *fixed* positive probability, one can get *as much* as one likes in *any* currency unit one *chooses* and with as *tiny* an initial investment as one likes; see Remark III.1.6 (a). Conversely, if \mathcal{S} satisfies NINA, there exist enough “good”,

i.e., strongly maximal, numéraire strategies η^* ; see Theorem VI.1.6. The only, but crucial caveat is that *one cannot ex ante choose an arbitrary numéraire strategy and expect to have good properties in the currency unit corresponding to this particular numéraire representative*. A moment of thought makes it clear that the last caveat is very natural. It is simply too much to assume that the particular numéraire strategy $e_1 = (1, 0, \dots, 0)$ is strongly maximal. One should better only assume NINA; then one can replace $e_1 = (1, 0, \dots, 0)$ by a dominating maximal numéraire strategy η^* ; see Theorem VI.1.6. The corresponding numéraire representative $S^{(\eta^*)}$ is “close” to $S^{(e_1)}$ in the sense of Theorem VI.1.1, and one has NFLVR $^{(\eta^*)}$ for $S^{(\eta^*)}$, see Proposition III.3.24. Thus, interpreting $V(\eta^*)(S^{(\eta^*)})$ as “riskless” asset and $S^{(\eta^*),1}, \dots, S^{(\eta^*),N}$ as “risky” assets, one can then use *in these units* all results from mathematical finance that are formulated in the standard framework.

2 The fair market framework

Our framework also ties up with the concepts of *fair markets* and *allowable strategies* developed by Yan [81, 82, 83].

[82] starts with a $(d + 1)$ -dimensional semimartingale $Y = (Y^0, \dots, Y^d)$ describing the evolution of $d + 1$ positive assets (where we write Y instead of the usual S for notational convenience).¹ He considers *discounted* asset prices, where one either discounts by one of the basic assets Y^0, \dots, Y^d or by their sum $\sum_{i=0}^d Y_i$. A market is called *fair* if (Y^0) -discounted prices admit a (true) equivalent martingale measure (EMM). It is shown that a market is fair if and only if the prices discounted by *any* of the assets Y^0, \dots, Y^d admits an EMM. A (self-financing) strategy is called *allowable* if it is bounded from below by $-a \sum_{i=0}^d Y_i$ for some $a > 0$, and it is shown that a market is fair if and only if Y satisfies NFLVR for allowable strategies.

From the perspective of numéraire-independent modelling, one considers the market \mathcal{S} generated by $S = (S^1, \dots, S^N) := (Y^0, \dots, Y^d)$, $N = d + 1$, and assumes that for each asset i , its buy-and-hold strategy e_i is a numéraire strategy. The market is fair if $S^{(e_1)}$ is a (true) equivalent martingale representative. By Corollaries VI.2.4 and VI.2.5, this is the case if and only if each $S^{(e_i)}$ and $S^{(\eta^S)}$ are equivalent martingale representatives, where $\eta^S = (1, \dots, 1)$ denotes the market portfolio. An *allowable strategy* is simply an $S^{(\eta^S)}$ -admissible investment process, and Proposition III.3.24 (c) and Corollary VI.2.5 imply that a market is fair if and only if $S^{(\eta^S)}$ satisfies NFLVR $^{(\eta^S)}$. So we recover the results of [82]. Note, however, that the concept of a fair market is not numéraire-independent: It depends on the original currency unit of \tilde{S} and on the sum $\sum_{i=1}^N \tilde{S}^i$ of the price processes of all assets—in the same way as the standard approach depends on the original currency unit of \tilde{S} and on the price process of the numéraire asset \tilde{S}^k ; cf. also the

¹In [83], it is only assumed that asset prices are nonnegative but that their sum $\sum_{i=0}^d Y_i$ is positive. (To be precise, the setup in [83] is more symmetric and there is no asset 0.) Otherwise, the concepts are the same as in [82].

discussion in Chapter I. So the framework in [83] is not numéraire-free and the title a misnomer.

3 The numéraire portfolio and the benchmark approach

Next, our approach sheds new light on the *numéraire portfolio* studied by various authors in increasing generality (see for instance [59, 2, 46]), and on the closely related *benchmark approach* of Platen [65].

There one usually starts out with a classic model $S = (1, X)$ and looks for a self-financing strategy ρ (parametrised in fractions of wealth) with positive wealth (=value) process W^ρ (and $W_0^\rho = 1$) such that the relative wealth process W^π/W^ρ for every self-financing strategy π is a \mathbb{P} -supermartingale. Such a strategy ρ (or alternatively its wealth process W^ρ) is called the *numéraire portfolio* or *growth optimal portfolio* (GOP) or *benchmark portfolio*. If it exists, it is unique and has several compelling features, e.g. in terms of utility maximisation; see [46, Remark 3.2]. Given the existence of the GOP, the *benchmark approach* consists in discounting S by the GOP; the so discounted assets or portfolios are called *benchmarked* assets or portfolios. This procedure is appealing for two reasons: First, benchmarked securities and portfolios are \mathbb{P} -supermartingales, and hedging and pricing can be done under the *physical* measure \mathbb{P} . This has a clear economic meaning, in contrast to pricing in the standard framework under a *risk-neutral* measure \mathbb{Q} , which often lacks a straightforward economic interpretation in incomplete markets. Second, the GOP can also exist in markets “admitting arbitrage” in the sense that X fails NFLVR. More precisely, the GOP exists if and only if X satisfies NUPBR; see [46, Theorem 4.12]. Thus, the benchmark approach is a bit more general than the standard framework.

From the perspective of numéraire-independent modelling, one considers the market \mathcal{S} generated by $S = (1, X)$. The numéraire portfolio corresponds to a numéraire strategy η^* with the property that for each other numéraire strategy η , the value process $V(\eta)(S^{(\eta^*)})$ is a \mathbb{P} -supermartingale. (Note that if we do not fix a currency unit ex ante, η^* is only unique up to multiplication with positive constants.) If such an η^* exists, a similar argument as in the proof of Theorem VI.2.1 shows that η^* is strongly maximal for \mathcal{U} . Thus, \mathcal{S} satisfies NINA by Proposition III.3.15. Conversely, NINA implies the existence of η^* as follows: NINA implies NUPBR^(e1) for $S^{(e1)}$ by Proposition III.3.24 (b), and the existence of η^* follows from [46].

Remark 3.1. Of course, it would be nice to deduce the last implication directly from our numéraire-independent framework. Using [2, Theorem 4.5], it would be sufficient to show that there exists a maximal numéraire strategy η^* such that

$$\sup_{\eta \text{ num.}} \mathbb{E} \left[\log \frac{V_T(\eta)(S^{(\eta^*)})}{V_0(\eta)(S^{(\eta^*)})} \right] < \infty, \quad (3.1)$$

where “num.” stands for “numéraire strategy”, and the denominator is needed for

normalisation. Moreover, by Theorem VI.1.6, it even suffices to consider in (3.1) only maximal numéraire strategies. Using (II.3.1), (3.1) is therefore equivalent to

$$\inf_{\tilde{\eta} \text{ max. num.}} \sup_{\eta \text{ max. num.}} \mathbb{E} \left[\log \left(\frac{V_T(\eta)(S) V_0(\tilde{\eta})(S)}{V_T(\tilde{\eta})(S) V_T(\eta)(S)} \right) \right] < \infty, \quad (3.2)$$

where “max. num.” stands for “maximal numéraire strategy” and $S \in \mathcal{S}$ is arbitrary. At present, however, we do not know whether (3.2) can be proved directly from NINA without excessive effort.

4 The equivalent local martingale deflator framework

Last but not least, our approach can also be linked to the concept of *equivalent local martingale deflators* [48, 73].

In the approach of [48] and [73], one starts out with a classic model $S = (1, X)$. An *equivalent local martingale deflator* (ELMD) is a positive local martingale Z with $Z_0 = 1$ such that ZW is a local martingale for any nonnegative wealth process W . Using semimartingale characteristics [48], or a change of numéraire argument [73],² it is shown that an ELMD exists if and only if S satisfies NA_1 .

From the perspective of numéraire-independent modelling, one considers the market \mathcal{S} generated by $S = (1, X)$ and works with the numéraire representative $S^{(e_1)}$ of e_1 . An ELMD is a (normalised) exchange rate process D with the property that $DS^{(e_1)}$ is a local (\mathbb{P} -)martingale representative. Since $S^{(e_1),1} \equiv 1$, this automatically implies that D is a local martingale. Since $\text{NA}_1^{(e_1)}$ for $S^{(e_1)}$ is equivalent to \mathcal{S} satisfying NINA (Proposition III.3.24 (b)), we recover the main results in [48, 73] from our Theorem VI.1.10; see also Remark VI.1.11.

²The result of [73] has been developed completely independent from and parallel to our work. Notwithstanding, since the *methods* are somewhat related, the proof of Proposition 2.7 (i) in [73] is similar to the proof of our Proposition VI.1.4, and the first part of the proof of Theorem 2.6 in [73] uses similar arguments as the proof of our Theorem VI.1.1.

Chapter VIII

Bubbles from a numéraire-independent perspective

In this final chapter, we develop a new approach for modelling financial bubbles using our numéraire-independent paradigm. Unlike most papers in the recent literature, e.g. [56, 8, 40, 41, 67], we do not define bubbles by a dual object, usually a *strict* local martingale measure, but start from primary notions that are economically motivated. After explaining the main concepts of static and dynamic *viability* and *efficiency* in Section 1, we illustrate them by several examples in Section 2 before deriving their dual characterisations in Section 3. In particular, we show that strict local martingale measures *arise naturally* in the context of modelling financial bubbles; see Theorem 3.22. After providing some further examples of what we call *nontrivial bubbly markets* in Section 4, we compare our definitions and results to the existing literature on bubbles in Section 5.

This chapter, which is joint work with Martin Schweizer, uses a somewhat different setup than the previous chapters. This is partly more and partly less general than the setup of Chapter II. For this reason, we give a largely self-contained presentation, even though this creates some redundancies.

In the sequel we also consider trading strategies which start at some stopping time $\sigma \in \mathcal{T}_{[0,T]}$ later than 0. To this end, we have to introduce some additional notation. For $\sigma \in \mathcal{T}_{[0,T]}$, we set $\mathcal{T}_{[\sigma,T]} := \{\tau \in \mathcal{T}_{[0,T]} : \tau \geq \sigma\}$. A product-measurable process $\xi = (\xi_t)_{t \in [0,T]}$ is called *predictable on* $[[\sigma, T]]$ if the random variable ξ_σ is \mathcal{F}_σ -measurable and the process $\xi \mathbb{1}_{[\sigma, T]}$ is predictable. So if ξ is predictable on $[[\sigma, T]]$ and $A_\sigma \in \mathcal{F}_\sigma$, also $\xi \mathbb{1}_{A_\sigma}$ is predictable on $[[\sigma, T]]$. For an \mathbb{R}^N -valued semimartingale $X = (X_t^1, \dots, X_t^N)_{t \in [0,T]}$ and $\sigma \in \mathcal{T}_{[0,T]}$, we denote by $L_\sigma(X)$ the set of all \mathbb{R}^N -valued processes $\zeta = (\zeta_t^1, \dots, \zeta_t^N)_{t \in [0,T]}$ which are predictable on $[[\sigma, T]]$ and for which the stochastic integral $\int_\sigma^T \zeta_s dX_s := \int_{(0,T]} \zeta_s \mathbb{1}_{[\sigma, T]}(s) dX_s$ is defined in the sense of N -dimensional stochastic integration (consult [37] for details).

1 Main concepts

As in the previous chapters, we consider a financial market consisting of $N > 1$ assets. We initially denote by $\tilde{S} = (\tilde{S}_t^1, \dots, \tilde{S}_t^N)_{t \in [0, T]}$ the price process of the N assets in some fixed but not specified currency unit. This unit may or may not be tradable (e.g. in the form of a bank account); we deliberately do not assume that one of the assets is identically 1, nor that there exists a “bank account” \tilde{S}^0 somewhere in the background. All we initially impose is that the process \tilde{S} is \mathbb{R}^N -valued, adapted and RCLL, that $\tilde{S}^i \geq 0$ \mathbb{P} -a.s. for each i , since we have primary assets in mind, and that the financial market is *nondegenerate* in the sense that

$$\inf_{t \in [0, T]} \sum_{i=1}^N \tilde{S}_t^i > 0 \text{ } \mathbb{P}\text{-a.s.} \quad (1.1)$$

This excludes the case that all assets default and we are left with a nonexistent market; see also the discussion after Definition II.1.3.

It is a folklore result in mathematical finance that in a reasonable financial market, relative prices should be semimartingales after some suitable discounting; see for example Kardaras and Platen [49] or Beiglböck and Schachermayer [3] and the references therein. To capture this idea in a more rigorous form, we first introduce the set $\tilde{\mathcal{D}}$ of all real-valued adapted RCLL processes $\tilde{D} = (\tilde{D}_t)_{t \in [0, T]}$ with

$$\inf_{t \in [0, T]} \tilde{D}_t > 0 \text{ } \mathbb{P}\text{-a.s.} \quad (1.2)$$

We call the elements of $\tilde{\mathcal{D}}$ *generalised exchange rate processes*. We assume that

$$\text{there exists some } \tilde{D} \in \tilde{\mathcal{D}} \text{ such that } \tilde{D}\tilde{S} \text{ is a semimartingale,} \quad (1.3)$$

and we choose and fix one such \tilde{D} and the corresponding process $S := \tilde{D}\tilde{S}$. We also call S a *semimartingale representative* of the market described by \tilde{S} .

From economic considerations, it is clear that all prices are relative, and that the basic qualitative properties of a model should not depend on the chosen currency unit. To make this precise, we call a process \tilde{S}' *economically equivalent* to \tilde{S} if \tilde{S}' is also \mathbb{R}^N -valued, adapted and RCLL, and if we can write $\tilde{S}' = \tilde{D}'\tilde{S}$ for some $\tilde{D}' \in \tilde{\mathcal{D}}$. In other words, two processes are economically equivalent if they describe the same assets in possibly different currency units.

Our first simple result shows that our modelling approach does not depend on the initial choice of \tilde{S} and has nice semimartingale properties, in the following sense.

Lemma 1.1. *Suppose that \tilde{S} and \tilde{S}' are economically equivalent. If \tilde{S} satisfies (1.1) or (1.3), then so does \tilde{S}' . If \tilde{S} satisfies both (1.1) and (1.3) and we choose a semimartingale representative $S = \tilde{D}\tilde{S}$, then each semimartingale representative $S' = \tilde{D}'\tilde{S}'$ is economically equivalent to S with an exchange rate process $D \in \tilde{\mathcal{D}}$ which is even a semimartingale.*

Proof. Since $\tilde{D}' > 0$ in the sense of (1.2), (1.1) directly transfers from \tilde{S} to $\tilde{S}' = \tilde{D}'\tilde{S}$. From (1.3) for \tilde{S} , we obtain $S = \tilde{D}\tilde{S}$ for some semimartingale S and

some $\tilde{D} \in \tilde{\mathcal{D}}$. So $S = \tilde{D}(\tilde{S}'/\tilde{D}') = (\tilde{D}/\tilde{D}')\tilde{S}'$ is a semimartingale and \tilde{D}/\tilde{D}' is in $\tilde{\mathcal{D}}$, and we see that \tilde{S}' also satisfies (1.3). If $S' = \bar{D}'\tilde{S}'$ is a semimartingale, we can use $\tilde{S}' = \tilde{D}'\tilde{S}$ to write $S' = DS$ with $D := \bar{D}'(\tilde{D}'/\bar{D})$ which is clearly in $\tilde{\mathcal{D}}$. But $S = \tilde{D}\tilde{S}$ and $S' = \tilde{D}'\tilde{S}'$ both also satisfy (1.1), and so we can write $D = (\sum_{i=1}^N S'^i)/(\sum_{i=1}^N S^i)$ to see that D is also a semimartingale. \square

In the sequel, we always assume that (1.1) and (1.3) are satisfied, and we choose a semimartingale representative S . All other semimartingale representatives are then economically equivalent to S with a (*semimartingale*) *exchange rate process*, and we introduce the set of exchange rate processes,

$$\begin{aligned} \mathcal{D} &:= \tilde{\mathcal{D}} \cap \{\text{semimartingales}\} \\ &= \{\text{all real-valued semimartingales } D = (D_t)_{t \in [0, T]} \text{ with } \inf_{t \in [0, T]} D_t > 0\}, \end{aligned}$$

and the *market generated by S* , which is

$$\mathcal{S} := \{S' = DS : D \in \mathcal{D}\}.$$

The key difference between S and \tilde{S} is that S is a semimartingale, and we exploit this when we formalise trading and self-financing strategies with the help of stochastic integrals. Up to a change of currency unit, however, S and \tilde{S} agree; so we can view the choice of working with S as merely dictated by convenience, and we could always rewrite everything back into the units of \tilde{S} if that is preferred for some reason; see also Remark 1.5 below for more details.

Example 1.2 (*Classic setup of mathematical finance*). One particular case is what we call the *classic setup* of mathematical finance. Suppose there is one asset which has for \mathbb{P} -almost all ω , a positive price. (A bit more precisely, we need that $\mathbb{P}[\inf_{t \in [0, T]} \tilde{S}_t^k > 0] = 1$ for some k , so that $\tilde{S}^k \in \tilde{\mathcal{D}}$ is a generalised exchange rate process.) Then we can express all other assets in units of that special asset by defining $X^i := \tilde{S}^i/\tilde{S}^k$ and then relabel the assets; we call that particular asset k now asset 0 or *bank account*, and we call the other $d := N - 1$ assets the *risky assets*, discounted by the bank account. Then we have $N = d + 1$ basic assets; but they are not symmetric because one of them is a riskless bank account which can never reach the value 0. Moreover, if there are initially several assets like \tilde{S}^k , there is arbitrariness in the choice of the one we use for discounting. As a consequence, if we define concepts in terms of X , they depend implicitly on the choice of the exchange rate process \tilde{S}^k , and it may become quite difficult to keep track of this all the time.

The vast majority of papers in mathematical finance works with the end result of the above setup. Usually, they start with an \mathbb{R}^d -valued process X and call this the (discounted) price process of d risky assets. Almost without exception, it is also assumed (but very often not mentioned explicitly) that there is in addition to X a riskless bank account whose price is identically 1—and this assumption is exploited in the standard problem formulations. (The papers also assume that X is a semimartingale, which corresponds to our choice of a semimartingale representative S .)

As one can see, the classic setup is intrinsically asymmetric. This hides or obscures a number of important phenomena, and we therefore want to start with a *symmetric* treatment of all assets. Since we make no assumptions on \tilde{D} in (1.2) except strict positivity, all our results include the classic setup with nonnegative prices; but they do not exploit its assumptions and asymmetry, and hence they are both more general and in our view more natural. The simplest example of a model which cannot be formulated in the classic setup is one with two assets ($N = 2$); they are both nonnegative, but both can default, i.e., become 0. One of them hits 0 at some (maybe random) time on a set A only; the other hits 0 on A^c only. If $0 < \mathbb{P}[A] < 1$, this cannot be put into the form of the classic setup. For a more detailed and intuitive formulation, see Example 2.1 below.

Remark 1.3. One could of course argue in the above example that adding a third asset of the form $S^3 = \alpha S^1 + (1 - \alpha)S^2$ with $\alpha \in (0, 1)$ would lead us back into the classic setup without changing the market because we have the same trading opportunities. However, this easy way out is an ad hoc problem fix, and it also raises the question how the resulting classic setup depends perhaps on the choice of α . Rather than trying to find a case-by-case approach, we prefer to deal with (1.1) and (1.3) in a general and systematic way.

Now let us return to our basic model. We want to describe (frictionless) continuous trading and work with self-financing strategies; so we need to use stochastic integrals, and therefore exploit below that S is a semimartingale. Again, this includes the classic setup.

One direct consequence of our symmetric formulation is as follows. If $D' \in \tilde{\mathcal{D}}$ is any generalised exchange rate process, the process $S' := D'S$ describes the same market, but in a different currency unit. (Like the original currency unit, the new unit induced by the “exchange rate” D may or may not be available for trade.) So our initially chosen $S \in \mathcal{S}$ is just one fixed semimartingale representative of the market \mathcal{S} generated by S or \tilde{S} .

In the sequel, we only want to work with notions which are independent of the choice of a specific semimartingale representative $S' \in \mathcal{S}$ (or a particular currency unit). More precisely, we want to have that a notion holds for our fixed semimartingale representative S if and only if it holds for each $S' \in \mathcal{S}$, in which case we also say that the notion holds for the market \mathcal{S} and call it “*numéraire-independent*”. Wherever this “*numéraire independence*” is not directly clear from the context or the definitions, we shall make a comment or give an explanation.

1.1 Self-financing strategies, numéraires and strategy cones

In this section, we introduce trading strategies. This is almost standard, with small (but important) differences because we are not in the classic setup. Recall that $S \in \mathcal{S}$ is a semimartingale representative of the market \mathcal{S} .

Definition 1.4. Fix a stopping time $\sigma \in \mathcal{T}_{[0, T]}$. A *self-financing strategy* (for \mathcal{S}) on $[[\sigma, T]]$ is an N -dimensional product-measurable process ϑ which is predictable

on $[[\sigma, T]]$, in $L_\sigma(S)$, and such that

$$V(\vartheta)(S) := \vartheta \cdot S = \vartheta_\sigma \cdot S_\sigma + \int_\sigma \vartheta_u \, dS_u \quad \mathbb{P}\text{-a.s. on } [[\sigma, T]] \quad (1.4)$$

We denote the space of all these strategies by $L_\sigma^{\text{sf}}(\mathcal{S})$ or just L_σ^{sf} , and we call $V(\vartheta)(S)$ the *value process* of ϑ (in the currency unit corresponding to S).

It is not immediately obvious but true that the concept of a self-financing strategy as above is “numéraire-independent”. Indeed, the result and the proof of Lemma II.2.5 for $\sigma \neq 0$ is mutatis mutandis the same as for $\sigma = 0$, and so if ϑ is in $L_\sigma(S)$ and satisfies (1.4), then it is also in $L_\sigma(S')$ and satisfies (1.4) for S' instead of S . In particular, writing $L_\sigma^{\text{sf}}(\mathcal{S})$ and not $L_\sigma^{\text{sf}}(S)$ is justified. Another consequence of the above argument is that the value process of a self-financing strategy ϑ satisfies the “*exchange rate consistency property*” (or *change-of-numéraire formula*)

$$V(\vartheta)(DS) = DV(\vartheta)(S) \quad \text{for every exchange rate process } D \in \mathcal{D}. \quad (1.5)$$

This means that when we change units from S to $S' = DS$, the wealth from ϑ in new units is simply the old wealth multiplied by D —as one expects from basic financial intuition.

Remark 1.5. (a) The exchange rate consistency (1.5) is economically completely natural. It has appeared, among others, in El Karoui et al. [17], Gouriéroux et al. [26] or Schweizer and Takaoka [73]. Since (1.5) *follows* from the definition (1.4) for any semimartingale representative S , it is natural to extend it *by definition* to all other representatives \tilde{S} as well. So if we want to work with self-financing strategies not for S , but the original (possibly non-semimartingale) $\tilde{S} = (1/\tilde{D})S$, we rewrite the self-financing condition (1.4) in the units corresponding to \tilde{S} as

$$\begin{aligned} V(\vartheta)(\tilde{S}) &:= \vartheta \cdot \tilde{S} = \vartheta_\sigma \cdot \tilde{S}_\sigma + \frac{1}{\tilde{D}} \int_\sigma \vartheta_u \, dS_u \\ &= \vartheta_\sigma \cdot \tilde{S}_\sigma + \frac{1}{\tilde{D}} \int_\sigma \vartheta_u \, d(\tilde{D}\tilde{S})_u \quad \mathbb{P}\text{-a.s. on } [[\sigma, T]], \end{aligned} \quad (1.6)$$

i.e., by multiplying everything by the exchange rate process $1/\tilde{D}$ at the appropriate time. This avoids the need of defining stochastic integrals with respect to \tilde{S} (which might even be impossible).

(b) In the classic setup with $N = d+1$, $S = (1, X)$ and discounted asset prices given by the \mathbb{R}^d -valued semimartingale X , self-financing strategies on $[[\sigma, T]]$ can be identified with pairs (v_σ, ψ) of \mathcal{F}_σ -measurable random variables v_σ and \mathbb{R}^d -valued predictable X -integrable processes ψ . This is because we can write (1.4) for a strategy $\vartheta = (\eta, \psi)$ in $S = (1, X)$ as

$$\eta = V(\vartheta)(S) - \psi \cdot X = v_\sigma + \int_\sigma \vartheta_u \, dX_u,$$

setting $v_\sigma := V_\sigma(\vartheta)(S)$ and using that asset 0 has a constant price of 1; see for example Bingham and Kiesel [5, Proposition 4.1.3] or Elliott and Kopp [18, Lemma 2.2.1]. Since it is so familiar, this identification of ϑ with (v_σ, ψ) is usually done even without mention in most papers. In our symmetric setup, such a simple identification is no longer possible; trading strategies must be treated as processes of dimension $N = d+1$, and the self-financing condition (1.4) imposes a nontrivial linear constraint on their coordinates.

Clearly, L_σ^{sf} is a vector space. It is also closed under multiplication with \mathcal{F}_σ -measurable random variables, which means that on $[[\sigma, T]]$, we can scale a strategy not only by a constant (as follows from the vector space structure), but even by a random factor if this is known at the beginning σ of the time period on which we trade.

To avoid doubling phenomena, we usually consider sub-cones of L_σ^{sf} for “allowed” trading. We first give the abstract definition.

Definition 1.6. For a stopping time $\sigma \in \mathcal{T}_{[0, T]}$, a *strategy cone (for \mathcal{S}) on $[[\sigma, T]]$* is a nonempty subset $\Gamma \subseteq L_\sigma^{\text{sf}}$ with the properties

- (a) if $\vartheta^{(1)}, \vartheta^{(2)} \in \Gamma$ and $c_\sigma^{(1)}, c_\sigma^{(2)} \in \mathbf{L}_+^0(\mathcal{F}_\sigma)$, then $c_\sigma^{(1)}\vartheta^{(1)} + c_\sigma^{(2)}\vartheta^{(2)} \in \Gamma$,
- (b) if $(\vartheta^{(n)})_{n \in \mathbb{N}}$ is a countable family in Γ and $(A_n)_{n \in \mathbb{N}}$ an \mathcal{F}_σ -measurable partition of Ω , then $\sum_{n=1}^\infty \mathbf{1}_{A_n} \vartheta^{(n)} \in \Gamma$.

A family of strategy cones $(\Gamma_\sigma)_{\sigma \in \mathcal{T}_{[0, T]}}$, where each Γ_σ is a strategy cone on $[[\sigma, T]]$, is called *time-consistent* if $\Gamma_{\sigma_1} \subseteq \Gamma_{\sigma_2}$ for $\sigma_1 \leq \sigma_2$ in $\mathcal{T}_{[0, T]}$.

The simplest example of a strategy cone on $[[\sigma, T]]$ is L_σ^{sf} itself. Moreover, the family $(L_\sigma^{\text{sf}})_{\sigma \in \mathcal{T}_{[0, T]}}$ is clearly time-consistent. Another example is given below in Definition 1.8.

If $\Gamma \subseteq L_\sigma^{\text{sf}}$ is a strategy cone on $[[\sigma, T]]$, we set, for any norm $\|\cdot\|$ in \mathbb{R}^N ,

$$\mathbf{b}\Gamma := \left\{ \vartheta \in \Gamma : \sup_{(\omega, t) \in \Omega \times [0, T]} \|\vartheta \mathbf{1}_{[[\sigma, T]]}\| \leq c_\sigma \text{ for some } c_\sigma \in \mathbf{L}_+^0(\mathcal{F}_\sigma) \right\},$$

$$\mathbf{h}\Gamma := \left\{ \vartheta \in \Gamma : \vartheta \mathbf{1}_{[[\sigma, T]]} = \vartheta_\sigma \mathbf{1}_{[[\sigma, T]]} \right\}.$$

It is straightforward to check that $\{0\} \subseteq \mathbf{h}\Gamma \subseteq \mathbf{b}\Gamma \subseteq \Gamma$, and that $\mathbf{h}\Gamma$ and $\mathbf{b}\Gamma$ are again strategy cones on $[[\sigma, T]]$. We call $\vartheta \in \mathbf{b}\Gamma$ a *bounded* strategy in Γ and $\vartheta \in \mathbf{h}\Gamma$ an *invest-and-keep* strategy in Γ . Note that if $(\Gamma_\sigma)_{\sigma \in \mathcal{T}_{[0, T]}}$ is a time-consistent family of strategy cones, then $(\mathbf{b}\Gamma_\sigma)_{\sigma \in \mathcal{T}_{[0, T]}}$ and $(\mathbf{h}\Gamma_\sigma)_{\sigma \in \mathcal{T}_{[0, T]}}$ are so, too.

Remark 1.7. (a) An invest-and-keep strategy is the simplest and most naive strategy, and requires from the investor only one decision: He should decide when to start his investment and how much of each asset he wants to have. After that, he just waits until the end T of the trading interval.

(b) Calling strategies in $\mathbf{b}\Gamma$ *bounded* may seem puzzling at first sight. However, each $\vartheta \in \mathbf{b}\Gamma$ is uniformly bounded on $[[\sigma, T]]$ by an \mathcal{F}_σ -measurable random

variable, and the latter play the role of “constants” on $[[\sigma, T]]$. (Recall that similarly, L_σ^{sf} is closed under multiplication with \mathcal{F}_σ -measurable random variables, and we also stipulate the cone structure in Definition 1.6 for nonnegative \mathcal{F}_σ -measurable random variables.) In particular, for $\sigma = 0$, we recover the usual concept of a bounded strategy.

It is well known that to avoid undesirable phenomena in a financial market, one must exclude doubling-type strategies. The usual way to do that is to impose solvency constraints, i.e., strategies are allowable for trading only if their value processes are bounded below by some quantity. If this should not depend on a specific currency unit, then the only possible choice for the lower bound is 0. This motivates the following definition.

Definition 1.8. Fix a stopping time $\sigma \in \mathcal{T}_{[0, T]}$. We call a strategy $\vartheta \in L_\sigma^{\text{sf}}$ an *undefaultable strategy on $[[\sigma, T]]$* and write $\vartheta \in \mathcal{U}_\sigma$ if

$$V(\vartheta)(S) \geq 0 \text{ P-a.s. on } [[\sigma, T]].$$

The notion of an undefaultable strategy is clearly numéraire-independent. Moreover, each \mathcal{U}_σ is a strategy cone, and $(\mathcal{U}_\sigma)_{\sigma \in \mathcal{T}_{[0, T]}}$ is a time-consistent family of strategy cones.

The next concept we need is a “numéraire”. We first give the definition and then some comments.

Definition 1.9. A strategy $\eta \in L_0^{\text{sf}}$ is called a *numéraire strategy* (for the market \mathcal{S}) if $\inf_{t \in [0, T]} V_t(\eta)(S) > 0$ P-a.s., i.e., if $V(\eta)(S)$ is an exchange rate process. If such an η exists, \mathcal{S} is called a *numéraire market*.

Note that the above concept is numéraire-independent since $V(\eta)(S) > 0$ holds for some $S \in \mathcal{S}$ if and only if it holds for all $S' \in \mathcal{S}$, due to (1.5). Note also that any numéraire strategy is automatically in \mathcal{U}_0 .

By our nondegeneracy assumption (1.1), the *market portfolio* $\eta^{\mathcal{S}} := (1, \dots, 1)$ of holding one unit of each asset is always a numéraire strategy, and lies in $\mathbf{h}\mathcal{U}_0$. Similarly, in the classic setup $S = (B, Y)$ with $N = d + 1$ assets, where B denotes an undiscounted “bank account” satisfying $\inf_{t \in [0, T]} B_t > 0$ and Y denotes d undiscounted “risky assets”, the buy-and-hold strategy $e_1 = (1, 0, \dots, 0)$ of the “bank account” is a numéraire strategy. For general η , in the classic setup, one calls $V(\eta)(S)$ (which equals B in the latter example) a “numéraire” or “tradable numéraire”. But doing this implies that we work in the currency unit corresponding to the particular representative S , and such a dependence is precisely what we want to avoid. We therefore describe “numéraires” not by their wealth, but in terms of their asset holdings, which do not depend on any currency unit.

For each numéraire strategy η , there exists a P-a.s. unique *numéraire representative* $S^{(\eta)} \in \mathcal{S}$ such that $V(\eta)(S^{(\eta)}) \equiv 1$. It is given explicitly by “ $V(\eta)$ -discounted prices”

$$S^{(\eta)} := \frac{S}{V(\eta)(S)}. \quad (1.7)$$

Note that $S^{(\eta)}$ is well defined because $V(\eta)$ satisfies the exchange rate consistency property (1.5) (or (1.6)); this ensures that the right-hand side of (1.7) yields the same result for any other representative $S' = DS$ of \mathcal{S} . In the classic setup as above with a bank account B and $e_1 = (1, 0, \dots, 0)$, (1.7) reduces to $S^{(e^1)} = S/B = (1, X)$ as in Example 1.2.

1.2 Maximal strategies

Suppose we are given a class Γ of possible strategies. A strategy $\vartheta \in \Gamma$ can be considered as a “reasonable investment” from that class only if it cannot be directly improved by another strategy from the same class. More precisely, using strategies in Γ with the same (or a lower) initial investment should not allow one to create more wealth at time T . It is natural to call such a strategy ϑ *maximal*; see Remark 1.12 below for more comments.

Definition 1.10. Let $\sigma \in \mathcal{T}_{[0, T]}$ be a stopping time and Γ a strategy cone on $[[\sigma, T]]$. Then a strategy $\vartheta \in \Gamma$ is called *weakly maximal for Γ* if there is no pair $(f, \bar{\vartheta})$ consisting of a nonzero random variable $f \in \mathbf{L}_+^0(\mathcal{F}_T) \setminus \{0\}$ and a strategy $\bar{\vartheta} \in \Gamma$ such that

$$V_\sigma(\bar{\vartheta})(S) \leq V_\sigma(\vartheta)(S) \text{ P-a.s.} \quad \text{and} \quad V_T(\bar{\vartheta})(S) \geq V_T(\vartheta)(S) + f \text{ P-a.s.} \quad (1.8)$$

Note above that f , which satisfies $f \geq 0$ P-a.s. and $\mathbb{P}[f > 0] > 0$, stands for the extra wealth at time T , on top of what we get from ϑ , that we generate by $\bar{\vartheta}$ without increasing the initial capital at time σ . If we require for maximality that (1.8) is impossible even in an approximate sense, we are led to a stronger notion of maximality.

Definition 1.11. Let $\sigma \in \mathcal{T}_{[0, T]}$ be a stopping time and Γ a strategy cone on $[[\sigma, T]]$. Then a strategy $\vartheta \in \Gamma$ is called *strongly maximal for Γ* if there is no nonzero random variable $f \in \mathbf{L}_+^0(\mathcal{F}_T) \setminus \{0\}$ such that for all $\varepsilon > 0$, there exists a strategy $\bar{\vartheta} \in \Gamma$ with

$$V_\sigma(\bar{\vartheta})(S) \leq V_\sigma(\vartheta)(S) + \varepsilon \text{ P-a.s.} \quad \text{and} \quad V_T(\bar{\vartheta})(S) \geq V_T(\vartheta)(S) + f \text{ P-a.s.} \quad (1.9)$$

Again, f stands for the nontrivial extra wealth we should like to achieve by changing ϑ to $\bar{\vartheta}$; and if ϑ is strongly maximal, this cannot be done, not even if we are allowed a small but strictly positive increase of initial capital at σ . Both concepts are clearly numéraire-independent.

Remark 1.12. (a) The terminology “maximal strategy” goes back at least to Delbaen and Schachermayer; see [12]. However, the setting there is different from here so that also maximality has a different meaning. More precisely, [12] is cast in the classic setup, uses for Γ the class \mathcal{A} (which is not numéraire-independent in our sense) of so-called admissible strategies on $[0, T]$, and a priori also imposes some absence-of-arbitrage conditions. In our terminology, a maximal strategy in the sense of [12] is then weakly maximal for \mathcal{A} .

(b) Both above definitions of maximality are slightly different from the Definitions III.3.1 and III.3.11 above. However, if Γ is a strategy cone (on $\llbracket 0, T \rrbracket$) which allows switching to numéraires (see Definition III.2.1), e.g. $\Gamma = \mathcal{U}_0$, using Propositions III.3.5, III.3.8, III.3.17 and III.3.18, it is not difficult to check that both definitions coincide (for strategy cones on $\llbracket 0, T \rrbracket$).

(c) It is clear that strong implies weak maximality, but the converse does not hold in general; indeed, Example III.3.14 shows that there are markets, where *every* $\vartheta \in \mathcal{U}_0$ is weakly maximal for \mathcal{U}_0 , but *no* $\vartheta \in \mathcal{U}_0$ is strongly maximal for \mathcal{U}_0 . Notwithstanding, if the zero strategy 0 is strongly maximal for \mathcal{U}_σ , $\sigma \in \mathcal{T}_{[0, T]}$, then weak implies strong maximality for \mathcal{U}_σ ; see Lemma 3.10 below.

1.3 Viability and efficiency criteria for markets

A financial market should behave in a reasonable manner, and this should be reflected in the properties of its model description. Let us formalise this and then explain the intuition.

Definition 1.13. A market \mathcal{S} is called

- *statically viable* if the zero strategy 0 is strongly maximal for $\mathbf{h}\mathcal{U}_\sigma$, for each $\sigma \in \mathcal{T}_{[0, T]}$.
- *dynamically viable* if the zero strategy 0 is strongly maximal for \mathcal{U}_σ , for each $\sigma \in \mathcal{T}_{[0, T]}$.

Static viability means that at every stopping time σ , just doing nothing cannot be improved by a naive invest-and-keep strategy. Dynamic viability is even stronger—one cannot improve on inactivity by trading, even if one trades continuously in time. Of course, in both cases, one must observe the constraint (from \mathcal{U}_σ) of keeping wealth nonnegative.

Dynamic viability by its definition implies static viability, but the converse is not true. Even in a finite-state discrete-time setup (with more than one time period), static viability is strictly weaker than dynamic viability; see Example 2.5 below. For finite discrete time, we show below in Lemma 1.18 that dynamic viability is equivalent to the classic *no-arbitrage* condition NA. In general, Corollary 3.9 below implies that a market \mathcal{S} is dynamically viable if and only if the zero strategy is strongly maximal for \mathcal{U}_0 ; in other words, it is enough to check maximality for the starting time 0 instead of all $\sigma \in \mathcal{T}_{[0, T]}$. Hence, dynamic viability is equivalent to numéraire-independent no-arbitrage (NINA); see Chapter III.

The next concept strengthens viability.

Definition 1.14. A market \mathcal{S} is called

- *statically efficient* if each strategy $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ is strongly maximal for $\mathbf{h}\mathcal{U}_\sigma$, for each $\sigma \in \mathcal{T}_{[0, T]}$.
- *dynamically efficient* if each strategy $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ is strongly maximal for \mathcal{U}_σ , for each $\sigma \in \mathcal{T}_{[0, T]}$.

Viability means that one cannot improve the zero strategy of doing nothing. Efficiency means that the market has even more structure—all naive invest-and-keep strategies are good in the sense that they cannot be improved, in a certain class, without risk. It is clear from the definition that dynamic efficiency implies static efficiency, and like for viability, the converse is not true; this is also illustrated below in Example 2.5.

The connection between viability and efficiency is more subtle. It is clear that efficiency (dynamic or static) implies viability (of the same kind). At first sight, one might expect that the converse holds as well—why should it make a difference whether one tries to improve zero or a general naive invest-and-keep strategy? But it turns out that there *is* a difference, and the reason behind this is that one must look for improvements in the class of undefaultable strategies. This is a cone, but not a linear space. If the strategy to be improved is changed from 0 to another strategy, this upsets the balance with the strategies we are allowed to use for improvement. In general, when we take differences to construct something better, this leads us outside a cone—except of course if we subtract zero.

Interestingly and notably, the above difference between efficiency and viability does not yet appear in finite discrete time. In fact, we show below in Lemmas 1.18 and 1.19 that static/dynamic efficiency is equivalent to static/dynamic viability in finite discrete time. This reflects the well-known fact that if one can achieve arbitrage in finite discrete time with a general strategy, one can also achieve arbitrage with an undefaultable strategy. This result is specific to finite discrete time because the proof relies on backward induction; see for example Elliott and Kopp [18, Section 2.2] or Lamberton and Lapeyre [53, Lemma 1.2.7].

In a market with infinitely many trading dates, things change. In continuous time, Example 4.1 shows that static/dynamic viability in general does not imply static/dynamic efficiency. So the next concept is meaningful.

Definition 1.15. A market \mathcal{S} is called a *bubbly market* if it fails to be dynamically efficient. It is called a *nontrivial bubbly market* if in addition \mathcal{S} is dynamically viable and statically efficient.

The idea behind this definition is simple. We call \mathcal{S} a bubbly market if there is some naive invest-and-keep strategy which can be improved (approximately) by dynamic trading. Of course, this will happen if \mathcal{S} “allows arbitrage”, i.e., if \mathcal{S} is not dynamically viable; but this is not the really interesting situation of a bubbly market. It is much more challenging to study what happens if \mathcal{S} does not admit arbitrage (i.e., is dynamically viable), and nevertheless, naive invest-and-keep strategies in the underlying assets can be improved—however, not by other naive invest-and-keep strategies (i.e., \mathcal{S} is statically efficient), but only by genuinely dynamic trading (i.e., \mathcal{S} is not dynamically efficient).

Note that a nontrivial bubbly market can only appear in a model with infinitely many trading dates. In finite discrete time (see Lemma 1.18), the equivalence of dynamic efficiency and the no-arbitrage property NA implies that every bubbly market allows arbitrage and hence cannot be nontrivially bubbly. We believe that this dichotomy is natural and that some interesting phenomena inherently hinge on an infinite set of trading dates.

Remark 1.16. (a) Throughout this paper, we consider a setting where there is a last trading date; we either work in continuous time on the (right-closed) interval $[0, T]$ or in discrete time on $\{0, 1, \dots, T\}$ (then with $T \in \mathbb{N}$). We believe that results like those for $[0, T]$ can also be developed for trading dates in $[0, \infty)$ or in $\mathbb{N}_0 = \{0, 1, 2, \dots\}$, but one must take some extra care as time goes to ∞ . This is left for future research.

(b) Note that in contrast to much of the existing work on bubbles, our definitions do not involve any martingale or strict local martingale property of some asset prices. Our approach is to start with an economically compelling notion of a bubbly market, and then to *prove* that there must be a close connection to strict local martingales. This is done below in Theorem 3.22.

(c) Another contrast to existing work is that we do not try to define, for one single (or multivariate) asset, whether or not this asset itself is a bubble. We look instead at the market as a whole and try to define (and then characterise) whether or not it contains a “bubble” somewhere. We discuss the connections to the literature in more detail in Section 5.

1.4 The case of finite discrete time

In this section, we show that many concepts and results simplify in a model with finite discrete time. This is of course no surprise since that setting is well known to be much easier than a situation with infinitely many trading dates.

Definition 1.17. We say that the market \mathcal{S} satisfies *no arbitrage (NA)* if there is no strategy $\vartheta \in \mathcal{U}_0$ satisfying, for some (or equivalently all) $S \in \mathcal{S}$,

$$V_0(\vartheta)(S) = 0, \quad V_T(\vartheta)(S) \geq 0 \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[V_T(\vartheta)(S) > 0] > 0. \quad (1.10)$$

For finite discrete time and the classic setup as in Example 1.2, the above is just the standard classic definition of absence of arbitrage; see for example Elliott and Kopp [18, Definition 2.2.3 and the subsequent section]. For a continuous-time model in the classic setup, the above condition was studied under the names NA^+ or NA_+ by Strasser [77] and Hulley [35]. Note also that Definition 1.17 is numéraire-independent, and that the explicit requirement $V_T(\vartheta)(S) \geq 0$ P-a.s. is actually redundant since $\vartheta \in \mathcal{U}_0$.

Lemma 1.18. *For a market \mathcal{S} in finite discrete time, the following are equivalent:*

- (a) \mathcal{S} satisfies NA.
- (b) \mathcal{S} is dynamically viable.
- (c) \mathcal{S} is dynamically efficient.
- (d) For each numéraire strategy η , there exists a probability measure $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that the $V(\eta)$ -discounted price process $S^{(\eta)} = \frac{S}{V(\eta)(S)}$ is a \mathbb{Q} -martingale.

Proof. We show below in Theorem 3.11 that (d) implies (c), and it is clear that (c) implies (b). Next, (b) implies that 0 is (strongly and a fortiori) weakly maximal for \mathcal{U}_0 , which is in turn equivalent to \mathcal{S} satisfying NA, and we obtain (a). So it only remains to argue that (a) implies (d), and this is where we exploit the setting of finite discrete time.

Let η be a numéraire strategy, e.g. the market portfolio $\eta^S = (1, \dots, 1)$. Then we have (1.10), with S replaced by $X := S^{(n)}$. We claim that for $t \in \{1, \dots, T\}$, there is no \mathcal{F}_{t-1} -measurable \mathbb{R}^N -valued random vector $\xi = (\xi^1, \dots, \xi^N)$ such that

$$\xi \cdot (X_t - X_{t-1}) \geq 0 \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[\xi \cdot (X_t - X_{t-1}) > 0] > 0. \quad (1.11)$$

Indeed, if we have such t and ξ , we can define an \mathbb{R}^N -valued predictable process ϑ by

$$\vartheta_k = \begin{cases} 0, & k \leq t-1, \\ \xi - (\xi \cdot X_{t-1})\eta_t, & k = t, \\ \xi \cdot (X_t - X_{t-1})\eta_k, & k > t. \end{cases}$$

It is easy to check that ϑ is in \mathcal{U}_0 due to (1.11), and because $V(\eta)(S^{(n)}) \equiv 1$, ϑ satisfies

$$V_0(\vartheta)(S^{(n)}) = 0 \quad \text{and} \quad \mathbb{P}[V_T(\vartheta)(S^{(n)}) > 0] = \mathbb{P}[\xi \cdot (X_t - X_{t-1}) > 0] > 0,$$

contradicting (1.10). Thus, applying Föllmer and Schied [25, Proposition 5.11 and Theorem 5.16] to the model $(1, X)$ gives $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $X = S^{(n)}$ is a \mathbb{Q} -martingale, and we have (d). \square

Lemma 1.19. *For a market \mathcal{S} in finite discrete time, the following are equivalent:*

- (a) \mathcal{S} is statically viable.
- (b) \mathcal{S} is statically efficient.

Proof. It is clear that (b) implies (a). For the converse, due to Lemma 6.2 below, it suffices to show that 0 is weakly maximal for $\mathbf{h}L_s^{\text{sf}}$, for each $s \in \{0, \dots, T\}$. Suppose to the contrary that we have $s \in \{0, \dots, T\}$ and $\vartheta \in \mathbf{h}L_s^{\text{sf}}$ with $V_s(\vartheta)(S) \leq 0$ P-a.s., $V_T(\vartheta)(S) \geq 0$ P-a.s. and $\mathbb{P}[V_T(\vartheta)(S) > 0] > 0$. Now if $\mathbb{P}[V_t(\vartheta)(S) < 0] = 0$ for $t \in \{s, \dots, T-1\}$, then ϑ is in $\mathbf{h}\mathcal{U}_s$ and $V_s(\vartheta)(S) = 0$ P-a.s., and so 0 fails to be weakly, hence also strongly maximal for $\mathbf{h}\mathcal{U}_s$, in contradiction to static viability. Thus, the set $A_t := \{V_t(\vartheta)(S) < 0\}$ has $\mathbb{P}[A_t] > 0$ for some $t \in \{s, \dots, T-1\}$. Let $t^* = \max\{t \in \{s, \dots, T-1\} : \mathbb{P}[A_t] > 0\}$ and take η to be a numéraire strategy in $\mathbf{h}\mathcal{U}_{t^*}$ (e.g. the market portfolio $\eta^S = (1, \dots, 1)$). From the definitions of t^* and A_{t^*} , it is not difficult to check that the strategy $\vartheta^* := \mathbb{1}_{A_{t^*}}(\vartheta - V_{t^*}(\vartheta)(S^{(n)})\eta)$ is in $\mathbf{h}\mathcal{U}_{t^*}$ and satisfies

$$V_{t^*}(\vartheta^*)(S^{(n)}) = 0 \text{ P-a.s.} \quad \text{and} \quad V_T(\vartheta^*)(S^{(n)}) \geq -\mathbb{1}_{A_{t^*}} V_{t^*}(\vartheta)(S^{(n)}) \text{ P-a.s.}$$

Because $V_{t^*}(\vartheta)(S^{(n)}) < 0$ on A_{t^*} and $\mathbb{P}[A_{t^*}] > 0$, this shows that 0 is not weakly, hence also not strongly maximal for $\mathbf{h}\mathcal{U}_{t^*}$, in contradiction to static viability of \mathcal{S} . \square

2 First examples

In this section, we give a number of examples to illustrate the ideas and concepts introduced so far, focussing mainly on the (sometimes subtle) differences between different notions. With the exception of Example 2.3, we do not yet need here the dual characterisations presented in Section 3.

We start with an explicit example to show that our approach is more general than the classic setup of mathematical finance discussed in Example 1.2.

Example 2.1 (*A market that does not fit into the classic setup*). Consider two independent Brownian motions $W^1 = (W_t^1)_{t \in [0, T]}$ and $W^2 = (W_t^2)_{t \in [0, T]}$ and two stopping times τ_1 and τ_2 with $\mathbb{P}[0 < \tau_i < T] > 0$, $i = 1, 2$, for the usual (augmented) filtration generated by W^1 and W^2 . In addition, let X_1 and X_2 be random variables which are independent of W^1 and W^2 and satisfy $\mathbb{P}[X_i = -1] = p_i$, $\mathbb{P}[X_i = \alpha_i] = 1 - p_i$ with $\alpha_i > 0$, $p_i \in (0, 1)$, $i = 1, 2$, and $\mathbb{P}[X_1 = -1, X_2 = -1] = 0$. Define the one-jump processes $N^i = (N_t^i)_{t \in [0, T]}$ by $N_t^i = X_i \mathbb{1}_{\{t \geq \tau_i\}}$, $i = 1, 2$, and let $(\mathcal{F}_t)_{t \in [0, T]}$ be the (augmented) filtration generated by W^1, W^2, N^1, N^2 . Let $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1, \sigma_2 > 0$, and define the process $S = (S_t^1, S_t^2)_{t \in [0, T]}$ by the SDEs

$$dS_t^i = S_{t-}^i (\mu_i dt + \sigma_i dW_t^i + dN_t^i), \quad S_0^i = s_i > 0, i = 1, 2.$$

It is straightforward to see that for $i = 1, 2$, prior to τ_i , S^i is a geometric Brownian motion with drift μ_i and volatility σ_i . At τ_i , it either jumps to zero (if $X_i = -1$) and stays there, or it jumps to $(1 + \alpha_i)S_{\tau_i-}^i$ (if $X_i = \alpha_i$) and evolves from there as a geometric Brownian motion with the same parameters μ_i, σ_i as before the jump.

Note that both S^1 and S^2 may jump to zero with positive probability; but because of $\mathbb{P}[X_1 = -1, X_2 = -1] = 0$, at least one of them stays positive until time T . This shows that (1.1) is satisfied for S , and (1.3) also holds (with $D \equiv 1$) since S is a semimartingale. It is clear that this example cannot be treated in the classic setup since no asset price process is guaranteed to remain positive with probability 1.

One intuitive (but somewhat extreme) situation where this model is natural is as follows. Suppose $\tau_1 = \tau_2 =: \tau$ and $X_1 = -X_2$, which forces $\alpha_1 = \alpha_2 = 1$. Then both S^1 and S^2 evolve as (independent) geometric Brownian motions up to the random time τ , which can be interpreted as the time that some important announcement is made. One of the two assets then drops to 0, while the other instantaneously doubles its price and then continues as a GBM. Which of the two assets “defaults” is determined by the value of X_1 , which can therefore be interpreted as a signal to the market at time τ . For instance, S^1, S^2 could be the market values of two firms competing for a monopoly in the same market sector, and $X_1 = +1$ means that firm 1 gets control of that sector—for example because it obtains the single licence available for a telecom market.

Remark 2.2. In the above example, we have processes that jump to zero when they default. One could also construct analogous examples where the processes

are continuous and creep down to zero. Our reason for using a jump process is just that this gives a simpler construction.

Because of its purpose, Example 2.1 cannot be formulated in the classic setup. However, all the subsequent examples in this section have one asset whose price is identically 1; this means that they can be realised in the familiar classic setup—however, we still work with our concepts which are based on undefaultable as opposed to classic admissible strategies.

Our second example, which is essentially taken from Delbaen and Schachermayer [11], exhibits a model and an explicit strategy which is not maximal. It also has other properties to which we shall return later.

Example 2.3 (*A strategy which is not maximal*). Let $W = (W_t)_{t \in [0, T]}$ be a Brownian motion with respect to a given filtration $(\mathcal{F}_t)_{t \in [0, T]}$; this need not be generated by W . Consider the market \mathcal{S} generated by $S = (S^1, S^2) = (1, X_t)_{t \in [0, T]}$, where X is the well-known three-dimensional Bessel process BES^3 , i.e., the unique strong solution of the SDE

$$dX_t = \frac{dt}{X_t} + dW_t, \quad X_0 = s_0 > 0. \quad (2.1)$$

We claim that the buy-and-hold strategy $e_1 = (1, 0)$ of the first asset fails to be weakly (and a fortiori strongly) maximal for \mathcal{U}_0 ; in other words, it is not a good idea in this market to put money into the bank account.

To see this, we first look at the buy-and-hold strategy $e_2 = (0, 1)$ of the second asset. Note that $S^{(e_2)} = (1/X, 1)$. Itô's formula shows that $Y := 1/X$ satisfies the SDE

$$dY_t = -|Y_t|^2 dW_t, \quad Y_0 = 1/s_0, \quad (2.2)$$

and so Y is a local \mathbb{P} -martingale. In fact, it is well known that Y is even a strict local \mathbb{P} -martingale; see e.g. Revuz and Yor [69, Proposition VI.3.3 and Exercise V.2.13]. Due to Theorem 3.1 below, the market \mathcal{S} is dynamically viable, and by Lemma 3.10 below, weak and strong maximality for each \mathcal{U}_σ are therefore equivalent.


In order to establish that e_1 is not maximal for \mathcal{U}_0 , we use Theorem VI.2.1. Since $S^{(e_2)} = (1/X, 1) = (Y, 1)$ is a local \mathbb{P} -martingale and $V(e_2)(S^{(e_2)}) \equiv 1$ is a (true) \mathbb{P} -martingale, the numéraire strategy e_2 is maximal for \mathcal{U}_0 by Theorem VI.2.1 (d). Therefore, by Theorem VI.2.1 (c), it suffices to show that for every $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T for which $S^{(e_2)}$ is a local \mathbb{Q} -martingale, $V(e_1)(S^{(e_2)}) = Y$ is a *strict* local \mathbb{Q} -martingale. So let $\mathbb{Q} \approx \mathbb{P}$ be such that $S^{(e_2)} = (Y, 1)$ is a local \mathbb{Q} -martingale. Then Y still satisfies the SDE (2.2), written for clarity as $dY_t = -|Y_t|^2 dW_t^{\mathbb{P}}$, so that $\langle Y \rangle_t = \int_0^t |Y_s|^4 ds$. Moreover, since $Y = 1/X$ is positive, we can write $W^{\mathbb{P}} = \int -\frac{1}{|Y|^2} dY$. The latter process is a continuous local \mathbb{Q} -martingale like Y , and it has quadratic variation $\langle W^{\mathbb{P}} \rangle_t = \int_0^t \frac{1}{|Y_s|^4} d\langle Y \rangle_s = t$, $t \in [0, T]$. Hence, by Lévy's characterisation of Brownian motion, $W^{\mathbb{P}}$ is also a \mathbb{Q} -Brownian motion $W^{\mathbb{Q}}$, and so Y has the same distribution under \mathbb{Q} as under \mathbb{P} . In particular, Y is also a strict local \mathbb{Q} -martingale, and so e_1 fails to be maximal for \mathcal{U}_0 .

Of course, this example is classic; it is well known that the BES^3 process can be used to construct counterexamples, and the non-maximality of e_1 is also already pointed out in [11, Remark after Corollary 5] (even if maximality there has a slightly different meaning, as discussed in Remark 1.12). One small but neat novelty of the present example is that we do not need to assume that the filtration is generated by W or by X .

Remark 2.4. Although the filtration in Example 2.3 is general, the model is nevertheless quite special because it is complete for its own filtration. More precisely, if we denote by $(\mathcal{F}_t^S)_{t \in [0, T]}$ the (augmented) filtration generated by S (or equivalently by W or by Y , due to (2.2)), then the local \mathbb{Q} -martingale $Y = 1/X$ has in $(\mathcal{F}_t^S)_{t \in [0, T]}$ the predictable representation property. We shall see below in Example 4.3 another natural example which is genuinely incomplete.

For both viability and efficiency, we have introduced a static and a dynamic version. The next, very simple example shows that the two concepts are different. It also clarifies where the difference comes from.

Example 2.5 (*Static versus dynamic viability/efficiency*). Let \mathcal{S} be the market generated by $S = (S^1, S^2) = (1, X_k)_{k \in \{0, 1, 2\}}$ with X given by the following event tree, where each branch is assumed to have positive probability.



The key feature of this model is that from time 0 directly to the final time 2 and from every node at time 1 to time 2, the asset price $S^2 = X$ can both go strictly up and down. Hence each such “one-step” model is arbitrage-free, naive invest-and-keep trading (which precisely corresponds to trading in such a one-step model) cannot improve the zero strategy, and so \mathcal{S} is statically efficient, and a fortiori statically viable. Note that this is because our invest-and-keep strategies must always be kept until the end ($T = 2$ here). However, from time 0 to time 1, X stays the same or goes up; so the one-step model in this sub-tree already admits arbitrage, then so does the whole model S , and \mathcal{S} is not dynamically viable (and a fortiori not dynamically efficient either). So both for viability and efficiency, the static version is weaker than the dynamic one.

3 Dual characterisation of nontrivial bubbly markets

Even though they are mathematically rigorous, the definitions of viability, efficiency and of a bubbly market are economic by nature and intuition. For working

with them and for proving further results, however, it is more useful to have an equivalent description in terms of dual objects. Our goal in this section is to provide such dual characterisations.

3.1 Dual objects in numéraire-independent modelling

To motivate the subsequent definitions, let us first recall the well-known primal and dual objects in the classic setup of mathematical finance; see Example 1.2. There one starts with an \mathbb{R}^d -valued process X and thinks of this as modelling the prices of d risky assets, *expressed in units of a further asset* labelled 0 and called bank account. Sometimes one also says that the bank account (or asset 0) is used as numéraire. (It is actually rare to find a precise definition of the term “numéraire”; the expression seems to be viewed as self-explaining.) With the above notations, *primal objects* are *self-financing strategies*, and these can be parametrised by pairs (v_0, ψ) of initial wealths $v_0 \in L^0(\mathcal{F}_0)$ and \mathbb{R}^d -valued predictable X -integrable processes ψ describing the holdings in the risky assets; see Remark 1.5. *Dual objects* are then *equivalent local martingale measures (ELMMs) for X* , i.e., probability measures \mathbb{Q} equivalent to \mathbb{P} on \mathcal{F}_T , $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that X is a local \mathbb{Q} -martingale. (We do not need σ -martingales when price processes are nonnegative.) Finally, the fundamental theorem of asset pricing (FTAP) says that absence of arbitrage for X under \mathbb{P} (in the sense that X satisfies the condition NFLVR of no free lunch with vanishing risk) is equivalent to the existence of an ELMM $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T for X .

In the above classic setup, the very first step is to choose and fix a “numéraire”, namely asset 0 (the “bank account”). All subsequent definitions and results depend on this, and some concepts even cannot be defined without this. (A more thorough discussion can be found in Chapter I.) In our more general setup, all $N = d + 1$ assets in S are treated symmetrically and we do not (and do not want to) choose a priori any particular numéraire. Primal objects are again self-financing strategies, which are now parametrised by \mathbb{R}^N -valued predictable S -integrable processes ϑ which satisfy the self-financing constraint (1.4). But dual objects, as can be seen in Chapter VI, are now *pairs* consisting of a numéraire η and an ELMM \mathbb{Q} for the corresponding discounted prices.

To make this precise, we first need a clear (and numéraire-independent) definition of “numéraire”. This has already been given above in Definition 1.9 where we have introduced the concept of a numéraire strategy. With that terminology, we can now precisely describe the *dual objects* in our numéraire-independent framework. They are *pairs* (η, \mathbb{Q}) , where η is a *numéraire strategy* and \mathbb{Q} is an *ELMM for the numéraire representative $S^{(\eta)}$* . To illustrate this, we recall from Chapter VI the numéraire-independent version of the FTAP in the setup of this chapter; i.e., for nonnegative (numéraire) markets.

Theorem 3.1. *The following are equivalent:*

- (a) \mathcal{S} is dynamically viable.

- (b) \mathcal{S} satisfies numéraire-independent no-arbitrage (NINA), i.e., the zero strategy 0 is strongly maximal for \mathcal{U}_0 .
- (c) There exists a pair (η, \mathbb{Q}) , where η is a numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that $S^{(\eta)}$ is a local \mathbb{Q} -martingale.
- (d) There exist a representative $\bar{S} \in \mathcal{S}$ and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that \bar{S} is a local \mathbb{Q} -martingale.

Note that in general $S^{(\eta)}$ or \bar{S} might fail to be true \mathbb{Q} -martingales.

Proof. The equivalence of (b), (c) and (d) follows from the equivalence of (a), (b) and (d) in Theorem VI.1.10. (a) implies (b) by the Definition 1.13 of dynamic viability, and that (b) implies (a) is shown below in Corollary 3.9. \square

Remark 3.2. As mentioned before, we get local martingales and not only σ -martingales because our prices are nonnegative. For a comparison of the above result to the classic FTAP, we refer to Chapter VI.

3.2 Contingent claims and superreplication prices

We proceed to recall the notions of contingent claims and superreplication prices from Chapter II and extend them to the setup of this chapter.

Definition 3.3. An *improper contingent claim* at time $\tau \in \mathcal{T}_{[0,T]}$ for the market \mathcal{S} is a map $F : \mathcal{S} \rightarrow \bar{\mathbf{L}}_+^0(\mathcal{F}_\tau)$ satisfying the *exchange rate consistency condition*

$$F(DS) = D_\tau F(S) \text{ P-a.s. for all } S \in \mathcal{S} \text{ and } D \in \mathcal{D}. \quad (3.1)$$

F is called a *contingent claim* at time τ if it is valued in $\mathbf{L}_+^0(\mathcal{F}_\tau)$, and *positive* if it is valued in $\mathbf{L}_{++}^0(\mathcal{F}_\tau)$.

A contingent claim in our abstract setup is a mapping which assigns to each semimartingale representative S (corresponding to one particular choice of currency unit) a payoff $F(S)$ at time τ (in the same unit), which is then just an \mathcal{F}_τ -measurable random variable. The simplest example is the value process $V_\tau(\vartheta)$ at time τ of any self-financing strategy ϑ . The property (3.1) here follows directly from (1.5). The canonical and most general example is obtained as follows: We fix a random variable $g \in \mathbf{L}_+^0(\mathcal{F}_\tau)$, choose a representative $S \in \mathcal{S}$ and define F by $F(S') = F(DS) := D_\tau g$ for any $S' = DS$ in \mathcal{S} . Then g represents a payoff in the currency units corresponding to S , and we could call F , which is well-defined due to (3.1), the contingent claim at time τ induced by g with respect to S ; see also Proposition II.5.3.

Remark 3.4. As for the self-financing condition in (1.6), we extend the relation (3.1) to arbitrary representatives $\tilde{S} = \tilde{D}S$ by setting

$$F(\tilde{S}) := \tilde{D}_\tau F(S).$$

Definition 3.5. Let $\sigma \leq \tau \in \mathcal{T}_{[0,T]}$ be stopping times, Γ a strategy cone on $[[\sigma, T]]$ and F a contingent claim at time τ . The function $\Pi_\sigma(F | \Gamma) : \mathcal{S} \rightarrow \bar{\mathbf{L}}_+^0(\mathcal{F}_\sigma)$ defined by

$$\Pi_\sigma(F | \Gamma)(S) := \text{ess inf} \left\{ v \in \bar{\mathbf{L}}_+^0(\mathcal{F}_\sigma) : \text{there is } \vartheta \in \Gamma \text{ such that on } \{v < \infty\}, \right. \\ \left. V_\sigma(\vartheta)(S) \leq v \text{ P-a.s. and } V_\tau(\vartheta)(S) \geq F(S) \text{ P-a.s.} \right\}$$

is called the *superreplication price of F at time σ for Γ* .

It is not difficult to check that $\Pi_\sigma(F | \Gamma)$ is an improper contingent claim at time σ . The following result lists some other basic properties. Note that these are properties of functions on \mathcal{S} , and that they are all numéraire-independent in the (usual) sense that they hold for some $S \in \mathcal{S}$ if and only if they hold for all $S' \in \mathcal{S}$; this is due to the exchange rate consistency property (3.1). The proofs are straightforward and hence omitted.

Proposition 3.6. *Let $\sigma \leq \tau \in \mathcal{T}_{[0,T]}$ be stopping times, Γ a strategy cone on $[[\sigma, T]]$ and F, F_1, F_2, G contingent claims at time τ with $F \leq G$ P-a.s. Let c_σ be a nonnegative \mathcal{F}_σ -measurable random variable. Then*

$$\begin{aligned} \Pi_\sigma(F | \Gamma) &\leq \Pi_\sigma(G | \Gamma) && \text{(monotonicity),} \\ \Pi_\sigma(c_\sigma F | \Gamma) &= c_\sigma \Pi_\sigma(F | \Gamma) && \text{(positive } \mathcal{F}_\sigma\text{-homogeneity),} \\ \Pi_\sigma(F_1 + F_2 | \Gamma) &\leq \Pi_\sigma(F_1 | \Gamma) + \Pi_\sigma(F_2 | \Gamma) && \text{(subadditivity).} \end{aligned}$$

Note that positive \mathcal{F}_σ -homogeneity implies that $\Pi_\sigma(\mathbf{1}_{A_\sigma} F | \Gamma) = \mathbf{1}_{A_\sigma} \Pi_\sigma(F | \Gamma)$ for $A_\sigma \in \mathcal{F}_\sigma$. In the context of conditional risk measures, this is sometimes called *locality* or the *local property*.

The next auxiliary technical result can be used to approximate superreplication prices.

Lemma 3.7. *Let $\sigma \leq \tau \in \mathcal{T}_{[0,T]}$ be stopping times, Γ a strategy cone on $[[\sigma, T]]$ and F a contingent claim at time τ with $\Pi_\sigma(F | \Gamma) < \infty$ P-a.s. Then for all $\delta > 0$ and all positive contingent claims C at time σ , there exists a strategy $\vartheta \in \Gamma$ satisfying*

$$V_\sigma(\vartheta) \leq \Pi_\sigma(F | \Gamma) + \delta C \text{ P-a.s.} \quad \text{and} \quad V_\tau(\vartheta) \geq F \text{ P-a.s.} \quad (3.2)$$

Proof. First, for any positive contingent claim C at time σ , there exists $S' \in \mathcal{S}$ with $C(S') = 1$ P-a.s. Indeed, set $D_T := 1/C(S) \in \mathbf{L}_{++}^0(\mathcal{F}_\sigma) \subset \mathbf{L}_{++}^0(\mathcal{F}_T)$ (for some fixed $S \in \mathcal{S}$), take $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T with $\mathbb{E}_{\mathbb{Q}}[D_T] < \infty$ and define the exchange rate process D as the RCLL version of the \mathbb{Q} -martingale $D_t = \mathbb{E}_{\mathbb{Q}}[D_T | \mathcal{F}_t]$, $t \in [0, T]$. Note that $D_\sigma = \mathbb{E}_{\mathbb{Q}}[D_T | \mathcal{F}_\sigma] = D_T$ P-a.s. since D_T is \mathcal{F}_σ -measurable. If we set $S' = DS$, the exchange rate consistency (3.1) for C gives

$$C(S') = C(DS) = D_\sigma C(S) = D_T \frac{1}{D_T} = 1 \text{ P-a.s.}$$

Now take $\delta > 0$ and note that $\Pi_\sigma(F | \Gamma)$ is a contingent claim at time σ and $C(S') = 1$. So by the exchange rate consistency (3.1), it suffices to show that there is $\vartheta \in \Gamma$ with

$$V_\sigma(\vartheta)(S') \leq \Pi_\sigma(F | \Gamma)(S') + \delta \text{ P-a.s.} \quad \text{and} \quad V_\tau(\vartheta)(S') \geq F(S') \text{ P-a.s.} \quad (3.3)$$

By assumption, the set

$$\mathcal{V} := \{v \in \mathbf{L}_+^0(\mathcal{F}_\sigma) : \text{there is } \vartheta \in \Gamma \text{ with} \\ V_\sigma(\vartheta)(S') \leq v \text{ P-a.s. and } V_\tau(\vartheta)(S') \geq F(S') \text{ P-a.s.}\}$$

is nonempty. We claim that it is also closed under taking minima. Indeed, if $v^{(1)}, v^{(2)} \in \mathcal{V}$ and $\vartheta^{(1)}, \vartheta^{(2)} \in \Gamma$ are such that $V_\sigma(\vartheta^{(i)})(S') \leq v^{(i)}$ P-a.s. and $V_\tau(\vartheta^{(i)})(S') \geq F(S')$ P-a.s., then $\hat{\vartheta} := \vartheta^{(1)}\mathbf{1}_{\{v^{(1)} \leq v^{(2)}\}} + \vartheta^{(2)}\mathbf{1}_{\{v^{(1)} > v^{(2)}\}} \in \Gamma$, $V_\sigma(\hat{\vartheta})(S') \leq v^{(1)} \wedge v^{(2)}$ P-a.s. and $V_\tau(\hat{\vartheta})(S') \geq F(S')$ P-a.s. So there is a nonincreasing sequence $(v_n)_{n \in \mathbb{N}}$ in \mathcal{V} with $\lim_{n \rightarrow \infty} v_n = \text{ess inf } \mathcal{V} = \Pi_\sigma(F | \Gamma)(S')$ P-a.s. If we set $B_n := \{v_n \leq \Pi_\sigma(F | \Gamma)(S') + \delta\}$, $B_0 := \emptyset$ and $A_n := B_n \setminus B_{n-1}$, then $(A_n)_{n \in \mathbb{N}}$ is an \mathcal{F}_σ -measurable partition of Ω . Let $(\vartheta^{(n)})_{n \in \mathbb{N}} \in \Gamma$ be such that $V_\tau(\vartheta^{(n)})(S') \geq F(S')$ P-a.s. and $V_\sigma(\vartheta^{(n)})(S') \leq v_n$ P-a.s. for each $n \in \mathbb{N}$. Then $\vartheta := \sum_{n=1}^{\infty} \mathbf{1}_{A_n} \vartheta^{(n)}$ is in Γ and satisfies (3.3). \square

Recall that for $\sigma \in \mathcal{T}_{[0, T]}$ and a strategy cone Γ on $[[\sigma, T]]$, a strategy $\vartheta \in \Gamma$ is strongly maximal for Γ if there is no random variable $f \in \mathbf{L}_+^0(\mathcal{F}_T) \setminus \{0\}$ such that for all $\varepsilon > 0$, there is a strategy $\bar{\vartheta} \in \Gamma$ with $V_\sigma(\bar{\vartheta})(S) \leq V_\sigma(\vartheta)(S) + \varepsilon$ P-a.s. and $V_T(\bar{\vartheta})(S) \geq V_T(\vartheta)(S) + f$ P-a.s. This can now be reformulated more compactly: $\vartheta \in \Gamma$ is strongly maximal for Γ if and only if there is no nonzero contingent claim F at time T such that

$$\Pi_\sigma(V_T(\vartheta) + F | \Gamma)(S) \leq V_\sigma(\vartheta)(S) \text{ P-a.s.} \quad (3.4)$$

The next result shows that superreplication prices for undefaultable strategies are time-consistent. This exploits that the family of all \mathcal{U}_σ is time-consistent; see Section 1.1.

Proposition 3.8. *Let $\sigma_1 \leq \sigma_2 \leq \tau \in \mathcal{T}_{[0, T]}$ be stopping times and F a contingent claim at time τ with $\Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) < \infty$ P-a.s. Then*

$$\Pi_{\sigma_1}(F | \mathcal{U}_{\sigma_1}) = \Pi_{\sigma_1}(\Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) | \mathcal{U}_{\sigma_1}) \text{ P-a.s.} \quad (3.5)$$

Proof. Denote the left- and right-hand sides of (3.5) by L and R , respectively. For “ \leq ”, it suffices to show the inequality on the set $A := \{R(S) < \infty\} \in \mathcal{F}_{\sigma_1} \subseteq \mathcal{F}_{\sigma_2}$ for some $S \in \mathcal{S}$. By positive \mathcal{F}_{σ_i} -homogeneity, we may thus replace F by $FI_{\{R < \infty\}}$, or equivalently assume without loss of generality that $R < \infty$ P-a.s. Analogously, for proving “ \geq ”, we may assume without loss of generality that $L < \infty$ P-a.s.

“ \leq ”. Fix a numéraire strategy η . Take $\delta > 0$ and denote by C^i the contingent claim at time σ_i satisfying $C^i(S^{(\eta)}) = 1$, $i = 1, 2$. In other words, C^i is induced at time σ_i by the constant 1 with respect to $S^{(\eta)}$. By Proposition 3.7, there exist $\vartheta^{(1)} \in \mathcal{U}_{\sigma_1}$ satisfying

$$V_{\sigma_1}(\vartheta^{(1)}) \leq R + \delta C^1 \text{ P-a.s.} \quad \text{and} \quad V_{\sigma_2}(\vartheta^{(1)}) \geq \Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) \text{ P-a.s.} \quad (3.6)$$

and $\vartheta^{(2)} \in \mathcal{U}_{\sigma_2}$ satisfying

$$V_{\sigma_2}(\vartheta^{(2)}) \leq \Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) + \delta C^2 \text{ P-a.s.} \quad \text{and} \quad V_\tau(\vartheta^{(2)}) \geq F \text{ P-a.s.} \quad (3.7)$$

By the choice of C^2 , (3.6) and (3.7),

$$V_{\sigma_2}(\vartheta^{(1)} + \delta\eta)(S^{(\eta)}) = V_{\sigma_2}(\vartheta^{(1)})(S^{(\eta)}) + \delta C^2(S^{(\eta)}) \geq V_{\sigma_2}(\vartheta^{(2)})(S^{(\eta)}) \text{ P-a.s.} \quad (3.8)$$

Set

$$\vartheta := (\vartheta^{(1)} + \delta\eta)\mathbf{1}_{\llbracket\sigma_1, \sigma_2\rrbracket} + (\vartheta^{(2)} + V_{\sigma_2}(\vartheta^{(1)} + \delta\eta - \vartheta^{(2)})(S^{(\eta)})\eta)\mathbf{1}_{\llbracket\sigma_2, T\rrbracket}.$$

From (3.8) and the fact that $\vartheta^{(i)} \in \mathcal{U}_{\sigma_i}$ and $\eta \in \mathcal{U}_0 \subseteq \mathcal{U}_{\sigma_i}$, it is easy to check that $\vartheta \in \mathcal{U}_{\sigma_1}$. Moreover, the definition of ϑ gives by (3.7), (3.6) and (3.8) that

$$V_{\sigma_1}(\vartheta) \leq R + 2\delta C^1 \text{ P-a.s.} \quad \text{and} \quad V_T(\vartheta) \geq F \text{ P-a.s.}$$

Thus $\Pi_{\sigma_1}(F | \mathcal{U}_{\sigma_1}) \leq R + 2\delta C^1$ by the Definition 3.5 of superreplication prices, and letting $\delta \searrow 0$ yields the claim.

“ \geq ”. Fix $\delta > 0$ and a positive contingent claim C at time σ_1 . By Proposition 3.7, there exists $\vartheta \in \mathcal{U}_{\sigma_1}$ satisfying

$$V_{\sigma_1}(\vartheta) \leq L + \delta C \text{ P-a.s.} \quad \text{and} \quad V_T(\vartheta) \geq F \text{ P-a.s.}$$

So the definition of superreplication prices gives first $\Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) \leq V_{\sigma_2}(\vartheta)$ P-a.s. and then

$$R = \Pi_{\sigma_1}(\Pi_{\sigma_2}(F | \mathcal{U}_{\sigma_2}) | \mathcal{U}_{\sigma_1}) \leq V_{\sigma_1}(\vartheta) \leq L + \delta C \text{ P-a.s.}$$

The claim follows by letting $\delta \searrow 0$. \square

Proposition 3.8 says that our valuation by superreplication is consistent over time. This is well known in the classic setup; see for instance Föllmer and Schied [25, Example 11.2.4] for discrete time or Klöppel and Schweizer [51, Theorem 5.1] for continuous time. One very useful consequence in our framework is that for undefaultable strategies, maximality needs only to be tested from time 0, i.e., on $\llbracket 0, T \rrbracket$.

Corollary 3.9. *Let $\sigma_1 \leq \sigma_2 \in \mathcal{T}_{[0, T]}$ be stopping times. If $\vartheta \in \mathcal{U}_{\sigma_1}$ is strongly maximal for \mathcal{U}_{σ_1} , it is also strongly maximal for \mathcal{U}_{σ_2} . Hence any $\vartheta \in \mathcal{U}_0$ is strongly maximal for each \mathcal{U}_σ , $\sigma \in \mathcal{T}_{[0, T]}$, if and only if it is strongly maximal for \mathcal{U}_0 .*

An analogous statement holds for “strong” replaced by “weak”.

Proof. Suppose by way of contradiction that $\vartheta \in \mathcal{U}_{\sigma_1} \subseteq \mathcal{U}_{\sigma_2}$ fails to be strongly maximal for \mathcal{U}_{σ_2} . Then there exists a nonzero contingent claim F at time T with $\Pi_{\sigma_2}(V_T(\vartheta) + F | \mathcal{U}_{\sigma_2}) \leq V_{\sigma_2}(\vartheta) < \infty$ P-a.s. Proposition 3.8, monotonicity and the definition of superreplication prices give

$$\begin{aligned} \Pi_{\sigma_1}(V_T(\vartheta) + F | \mathcal{U}_{\sigma_1}) &= \Pi_{\sigma_1}(\Pi_{\sigma_2}(V_T(\vartheta) + F | \mathcal{U}_{\sigma_2}) | \mathcal{U}_{\sigma_1}) \\ &\leq \Pi_{\sigma_1}(V_{\sigma_2}(\vartheta) | \mathcal{U}_{\sigma_1}) \leq V_{\sigma_1}(\vartheta) \text{ P-a.s.,} \end{aligned}$$

and so ϑ fails to be strongly maximal for \mathcal{U}_{σ_1} , and we arrive at a contradiction. \square

3.3 Dual characterisation of dynamic efficiency

Recall that \mathcal{S} is a bubbly market by definition if it is not dynamically efficient, i.e., if some naive invest-and-keep strategy can be improved (approximately) by dynamic trading. We want to characterise this via dual objects and martingale properties, and so we first give a dual characterisation of dynamic efficiency.

In preparation, we need the following useful result already announced in Remark 1.12. It says that under dynamic viability, the notions of weak and strong maximality for the class \mathcal{U} of undefaultable strategies are equivalent. This is a generalisation of Proposition III.3.19 from an initial trading time 0 to a general stopping time $\sigma \in \mathcal{T}_{[0,T]}$.

Lemma 3.10. *Let \mathcal{S} be dynamically viable. Fix $\sigma \in \mathcal{T}_{[0,T]}$. Then $\vartheta \in \mathcal{U}_\sigma$ is weakly maximal for \mathcal{U}_σ if and only if it is strongly maximal for \mathcal{U}_σ .*

Proof. Strong clearly implies weak maximality; so we prove the “only if” part. Let $\vartheta \in \mathcal{U}_\sigma$ be weakly maximal.

In a first step, we show that for each $\bar{\vartheta} \in \mathcal{U}_\sigma$ with $V_T(\bar{\vartheta}) \geq V_T(\vartheta)$ P-a.s., we have $V(\bar{\vartheta}) \geq V(\vartheta)$ P-a.s. on $[[\sigma, T]]$, so that $\bar{\vartheta} - \vartheta \in \mathcal{U}_\sigma$. Indeed, if $\tau \in \mathcal{T}_{[\sigma, T]}$ is a stopping time such that $A := \{V_\tau(\bar{\vartheta}) < V_\tau(\vartheta)\}$ has $\mathbb{P}[A] > 0$, we take a numéraire strategy η and set

$$\hat{\vartheta} := \vartheta \mathbf{1}_{[[\sigma, \tau]]} + \left(\mathbf{1}_{A^c} \vartheta + \mathbf{1}_A (\bar{\vartheta} + V_\tau(\vartheta - \bar{\vartheta})(S^{(\eta)})) \right) \mathbf{1}_{] \tau, T]}.$$

Then $\hat{\vartheta} \in \mathcal{U}_\sigma$, we have $V_\sigma(\hat{\vartheta}) = V_\sigma(\vartheta)$ P-a.s., and using that $V_T(\bar{\vartheta}) \geq V_T(\vartheta)$ P-a.s. gives

$$\begin{aligned} V_T(\hat{\vartheta}) &= \mathbf{1}_{A^c} V_T(\vartheta) + \mathbf{1}_A (V_T(\bar{\vartheta}) + V_\tau(\vartheta - \bar{\vartheta})(S^{(\eta)})) V_T(\eta) \\ &\geq \mathbf{1}_{A^c} V_T(\vartheta) + \mathbf{1}_A (V_T(\vartheta) + V_\tau(\vartheta - \bar{\vartheta})(S^{(\eta)})) V_T(\eta) \\ &= V_T(\vartheta) + \mathbf{1}_A V_\tau(\vartheta - \bar{\vartheta})(S^{(\eta)}) V_T(\eta) \text{ P-a.s.} \end{aligned}$$

Since $\mathbb{P}[A] > 0$ and $V_T(\hat{\vartheta}) > V_T(\vartheta)$ on A , this shows that ϑ fails to be weakly maximal, and we arrive at a contradiction.

We proceed to establish that ϑ is strongly maximal for \mathcal{U}_σ . Seeking a contradiction, suppose there exists a nonzero contingent claim F at time T satisfying $\Pi_\sigma(V_T(\vartheta) + F | \mathcal{U}_\sigma) \leq V_\sigma(\vartheta)$ P-a.s. Take $\delta > 0$ and a positive contingent claim C at time σ . By Proposition 3.7, there is $\bar{\vartheta} \in \mathcal{U}_\sigma$ with $V_T(\bar{\vartheta}) \geq V_T(\vartheta) + F$ P-a.s. and

$$V_\sigma(\bar{\vartheta}) \leq \Pi_\sigma(V_T(\vartheta) + F | \mathcal{U}_\sigma) + \delta C \leq V_\sigma(\vartheta) + \delta C \text{ P-a.s.}$$

By the first step, $\vartheta' := \vartheta - \bar{\vartheta}$ is in \mathcal{U}_σ . Moreover, we have $V_T(\vartheta') \geq F$ P-a.s. and $V_\sigma(\vartheta') \leq \delta C$ P-a.s. so that we get $\Pi_\sigma(F | \mathcal{U}_\sigma) \leq \delta C$ P-a.s. Letting $\delta \searrow 0$ gives $\Pi_\sigma(F | \mathcal{U}_\sigma) = 0$ P-a.s., and so 0 is not strongly maximal for \mathcal{U}_σ , in contradiction to dynamic viability of \mathcal{S} . \square

Our first characterisation of dynamic efficiency now follows by combining two results from Chapter VI. Recall that due to (1.1), our \mathcal{S} is a numéraire market.

Theorem 3.11. *The following are equivalent:*

- (a) \mathcal{S} is dynamically efficient.
- (b) For each bounded numéraire strategy η there exists $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)}$ is a (true) \mathbb{Q} -martingale.
- (c) There exists a pair (η, \mathbb{Q}) , where η is a bounded numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , such that $S^{(\eta)}$ is a (true) \mathbb{Q} -martingale.
- (d) There exists a representative $\bar{S} \in \mathcal{S}$ and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that \bar{S} is a (true) \mathbb{Q} -martingale.

Proof. Since S is nonnegative, both the market portfolio $\eta^S = (1, \dots, 1)$ and the corresponding representative $S^{(\eta^S)} = S / (\sum_{i=1}^N S^i)$ are bounded, as required for Corollaries VI.2.4 and VI.2.5. Now if we have (a), then η^S is strongly maximal for \mathcal{U}_0 and (b) follows from Corollaries VI.2.5, (c) \Rightarrow (a), and VI.2.4, (a) \Rightarrow (c). It is clear that (b) implies (c), and that (c) implies (d).

Suppose we have (d). Fix any stopping time $\sigma \in \mathcal{T}_{[0, T]}$ and any $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$. As \bar{S} is a \mathbb{Q} -martingale on $[[\sigma, T]]$, so is $V(\vartheta)(\bar{S})$. If we take any $\bar{\vartheta} \in \mathcal{U}_\sigma$, then $V(\bar{\vartheta})(\bar{S})$ is a nonnegative stochastic integral of a \mathbb{Q} -martingale and hence a \mathbb{Q} -supermartingale. So if $V_T(\bar{\vartheta}) \geq V_T(\vartheta)$, we get

$$V_\sigma(\vartheta)(\bar{S}) = \mathbb{E}_{\mathbb{Q}}[V_T(\vartheta)(\bar{S}) | \mathcal{F}_\sigma] \leq \mathbb{E}_{\mathbb{Q}}[V_T(\bar{\vartheta})(\bar{S}) | \mathcal{F}_\sigma] \leq V_\sigma(\bar{\vartheta})(\bar{S}),$$

and this shows that ϑ is weakly maximal for \mathcal{U}_σ . But (d) implies by Theorem 3.1 also that \mathcal{S} is dynamically viable, and so weak maximality in \mathcal{U}_σ is equivalent to strong maximality in \mathcal{U}_σ , by Lemma 3.10. So we get (a) and the proof is complete. \square

With the help of the above dual characterisation, we can give some equivalent primal descriptions of dynamically efficient markets.

Corollary 3.12. *The following are equivalent:*

- (a) \mathcal{S} is dynamically efficient.
- (b) The market portfolio $\eta^S = (1, \dots, 1)$ (buy and hold one unit of each asset) is strongly maximal for \mathcal{U}_0 .
- (c) For each $i = 1, \dots, N$, the strategy $e_i := (0, \dots, 0, 1, 0, \dots, 0)$, where the 1 is at position i , (buy and hold one unit of asset i) is strongly maximal for \mathcal{U}_0 .
- (d) For each stopping time $\sigma \in \mathcal{T}_{[0, T]}$, each $\vartheta \in \mathbf{b}\mathcal{U}_\sigma$ is strongly maximal for \mathcal{U}_σ .

Proof. Because $\mathbf{h}\mathcal{U} \subseteq \mathbf{b}\mathcal{U}$, it is clear that (d) implies (a). Next, (a) trivially implies (b), which is in turn equivalent to (c) by Corollary VI.2.5. Finally, to establish (b) \Rightarrow (d), we first show that $V(\vartheta)(S^{(\eta^S)})$ is a true martingale on $[[\sigma, T]]$ for each $\vartheta \in \mathbf{b}\mathcal{U}_\sigma$ by arguing as in Lemma VI.2.2 and then proceed as in the second part of the proof of Theorem 3.11, with $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ there replaced by $\vartheta \in \mathbf{b}\mathcal{U}_\sigma$. \square

The following characterisation is immediate from Corollary 3.12. It clarifies precisely what makes a market \mathcal{S} a bubbly market.

Corollary 3.13. *\mathcal{S} is a bubbly market if and only if there exists at least one index $i \in \{1, \dots, N\}$ such that the buy-and-hold strategy e_i of asset i is not strongly maximal for \mathcal{U}_0 .*

In other words, \mathcal{S} is a bubbly market if and only if it contains an asset which is bad (or stupid) enough that its evolution can be beaten by dynamic trading in the whole market, without running into debt.

If we specialise Corollary 3.13 to the classic setup with $S = (1, X)$, we see that a bubbly market in our sense can arise in two ways. It may happen that one of the *risky* assets in X can be dominated by dynamic trading in the other risky assets and the bank account; then that asset might be called a bubble. But it may also happen that the *bank account* itself can be dominated by trading in the other (risky assets)—and this would mean that our initial choice of the bank account as numéraire was rash, because it does not behave well in comparison to the other assets. Put differently, discounting with such a bank account is maybe not a good starting point from an economic perspective.

3.4 Dual characterisation of static efficiency

By definition, a nontrivial bubbly market must be statically efficient. So it is also of interest to characterise static efficiency by dual objects and martingale properties. Our first result in this section gives such a characterisation, but it is for most concrete models too complicated to be used in practice. We therefore provide a second, more tractable description under the additional assumption that \mathcal{S} is statically viable. This combination is almost tailor-made for our applications because a nontrivial bubbly market must be dynamically (and hence a fortiori statically) viable.

We first introduce some notation.

Definition 3.14. Let $\sigma \in \mathcal{T}_{[0, T]}$ be a stopping time and $\eta \in \mathbf{h}\mathcal{U}_\sigma$ an (invest-and-keep) numéraire strategy. A *one-step equivalent martingale measure (EMM)* for $S^{(\eta)}$ on $\{\sigma, T\}$ is a probability measure $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T with $S_T^{(\eta)} \in L^1(\mathbb{Q})$ and $\mathbb{E}_{\mathbb{Q}}[S_T^{(\eta)} | \mathcal{F}_\sigma] = S_\sigma^{(\eta)}$ \mathbb{Q} -a.s.

In other words, \mathbb{Q} is simply an equivalent martingale measure in the classic sense for the one-period model with $\mathcal{G}_0 = \mathcal{F}_\sigma, \mathcal{G}_1 = \mathcal{F}_T$ and $X_0 = S_\sigma^{(\eta)}, X_1 = S_T^{(\eta)}$. We denote below by $\mathcal{L}(S_T^{(\eta)} | \mathcal{F}_\sigma)$ a regular conditional distribution of $S_T^{(\eta)}$ given \mathcal{F}_σ , and by $\text{ri conv supp } \mu$, for a probability measure μ on \mathbb{R}^N , the relative interior of the convex hull of the topological support of μ .

With this, we can formulate a dual characterisation of statically efficient markets.

Theorem 3.15. *The following are equivalent:*

- (a) \mathcal{S} is statically efficient.

- (b) For each stopping time $\sigma \in \mathcal{T}_{[0,T]}$, there exists a pair (η, \mathbb{Q}) such that η is a numéraire strategy in $\mathbf{h}\mathcal{U}_\sigma$ and \mathbb{Q} is a one-step EMM for $S^{(\eta)}$ on $\{\sigma, T\}$.
- (c) For each stopping time $\sigma \in \mathcal{T}_{[0,T]}$, there exists a numéraire strategy η in $\mathbf{h}\mathcal{U}_\sigma$ such that $S_\sigma^{(\eta)} \in \text{ri conv supp } \mathcal{L}(S_T^{(\eta)} | \mathcal{F}_\sigma)$ \mathbb{P} -a.s.

Proof. Lemma 6.2 below shows that (a) is equivalent to saying that each $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ is strongly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$, for each $\sigma \in \mathcal{T}_{[0,T]}$. The latter is then shown to be equivalent to (b) as well as to (c) in Lemma 6.1 below. \square

Theorem 3.15 can be interpreted as follows. Static efficiency means that for every starting date $\sigma \in \mathcal{T}_{[0,T]}$, the one-period model with trading dates σ and T must be arbitrage-free. (Note that while the starting time can vary, the end date is always fixed at T .) From a theoretical point of view, this is very illuminating and reflects in a natural way the structure of the static efficiency condition. But since we must verify part (b) or (c) for all stopping times σ , it seems almost hopeless to use this result in practice.

Our second characterisation below shows that if we already know that our market is statically viable, it is enough to look only at *deterministic* times. This is a significant improvement if we consider for example markets with diffusion dynamics. There one can first use Theorem 3.1 to check for dynamic viability, because this implies static viability. For checking Theorem 3.15, in particular its condition (c), one can then exploit the many known results about transition densities for diffusion processes—and since these results are only for deterministic times, it is important that we do not need to check the behaviour at stopping times. An explicit example to illustrate this procedure can be found in Example 4.2.

Theorem 3.16. *Suppose \mathcal{S} is statically viable. Then the following are equivalent:*

- (a) \mathcal{S} is statically efficient.
- (b) For each deterministic time $s \in [0, T]$, there exists a pair (η, \mathbb{Q}) such that η is a numéraire strategy in $\mathbf{h}\mathcal{U}_s(\mathcal{S})$ and \mathbb{Q} is a one-step EMM for $S^{(\eta)}$ on $\{s, T\}$.
- (c) For each deterministic time $s \in [0, T]$, there exists a numéraire strategy η in $\mathbf{h}\mathcal{U}_s(\mathcal{S})$ such that $S_s^{(\eta)} \in \text{ri conv supp } \mathcal{L}(S_T^{(\eta)} | \mathcal{F}_s)$ \mathbb{P} -a.s.

Proof. The assumption of static viability of \mathcal{S} allows us to use Lemma 6.2 below and obtain that (a) is equivalent to weak maximality of 0 for $\mathbf{h}L_s^{\text{sf}}$, for each $s \in [0, T]$. By Lemma 6.1 below, with $\sigma := s$, the latter is equivalent to (b) as well as to (c). \square

In the classic setup, it is well known for finite discrete time and in particular for one-period models that for absence-of-arbitrage and valuation questions, it does not matter if one uses all self-financing strategies or only those with nonnegative wealth; see for instance Elliott and Kopp [18, Section 2.2] or Lambertson and Lapeyre [53, Lemma 1.2.7]. The next result is a generalisation of this (basically)

one-period property to invest-and-keep strategies in continuous time. The key ingredient is static efficiency which is equivalent to absence of arbitrage in a one-period model; see Lemma 1.18 and note that dynamic and static efficiency coincide for one-period models.

Lemma 3.17. *Suppose \mathcal{S} is statically efficient and let $\sigma \in \mathcal{T}_{[0,T]}$ be a stopping time.*

- (a) *If $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ satisfies $V_T(\vartheta) \geq 0$ P-a.s., then $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$: An invest-and-keep strategy with nonnegative final wealth has nonnegative wealth over its entire lifetime.*
- (b) *If F is a contingent claim at time T , then*

$$\Pi_\sigma(F | \mathbf{h}\mathcal{U}_\sigma) = \Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) \text{ P-a.s.} \quad (3.9)$$

If moreover $\Pi_\sigma(F | \mathbf{h}\mathcal{U}_\sigma) < \infty$ P-a.s., there exists a strategy $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ with

$$V_T(\vartheta) \geq F \text{ P-a.s.} \quad \text{and} \quad V_\sigma(\vartheta) = \Pi_\sigma(F | \mathbf{h}\mathcal{U}_\sigma) \text{ P-a.s.}$$

Valuation with undefaultable or with arbitrary self-financing invest-and-keep strategies yields the same result; and if the result from superreplication is finite, the essential infimum from the definition is attained as a minimum.

Proof. (a) is shown in Lemma 6.2 below. For (b), the inequality “ \geq ” in (3.9) is clear since $\mathbf{h}\mathcal{U}_\sigma \subseteq \mathbf{h}L_\sigma^{\text{sf}}$. For “ \leq ”, we may assume without loss of generality that the right-hand side is finite P-a.s.; this is argued as in the proof of Proposition 3.8 via positive \mathcal{F}_σ -homogeneity. From static efficiency, Theorem 3.15 gives the existence of a one-step EMM for $S^{(\eta)}$ on $\{\sigma, T\}$, for some numéraire strategy $\eta \in \mathbf{h}\mathcal{U}_\sigma \subseteq \mathbf{h}L_\sigma^{\text{sf}}$, and so Lemma 6.1 below yields the existence of some $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ with

$$V_T(\vartheta) \geq F \geq 0 \text{ P-a.s.} \quad \text{and} \quad V_\sigma(\vartheta) = \Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) \text{ P-a.s.}$$

But now $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ by part (a), and so $\Pi_\sigma(F | \mathbf{h}\mathcal{U}_\sigma) \leq V_\sigma(\vartheta)$ P-a.s. yields “ \leq ” in (3.9). \square

3.5 No dominance

It is folklore in mathematical finance that for simple risk-neutral valuation results, one needs something extra in addition to absence of arbitrage. Like many important insights, this can be traced back to work of R. Merton who introduced for this purpose the notion of “no dominance”. In Merton’s words [62], “security (portfolio) A is *dominant* over security (portfolio) B , if on some known date in the future, the return on A will exceed the return on B for some possible states of the world, and will be at least as large as on B , in all possible states of the world”, and “a necessary condition for a rational option pricing theory is that the option be priced such that it is neither a dominant nor a dominated security”.

The above formulation is intuitive, but not very precise. It does not tell us exactly what a “security” or a “portfolio” is, and it also does not tell us exactly

what “the return” is. Subsequent papers have therefore developed different mathematical formulations for the idea; the key difference lies precisely in the above two terms.

The works of Protter and co-authors [40, 41, 67] incorporate “return” by the assumption that each financial product (including basic assets and dynamic trading strategies) has a market price at each time. It is not explained in detail where market values come from; results are obtained by imposing certain structural assumptions on market prices, including “no dominance”. In contrast, Jarrow and Larsson [39] only talk about basic assets and compute the “return” from the value processes of self-financing strategies. This is more specific than the approach in [40, 41, 67], but it also gives in our view potentially sharper results with weaker assumptions on the underlying market. In particular, one can try to impose “no dominance” only on basic assets and then try to deduce analogous properties for suitable valuations applied to complex assets, portfolios or derivatives. We therefore follow [39] in spirit when we introduce our numéraire-independent versions of no dominance.

Definition 3.18. The market \mathcal{S} is said to satisfy

- *static no dominance* (static ND) if the market portfolio $\eta^{\mathcal{S}} = (1, \dots, 1)$ is weakly maximal for $\mathbf{h}\mathcal{U}_{\sigma}$, for each $\sigma \in \mathcal{T}_{[0,T]}$.
- *dynamic no dominance* (dynamic ND) if the market portfolio $\eta^{\mathcal{S}}$ is weakly maximal for \mathcal{U}_{σ} , for each $\sigma \in \mathcal{T}_{[0,T]}$.

Due to Corollary 3.9, dynamic ND is equivalent to requiring that the market portfolio $\eta^{\mathcal{S}}$ is weakly maximal for \mathcal{U}_0 . Moreover, one can show that $\eta^{\mathcal{S}}$ is weakly maximal for \mathcal{U}_0 if and only if for $i = 1, \dots, N$, the buy-and-hold strategies $e_i = (0, \dots, 0, 1, 0, \dots, 0)$ of each basic asset i are weakly maximal for \mathcal{U}_0 . In fact, the “only if” part is clear since any improvement of an e_i will also improve $\eta^{\mathcal{S}}$, and the “if” part follows from Corollary III.3.10, where it is shown that the weakly maximal strategies in \mathcal{U}_0 form a convex cone. This shows that our definition of dynamic ND is very close in spirit to the concept of no dominance used in [39]. On the other hand, the concept of static ND seems to be new. It is more delicate to analyse; we mention for example that static ND is *not* equivalent to weak maximality of $\eta^{\mathcal{S}}$ (or of e_i , for $i = 1, \dots, N$) for $\mathbf{h}\mathcal{U}_0$. This can be easily seen if we take Example 2.5 and modify it slightly so that the two possible values of S_1^2 at time 1 are no longer 2 and 1, but 2 and $3/2$. We leave the details to the reader.

Our next result connects the notions introduced so far. It shows that no dominance is precisely the extra ingredient that distinguishes efficiency from viability.

Proposition 3.19. *The market \mathcal{S} is statically/dynamically efficient if and only if it is statically/dynamically viable and satisfies static/dynamic ND.*

Proof. Static/dynamic efficiency trivially implies static/dynamic viability and yields that $\eta^{\mathcal{S}}$ is strongly (and a fortiori weakly) maximal for $\mathbf{h}\mathcal{U}_{\sigma}/\mathcal{U}_{\sigma}$, for each $\sigma \in \mathcal{T}_{[0,T]}$.

For the converse direction, we first consider the static case. By Lemma 6.2, it suffices to show that under static ND, 0 is weakly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$, for each $\sigma \in \mathcal{T}_{[0,T]}$. Suppose to the contrary that there are $\sigma \in \mathcal{T}_{[0,T]}$ and $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ with

$$V_\sigma(\vartheta) \leq 0 \text{ P-a.s.}, \quad V_T(\vartheta) \geq 0 \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[V_T(\vartheta) > 0] > 0.$$

Then $c_\sigma := 1/(1 + \sum_{i=1}^N |\vartheta_\sigma^i|)$ is in $\mathbf{L}_{++}^0(\mathcal{F}_\sigma)$, and $\bar{\vartheta} := c_\sigma \vartheta + \eta^S$ is in $\mathbf{h}\mathcal{U}_\sigma$ and satisfies

$$V_\sigma(\bar{\vartheta}) \leq V_\sigma(\eta^S) \text{ P-a.s.}, \quad V_T(\bar{\vartheta}) \geq V_T(\eta^S) \text{ P-a.s.} \quad \text{and} \quad \mathbb{P}[V_T(\bar{\vartheta}) > V_T(\eta^S)] > 0.$$

Thus η^S fails to be weakly maximal for $\mathbf{h}\mathcal{U}_\sigma$, in contradiction to static ND.

For the dynamic case, we know from Lemma 3.10 that under dynamic viability, weak is equivalent to strong maximality of η^S for \mathcal{U}_0 . So dynamic efficiency follows from Corollary 3.12. \square

Remark 3.20. If we look a bit more carefully at the proof of Proposition 3.19, we see that we have actually proved that static efficiency and static ND are equivalent.

One of the main results on no dominance in the classic setup is that it is the extra strengthening of “absence of arbitrage” that is required to obtain the existence not only of an equivalent local (or σ -)martingale measure, but of a true martingale measure; see Theorem 3.2 of [39]. Our next result shows that this connection also holds in our numéraire-independent framework.

Corollary 3.21. *The following are equivalent:*

- (a) \mathcal{S} satisfies NINA and dynamic ND.
- (b) There exists a pair (η, \mathbb{Q}) , where η is a bounded numéraire strategy and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T is such that the $V(\eta)$ -discounted price process $S^{(\eta)} = \frac{S}{V(\eta)(S)}$ is a \mathbb{Q} -martingale.

Proof. By Theorem 3.1, NINA or strong maximality of 0 in \mathcal{U}_0 is equivalent to dynamic viability of \mathcal{S} . Together with dynamic ND, this is by Proposition 3.19 equivalent to dynamic efficiency of \mathcal{S} , and this in turn is equivalent to (b) by Theorem 3.11. \square

3.6 Bubbly markets and strict local martingales

In this section, we derive the promised connections between nontrivial bubbly markets and strict local martingales. A large part of the literature on bubbles in financial markets starts, in the classic setup of Example 1.2, with the *assumption* that the discounted price process X is a strict local martingale (sometimes under \mathbb{P} itself, sometimes under a chosen risk-neutral or valuation measure \mathbb{Q}). In distinct contrast, we show here that our definition of a nontrivial bubbly market leads to the *conclusion* that we must have strict local martingale properties. Moreover, this result is robust in the sense that we have it simultaneously under

all possible valuation measures \mathbb{Q} —it cannot happen that we “see” a bubble under one measure and no bubble under another. In other words, our definition of a nontrivial bubbly market does not depend on the a priori choice of a valuation or risk-neutral or martingale measure \mathbb{Q} . This is in marked contrast to the approach of Protter et al. [40, 41, 67]; see also Section 5. To obtain these connections, we combine our results so far with a result from Chapter VI.

Theorem 3.22. *Suppose \mathcal{S} is a nontrivial bubbly market. Then*

- (a) *There exist a numéraire strategy η and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $S^{(\eta)}$ is a local \mathbb{Q} -martingale; and for any such pair (η, \mathbb{Q}) , the process $S^{(\eta)}$ is a strict local \mathbb{Q} -martingale.*
- (b) *There exists $\bar{S} \in \mathcal{S}$ and $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that \bar{S} is a local \mathbb{Q} -martingale; and for any such pair (\bar{S}, \mathbb{Q}) , the process \bar{S} is a strict local \mathbb{Q} -martingale.*

Proof. Because \mathcal{S} is dynamically viable, the existence of (η, \mathbb{Q}) as in (a) and (\bar{S}, \mathbb{Q}) as in (b) follows from Theorem 3.1. Because \mathcal{S} is not dynamically efficient, we know from Corollary 3.12 that the market portfolio $\eta^{\mathcal{S}} = (1, \dots, 1)$ is not strongly maximal for \mathcal{U}_0 , and so Corollary VI.2.5 implies that there cannot be a pair (\bar{S}, \mathbb{Q}) with $\bar{S} \in \mathcal{S}$, $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T and \bar{S} a true \mathbb{Q} -martingale. This gives the second part of (b), and with $\bar{S} := S^{(\eta)}$ also the second part of (a). \square

Note that we are in a market with $N > 1$ traded primary assets. Saying that a representative \bar{S} , which is an \mathbb{R}^N -valued process, is a strict local \mathbb{Q} -martingale means that there is at least one coordinate \bar{S}^i with $i \in \{1, \dots, N\}$ which has the local, but not the true \mathbb{Q} -martingale property. This reflects Corollary 3.12 which says that the market fails to be dynamically efficient if and only if at least one of the buy-and-hold strategies $e_i = (0, \dots, 0, 1, 0, \dots, 0)$ is not strongly maximal for \mathcal{U}_0 .

Remark 3.23. It follows immediately from Theorems 3.1 and 3.11 that (a) or (b) in Theorem 3.22 are equivalent to \mathcal{S} being a dynamically viable bubbly market. What is still missing for a *nontrivial* bubbly market, however, is static efficiency. So not every market satisfying (a) or (b) in Theorem 3.22 is a nontrivial bubbly market.

4 Further examples

Our next goal is to relate viability and efficiency. We know from the definitions that efficiency implies viability, and we have shown in Lemmas 1.18 and 1.19 that the converse is also true for a market in finite discrete time.

The next example illustrates how the picture changes if we go to continuous time. We first keep the model simple but abstract, so that one easily sees which basic properties drive its behaviour, and then give concrete examples of such models.

Example 4.1 (*Dynamic viability and static/dynamic efficiency*). Consider the market \mathcal{S} generated by $S = (S^1, S^2) = (X, 1)$, where X is a strict local \mathbb{P} -martingale. (So in these units, the bank account is asset 2, for a change.) We also suppose that \mathcal{S} is complete (which means that X has the predictable representation property in the filtration $(\mathcal{F}_t)_{t \in [0, T]}$ we work with). We claim that \mathcal{S} is dynamically viable, but not dynamically efficient.

To see this, note first that dynamic viability follows directly from Theorem 3.1 when we choose for η the buy-and-hold strategy $e_2 = (0, 1)$ of the second asset. Next, completeness and continuity of X imply that the only \mathbb{P} -martingales strongly \mathbb{P} -orthogonal to X are constants, since \mathcal{F}_0 is \mathbb{P} -trivial. Thus the density process Z of any ELMM \mathbb{Q} for X must be constant, hence 1, so that $\mathbb{Q} \equiv \mathbb{P}$. So because X is a strict local \mathbb{P} -martingale, there cannot be any $\mathbb{Q} \approx \mathbb{P}$ which makes X a true \mathbb{Q} -martingale, and since this means that (c) in Theorem 3.11 with $\eta = e_2$ fails, we conclude that \mathcal{S} is not dynamically efficient. In particular, \mathcal{S} is a bubbly market.

For \mathcal{S} to be a nontrivial bubbly market, since we already know that it is dynamically viable, we only need to show that it is addition statically efficient. By Theorem 3.16, this is the case if and only if

$$X_s \in \text{ri conv supp } \mathcal{L}(X_T | \mathcal{F}_s) \text{ } \mathbb{P}\text{-a.s.} \quad \text{for each } s \in [0, T), \quad (4.1)$$

where $\text{supp } \mathcal{L}(S_T^1 | \mathcal{F}_s)$ is the (ω -dependent) support of the regular conditional distribution of X_T given \mathcal{F}_s , conv denotes the convex hull, and ri the relative interior.

Example 4.2 (*Complete nontrivial bubbly markets*). For a concrete example of a strict local \mathbb{P} -martingale which has the predictable representation property, we can go back to Example 2.3 where $S = (1, X)$ and X is under \mathbb{P} a BES^3 process. The representative $S^{(e_2)} = (1/X, 1) := (Y, 1)$ is then of the form that we want. More generally, we could assume that Y is a *constant elasticity of variance (CEV)* process, i.e., satisfies the SDE

$$dY_t = \sigma |Y_t|^\beta dW_t^\mathbb{P}, \quad Y_0 = y_0 > 0, \quad (4.2)$$

with $\sigma > 0$ and $\beta > 1$. It is well known that the SDE (4.2) has a unique strong solution Y which is a positive continuous strict local \mathbb{P} -martingale; see [55, Section 9.8] for a detailed discussion of the CEV model.¹ As in Example 2.3 and Remark 2.4, one can also see that the CEV process has the predictable representation property for its own filtration (which can equivalently be generated by $W^\mathbb{P}$). Note that the BES^3 process in (2.2) is obtained as the special case where $\beta = 2$ and $\sigma = 1$.

If we want to check for the above examples the relative interior condition (4.1), we can use the transition densities $f(T, y; s, x)$ for the conditional distribution at

¹It seems to be difficult to find an exact single reference, where it is *rigorously* shown that the SDE (4.2) has a unique strong solution and that this solution is a strict local \mathbb{P} -martingale. But of course this follows easily from the general theory of one-dimensional SDEs and the explicit transition density computed by Emanuel and MacBeth [19, Equation (7)].

time T , given that we are in x at time s . The explicit formula for the CEV model can be found for instance in Emanuel and MacBeth [19, Equation (7)]. One can see from that expression that $\mathcal{L}(Y_T | Y_s)$ has all of $(0, \infty)$ as its support; so (4.1) is clearly satisfied and hence \mathcal{S} is statically efficient.

We note in passing that Emanuel and MacBeth [19] also say implicitly that Y is a strict local \mathbb{P} -martingale—they mention that computing the mean of the transition density yields $\mathbb{E}_{\mathbb{P}}[Y_T | Y_t] \neq Y_t$.

In summary, we see from this example that both the CEV process in (4.2) and the BES^3 process in (2.2) lead to a nontrivial bubbly market.

The preceding example is set in a complete financial markets. The next example provides a situation where we have a genuinely incomplete market, and we comment after the example why this is new and of interest.

Example 4.3 (*An incomplete nontrivial bubbly market*). Consider two independent \mathbb{P} -Brownian motions $W^{\mathbb{P}} = (W_t^{\mathbb{P}})_{t \in [0, T]}$ and $B = (B_t)_{t \in [0, T]}$ with respect to a given filtration $(\mathcal{F}_t)_{t \in [0, T]}$; the latter need not be generated by $W^{\mathbb{P}}$ and B . Let \mathcal{S} be the market generated by $S = (S^1, S^2) = (X_t, 1)_{t \in [0, T]}$, where X satisfies the SDE

$$dX_t = V_t |X_t|^{\beta} dW_t^{\mathbb{P}}, \quad X_0 = x_0 > 0. \quad (4.3)$$

Here $\beta > 1$ is a constant and the volatility $V = (V_t)_{t \in [0, T]}$ is stochastic and satisfies the SDE

$$dV_t = \alpha(V_t - \bar{\sigma})(V_t - \underline{\sigma}) dB_t, \quad V_0 = v_0 \in (\underline{\sigma}, \bar{\sigma}), \quad (4.4)$$

for some constants $\alpha > 0$ and $\bar{\sigma} > \underline{\sigma} > 0$. The SDE (4.3) can be interpreted as a CEV model (see Example 4.2) with stochastic volatility V and elasticity of variance $\beta > 1$. It is not difficult to check that (4.4) has a unique strong solution satisfying $\underline{\sigma} < V < \bar{\sigma}$ \mathbb{P} -a.s.; see e.g. Rady [68, Section 3]. We point out that the exact form of the volatility process V is not important for the argument that follows; we only use that V is a continuous (\mathcal{F}_t) -adapted strong Markov process that is uniformly bounded from above and below by positive constants.

We proceed to argue that (4.3) has a unique strong solution. More precisely, we show a more general result, which will be used for other purposes as well: Let $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T be such that $W^{\mathbb{P}}$ is a \mathbb{Q} -Brownian motion (on $[0, T]$). Then (4.3) has a unique strong solution X satisfying $\mathbb{E}_{\mathbb{Q}}[X_t] < x_0$, $t \in (0, T]$, i.e., X is a strict local \mathbb{Q} -martingale. Moreover, there exists $\varepsilon \in (0, T]$ which depends on x_0 , but not on v_0 , such that

$$\mathbb{E}_{\mathbb{Q}}[X_{\varepsilon}] > \frac{x_0}{2}. \quad (4.5)$$

Let us argue the above claims, using Protter [66, Chapter V] as reference. Uniqueness of a solution to (4.3) under \mathbb{Q} (up to a possible explosion time) holds because the function $f : \mathbb{R}_+ \times \Omega \times \mathbb{R}_+ \rightarrow \mathbb{R}$ given by $f(t, \omega, x) := V_t(\omega)|x|^{\beta}$ is uniformly in t locally random Lipschitz in x , i.e., for each $n \in \mathbb{N}$, there exists a finite random variable K_n such that $\sup_{t \in [0, T]} |f(t, \omega, x) - f(t, \omega, y)| \leq K_n(\omega)|x - y|$ for all $x, y \in [0, n]$. To establish existence of a solution to (4.3) under \mathbb{Q} and to

prove the remaining assertions, we use a time-change argument reducing (4.3) to the SDE of the standard CEV model. To simplify the presentation, we may assume that after possibly enlarging the original probability space, there exists a \mathbb{Q} -Brownian motion $W^{\mathbb{Q}} = (W_t^{\mathbb{Q}})_{t \geq 0}$ such that $(W_t^{\mathbb{Q}})_{t \in [0, T]} = W^{\mathbb{P}}$. Define the filtration $(\tilde{\mathcal{F}}_t)_{t \geq 0}$ by $\tilde{\mathcal{F}}_t = \mathcal{F}_{t \wedge T} \vee \sigma(W_s^{\mathbb{Q}}; s \leq t) \vee \mathcal{N}$, where \mathcal{N} are the \mathbb{Q} -nullsets in $\mathcal{F}_T \vee \sigma(W_s^{\mathbb{Q}}; s \geq 0)$. Then the process $\tilde{V} = (\tilde{V}_t)_{t \geq 0}$ defined by $\tilde{V}_t = V_{t \wedge T}$ is a continuous $(\tilde{\mathcal{F}}_t)$ -adapted process which satisfies $(\tilde{V}_t)_{t \in [0, T]} = V$ and takes values in $(\underline{\sigma}, \bar{\sigma})$ \mathbb{Q} -a.s. We are going to construct a strong solution on $[0, \infty)$ of the SDE

$$d\tilde{X}_t = \tilde{V}_t |\tilde{X}_t|^\beta dW_t^{\mathbb{Q}}, \quad \tilde{X}_0 = x_0 > 0, \quad (4.6)$$

and it is clear that $X = (X_t)_{t \in [0, T]}$ defined by $X_t = \tilde{X}_t$ for $t \in [0, T]$ is then a strong solution to (4.3).

Define the process $\tilde{M} = (\tilde{M}_t)_{t \geq 0}$ by $\tilde{M}_t = \int_0^t \tilde{V}_s dW_s^{\mathbb{Q}}$ and the process $(\Lambda_t)_{t \geq 0}$ by $\Lambda_t = \int_0^t |\tilde{V}_s|^2 ds$. Then \tilde{M} is under \mathbb{Q} a continuous local $(\tilde{\mathcal{F}}_t)$ -martingale null at 0, and Λ has \mathbb{Q} -a.s. continuous trajectories, is null at 0, strictly increasing, and satisfies \mathbb{Q} -a.s.

$$\bar{\sigma}^2 t > \Lambda_t > \underline{\sigma}^2 t \quad \text{for } t \geq 0. \quad (4.7)$$

Define the process $\tau = (\tau_t)_{t \geq 0}$ by $\tau_t = \inf\{s \geq 0 : \Lambda_s \geq t\}$. Then τ is an increasing continuous time change for the filtration $(\tilde{\mathcal{F}}_t)_{t \geq 0}$. Define the time-changed filtration $(\hat{\mathcal{F}}_t)_{t \geq 0}$ by $\hat{\mathcal{F}}_t := \tilde{\mathcal{F}}_{\tau_t}$ and the time-changed process $\hat{W} = (\hat{W}_t)_{t \geq 0}$ by $\hat{W}_t := \tilde{M}_{\tau_t}$. Then \hat{W} is under \mathbb{Q} a continuous local $(\hat{\mathcal{F}}_t)$ -martingale with $\langle \hat{W} \rangle_t = \langle \tilde{M} \rangle_{\tau_t} = \Lambda_{\tau_t} = t$ \mathbb{Q} -a.s. and therefore a \mathbb{Q} -Brownian motion for the filtration $(\hat{\mathcal{F}}_t)_{t \geq 0}$ by Lévy's characterisation of Brownian motion. In the time-changed filtration, consider the SDE for the standard CEV model,

$$d\hat{X}_t = |\hat{X}_t|^\beta d\hat{W}_t, \quad \hat{X}_0 = x_0 > 0. \quad (4.8)$$

This SDE has a unique strong solution \hat{X} which is a positive continuous strict local \mathbb{Q} -martingale (cf. Example 4.2). Moreover, it follows from the explicit formula for the transition density that $\lim_{t \searrow 0} \mathbb{E}_{\mathbb{Q}}[\hat{X}_t] = x_0$; see [19, Equation (7)]. Define the process $(\tilde{X}_t)_{t \in [0, T]}$ by $\tilde{X}_t = \hat{X}_{\Lambda_t}$, and note that $\tilde{M}_t = \hat{W}_{\Lambda_t}$, $t \geq 0$. Then $(\tilde{X}_t)_{t \in [0, T]}$ is a positive continuous local \mathbb{Q} -martingale for the filtration $(\tilde{\mathcal{F}}_t)_{t \in [0, T]}$ and plugging in the definitions and using (4.8) shows that \tilde{X} satisfies the SDE

$$d\tilde{X}_t = |\tilde{X}_t|^\beta d\tilde{M}_t = \tilde{V}_t |\tilde{X}_t|^\beta dW_t^{\mathbb{Q}}, \quad \tilde{X}_0 = x_0.$$

Moreover, \tilde{X} is under \mathbb{Q} a positive $(\tilde{\mathcal{F}}_t)$ -supermartingale by Fatou's lemma, and so by (4.7) and the properties of \tilde{X} ,

$$\mathbb{E}_{\mathbb{Q}}[\tilde{X}_t] = \mathbb{E}_{\mathbb{Q}}[\hat{X}_{\Lambda_t}] \leq \mathbb{E}_{\mathbb{Q}}[\hat{X}_{\underline{\sigma}^2 t}] < x_0.$$

By the same argument, $\mathbb{E}_{\mathbb{Q}}[\tilde{X}_t] \geq \mathbb{E}_{\mathbb{Q}}[\hat{X}_{\bar{\sigma}^2 t}]$, and since the right-hand side does not depend on v_0 , this together with $\lim_{t \searrow 0} \mathbb{E}_{\mathbb{Q}}[\tilde{X}_t] = x_0$ establishes (4.5).

Next, we show that \mathcal{S} is dynamically viable but fails to be dynamically efficient. To this end, note that $S = S^{(e_2)}$ and $X = V(e_1)(S^{(e_2)})$, where $e_1 = (1, 0)$

and $e_2 = (0, 1)$ are the buy-and-hold strategies of the first and second asset. By the above result for $\mathbb{Q} = \mathbb{P}$, $X = V(e_1)(S^{(e_2)})$ is a local \mathbb{P} -martingale and so \mathcal{S} is dynamically viable by Theorem 3.1. To establish that \mathcal{S} is not dynamically efficient, by Theorem 3.1, it suffices to show that there is no $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that $X = V(e_1)(S^{(e_2)})$ is a (true) \mathbb{Q} -martingale. But if X is a local \mathbb{Q} -martingale under $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , then $W^{\mathbb{P}} = \int V^{-1}|X|^{-\beta} dX$, by (4.3) and strict positivity of X and V , is a continuous local \mathbb{Q} -martingale with quadratic variation $\langle W^{\mathbb{P}} \rangle_t = \int_0^t V_s^{-2}|X_s|^{-2\beta} d\langle X \rangle_s = t$, $t \in [0, T]$, and so $W^{\mathbb{P}}$ is also a \mathbb{Q} -Brownian motion. Again by the above result, X is therefore a strict local \mathbb{Q} -martingale. So \mathcal{S} is not dynamically efficient.

Finally, we show that \mathcal{S} is statically efficient, which together with the above yields that \mathcal{S} is a nontrivial bubbly market. We already know that \mathcal{S} is dynamically viable, and for each $s \in [0, T)$, e_2 is a numéraire strategy in $\mathbf{h}\mathcal{U}_s$ and $S = S^{(e_2)}$. By Theorem 3.16, it is thus enough to find some $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T such that \mathbb{Q} is a one-step EMM for $S^{(\eta)}$ on $\{s, T\}$, i.e., $\mathbb{E}_{\mathbb{Q}}[X_T | \mathcal{F}_s] = X_s$ \mathbb{Q} -a.s., and for that, it even suffices to show $\mathbb{P}[X_T > X_s | \mathcal{F}_s] > 0$ \mathbb{P} -a.s. and $\mathbb{P}[X_T < X_s | \mathcal{F}_s] > 0$ \mathbb{P} -a.s. Now the pair (S, V) is a strong Markov process for the filtration $(\mathcal{F}_t)_{t \in [0, T]}$; so it suffices to show $\mathbb{P}_{x,v}[X_{T-s} > x] > 0$ and $\mathbb{P}_{x,v}[X_{T-s} < x] > 0$ for all $x > 0$ and $v \in (\underline{\sigma}, \bar{\sigma})$, where we write $\mathbb{P}_{x,v}$ for the distribution of the solution (X, V) of (4.3) and (4.4) with initial value (x, v) .

To make the presentation more transparent, we work from now on without loss of generality on the canonical space and denote by (S, V) the canonical Markov process and by ϑ the shift operator. Fix $x > 0$ and $v \in (\underline{\sigma}, \bar{\sigma})$. Then $\mathbb{E}_{x,v}[X_{T-s}] < x$ by the above general result, and so $\mathbb{P}_{x,v}[X_{T-s} < x] > 0$. For the other inequality, choose $\varepsilon \in (0, T - s)$ small enough that $\mathbb{E}_{2x, \tilde{v}}[X_\varepsilon] > x$ for all $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$; see (4.5). Define the stopping times $\tau_x^\uparrow := \inf\{t \geq 0 : X_t \geq x\}$ and $\tau_x^\downarrow := \inf\{t \geq 0 : X_t \leq x\}$. We claim that $\mathbb{P}_{\tilde{x}, \tilde{v}}[\tau_{2x}^\uparrow < \varepsilon] > 0$ for $\tilde{x} > 0$, $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$ and $\mathbb{P}_{2x, \tilde{v}}[\tau_x^\downarrow > \varepsilon] > 0$ for $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$. The first claim follows immediately from the fact that X is a strict local $\mathbb{P}_{\tilde{x}, \tilde{v}}$ -martingale on $[0, \varepsilon]$ and hence cannot be uniformly bounded. The second claim holds because if we had $\mathbb{P}_{2x, \tilde{v}}[\tau_x^\downarrow \leq \varepsilon] = 1$ for some $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$, then the choice of ε and the supermartingale property of X under $\mathbb{P}_{2x, \tilde{v}}$ would yield a contradiction via $x < \mathbb{E}_{2x, \tilde{v}}[X_\varepsilon] \leq \mathbb{E}_{2x, \tilde{v}}[X_{\tau_x^\downarrow}] = \mathbb{E}_{2x, \tilde{v}}[x] = x$. Note also that $\mathbb{P}_{x, \tilde{v}}[X_{T-s-\varepsilon} \leq x] > 0$ for $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$ by the fact that $\mathbb{E}_{x, \tilde{v}}[X_{T-s-\varepsilon}] < x$. Combining everything then yields

$$\begin{aligned}
\mathbb{P}_{x,v}[X_{T-s} > x] &\geq \mathbb{E}_{x,v} \left[\mathbf{1}_{\{X_\varepsilon \circ \vartheta_{T-s-\varepsilon} > x\}} \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] \\
&= \mathbb{E}_{x,v} \left[\mathbb{E}_{X_{T-s-\varepsilon}, V_{T-s-\varepsilon}} \left[\mathbf{1}_{\{X_\varepsilon > x\}} \right] \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] \\
&\geq \mathbb{E}_{x,v} \left[\mathbb{E}_{X_{T-s-\varepsilon}, V_{T-s-\varepsilon}} \left[\mathbf{1}_{\{X_\varepsilon > x\}} \mathbf{1}_{\{\tau_{2x}^\uparrow < \varepsilon\}} \right] \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] \\
&\geq \mathbb{E}_{x,v} \left[\mathbb{E}_{X_{T-s-\varepsilon}, V_{T-s-\varepsilon}} \left[\mathbf{1}_{\{(\inf_{0 \leq u \leq \varepsilon} X_u) \circ \vartheta_{\tau_{2x}^\uparrow} > x\}} \mathbf{1}_{\{\tau_{2x}^\uparrow < \varepsilon\}} \right] \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] \quad (4.9) \\
&= \mathbb{E}_{x,v} \left[\mathbb{E}_{X_{T-s-\varepsilon}, V_{T-s-\varepsilon}} \left[\mathbb{E}_{2x, V_{\tau_{2x}^\uparrow}} \left[\mathbf{1}_{\{\inf_{0 \leq u \leq \varepsilon} X_u > x\}} \right] \mathbf{1}_{\{\tau_{2x}^\uparrow < \varepsilon\}} \right] \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] \\
&= \mathbb{E}_{x,v} \left[\mathbb{E}_{X_{T-s-\varepsilon}, V_{T-s-\varepsilon}} \left[\mathbb{E}_{2x, V_{\tau_{2x}^\uparrow}} \left[\mathbf{1}_{\{\tau_x^\downarrow > \varepsilon\}} \right] \mathbf{1}_{\{\tau_{2x}^\uparrow < \varepsilon\}} \right] \mathbf{1}_{\{X_{T-s-\varepsilon} \leq x\}} \right] > 0,
\end{aligned}$$

where we use at the start a trivial inclusion, the Markov property, another trivial inclusion, and after (4.9) the strong Markov property and the definition of τ_x^\downarrow , and finally that $\mathbb{P}_{2x, \tilde{v}}[\tau_x^\downarrow > \varepsilon] > 0$ for $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$, $\mathbb{P}_{\tilde{x}, \tilde{v}}[\tau_{2x}^\uparrow < \varepsilon] > 0$ for $\tilde{x} > 0$, $\tilde{v} \in (\underline{\sigma}, \bar{\sigma})$ and $\mathbb{P}_{x, v}[X_{T-s-\varepsilon} \leq x] > 0$. For (4.9), we use that if after time $\tau_{2x}^\uparrow < \varepsilon$, we stay above x for at least ε units of time, then we must be above x at time ε . This ends the example.

Example 4.3 is of interest for several reasons. First of all, it is a CEV model with stochastic volatility and therefore quite realistic from a practical perspective. In fact, if we replace the volatility process V from the SDE (4.4) by a geometric Brownian motion, we get the well-known SABR model (see [29]). Next, because $X = S^1$ is a strict local \mathbb{P} -martingale, we have a bubble model in the sense of Loewenstein and Willard [57], Protter et al. [40, 41, 67] or Cox and Hobson [8], among others; see Section 5. However, we actually have more. In the usual approaches to bubble modelling in incomplete markets (e.g. [41, 67]), one fixes an ELMM or risk-neutral measure \mathbb{Q} and assumes that the asset price is under \mathbb{Q} a strict local martingale. This should more accurately be called a \mathbb{Q} -bubble, because there might well be another ELMM \mathbb{Q}' under which the asset price is a true martingale, and so the above notion of a bubble possibly depends in a crucial way on the choice of the risk-neutral measure one works with. For a more thorough discussion of that issue, we refer to Section 5. In Example 4.3, this issue does not arise at all. We have a nontrivial bubbly market, and this means by Theorem 3.22 that for all possible representatives $S \in \mathcal{S}$ and all ELMMs \mathbb{Q} , we always have for S under \mathbb{Q} a strict local martingale. In other words, Example 4.3 gives us a market for a bubble which is robust with regard to the choice of the ELMM one wants to work with. This can also be seen from the above arguments—we show that whenever we have (4.3) and (4.4) under some $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T , we have for X a strict local \mathbb{Q} -martingale. Apart from Jarrow and Larsson [39, Theorem 5.7], such a robust bubble model has not been presented in the literature so far. (We also point out that in [39, Theorem 5.7] it is not checked that we have strict local \mathbb{Q} -martingales on (any) finite time horizon.) Last but not least, the market in Example 4.3 is statically efficient and so naive invest-and-keep trading is still optimal in its own class. Even though this is not discussed in this chapter, we remark that this property is crucial when studying economically consistent valuation of contingent claims.

Example 4.3 gives an incomplete market model which is a nontrivial bubbly market and therefore robust in the sense that the asset price is a strict local martingale under *every* ELMM \mathbb{Q} . However, it can also happen that we have for S under some ELMM \mathbb{Q} a strict local martingale, but under another ELMM \mathbb{Q}' a true martingale. The next example gives a concrete model where this happens. Of course, by Theorem 3.22, the market generated by this model is then not (nontrivial) bubbly in our sense.

Example 4.4 (*A \mathbb{Q} -bubble which is not a \mathbb{Q}' -bubble*). Take two \mathbb{Q} -Brownian motions $W^1 = (W_t^1)_{t \in [0, T]}$ and $W^2 = (W_t^2)_{t \in [0, T]}$ with respect to a given filtration $(\mathcal{F}_t)_{t \in [0, T]}$; this need not be generated by (W^1, W^2) . We assume that W^1 and

W^2 are positively but not perfectly correlated: there exists a constant $\lambda \in (0, 1)$ such that $d\langle W^1, W^2 \rangle_t = \lambda dt$ (we use ρ for something else below). Consider the market \mathcal{S} generated by the process $S = (S^1, S^2) = (1, X_t)_{t \in [0, T]}$, where X satisfies the SDE, for some constant $\xi > 0$,

$$dX_t = \xi X_t V_t dW_t^1, \quad X_0 = x_0 > 0. \quad (4.10)$$

The volatility process $V = (V_t)_{t \in [0, T]}$ is itself stochastic and satisfies the SDE

$$dV_t = bV_t dW_t^2, \quad V_0 = 1, \quad (4.11)$$

for some constant $b > 0$. It is clear that (4.11) and (4.10) have unique strong solutions V and X , and we claim that X is a strict local \mathbb{Q} -martingale on $[0, T]$ and that there exists a probability measure $\mathbb{Q}' \approx \mathbb{Q}$ on \mathcal{F}_T such that X is a true \mathbb{Q}' -martingale.

To prove this, we use the results of Sin [76]. Setting $a := (b\lambda, b\sqrt{1-\lambda^2})$, $\sigma := (\xi, 0)$ and $\rho = 0$, we are exactly in the setup of [76, Theorems 3.3 and 3.9] with $\alpha = 1$. Note that $a \cdot \sigma = \xi b\lambda > 0$, and a, σ are not parallel. So we immediately get the existence of \mathbb{Q}' (called \mathbb{Q}^a in [76, Theorem 3.9]). However, to show that X is a strict local \mathbb{Q} -martingale on $[0, T]$, we cannot directly rely on [76, Theorem 3.3], since a strict local martingale on $[0, \infty)$ might still be a true martingale on a given finite interval. But S^2 is a positive local \mathbb{Q} -martingale, hence a \mathbb{Q} -supermartingale, and so it suffices to show that $\mathbb{E}[X_T] < x_0$. For that, by [76, Lemma 4.2], it is enough to show that $\mathbb{Q}[\hat{\tau} < T] > 0$, where $\hat{\tau}$ is the explosion time of the SDE

$$d\hat{V}_t = b\hat{V}_t d\hat{W}_t + b\xi\lambda\hat{V}_t^2 dt, \quad \hat{V}_0 = 1, \quad (4.12)$$

with a generic \mathbb{Q} -Brownian motion $\hat{W} = (\hat{W}_t)_{t \geq 0}$. For the rest of the example, denote by \hat{V} the canonical process on the path space $C([0, \infty); (0, \infty) \cup \{\Delta\})$, where Δ is an absorbing cemetery state, by \mathbb{P}_v the distribution on the path space of the solution of (4.12) with initial value $v > 0$, and by ϑ the shift operator. It follows from [76, Lemma 4.3] that under each \mathbb{P}_v , \hat{V} explodes in finite time with positive probability, and is valued in $(0, \infty)$ before explosion (the argument in [76] does not depend on the initial value v). With $T_v := \inf\{T \geq 0 : \mathbb{P}_v[\hat{\tau} < T] > 0\}$ for $v > 0$, this means that $T_v < \infty$. We claim that in fact $T_v = 0$ for all $v > 0$, and this will complete the proof, because we then have $\mathbb{P}_v[\hat{\tau} < T] > 0$ for all $T > 0$, as desired.

We first establish that $v \mapsto T_v$ is decreasing. Suppose to the contrary that $T_{v_1} < T_{v_2}$ for $0 < v_1 < v_2$. Then there is $\varepsilon \in (0, T_{v_2} - T_{v_1})$ with $\mathbb{P}_{v_1}[\hat{\tau} < T_{v_1} + \varepsilon] > 0$ and $\mathbb{P}_{v_2}[\hat{\tau} < T_{v_1} + \varepsilon] = 0$. With $\tau_y^\uparrow := \inf\{t \geq 0 : \hat{V}_t \geq y\}$, we can then use $\hat{\tau} = \hat{\tau} \circ \vartheta_{\tau_{v_2}^\uparrow}$ and the strong Markov property to obtain that

$$0 < \mathbb{P}_{v_1}[\hat{\tau} < T_{v_1} + \varepsilon] = \mathbb{P}_{v_1}[\hat{\tau} \circ \vartheta_{\tau_{v_2}^\uparrow} < T_{v_1} + \varepsilon] = \mathbb{E}_{v_1}[\mathbb{P}_{v_2}[\hat{\tau} < T_{v_1} + \varepsilon]] = 0,$$

which is a contradiction. So T_v is decreasing in v and $T_\infty := \lim_{v \rightarrow \infty} T_v$ exists in $[0, \infty)$. If $T_\infty > 0$, then there is $\varepsilon > 0$ such that $\mathbb{P}_v[\hat{\tau} \leq \varepsilon] = 0$ for all $v \in (0, \infty)$. Then the Markov property gives for all $v \in (0, \infty)$ that

$$\mathbb{P}_v[\hat{\tau} \leq 2\varepsilon] = \mathbb{P}_v[\hat{\tau} \leq 2\varepsilon, \hat{\tau} > \varepsilon] = \mathbb{P}_v[\hat{\tau} \circ \vartheta_\varepsilon \leq 2\varepsilon] = \mathbb{E}_v[\mathbb{P}_{\hat{V}_\varepsilon}[\hat{\tau} \leq \varepsilon]] = 0.$$

Iterating this argument yields $\mathbb{P}_v[\hat{\tau} \leq n\varepsilon] = 0$ for all $n \in \mathbb{N}$, $v \in (0, \infty)$, and we arrive at a contradiction. So $T_\infty = 0$. Finally, we show that $T_v = 0$ for all $v > 0$. If this fails, there exists $v_0 \in (0, \infty)$ such that $T_{v_0} > 0$, and then there exists $\varepsilon > 0$ such that $\mathbb{P}_{v_0}[\hat{\tau} \leq 2\varepsilon] = 0$. Using that $v \mapsto T_v$ is decreasing and $T_\infty = 0$, pick $v_1 > v_0$ large enough that $T_{v_1} < \varepsilon$; then $\mathbb{P}_v[\hat{\tau} \leq \varepsilon] > 0$ for all $v \geq v_1$ since T_v is decreasing in v . Because $b\xi\lambda > 0$, a standard comparison argument for SDEs (see e.g. [44, Theorem 23.5]) yields $\hat{V} \geq \tilde{V}$ \mathbb{P}_{v_0} -a.s., where $\tilde{V} = v_0\mathcal{E}(b\hat{W})$ satisfies $d\tilde{V}_t = b\tilde{V}_t d\hat{W}_t$, and so $\mathbb{P}_{v_0}[\hat{V}_\varepsilon \geq v_1] \geq \mathbb{P}_{v_0}[\tilde{V}_\varepsilon \geq v_1] > 0$ since \tilde{V}_ε has a lognormal distribution. Using the Markov property then gives the contradiction

$$\begin{aligned} 0 &= \mathbb{P}_{v_0}[\hat{\tau} \leq 2\varepsilon] \geq \mathbb{P}_{v_0}[\hat{\tau} \leq 2\varepsilon, \hat{V}_\varepsilon \geq v_1, \hat{\tau} > \varepsilon] = \mathbb{P}_{v_0}[\hat{\tau} \circ \vartheta_\varepsilon \leq \varepsilon, \hat{V}_\varepsilon \geq v_1] \\ &= \mathbb{E}_{v_0}[\mathbb{E}_{V_\varepsilon}[\hat{\tau} \leq \varepsilon] \mathbf{1}_{\{V_\varepsilon \geq v_1\}}] > 0. \end{aligned}$$

So we have $T_v = 0$ for all $v > 0$, and X is a strict local \mathbb{Q} -martingale on $[0, T]$, for each $T > 0$. This ends the example.

5 Comparison to the literature

In this section, we point out some connections of our work to the existing literature.

5.1 General comments and framework

The literature on bubbles is vast, and it is impossible to survey this here, even only approximately. The Encyclopedia of Quantitative Finance [43] has for example a 15-page entry “Bubbles and crashes”, with a list of more than 100 references. A recent survey article by Scherbina and Schlusche [72] puts more emphasis on behavioural models and rational models with frictions, and also provides a brief overview on the history of bubbles. The books of Brunnermeier [6] or Shiller [75] are often quoted as early classics; and the recent paper by Protter [67] entitled “A mathematical theory of financial bubbles” also contains around 160 references plus some discussions of literature.

Our aim in this section is much more modest. We try to compare our definitions and results to some seminal recent papers from the mathematical finance body of the literature. To that end, it is helpful to first provide a unified framework within which different approaches can be analysed.

In discussions with economists or in the financial economics literature, the standard description of a bubble says that this is (or is linked to) an asset whose *market value* exceeds (or differs from) its *fundamental value*. To make this more formal, we fix a time horizon $T > 0$ and a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$ with $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ and describe a dividend-paying asset (D, S) by its cumulative dividend process $D = (D_t)_{t \in [0, T]}$ and its ex-dividend price process $S = (S_t)_{t \in [0, T]}$, both in the same fixed units. We also include a bank account $B = (B_t)_{t \in [0, T]}$; so if we hold the asset over a time interval $(t, u]$, meaning that we purchase it at

time t and sell it at u , we obtain at time u a total cashflow or gain of

$$S_u - \frac{B_u}{B_t} S_t + B_u \int_t^u \frac{1}{B_s} dD_s,$$

whose equivalent discounted back to time t is

$$B_t \left(\frac{S_u}{B_u} - \frac{S_t}{B_t} + \int_t^u \frac{1}{B_s} dD_s \right) =: B_t(W_u - W_t).$$

In discounted units, we therefore have

$$W_t = \frac{S_t}{B_t} + \int_0^t \frac{1}{B_s} dD_s \quad (5.1)$$

for the discounted gains process from holding one unit of the asset. Of course, both S and D could be multidimensional, of the same dimension; then we add a superscript i for S , D and W .

If we now denote by S_t^* the (undiscounted) fundamental value of our asset at time t , then $S_t^* \neq S_t$ (or $S_t^* < S_t$) would mean that the asset has a bubble, and the difference $\beta_t := S_t - S_t^*$ is usually called the (size of the) bubble or the bubble component of the asset.

The big snag with the above description is that it does not tell us what the “fundamental value” is, nor where it comes from. Axiomatically, it is reasonable to impose that we have a fundamental value operator which assigns fundamental values to assets or general financial products. Such an operator is usually monotone and linear, and if we add some mild continuity conditions, it is reasonable to assume that it has the form, for $0 \leq t \leq u \leq T$,

$$\begin{aligned} \Phi_t(S, D) &= B_t \mathbb{E} \left[\frac{Z_u}{Z_t} \frac{S_u}{B_u} + \int_t^u \frac{Z_s}{Z_t} \frac{1}{B_s} dD_s \mid \mathcal{F}_t \right] \\ &= \frac{\mathbb{E}[(Z_u/B_u)S_u + \int_t^u (Z_s/B_s) dD_s \mid \mathcal{F}_t]}{Z_t/B_t}, \end{aligned} \quad (5.2)$$

where the positive adapted processes Z and Z/B are often called a deflator or a state price density, respectively. If we consider the asset “bank account”, we get with S replaced by B and D replaced by 0 the fundamental value

$$\Phi_t(B, 0) = B_t \mathbb{E} \left[\frac{Z_u}{Z_t} \mid \mathcal{F}_t \right].$$

So if the bank account has no bubble, it is reasonable to assume that Z is a positive \mathbb{P} -martingale. More generally, it is only imposed that Z is a positive local \mathbb{P} -martingale with $Z_0 = 1$. In the martingale case, we can define a probability measure \mathbb{Q} equivalent to \mathbb{P} by $d\mathbb{Q} = Z_T d\mathbb{P}$ (assuming $Z_0 = 1$ so that $\mathbb{Q} = \mathbb{P}$ on \mathcal{F}_0) and then rewrite (5.2) as

$$\Phi_t(S, D) = B_t \mathbb{E}_{\mathbb{Q}} \left[\frac{S_u}{B_u} + \int_t^u \frac{1}{B_s} dD_s \mid \mathcal{F}_t \right]$$

by the Bayes rule. (To be accurate, this assumes that D is of finite variation.) Using (5.1), we can also reformulate this as

$$\Phi_t(S, D) = B_t \mathbb{E}_{\mathbb{Q}}[W_u - W_t | \mathcal{F}_t] + S_t. \quad (5.3)$$

So if we decide that the fundamental value of (S, D) at time t is

$$S_t^* := \Phi_t(S, D),$$

then having a bubble means that W is not a \mathbb{Q} -martingale. We emphasise, however, that this notion of fundamental value and hence also of a bubble depends on Z or \mathbb{Q} ; so it would be more accurate to talk about a \mathbb{Q} -bubble or a Z -bubble.

With the above terminology, we are now ready to discuss a number of important papers from the literature.

5.2 Loewenstein and Willard

In a seminal paper [57], these authors start with a financial market (B, S, D) as described above. All their processes are continuous, $B > 0$ is of finite variation, $D \geq 0$ is increasing, and $S \geq 0$. (Actually, both S and D are \mathbb{R}^K -valued in their setup.) They assume that there is a local \mathbb{P} -martingale $Z > 0$ with $Z_0 = 1$ such that deflated prices $\frac{Z}{B}S + \int \frac{Z}{B} dD$ form a local \mathbb{P} -martingale. They also impose completeness of the market by assuming that the above Z is unique. Then they say that asset i has a bubble if its deflated price process $\frac{Z}{B}S^i + \int \frac{Z}{B} dD^i$ is a strict local \mathbb{P} -martingale, and that the bank account has a bubble if $Z = \frac{Z}{B}B$ itself is a strict local \mathbb{P} -martingale. As in (5.2), the fundamental values are defined as

$$S_t^* := \Phi_t(S, D) := B_t \mathbb{E} \left[\frac{Z_T}{Z_t} \frac{S_T}{B_T} + \int_t^T \frac{Z_s}{Z_t} \frac{1}{B_s} dD_s \mid \mathcal{F}_t \right],$$

$$B_t^* := \Phi_t(B, 0) := B_t \mathbb{E} \left[\frac{Z_T}{Z_t} \mid \mathcal{F}_t \right],$$

and the bubble components of asset i and the bank account are then, for $t \in [0, T]$,

$$\beta_t^i := S_t^i - S_t^{*,i},$$

$$\beta_t^B := B_t - B_t^*.$$

The authors emphasise that their definition is in line with the previous economic literature (e.g. Diba and Grossman [15], Tirole [79], or Santos and Woodford [71]).

How does this compare with our approach? Suppose for simplicity that there are no dividends so that $D \equiv 0$. To avoid confusion in the notation, we write in the sequel Y instead of S for the stocks in the sense of Loewenstein and Willard, and reserve S for a representative of a market \mathcal{S} in our sense. So let \mathcal{S} be the market generated by $S = (B, Y)$, where B is a bank account and Y a stock in the sense of Loewenstein and Willard. In our setup, completeness translates into saying that for any numéraire strategy η , there is at most one equivalent local

martingale measure (ELMM) \mathbb{Q} for $V(\eta)$ -discounted prices $S^{(\eta)} = S/V(\eta)(S)$. If \mathcal{S} is dynamically viable, there exists a pair (η, \mathbb{Q}) such that \mathbb{Q} is an ELMM for $S^{(\eta)}$; see Theorem 3.1. If \mathcal{S} is in addition a bubbly market in our sense, the (unique) ELMM \mathbb{Q} for $S^{(\eta)}$ is such that for at least one index j , $S^{(\eta),j}$ is a strict local \mathbb{Q} -martingale; see Theorem 3.11 and Corollary 3.13. Let $Z^{\mathbb{Q}}$ be the density process of \mathbb{Q} with respect to \mathbb{P} . By Bayes' theorem, $Z^{\mathbb{Q}}S^{(\eta)}$ is a local \mathbb{P} -martingale and $Z^{\mathbb{Q}}S^{(\eta),j}$ is a strict local \mathbb{P} -martingale. Setting now $Z := Z^{\mathbb{Q}}B/V(\eta)(S)$, we get that Z and $\frac{Z}{B}Y$ are local \mathbb{P} -martingales, and either Z (if $j = 1$) or $\frac{Z}{B}Y^i$ (for $i = j - 1$ if $j > 1$) is a strict local \mathbb{P} -martingale. So under completeness, if the market \mathcal{S} generated by $S = (B, Y)$ is a dynamically viable bubbly market, either one of the stocks Y^i or the bank account B has a bubble in the sense of Loewenstein and Willard [57]. It is not difficult to check that also the converse is true. In this sense, [57] is closest to our setup in the existing literature. Note, however, that the tricky issue of bubbles in an *incomplete* financial market is not addressed there.

We remark in passing that a process Z as above has been called a *local martingale density* or a *local martingale deflator* in the recent literature, and that its existence, for a setting $D \equiv 0$ without dividends, has been shown to be equivalent to the absence-of-arbitrage condition of *no unbounded profit with bounded risk* (NUPBR); see e.g. Kardaras [48] or Takaoka and Schweizer [73]. This condition, as shown by Kardaras [48], is in turn equivalent to the absence of arbitrage of the first kind (NA₁) or absence of cheap thrills, and it is this latter condition that appears also in Loewenstein and Willard [56]; see also Chapters III and VI for a discussion of those concepts from a numéraire-independent perspective. For our discussion of bubble modelling, these remarks are at present mostly tangential, but we come back to them a bit later.

5.3 Jarrow, Protter and Shimbo

These authors provide a detailed study of asset price bubbles in two papers—one for complete [40] and one for incomplete markets [41]. Their basic setup is a financial market $(1, S, D)$ as above, with $B \equiv 1$ and $S, D \geq 0$ one-dimensional semimartingales. They work on a right-open stochastic interval $\llbracket 0, \tau \llbracket$ with a stopping time τ and add a liquidation value X_τ at τ to the stock S , but this is just a minor technical detail; we can replace X_τ by the final stock price S_T without changing the essence of the model. (We also point out that the dividend process D should actually be increasing for some of the arguments to work.) Instead of the existence of (a local martingale density) Z (or equivalently NUPBR, as discussed above in Section 5.2), [40, 41] impose the stronger condition NFLVR for the gains process $W = S + D$; so there exists an ELMM \mathbb{Q} for W by the fundamental theorem of asset pricing. The first paper [40] on complete markets assumes that \mathbb{Q} is unique; the second [41] does not, and we denote by $\mathcal{M}_e(W)$ the nonempty set of ELMMs \mathbb{Q} for W .

For the complete case [40], the fundamental value is defined as

$$S_t^* := \Phi_t(S, D) := \mathbb{E}_{\mathbb{Q}}[S_T + (D_T - D_t) \mid \mathcal{F}_t] = \mathbb{E}_{\mathbb{Q}}[W_T - W_t \mid \mathcal{F}_t] + S_t, \quad (5.4)$$

and so an asset price bubble

$$\beta_t := S_t - S_t^* = W_t - \mathbb{E}_{\mathbb{Q}}[W_T | \mathcal{F}_t]$$

appears if and only if the local \mathbb{Q} -martingale W is not a true \mathbb{Q} -martingale. However, unlike in the case of Loewenstein and Willard [57], a bubble in the bank account is not possible in [40] (and in [41], for that matter) due to the more restrictive assumption of NFLVR (instead of NUPBR). So similar to [57], the link between bubbles and strict local martingales is due to the definition (5.4) of the fundamental value. [40] also introduce different types of bubbles (depending on the time-horizon), provide a decomposition of bubbles and discuss the valuation of contingent claims.

The incomplete case in [41] is more challenging. Since $\mathcal{M}_e(W)$ is no longer a singleton, it is not immediately clear which ELMM \mathbb{Q} one should use to define a fundamental value as in (5.4). Just picking one \mathbb{Q} and using that throughout would be ad hoc and would also just lead back to the results from the complete case. To overcome this problem, [41] propose a mechanism where “the market” chooses and sometimes (at random times σ_i) changes the measure used in (5.4), so that one works with $\mathbb{Q}^i \in \mathcal{M}_e(W)$ for times t between σ_i and σ_{i+1} . In effect, this means that one uses a fundamental value of the form

$$S_t^* := \Phi_t(S, D) := \mathbb{E}_{\mathbb{Q}_t}[W_T - W_t | \mathcal{F}_t] + S_t, \quad (5.5)$$

where the measure \mathbb{Q}_t used at time t now depends on t as well, and so the analysis of bubbles becomes more involved. In the same spirit, but in a different setup², Biagini et al. [4] study the case where \mathbb{Q}_t moves smoothly from one \mathbb{Q} to another R . In both cases, however, the actual choice of \mathbb{Q}^i , or \mathbb{Q} and R , is not fully convincing from an economic perspective. For example, [41] makes the assumptions that there are enough liquidly traded derivatives in the market to determine the ELMM \mathbb{Q} , and that \mathbb{Q} can actually be identified from market prices. In effect, this practically leads us back to the complete case studied in [40]. Moreover, we are not aware of any well-established procedures to implement such an identification of \mathbb{Q} from market prices, and it seems also conceptually difficult to reconcile a \mathbb{Q} determined from liquid derivative prices with possible violations, due to bubbles, of e.g. put-call parity. In any case, the resulting object should more accurately be called a \mathbb{Q} -bubble, because there might well be another ELMM \mathbb{Q}' under which the asset price is a true martingale (see Example 4.4), and so the above notion of a bubble depends in a crucial way on the choice of the risk-neutral measure one works with. For all these reasons, we prefer an approach which only uses basic assets as given and does not need partly exogenous inputs to define bubbles. Hence, our notion of a bubbly market is more restrictive than the notion of a \mathbb{Q} -bubble in the sense of Protter et al. [41], i.e., if $S = (1, Y)$ has a \mathbb{Q} -bubble (writing as above Y instead of S for a stock in the sense of [41]), the market \mathcal{S} generated by S might not be (nontrivial) bubbly in our sense.

²[41] needs a bigger filtration \mathbb{G} to accommodate the σ_i (which are independent of \mathbb{F}), whereas [4] always stays within \mathbb{F} .

A major part of the analysis in [40, 41] is to study issues of valuation in markets with bubbles, and we comment on this below in Section 5.4. Protter [67] also discusses ideas to identify a bubble by statistical methods and gives in Section 11 an overview with discussion of some other approaches in the literature. We refer to that instead of repeating it here.

5.4 Bubbles and derivative pricing

Bubbles and strict local martingales have come up in mathematical finance with some prominence in the area of option pricing, in particular with relation to violations of put-call parity. Early work on that topic appears in Lewis [55], and this has been taken up in two seminal papers by Cox and Hobson [8] and Heston, Loewenstein and Willard [34].

The setup of [8] is like [40] similar to [57] but a bit more restrictive; they have a model $(1, S, 0)$ without dividends, where $S \geq 0$ is a continuous semimartingale, and they assume NFLVR and completeness so that they have a unique ELMM \mathbb{Q} for S ($= W$ here). They say that S has a bubble if it is a strict local \mathbb{Q} -martingale; so the definition is the same as in [40]. The main focus of [8] is then on valuation of options in the presence of bubbles, and in particular on violation of put-call parity.

Heston et al. [34] consider a setup $(1, S, 0)$ without dividends, where S is a one-dimensional local or stochastic volatility model; so they allow in particular incomplete markets. They say that “[a]n asset with a nonnegative price has a “bubble” if there is a self-financing portfolio with pathwise nonnegative wealth that costs less than the asset and replicates the asset’s price at a fixed future date. The bubble’s value is the difference between the asset’s price and the lowest cost replicating strategy” [34, Definition 2.1]. In terms of the discussion in Section 5.1, this means that they use as fundamental value the superreplication price. We note in passing that this kind of definition has also appeared in a recent paper by Loewenstein and Willard [58]. The approach of [34] is closely related to ours, but their notion of a bubble is more restrictive than ours in the more relevant case of incomplete markets. Indeed, let \mathcal{S} be the market generated by $S = (1, Y_t)_{t \in [0, T]}$ (writing as above Y instead of S for a stock in the sense of [34]). If there exists an ELMM \mathbb{Q} for S (or for Y , for that matter), then \mathcal{S} is dynamically viable (see Theorem 3.1), and if for each such \mathbb{Q} , Y is a strict local \mathbb{Q} -martingale, then \mathcal{S} is a bubbly market in our sense (see Theorem 3.22). Denote by $\mathcal{M}_e(Y)$ the nonempty set of all ELMMs for Y . Then by the classic dual characterisation of superreplication prices by Kramkov [52] and Föllmer and Kabanov [24], the classic superreplication price of Y_t is given by $\sup_{\mathbb{Q} \in \mathcal{M}_e(Y)} \mathbb{E}_{\mathbb{Q}}[Y_t]$, $t \in [0, T]$. Now it may happen that $Y_0 = \sup_{\mathbb{Q} \in \mathcal{M}_e(Y)} \mathbb{E}_{\mathbb{Q}}[Y_t]$ for every t even though Y is a strict local \mathbb{Q} -martingale for all $\mathbb{Q} \in \mathcal{M}_e(Y)$. Then \mathcal{S} is a bubbly market in our sense but not in sense of Heston et al. [34]. For complete markets, however, both concepts coincide; for if there exists only one ELMM for Y , then Y is a strict local \mathbb{Q} -martingale if and only if there exists $t \in [0, T]$ such that $Y_0 > \mathbb{E}_{\mathbb{Q}}[Y_t]$.

The main focus of [34] is to relate the existence of bubbles to multiplicity

(nonuniqueness) of solutions of the valuation PDEs of call and put options, but they also provide in their specific SV framework necessary and sufficient conditions for various bubbles (on the money market account or on the stock).

Other papers that study failures of option pricing properties in models with bubbles or strict local martingales are Ekström and Tysk [16] (who use PDE techniques), Pal and Protter [64] (who use h -transforms) or Madan and Yor [60] (who connect this to an extension of Itô's formula), among others.

5.5 Arbitrage aspects

In recent years, there has been a lot of interest in models which do not satisfy the classic strong absence-of-arbitrage condition NFLVR of Delbaen and Schachermayer [9, 13]. A major motivation has been that a number of empirical observations do not fit well with the stringent properties imposed by NFLVR, with prominent examples being stochastic portfolio theory (see Fernholz [21]) or the benchmark approach (see Platen [65] and Platen and Heath [28]). In most papers on these subjects, however, the models still satisfy NUPBR, and as discussed in Section 5.2, this brings us close again to the general setup presented in Section 5.1.

One typical question to ask is if or how hedging still works and if or how one could exploit the presence of potential arbitrages. The latter aspect is for example studied in Fernholz et al. [22]. They consider an Itô process model $(1, S, 0)$ where each asset is positive and do not assume NFLVR, but (implicitly) that the market price of risk is \mathbb{P} -a.s. square-integrable, which implies the existence of a local martingale deflator and hence NUPBR. Their goal is to explicitly construct portfolios which are better than the market portfolio; this is called “relative arbitrage”. In our terminology, they assume that the market is dynamically viable, but not dynamically efficient. Hence the market portfolio is not maximal, and so it is not surprising that there exist (maximal) strategies which “improve” it. But of course the main contribution of stochastic portfolio theory is to provide explicit formulas for such strategies.

In a similar vein, Ruf [70] discusses “hedging under arbitrage”, i.e., when NFLVR fails but we still have NUPBR. His setup is very similar to Loewenstein and Willard [57], but the market is not assumed to be complete. In a Markovian framework, [70] then uses Feynman–Kac type results to obtain and compute optimal strategies which superreplicate a given contingent claim with minimal initial capital. The link to bubbles here is mainly that by definition, bubbles are identified with assets whose prices follow strict local martingales.

In a quite different direction, one can ask what extra ingredient is needed to avoid having strict local martingales for asset prices, because many pricing anomalies turn out to be consequences of bubbles in that sense. As we have seen in Corollary 3.21, what is needed is a concept of no dominance. This idea goes back to Merton [62] and was formalised mathematically by Jarrow et al. [40, 41] in two different ways; related work was done apparently in parallel by Heston et al. [34]. A detailed study with very clear definitions was then given by Jarrow and Larsson [39]. They consider a market $(1, S, 0)$, where S is a d -dimensional

semimartingale for the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$. They say that “the market (\mathbb{F}, S) is [...] efficient on $[0, T]$ with respect to \mathbb{F} if there exist a consumption good price index ψ and an economy $((P_k)_{k=1}^K, \mathbb{F}, (\varepsilon_k)_{k=1}^K, (U_k)_{k=1}^K)$ for which (ψ, S) is an equilibrium price process S on $[0, T]$ ”. Here $k = 1, \dots, K$ denote different investors with beliefs $P_k \approx \mathbb{P}$ (subjective probability measures), endowment streams ε_k and (time-dependent) utility functions U_k . [39] show that (\mathbb{F}, S) is efficient on $[0, T]$ if and only if S satisfies NFLVR and no dominance (ND), i.e., the buy-and-hold strategy of each risky asset is maximal, or equivalently if and only if there exists a equivalent (true) martingale measure $\mathbb{Q} \approx \mathbb{P}$ on \mathcal{F}_T . Moreover, they consider the case of different information sets and finally provide examples of efficient and inefficient markets, namely local and stochastic volatility models.

Our definition of dynamic efficiency and dynamic no dominance is directly inspired by [39], and our Corollary 3.21 is a numéraire-independent version of (part of) their key result [39, Theorem 3.2]. In particular, that result justifies our terminology of dynamic efficiency and also motivates our notion of static efficiency. Nevertheless, our definition of dynamic no dominance is (formally) a bit weaker since it only imposes maximality for the market portfolio, not for all individual buy-and-hold strategies in all assets.

5.6 Delbaen and Schachermayer

Despite its drawbacks as briefly discussed in Remark 1.12 (a), one important inspiration for many of our concepts has been the work of Delbaen and Schachermayer, especially [11] for numéraire changes and related topics and [12] for maximality. We emphasise again that a direct comparison is delicate because we work with a different notion of admissible strategies. But there is no doubt that F. Delbaen is also well aware of the close connections between maximal elements or strategies, bubbles, and strict local martingales. This is illustrated by a presentation given in June 2012 at the QMF conference in Cairns, Australia. We quote from these slides that “[a] bubble is something that has a price that is too high or for the same amount of money you can get something better” and that “ $H \cdot S$, acceptable, could be called a bubble if the price of $f = (H \cdot S)_\infty$ is strictly lower than 0”. Delbaen also proposes some ideas to define non-bubbles; however, we have not seen any published work or preprint so far.

6 Appendix

In this appendix, we formulate and prove two technical results needed mainly for the dual characterisation of static efficiency in Section 3.4.

The first technical result gives a dual characterisation of weakly and strongly maximal strategies in $\mathbf{h}L_\sigma^{\text{sf}}$. The key ingredients are the Dalang–Morton–Willinger theorem, the one-step superhedging duality or optional decomposition, and the geometric characterisation of absence of arbitrage (all with contingent initial data). For the last point, for an \mathbb{R}^N -valued random vector X and a sub- σ -field \mathcal{G} , we write $\text{ri conv supp } \mathcal{L}(X | \mathcal{G})$ for the relative interior of the convex hull of the

(ω -dependent) topological support of the regular conditional distribution of X given \mathcal{G} .

Lemma 6.1. *For any stopping time $\sigma \in \mathcal{T}_{[0,T]}$, the following are equivalent:*

- (a) 0 is weakly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$.
- (b) 0 is strongly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$.
- (c) Each $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ is strongly (and hence weakly) maximal for $\mathbf{h}L_\sigma^{\text{sf}}$.
- (d) For each numéraire strategy η in $\mathbf{h}L_\sigma^{\text{sf}}$, there exists a one-step EMM \mathbb{Q} on $\{\sigma, T\}$ for $S^{(\eta)}$.
- (d') For each numéraire strategy η in $\mathbf{h}L_\sigma^{\text{sf}}$, we have

$$S_\sigma^{(\eta)} \in \text{ri conv supp } \mathcal{L}(S_T^{(\eta)} | \mathcal{F}_\sigma) \text{ P-a.s.}$$

- (e) There exist a numéraire strategy η in $\mathbf{h}L_\sigma^{\text{sf}}$ and a one-step EMM \mathbb{Q} on $\{\sigma, T\}$ for $S^{(\eta)}$.
- (e') There exists a numéraire strategy η in $\mathbf{h}L_\sigma^{\text{sf}}$ such that

$$S_\sigma^{(\eta)} \in \text{ri conv supp } \mathcal{L}(S_T^{(\eta)} | \mathcal{F}_\sigma) \text{ P-a.s.}$$

Moreover, if one of the above equivalent conditions is satisfied, then for any contingent claim F at time T with $\Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) < \infty$ P-a.s., there is $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ with

$$V_\sigma(\vartheta) = \Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) \text{ P-a.s.} \quad \text{and} \quad V_T(\vartheta) \geq F \text{ P-a.s.}$$

Proof. “(c) \Rightarrow (b) \Rightarrow (a)”. This is trivial.

“(b) \Rightarrow (c)”. Seeking a contradiction, suppose there is $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ which is not strongly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$. By the characterisation of strong maximality after Lemma 3.7, there exists a nonzero contingent claim F at time T such that

$$\Pi_\sigma(V_T(\vartheta) + F | L_\sigma^{\text{sf}}) \leq V_\sigma(\vartheta) < \infty \text{ P-a.s.} \quad (6.1)$$

Let $\delta > 0$ be arbitrary and C a positive contingent claim at time σ . By Lemma 3.7 applied to $F + V_T(\vartheta)$ and by (6.1), there exists $\bar{\vartheta} \in L_\sigma^{\text{sf}}$ such that

$$V_\sigma(\bar{\vartheta}) \leq V_\sigma(\vartheta) + \delta C \text{ P-a.s.} \quad \text{and} \quad V_T(\bar{\vartheta}) \geq V_T(\vartheta) + F \text{ P-a.s.}$$

Set $\hat{\vartheta} := \bar{\vartheta} - \vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$. Then

$$V_\sigma(\hat{\vartheta}) \leq \delta C \text{ P-a.s.} \quad \text{and} \quad V_T(\hat{\vartheta}) \geq F \text{ P-a.s.}$$

This implies $\Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) \leq \delta C$, and letting $\delta \searrow 0$ gives $\Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) = 0$. This contradicts the hypothesis of strong maximality of 0 for $\mathbf{h}L_\sigma^{\text{sf}}$.

“(a) \Rightarrow (d)”. Take a numéraire strategy η in $\mathbf{h}L_\sigma^{\text{sf}}$ and set $\mathcal{G}_0 := \mathcal{F}_\sigma$, $\mathcal{G}_1 := \mathcal{F}_T$, $X_0 := S_\sigma^{(\eta)}$, $X_1 := S_T^{(\eta)}$. Then (a) implies that the classic discounted one-period model $(1, X)$ is arbitrage-free in the standard sense that there is no \mathcal{G}_0 -measurable \mathbb{R}^N -valued random vector ξ with $\xi \cdot (X_1 - X_0) \geq 0$ P-a.s. and $\mathbb{P}[\xi \cdot (X_1 - X_0) > 0] > 0$. Indeed, if such a ξ exists, then $\vartheta := \xi \mathbf{1}_{[\sigma, T]} - (\xi \cdot S_\sigma^{(\eta)})\eta$ is in $\mathbf{h}L_\sigma^{\text{sf}}$, and noting that $V(\eta)(S^{(\eta)}) \equiv 1$ due to (1.7), we have

$$\begin{aligned} V_\sigma(\vartheta)(S^{(\eta)}) &= \xi \cdot S_\sigma^{(\eta)} - (\xi \cdot S_\sigma^{(\eta)}) = 0 \text{ P-a.s.}, \\ V_T(\vartheta)(S^{(\eta)}) &= \xi \cdot S_T^{(\eta)} - (\xi \cdot S_\sigma^{(\eta)}) = \xi \cdot (X_1 - X_0) \geq 0 \text{ P-a.s.}, \\ \mathbb{P}[V_T(\vartheta)(S^{(\eta)}) > 0] &= \mathbb{P}[\xi \cdot (X_1 - X_0) > 0] > 0, \end{aligned}$$

which contradicts (a). So by the Dalang–Morton–Willinger theorem, e.g. in the form of [25, Theorem 1.54]), there exists an EMM \mathbb{Q} for the above one-period model, and translating everything back to our setup, we see that \mathbb{Q} is a one-step EMM on $\{\sigma, T\}$ for $S^{(\eta)}$.

“(d) \Leftrightarrow (d’) & (e) \Leftrightarrow (e’)”. This follows directly from Jacod and Shiryaev [36, Theorem 3].

“(d) \Rightarrow (e)”. This is trivial as the market portfolio $\eta^S = (1, \dots, 1)$ is a numéraire strategy in $\mathbf{h}L_\sigma^{\text{sf}}$.

“(e) \Rightarrow (b) & additional assertion”. Let η and \mathbb{Q} be as in (e) and F a contingent claim at time T with $\Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}}) < \infty$ P-a.s. Denote by $\mathcal{Q} \neq \emptyset$ the set of all one-step EMMs on $\{\sigma, T\}$ for $S^{(\eta)}$. For each $\bar{\vartheta} \in \mathbf{h}L_\sigma^{\text{sf}}$ with $V_T(\bar{\vartheta}) \geq F$ and each $\tilde{\mathbb{Q}} \in \mathcal{Q}$, we then have

$$\begin{aligned} V_\sigma(\bar{\vartheta})(S^{(\eta)}) &= \bar{\vartheta}_\sigma \cdot S_\sigma^{(\eta)} = \mathbb{E}_{\tilde{\mathbb{Q}}}[\bar{\vartheta}_\sigma \cdot S_T^{(\eta)} | \mathcal{F}_\sigma] \\ &= \mathbb{E}_{\tilde{\mathbb{Q}}}[V_T(\bar{\vartheta})(S^{(\eta)}) | \mathcal{F}_\sigma] \geq \mathbb{E}_{\tilde{\mathbb{Q}}}[F(S^{(\eta)}) | \mathcal{F}_\sigma] \text{ P-a.s.} \end{aligned}$$

Thus, by the definition of superreplication prices,

$$\operatorname{ess\,sup}_{\tilde{\mathbb{Q}} \in \mathcal{Q}} \mathbb{E}_{\tilde{\mathbb{Q}}}[F(S^{(\eta)}) | \mathcal{F}_\sigma] \leq \Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}})(S^{(\eta)}) < \infty \text{ P-a.s.} \quad (6.2)$$

In particular, $\mathbb{P}[\Pi_\sigma(F | \mathbf{h}L_\sigma^{\text{sf}})(S^{(\eta)}) > 0] > 0$ if $\mathbb{P}[F(S^{(\eta)}) > 0] > 0$, which gives (b) in view of the characterisation of strong maximality after Lemma 3.7. Now define a one-period market as in “(a) \Rightarrow (d)” and set $\gamma_1 := F(S^{(\eta)})$ and $\gamma_0 := \operatorname{ess\,sup}_{\tilde{\mathbb{Q}} \in \mathcal{Q}} \mathbb{E}_{\tilde{\mathbb{Q}}}[\gamma_1 | \mathcal{G}_0] < \infty$. Then \mathcal{Q} is the set of all EMMs for $(1, X)$ and γ_1 is a contingent claim in this market, all in the classic sense. By the superhedging duality or optional decomposition (see e.g. [25, Corollary 7.15]), there exists an \mathbb{R}^N -valued \mathcal{F}_σ -measurable random vector ξ with $\gamma_0 + \xi \cdot (X_1 - X_0) \geq \gamma_1$ P-a.s. Set $\vartheta := \xi \mathbf{1}_{[\sigma, T]} + (\gamma_0 - \xi \cdot S_\sigma^{(\eta)})\eta$, which is in $\mathbf{h}L_\sigma^{\text{sf}}$. Then as in “(a) \Rightarrow (d)”,

$$\begin{aligned} V_\sigma(\vartheta)(S^{(\eta)}) &= \xi \cdot S_\sigma^{(\eta)} + (\gamma_0 - \xi \cdot S_\sigma^{(\eta)}) = \gamma_0 = \operatorname{ess\,sup}_{\tilde{\mathbb{Q}} \in \mathcal{Q}} \mathbb{E}_{\tilde{\mathbb{Q}}}[F(S^{(\eta)}) | \mathcal{F}_\sigma] \text{ P-a.s.}, \\ V_T(\vartheta)(S^{(\eta)}) &= \xi \cdot S_T^{(\eta)} + (\gamma_0 - \xi \cdot S_\sigma^{(\eta)}) = \gamma_0 + \xi \cdot (X_1 - X_0) \geq \gamma_1 = F(S^{(\eta)}) \text{ P-a.s.} \end{aligned}$$

Together with (6.2) and the definition of superreplication prices, this establishes the additional assertion. \square

The second technical result provides equivalent characterisations of static efficiency.

Lemma 6.2. *The following are equivalent:*

- (a) \mathcal{S} is statically efficient.
- (b) \mathcal{S} is statically viable, and for each stopping time $\sigma \in \mathcal{T}_{[0,T]}$, every $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ which satisfies $V_T(\vartheta) \geq 0$ P-a.s. is in $\mathbf{h}\mathcal{U}_\sigma$.
- (b') \mathcal{S} is statically viable, and for each deterministic time $s \in [0, T)$, every $\vartheta \in \mathbf{h}L_s^{\text{sf}}$ which satisfies $V_T(\vartheta) \geq 0$ P-a.s. is in $\mathbf{h}\mathcal{U}_s$.
- (c) \mathcal{S} is statically viable, and 0 is weakly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$, for each stopping time $\sigma \in \mathcal{T}_{[0,T]}$.
- (c') \mathcal{S} is statically viable, and 0 is weakly maximal for $\mathbf{h}L_s^{\text{sf}}$, for each deterministic time $s \in [0, T)$.
- (d) The zero strategy 0 is weakly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$, for each stopping time $\sigma \in \mathcal{T}_{[0,T]}$.
- (e) Every strategy $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ is strongly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$, for each stopping time $\sigma \in \mathcal{T}_{[0,T]}$.

Proof. “(a) \Rightarrow (b)”. Static viability follows from static efficiency. Fix $\sigma \in \mathcal{T}_{[0,T]}$ and take $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ with $V_T(\vartheta) \geq 0$ P-a.s. Set $|\vartheta| := (|\vartheta^1|, \dots, |\vartheta^d|)$. Then both $|\vartheta|$ and $|\vartheta| + \vartheta$ are in $\mathbf{h}\mathcal{U}_\sigma$, and $V_T(|\vartheta| + \vartheta) \geq V_T(|\vartheta|)$ P-a.s. For any $\tau \in \mathcal{T}_{[\sigma, T]}$, we have $|\vartheta|, |\vartheta| + \vartheta \in \mathbf{h}\mathcal{U}_\tau$. Since $|\vartheta|$ is (strongly and hence) weakly maximal for $\mathbf{h}\mathcal{U}_\tau$ by static efficiency, we first conclude that $V_\tau(\vartheta + |\vartheta|) \geq V_\tau(|\vartheta|)$ P-a.s. and hence $V_\tau(\vartheta) \geq 0$ P-a.s. So $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$ because $\tau \in \mathcal{T}_{[\sigma, T]}$ was arbitrary.

“(b) \Rightarrow (c)”. Seeking a contradiction, suppose there is $\sigma \in \mathcal{T}_{[0,T]}$ such that 0 is not strongly maximal for $\mathbf{h}L_\sigma^{\text{sf}}$. Then there is $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ such that $V_\sigma(\vartheta) \leq 0$ P-a.s. and $V_T(\vartheta) \geq 0$ P-a.s., where the second inequality is strict with positive probability. Now (b) gives that $\vartheta \in \mathbf{h}\mathcal{U}_\sigma$, and so 0 fails to be weakly (and a fortiori strongly) maximal for $\mathbf{h}\mathcal{U}_\sigma$, in contradiction to static viability of \mathcal{S} .

“(c) \Rightarrow (c’)”. This is trivial.

“(c’) \Rightarrow (b’)”. Fix $s \in [0, T)$ and take $\vartheta \in \mathbf{h}L_s^{\text{sf}}$ with $V_T(\vartheta) \geq 0$ P-a.s. If ϑ is not in $\mathbf{h}\mathcal{U}_s$, by right-continuity of the paths of $V(\vartheta)(S)$, there exists $r \in (s, T)$ such that $\mathbb{P}[V_r(\vartheta) < 0] > 0$. Let η be a numéraire strategy in $\mathbf{h}L_r^{\text{sf}}$, e.g. $\eta = \eta^S$, and set $\tilde{\vartheta} := (\vartheta + |V_r(\vartheta)(S^{(\eta)})|\eta)\mathbf{1}_{\{V_r(\vartheta)(S^{(\eta)}) < 0\}}\mathbf{1}_{[r, T]} \in \mathbf{h}L_r^{\text{sf}}$. Using that $V_T(\vartheta) \geq 0$ P-a.s., we get

$$V_r(\tilde{\vartheta}) = 0 \text{ P-a.s.} \quad \text{and} \quad V_T(\tilde{\vartheta}) \geq |V_r(\vartheta)(S^{(\eta)})|V_T(\eta)\mathbf{1}_{\{V_r(\vartheta)(S^{(\eta)}) < 0\}} \text{ P-a.s.}$$

Thus, $V_T(\tilde{\vartheta}) \geq 0$ P-a.s. and $\mathbb{P}[V_T(\tilde{\vartheta}) > 0] > 0$, in contradiction to strong maximality of 0 for $\mathbf{h}L_r^{\text{sf}}$.

“(b’) \Rightarrow (b)”. Note that the second statement in (b’) trivially also holds for $s = T$. Fix $\sigma \in \mathcal{T}_{[0,T]}$ and take $\vartheta \in \mathbf{h}L_\sigma^{\text{sf}}$ with $V_T(\vartheta) \geq 0$ P-a.s. First, let

$\tau \in \mathcal{T}_{[\sigma, T]}$ be of the form $\tau = \sum_{i=1}^n t_i \mathbf{1}_{A_i}$, where $n \in \mathbb{N}$, $0 \leq t_1 < \dots < t_n \leq T$, and $(A_i)_{i \in \{1, \dots, n\}}$ is a partition of Ω with $A_i \in \mathcal{F}_{t_i}$. For $i \in \{1, \dots, n\}$, set $\vartheta^{(i)} := \mathbf{1}_{A_i} \vartheta \in \mathbf{h}L_{\tau}^{\text{sf}} \cap \mathbf{h}L_{t_i}^{\text{sf}}$ so that $\vartheta = \sum_{i=1}^n \vartheta^{(i)}$. Since each $V_{t_i}(\vartheta^{(i)}) \geq 0$ \mathbb{P} -a.s. by (b'), we have $V_{\tau}(\vartheta) \geq 0$ \mathbb{P} -a.s. For general $\tau \in \mathcal{T}_{[\sigma, T]}$, there is a nonincreasing sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times in $\mathcal{T}_{[0, T]}$ each taking only finitely many values such that $\lim_{n \rightarrow \infty} \tau_n = \tau$. Since each $V_{\tau_n}(\vartheta) \geq 0$ \mathbb{P} -a.s. by the first part of the argument, right-continuity of the paths of $V(\vartheta)(S)$ yields $V_{\tau}(\vartheta) \geq 0$ \mathbb{P} -a.s. So $\vartheta \in \mathbf{h}\mathcal{U}_{\sigma}$ because $\tau \in \mathcal{T}_{[\sigma, T]}$ was arbitrary.

“(c) \Rightarrow (d)”. This is trivial.

“(d) \Rightarrow (e)”. This follows from “(a) \Rightarrow (c)” in Lemma 6.1.

“(e) \Rightarrow (a)”. This is clear from $\mathbf{h}\mathcal{U}_{\sigma} \subseteq \mathbf{h}L_{\sigma}^{\text{sf}}$, for each $\sigma \in \mathcal{T}_{[0, T]}$. □

Remark 6.3. The equivalence of (c) and (d) in Lemma 6.2 shows that (c) is in fact equivalent to the same statement *without* assuming static viability of \mathcal{S} . However, this is not true for statement (c').

Bibliography

- [1] J.-P. Ansel and C. Stricker, *Couverture des Actifs Contingents et Prix Maximum*, Ann. Inst. H. Poincaré Probab. Statist. **30** (1994), 303–315.
- [2] D. Becherer, *The Numeraire Portfolio for Unbounded Semimartingales*, Finance Stoch. **5** (2001), 327–341.
- [3] M. Beiglböck, W. Schachermayer, and B. Veliyev, *A Direct Proof of the Bichteler–Dellacherie Theorem and Connections to Arbitrage*, Ann. Prob. **39** (2011), 2424–2440.
- [4] F. Biagini, H. Föllmer, and S. Nedelcu, *Shifting Martingale Measures and the Birth of a Bubble as a Submartingale*, Finance Stoch. **18** (2014), 297–326.
- [5] N. H. Bingham and R. Kiesel, *Risk-Neutral Valuation*, second ed., Springer Finance, Springer, London, 2004.
- [6] M. K. Brunnermeier, *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis, and Herding*, Oxford University Press, 2001.
- [7] P. Carr, T. Fisher, and J. Ruf, *On the Hedging of Options on Exploding Exchange Rates*, Finance Stoch. **18** (2014), 115–144.
- [8] A. M. G. Cox and D. G. Hobson, *Local Martingales, Bubbles and Option Prices*, Finance Stoch. **9** (2005), 477–492.
- [9] F. Delbaen and W. Schachermayer, *A General Version of the Fundamental Theorem of Asset Pricing*, Math. Ann. **300** (1994), 463–520.
- [10] ———, *The Existence of Absolutely Continuous Local Martingale Measures*, Ann. Appl. Probab. **5** (1995), 926–945.
- [11] ———, *The No-Arbitrage Property under a Change of Numéraire*, Stoch. Stoch. Rep. **53** (1995), 213–226.
- [12] ———, *The Banach Space of Workable Contingent Claims in Arbitrage Theory*, Ann. Inst. H. Poincaré Probab. Statist. **33** (1997), 113–144.
- [13] ———, *The Fundamental Theorem of Asset Pricing for Unbounded Stochastic Processes*, Math. Ann. **312** (1998), 215–250.

- [14] ———, *The Mathematics of Arbitrage*, Springer Finance, Springer, Berlin, 2006.
- [15] B. T. Diba and H. I. Grossman, *On the Inception of Rational Bubbles*, Q. J. Econ. **102** (1987), 697–700.
- [16] E. Ekström and J. Tysk, *Bubbles, Convexity and the Black–Scholes Equation*, Ann. Appl. Probab. **19** (2009), 1369–1384.
- [17] N. El Karoui, H. Geman, and J.-C. Rochet, *Changes of Numéraire, Changes of Probability Measure and Option Pricing*, J. Appl. Probab. **32** (1995), 443–458.
- [18] R. J. Elliott and P. E. Kopp, *Mathematics of Financial Markets*, second ed., Springer Finance, Springer-Verlag, New York, 2005.
- [19] D. C. Emanuel and J. D. MacBeth, *Further Results on the Constant Elasticity of Variance Call Option Pricing Model*, J. Financ. Quant. Anal. **17** (1982), 533–554.
- [20] M. Emery, *Une Topologie sur l'Espace des Semimartingales*, Séminaire de Probabilités, XIII, Lecture Notes in Math., vol. 721, Springer-Verlag, Berlin, 1979, p. 260–280.
- [21] R. Fernholz, *Stochastic Portfolio Theory*, Applications of Mathematics, vol. 48, Springer-Verlag, New York, 2002.
- [22] R. Fernholz, I. Karatzas, and C. Kardaras, *Diversity and Relative Arbitrage in Equity Markets*, Finance Stoch. **9** (2005), 1–27.
- [23] M. Frittelli and G. Scandolo, *Risk Measures and Capital Requirements for Processes*, Math. Finance **16** (2006), 589–612.
- [24] H. Föllmer and Y. M. Kabanov, *Optional Decomposition and Lagrange Multipliers*, Finance Stoch. **2** (1998), 69–81.
- [25] H. Föllmer and A. Schied, *Stochastic Finance*, second ext. ed., de Gruyter Studies in Mathematics, vol. 27, Walter de Gruyter & Co., Berlin, 2004.
- [26] C. Gourieroux, J. P. Laurent, and H. Pham, *Mean-Variance Hedging and Numéraire*, Math. Finance **8** (1998), 179–200.
- [27] P. Guasoni and M. Rasonyi, *Fragility of Arbitrage and Bubbles in Diffusion Models*, Forthcoming in Finance Stoch.
- [28] D. Heath and E. Platen, *A Benchmark Approach to Quantitative Finance*, vol. 13, Springer, 2006.
- [29] P. Henry-Labordère, *SABR model*, In Encyclopedia of Quantitative Finance, R. Cont, ed., Wiley, 2010.

- [30] M. Herdegen, *A Numéraire Independent Modelling Framework for Financial Markets*, NCCR FINRISK working paper No. 741, ETH Zürich (2012), Available at: http://www.nccr-finrisk.uzh.ch/media/pdf/wp/WP741_D1.pdf.
- [31] ———, *No-Arbitrage in a Numéraire Independent Modelling Framework*, NCCR FINRISK working paper No. 775, ETH Zürich (2012), Available at: http://www.nccr-finrisk.uzh.ch/media/pdf/wp/WP775_D1.pdf.
- [32] ———, *No-Arbitrage in a Numéraire-Independent Modeling Framework*, Forthcoming in *Math. Finance*.
- [33] M. Herdegen and S. Herrmann, *A Class of Strict Local Martingales*, Swiss Finance Institute Research Paper No. 14-18 (2014), Available at SSRN: <http://ssrn.com/abstract=2402248>.
- [34] S. L. Heston, M. Loewenstein, and G. A. Willard, *Options and Bubbles*, *Rev. Financ. Stud.* **20** (2007), 359–390.
- [35] H. Hulley, *Strict Local Martingales in Continuous Financial Market Models*, Ph.D. thesis, University of Technology, Sydney, 2009.
- [36] J. Jacod and A. N. Shiryaev, *Local Martingales and the Fundamental Asset Pricing Theorems in the Discrete-Time Case*, *Finance Stoch.* **2** (1998), 259–273.
- [37] ———, *Limit Theorems for Stochastic Processes*, second ed., *Grundlehren der Mathematischen Wissenschaften*, vol. 288, Springer, Berlin, 2003.
- [38] J. Jacod, *Calcul Stochastique et Problèmes de Martingales*, *Lecture Notes in Mathematics*, vol. 714, Springer, Berlin, 1979.
- [39] R. A. Jarrow and M. Larsson, *The Meaning of Market Efficiency*, *Math. Finance* **22** (2012), 1–30.
- [40] R. A. Jarrow, P. Protter, and K. Shimbo, *Asset Price Bubbles in Complete Markets*, In *Advances in Mathematical Finance*, *Appl. Numer. Harmon. Anal.*, M. C. Fu, R. A. Jarrow, J.-Y. J. Yen, and R. J. Elliott, eds., Birkhäuser, Boston, 2007, p. 97–121.
- [41] ———, *Asset Price Bubbles in Incomplete Markets*, *Math. Finance* **20** (2010), 145–185.
- [42] Y. M. Kabanov, *On the FTAP of Kreps-Delbaen-Schachermayer*, In *Statistics and Control of Stochastic Processes: The Liptser Festschrift*, Y. M. Kabanov, B. L. Rozovskii, and A. N. Shiryaev, eds., World Sci. Publ., River Edge, NJ, 1997, p. 191–203.
- [43] T. Kaizoji and D. Sornette, *Bubbles and crashes*, In *Encyclopedia of Quantitative Finance*, R. Cont, ed., Wiley, 2010.

- [44] O. Kallenberg, *Foundations of Modern Probability*, second ed., Probability and its Applications, Springer, New York, 2002.
- [45] J. Kallsen, σ -Localization and σ -Martingales, *Theory Probab. Appl.* **48** (2004), 152–163.
- [46] I. Karatzas and C. Kardaras, *The Numéraire Portfolio in Semimartingale Financial Models*, *Finance Stoch.* **11** (2007), 447–493.
- [47] C. Kardaras, *Finitely Additive Probabilities and the Fundamental Theorem of Asset Pricing*, In *Contemporary Quantitative Finance*, C. Chiarella, and A. Novikov, eds., Springer, Berlin, 2010, p. 19–34.
- [48] ———, *Market Viability via Absence of Arbitrage of the First Kind*, *Finance Stoch.* **16** (2012), 651–667.
- [49] C. Kardaras and E. Platen, *On the Semimartingale Property of Discounted Asset-Price Processes*, *Stoch. Proc. Appl.* **121** (2011), 2678–2691.
- [50] A. Klenke, *Probability Theory*, Springer, London, 2008.
- [51] S. Klöppel and M. Schweizer, *Dynamic Indifference Valuation via Convex Risk Measures*, *Math. Finance* **17** (2007), 599–627.
- [52] D. O. Kramkov, *Optional Decomposition of Supermartingales and Hedging Contingent Claims in Incomplete Security Markets*, *Probab. Theory Relat. Fields* **105** (1996), 459–479.
- [53] D. Lamberton and B. Lapeyre, *Introduction to Stochastic Calculus Applied to Finance*, second ed., Chapman & Hall/CRC Financial Mathematics Series, CRC Press, Boca Raton, FL, 2008.
- [54] S. Levental and A. V. Skorokhod, *A Necessary and Sufficient Condition for Absence of Arbitrage with Tame Portfolios*, *Ann. Appl. Probab.* **5** (1995), 906–925.
- [55] A. L. Lewis, *Option Valuation under Stochastic Volatility*, Finance Press, Newport Beach, CA, 2000.
- [56] M. Loewenstein and G. A. Willard, *Local Martingales, Arbitrage, and Viability. Free Snacks and Cheap Thrills*, *Econ. Theory* **16** (2000), 135–161.
- [57] ———, *Rational Equilibrium Asset-Pricing Bubbles in Continuous Trading Models*, *J. Econ. Theory* **91** (2000), 17–58.
- [58] ———, *Consumption and Bubbles*, *J. Econ. Theory* **148** (2013), 563–600.
- [59] J. B. Long, Jr., *The Numeraire Portfolio*, *J. Financ. Econ.* **26** (1990), 29–69.

- [60] D. B. Madan and M. Yor, *Ito's Integrated Formula for Strict Local Martingales*, In memoriam Paul-André Meyer: Séminaire de Probabilités XXXIX, Lecture Notes in Math., vol. 1874, Springer, Berlin, 2006, p. 157–170.
- [61] W. Margrabe, *The Value of an Option to Exchange one Asset for Another*, J. Finance **33** (1978), 177–186.
- [62] R. C. Merton, *Theory of Rational Option Pricing*, Bell J. Econ. Manag. Sci. **4** (1973), 141–183.
- [63] J. Mémin, *Espaces de Semi Martingales et Changement de Probabilité*, Z. Wahrsch. verw. Gebiete **52** (1980), 9–39.
- [64] S. Pal and P. Protter, *Analysis of Continuous Strict Local Martingales via h -Transforms*, Stoch. Proc. Appl. **120** (2010), 1424–1443.
- [65] E. Platen, *A Benchmark Approach to Finance*, Math. Finance **16** (2006), 131–151.
- [66] P. Protter, *Stochastic Integration and Differential Equations*, second ed., Applications of Mathematics, vol. 21, Springer, New York, 1990.
- [67] ———, *A Mathematical Theory of Financial Bubbles*, Lecture Notes in Mathematics, vol. 2081, Springer, New York, 2013, p. 1–108.
- [68] S. Rady, *Option Pricing in the Presence of Natural Boundaries and a Quadratic Diffusion Term*, Finance Stoch. **1** (1997), 331–344.
- [69] D. Revuz and M. Yor, *Continuous Martingales and Brownian Motion*, correct. third print. of the third ed., Grundlehren der mathematischen Wissenschaften, vol. 293, Springer, Berlin, 2005.
- [70] J. Ruf, *Hedging under Arbitrage*, Math. Finance **23** (2013), 297–317.
- [71] M. S. Santos and M. Woodford, *Rational Asset Pricing Bubbles*, Econometrica **65** (1997), 19–57.
- [72] A. Scherbina and B. Schlusche, *Asset Price Bubbles: A Survey*, Quant. Finance **14** (2014), 589–604.
- [73] M. Schweizer and K. Takaoka, *A Note on the Condition of No Unbounded Profit with Bounded Risk*, Finance Stoch. **18** (2014), 393–405.
- [74] G. R. Sell and Y. You, *Dynamics of Evolutionary Equations*, Applied Mathematical Sciences, vol. 143, Springer, New York, 2002.
- [75] R. J. Shiller, *Irrational Exuberance*, Princeton Univ. Press, Princeton, NJ, 2000.
- [76] C. A. Sin, *Complications with Stochastic Volatility Models*, Adv. Appl. Probab. **30** (1998), 256–268.

- [77] E. Strasser, *Characterization of Arbitrage-Free Markets*, Ann. Appl. Probab. **15** (2005), 116–124.
- [78] C. Stricker, *Simple Strategies in Exponential Utility Maximization*, Séminaire de Probabilités, XXXVI, Lecture Notes in Math., vol. 1801, Springer, Berlin, 2003, p. 415–418.
- [79] J. Tirole, *On the Possibility of Speculation under Rational Expectations*, Econometrica **50** (1982), 1163–1181.
- [80] J. Vecer, *Stochastic Finance: A Numeraire Approach*, Chapman & Hall/CRC Financial Mathematics Series, CRC Press, Boca Raton, FL, 2011.
- [81] J. Xia and J.-A. Yan, *A New Look at Some Basic Concepts in Arbitrage Pricing Theory*, Sci. China Ser. A **46** (2003), 764–774.
- [82] J.-A. Yan, *A New Look at the Fundamental Theorem of Asset Pricing*, J. Korean Math. Soc. **35** (1998), 659–673.
- [83] ———, *A Numeraire-Free and Original Probability Based Framework for Financial Markets*, In Proceedings of the International Congress of Mathematicians, Vol. III, T. Li, ed., Higher Ed. Press, Beijing, 2002, p. 861–871.

Curriculum Vitae

Martin Peter Gerhard Herdegen
born April 10, 1986
German citizen

Education

PhD in Mathematics ETH Zürich, Switzerland	11/2009–07/2014
MSc in Mathematical Sciences University of Bath, United Kingdom	09/2008–09/2009
Vordiplom in Wirtschaftsmathematik LMU München, Germany	04/2006–08/2008
Abitur at Franz-Ludwig-Gymnasium Bamberg, Germany	09/1996–06/2005

Academic Employment

Exercise coordinator in Mathematics ETH Zürich, Switzerland	09/2010–07/2014
Tutorial assistant in Mathematics University of Bath, United Kingdom	09/2008–05/2009
Tutorial assistant in Mathematics LMU München, Germany	04/2008–07/2008