TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

DISSERTATION

Superreplication and Arbitrage in Multiasset Models under Proportional Transaction Costs

ausgeführt zum Zwecke der Erlangung des akademischen Grades eines Doktors der Naturwissenschaften unter der Leitung von

O. Univ.-Prof. Mag. Dr. Walter Schachermayer

Fakultät für Mathematik
Universität Wien

eingereicht an der Technischen Universität Wien
bei der Fakultät für Mathematik und Geoinformation

von

Dipl.-Math. Oec. Benedikt Frank Korbinian Blum
Matrikelnummer: 0827358
Krongasse 4/5, 1050 Wien

Wien, im Dezember 2010

Kurzfassung

Diese Dissertation widmet sich der Untersuchung von Transaktionskosten in zeitstetigen Finanzmarktmodellen und ihrem Einfluss auf bekannte Ergebnisse aus der Finanzmathematik. Abgesehen von der bloßen Tatsache, dass sie in praktisch jedem existenten Finanzmarkt beobachtet werden können, ist die Berücksichtigung von Transaktionskosten auch aus indirekten Gründen lohnenswert: Zum einen gibt es Modelle, die an sich Arbitrage zulassen, durch die Hinzunahme von Transaktionskosten aber arbitragefrei werden. Auf diese Weise erweitern Transaktionskosten also den Spielraum für die Modellierung von Preisprozessen. Zum anderen verbieten sich besonders aufwändige Strategien, die in der Praxis unmöglich umgesetzt werden können, durch Transaktionskosten von selbst, was bei den üblichen Hedging- und Investitionsproblemen zu realistischeren Lösungen führen kann.

Wir präsentieren ein Modell für den Handel auf mehrdimensionalen Märkten mit proportionalen Transaktionskosten. Dieses Modell kann als Verallgemeinerung von eindimensionalen Modellen gesehen werden, die in den letzten Jahren vorgeschlagen wurden. Der Preisprozess muss hierbei nur rechtsseitig stetig sein und linke Grenzwerte aufweisen (càdlàg), aber kein Semimartingal sein. Wir beweisen grundlegende Eigenschaften des Modells, beleuchten die Gemeinsamkeiten mit alternativen Ansätzen und gelangen so zu einem Satz über Superreplikation, der es erlaubt zu entscheiden, welche Contingent Claims trotz Transaktionskosten ohne Startkapital repliziert werden können. Wie bereits in früheren Veröffentlichungen gezeigt wurde, spielen sogenannte *Consistent Price Systems* hier eine zentrale Rolle.

Wir fahren mit dem Beweis einer mehrdimensionalen Version des sogenannten *Face-lifting Theorem* fort. Dieser Satz stellt eine Verbindung zwischen Superreplikationspreisen pfadunabhängiger Optionen und der Gestalt der zugrundeliegenden Payoff-Funktion her. Abgesehen von einer Verallgemeinerung eines entsprechenden eindimensionalen Ergebnisses sind wir in der Lage, dort getroffene Annahmen an den Preisprozess und die Payoff-Funktion deutlich abzuschwächen. Dieses Ergebnis verdeutlicht eine entscheidende Eigenschaft von

Transaktionskostenmodellen: Auch wenn der Grad der Kosten gegen Null tendiert, wie es bei großen Marktteilnehmern auf besonders liquiden Märkten der Fall ist, konvergieren manche Ergebnisse nicht gegen die entsprechenden transaktionskostenfreien Ergebnisse.

Das wichtigste Resultat dieser Dissertation ist eine mehrdimensionale Version des *Fundamental Theorem of Asset Pricing* für zeitstetige Modelle mit kleinen, proportinalen Transaktionskosten. Wir zeigen, dass solche Modelle genau dann für alle Kostenlevels arbitragefrei sind, wenn sie auch für alle Levels Consistent Price Systems zulassen. Die dabei verwendete Arbitrage-Definition ist explizit, also nicht von der Art eines *free lunch*.

Sowohl das Face-lifting- als auch das Fundamental Theorem werden im Spezialfall eindimensionaler exponentieller Lévy-Prozesse noch genauer untersucht. In diesem Fall ist nicht nur eine komplette Charakterisierung möglich, welche Prozesse in Abhängigkeit von ihrem Lévy-Chintschin-Tripel welche Consistent Price Systems zulassen, sondern es kann auch eine Version des Fundamental Theorem ohne die „für alle"-Quantoren bewiesen werden.

Zuletzt wird noch eine Verallgemeinerung des Transaktionskostenmodells vorgestellt, die sich hauptsächlich, aber nicht ausschließlich an Devisenmärkte richtet: Hier ist ein direkter Kaptialtransfer zwischen den Assets möglich, anstatt wie sonst das Kapital stets über ein Barkonto zu leiten. Auch können hier verschiedene Kostenlevels für Kauf und Verkauf oder auch nur einseitige Transaktionskosten berücksichtigt werden. Das Fundamental Theorem of Asset Pricing wird auf dieses Modell verallgemeinert.

Abstract

This thesis is dedicated to the study of transaction costs in continuous-time financial market models and their impact on well-known results from mathematical finance. Apart from the mere fact that they can be observed in virtually every financial market, transaction costs deserve consideration also for indirect reasons: First of all, certain models that allow for arbitrage in the frictionless case become arbitrage-free through the introduction of costs, thus we can widen the scope of processes to be considered in modelling. Furthermore, the presence of costs disallows especially extensive trading strategies that can impossibly be reproduced in practice, which can lead to more realistic solutions of common hedging and investment problems.

We present a model for trading on multidimensional continuous-time markets under proportional transaction costs. This model can be seen as generalization of one-dimensional models proposed in recent years. Our model is open for all right-continuous processes with finite left limits (càdlàg), hence does not require semimartingales. We prove basic properties of our model, highlight the parallels to other approaches and arrive at a crucial super-replication theorem, that allows us to decide which contingent claims can be superhedged without initial endowment by trading on the market under transaction costs. For this, the notion of *consistent price systems* plays a crucial role, as was shown in several earlier publications.

We resume by proving a multiasset version of the so-called *face-lifting theorem*, a result that links superreplication prices of path-independent options to the shape of the underlying payoff function. Apart from generalizing a predecessing result to multiple dimensions, we are able to significantly weaken the assumptions made on the process as well as the payoff function. This result highlights a critical feature of transaction costs models: Even if the level of costs tends to zero, as is the case for very large traders on highly liquid markets, some results do not converge to the corresponding results from the frictionless case.

The most important result of this thesis is a multidimensional version of the *Fundamental Theorem of Asset Pricing* for continuous models under small proportional transaction costs. We prove that such a model is arbitrage-free for all cost levels if and only if it admits consistent price systems for all levels. The notion of arbitrage used here is explicit, i. e. not of *free-lunch*-type.

Both the face-lifting theorem and the Fundamental Theorem of Asset Pricing are revisited in the special case of one-dimensional exponential Lévy processes. In this case we can give a complete characterization of which processes admit which consistent price systems, depending only on the Lévy-Khinchine triplet of the process. We also prove a version of the Fundamental Theorem without the “for all”-quantors.

Finally, we propose a generalized transaction costs regime for multiasset models that is aimed mainly, but not exclusively towards currency markets: Here, direct exchange of capital between the assets is possible instead of always funneling one’s capital through the cash account. The generalized model furthermore allows to assign different cost levels to buying and selling or to consider one-sided transaction costs. The Fundamental Theorem of Asset Pricing is extended to our generalized model.

Acknowledgements

Over more than the last decade, I spent a great portion of my time reading, writing, learning and teaching mathematics. These years have been full of challenging, but also satisfying moments. Not always seemed there to be a light at the end of the tunnel, and I know I wouldn't be where I am now if it hadn't been for many great teachers and mentors. Out of these, three extraordinary persons stand out, which happen to be the advisors of my three theses. I am thankful for their guidance, their patience and for their giant shoulders to stand on.

The first is StDir. i. R. Hans Andraschko, who was my teacher from 1997 to 2000 and the first to show me the beauty behind all the annoying formulas. Without him I hardly would have decided to study mathematics in the first place, and also my interest for stochastics is owed to him.

The second is Prof. Dr. Jan Kallsen, from whom I learned most of what I know about probability theory and stochastic analysis (actually, I think I learned more from him than I know today). He supervised my diploma thesis and reminded me that before enjoying the beauty behind, the annoying formulas have to be dealt with.

Finally, I am deeply indebted to Prof. Dr. Walter Schachermayer, the advisor of this thesis and my mentor for the last three and a half years. He introduced me to the exciting topic of arbitrage, shared his profound knowledge with me and was a good counselor in good times as well as in bad times. If you can answer all questions that I have been thinking about for weeks in less than a minute, you get a pass on not reading too much of what I sent you for years.

Also, I owe a lot to all my colleagues at the research unit for financial and actuarial mathematics (FAM) at the Vienna University of Technology. They have showed me that the self-denomination "FAMily" is not just a lame pun. It is an attitude. And a lame pun. Through their help an guidance inside and outside of our university, they have become true friends of mine. And in one instance, even more than that. Barbara, this one's for you.

I gratefully acknowledge the financial support from UniCredit Bank Austria and the Christian Doppler Research Association via the Christian Doppler Laboratory for Portfolio Risk Management, and the fruitful collaboration with my partners at Bank Austria. Without both this work would not have been possible.

I would also like to thank the facility maintenance unit (Gebäude und Technik) of the technical university, for the rare opportunity to disprove the great David Hilbert: In three years, they have not managed to repair the leaking ceiling above my desk, thereby verifying that, in fact, sometimes a math department *is* a swimming pool.

Last but foremost, I am deeply indebted to my family for their love, their endless support and for all those fancy first names. Having a mathematician in your family is certainly not the easiest fate, but my parents, my sister and my brother have always given me much more than what you can expect of a family, and more than I can ever pay back. The same holds for my friends, but at least them I might be able to pay back. As usual, the question is not how, but rather where.

Contents

Optima autem aleae commoditas est non ludere

Gerolamo Cardano (1501-1576)
Liber de ludo aleae, caput IV

Daß aber ein Mathematiker, aus dem Hexengewirre seiner Formeln heraus, zur Anschauung der Natur käme und Sinn und Verstand, unabhängig, wie ein gesunder Mensch brauchte, werd ich wohl nicht erleben.

Johann Wolfgang von Goethe (1749-1832)
Brief an Carl Friedrich Zelter, 17. Mai 1829

Chapter 1

Introduction

As long as mathematical models have were established to describe and predict real-world phenomena, they were criticized for being over-simplified and over-idealized. While a certain degree of idealization lies in the nature of modelling and a perfect model that considered each imaginable feature would be as useless as a world map of scale 1:1, every such criticism is justified as long as it may lead to an improvement of the model's predictions without making it unfeasible. In physics and engineering for instance, the incorporation of friction improves a model's prediction quality, but usually drastically complicates the model.

In continuous-time financial mathematics, this can be best observed in the case of the *Black-Scholes* model, which has, since it was first proposed by Samuelson [Sam65], established itself as the undoubted benchmark model and changed the whole industry, as was pointed out in the justification of the Nobel Foundation for awarding the 1997 Nobel prize in economics to F. Black and M. Scholes, two of the main contributors (along with R. Merton) of the theory behind the model [N98]. The model was the first to offer a convincing principle to find unique option prices based on the argument of *no arbitrage*.

The standard criticisms on the Black-Scholes lie in its central assumptions: Volatility is in practice neither deterministic nor constant, rather is it random and depends on the stock price. Normally distributed logarithmic returns and the absence of jumps were criticized for underestimating rare extreme events as early as 1965 [F65]. To counteract these shortcomings, processes with jumps, stochastic volatility and stochastic time change were proposed. However, further development of the theory soon made it clear that any non-trivial extension of Black-Scholes would mean to give up the crucial feature of completeness: In no model with more degrees of freedom all options can be replicated by an adequate trading strategy, hence there no more exist unique option prices, one of the main reasons for the success of Black-Scholes. However, the *Fundamental Theorem of Asset Pricing* (FTAP), first introduced in three seminal works of Harrison, Kreps and Pliska ([HK79], [HP81], [Kr81]) and further refined and generalized over the years (see [DS06]), still tells us exactly under which conditions the absence of arbitrage can be retrieved – as sort of a minimal standard for a reasonable market model.

Another set of criticism addresses not the way the price process is designed, but rather the whole fashion in which trading is performed: In classical finance one assumes that each market agent has the same information and can buy arbitrarily large and small amounts of the stock at each time at a fixed price, without paying any kind of transaction costs. In practice, asymmetric information and all kinds of trading restrictions are common. Prohibition of short sales, the restriction to integer quantities of assets and time-lags in order processing are common and can be ignored, if by anyone, only by large traders with the adequate infrastructure, acting on large markets. Transaction costs, commonly referred to as *market friction* due to their similarity to the corresponding physical phenomenon (as in: loss of energy/value caused by movement), are arguably the most convincing restriction, for a whole set of reasons.

First, they affect virtually every financial market there is, be it equity, fixed income, currencies or bonds. In every market, the trader faces not one fixed price, but rather a bid-ask spread: a lower price that one revenues for selling, and a higher price one has to pay for buying. This effect is smaller in highly liquid markets and bigger in markets with fewer frequent transactions, but never fully vanishes. Especially large agents that can neglect the restrictions on trading small quantities face this problem when buying large portions: If the supply is increasingly bought up, then the spread widens. The same happens when large quantities are sold and the market is flooded. This effect is called impact or liquidity costs.

Second, a typical trader also faces actual monetary costs of trading. Private or small investors have to pay a fee to their broker, while institutional traders still have to pay for their staff and the trading infrastructure, that they could have avoided by trading less ot not at all. These costs may become very small for large, highly automatized agents, but never really go down to zero.

A third source for transaction costs are all kinds of taxes on financial transactions. These taxes, that were common worldwide throughout most of the twentieth century, were successively reduced or abolished in most countries in the 1980s and 1990s, but never completely vanished and have even made a comeback to several political agendas in times where many national economies face severe budget crises. In particular, the idea of a transaction tax on the especially liquid currency market, coined *Tobin tax* after its most prominent advocate J. Tobin [Tob78], gains popularity among scholars, non governmental organizations and also governments. It was proposed as early as 1936 by J. M. Keynes to stabilize exchange rates and discourage undesired speculations [Ke36]. Such a tax, once introduced to a market, would halt any argument that transaction costs vanish for large traders and can therefore be neglected.

Apart from the mere fact that they exist in practice, there are other, rather mathematical arguments that suggest considering transaction costs in financial market models. As was pointed out by Magill and Constantinides [MC76], a typical solution to a hedging or optimal investment problem in the Back-Scholes or a similar continuous time model is a trading strategy of infinite variation. So to follow this strategy, the trader would have to make infinitely many trades and move around an infinite amount of money in each time interval – obviously an absurd conclusion that would make continuous-time models unfeasible altogether. By introducing even the smallest transaction costs, they conclude, one can save these models, since infinite-variation strategies are automatically dismissed as they would bankrupt the trader immediately.

Another argument deals with processes that are candidates for price processes because of favourable modelling properties but fail to be arbitrage free and are therefore dismissed. The most prominent example is fractional Brownian motion, proposed by Mandelbrot [Ma68]. The process fails to be a semimartingale and therefore admits free lunch [DS06, Theorem 9.7.2]. However, it was shown by Guasoni [Gu06] that, if arbitrarily small transaction costs are introduced, the arbitrage opportunity disappears and the process becomes eligible for modelling.

More general: Outside the realm of semimartingales, there lacks a consistent theory of stochastic integration, which is indispensable for mathematical finance. But with transaction costs restricting us to finite variation strategies, integration is possible even in a pathwise Lebesgue-Stieltjes sense.

As a related argument, it was pointed out by Guasoni et al. [GRS08a] that without transaction costs, important issues like arbitrage and option hedging depend critically on high frequency properties of the process such as quadratic variation and other small scale path properties that can hardly be observed and play no role in actual trading. Instead, *coarse* properties like supports become more relevant, that can be properly observed and tested. So incorporation of transaction costs can somehow decrease model risk, even to the extent that some results become virtually model-independent, like our version of the face-lifting theorem, see Chapter 3.

1.1 Outline

Much effort has been put into translating fundamental results from frictionless finance to markets with transaction costs. Originally the main focus was on problems of optimal consumption and investment ([MC76], [DN90]) and option hedging ([Le85], [DC94], [LS97]), while the question of no-arbitrage criteria was not addressed until 1995, presumably for a simple reason: The Fundamental Theorem of Asset Pricing [DS06] delivers the answer in the frictionless case, so simply assuming an equivalent martingale measure leads to sufficient conditions in the transaction costs case, which is enough to work with in most cases. This way, the question of necessary conditions for a market to be arbitrage-free under transaction costs remained unanswered. The first attempt by Jouini and Kallal [JK95] relied heavily on an L^2 framework and was limited to simple strategies, while Schachermayer [Scha04] and Guasoni et al. [GRS08b] address the multidimensional case in discrete time and the one-dimensional case in continuous time, respectively. One goal of this work is to establish a Fundamental Theorem of Asset Pricing (that is: necessary *and* sufficient conditions for a model to be arbitrage free) for continuous d-dimensional processes in continuous time.

Another subject that this work tries to cover originates in a conjecture made by Davis and Clark [DC94] about the superreplication price of a European call option. They suspected that under transaction costs, this price could not be lower than the price of the stock. This conjecture has been proven in increasing generality ([CSS95], [LS97], [CPT99], [CK96]) and subsequently led to a general pricing principle now known as *face-lifting pricing*. We attempt to generalize the corresponding theorem from Guasoni et al. [GRS08a] to multiple dimensions and additionally weaken the assumptions made to both the price process and the option to be superreplicated.

More general, this work attempts to establish a uniform and consistent framework for trading on multidimensional asset processes under proportional transaction costs in continuous time. We try to bring together models proposed so far, highlight their parallels, but also the differences. We explicitly do not assume the price process to be a semimartingale, therefore also covering fractional Brownian motion and related processes.

This work is organized as follows: In Section 2 we present a simple multidimensional trading framework, along with the basic definitions of trading, admissibility and consistent price systems. We further highlight an alternative model proposed by Kabanov [Ka99] and develop the connections between the models necessary to translate a crucial superreplication theorem into our framework.

Section 3 is devoted to the above-mentioned face-lifting theorem. We prove a general multidimensional version and present some examples to illustrate its implications, but also its limitations.

In Section 4 we highlight the special case of one-dimensional exponential Lévy models. Here we present sufficient conditions for the face-lifting theorem to hold and fully answer the question of the existence of consistent price systems for these processes. This serves as preparatory work for a Lévy version of the Fundamental Theorem of Asset Pricing.

Section 5 is devoted to a multidimensional version of the Fundamental Theorem of Asset Pricing in the spirit of [GRS08b]. We present several notions of arbitrage and prove the Fundamental Theorem both for d-dimensional continuous and for one-dimensional exponential Lévy processes. As a side result, we prove the equivalence between arbitrage and free lunch with bounded risk, as was done in [GRS08b] in the one-dimensional case.

Finally, we extend our framework to a more general principle of transaction costs, similar to [Ka99]. The Fundamental Theorem is also extended to this kind of transaction costs.

A general result on the d-dimensional Esscher transform is proved in the appendix. It is used for the proof of the Fundamental Theorem, but may also be of independent interest.

1.2 Historical overview

In this section we want to give a short overview over the development of transaction cost models in finance. It is not possible to fully trace back the term *transaction costs* to where it first evolved, but the idea seems to be present already in the earliest post-depression economics literature, where market friction is used to explain why perfect-economy models fail to predict real-world outcomes. One of the first backed sources where this phenomenon is studied is the later Nobel laureate R. Coase, who described *costs of using the price mechanism* [Co37] and later discussed *costs of market transactions* in greater detail [Co60]. In the analysis of the cash demand of economies, and in particular in the transactions demand for money, transaction costs have played a major role since the 1950s, as seen in seminal works of W. Baumol and J. Tobin ([Bau52], [Tob56]).

The first in-depth analysis of a transaction costs model in a post-Black-Scholes mathematical framework comes from Magill and Constantinides [MC76]. There the authors already mention that both a continuous-time market and the absence of transaction costs alone can be justified, while the combination of both leads to portfolio strategies that are unrealistic. They conjecture that once transaction costs are introduced, rational traders will only change their portfolio at finitely many random times. They consider an m-dimensional Black-Scholes market with fixed correlation and cost rates $\chi_i, \chi^i \in [0,1)$ for buying/selling the i-th asset. So the total costs for each rebalancing equal

$$T(v_1, v_2, \dots, v_m) = \sum_{i=1}^{m} (\chi^i \mathbb{1}_{v_i>0} + \chi_i \mathbb{1}_{v_i<0}) v_i,$$

where v_i denotes the change in the i-th asset, measured in cash rather than in physical units. In this model, an optimal investment and consumption problem (also called Merton problem [Me69]) is solved and the existence of a *no-trade region* around Merton's optimal proportion is proved. However, the restriction to piecewise constant trading strategies is imposed rather than proved. The possibility of continuous-time, finite-variation trading is not mentioned.

Leland [Le85] addresses the hedging problem in the Black-Scholes model under transaction costs by a heuristic discretization approach: If trading is only allowed on a fixed equidistant time grid, then the problem of infinite transaction costs vanishes and the transaction costs just contribute to the usual discretization error. Here the costs are also proportional, indicated as *round trip costs*, i. e. for buying and immediately selling a proportion k has to be paid. The transaction costs for portfolio rebalancing therefore equal

$$TC = \frac{k}{2} |\Delta D (S + \Delta S)|,$$

where S denotes the price of the stock and ΔD is the number of units purchased or sold. By empirical studies the author concludes that while the total amount of transaction costs paid is roughly proportional to k (i. e. higher transaction costs to not prevent the trader from trading), they increase quadratically as the distance between trades decreases. So a trade-off between good approximation and small transaction costs has to be made.

The model used in this work is in the tradition of Davis and Norman [DN90] who were among the first to offer a model that allowed for actual continuous-time trading. Their work also premiered the notion of a *solvency region*, see Definitions 2.2.1 and 6.1.5. They revisit the Merton problem in a one-dimensional Black-Scholes framework and introduce proportional transaction costs by proportions $\lambda, \mu \geq 0$ for buying and selling. The trading strategy is modelled by two separate increasing, right-continuous processes L, U counting the accumulated purchases and sales, again measured in cash. The cash-account s_0 therefore has the dynamics

$$ds_0(t) = \big(rs_0(t) - c(t)\big)dt - (1+\lambda)dL_t + (1-\mu)dU_t,$$

where r is the rate of interest and c the consumption function. Again, the existence of a no-trade region is proved. In a side remark, they also address the idea of fixed transaction costs and link this approach to impulse control problems.

This model is reused (in the special case $\lambda = \mu$) by Davis and Clark [DC94], who are among the first to tackle the problem of superreplication under transaction costs in continuous time, and by Soner et al. [CSS95], Leventhal and Skorohod [LS97] and Cvitanić et al. [CPT99], [CK96] for different diffusion models.

The work of Jouini and Kallal [JK95] can be seen as pioneer in multiple ways. Instead of modelling a price process and then wrapping it into a bid-ask spread, they model the bid-ask spread directly via two different price processes $Z' \leq Z$. Then, they were the first to offer framework of no-arbitrage pricing, thereby premiering what is today known as consistent price system (see our Definitions 2.1.3 and 6.1.10), namely a process Z^*, which admits a martingale measure equivalent to $\mathbf{P}$ and satisfies $Z'_t \leq Z^*_t \leq Z_t$ almost surely for all t.

All the models in the tradition of [DN90] follow the classic post-Black-Scholes literature in the sense that the investments in the risky assets are accounted in stochastic processes, usually decomposed in two increasing processes for purchase and sale, while the amount in the cash account follows from stochastic integration, therefore trading strategies are automatically self-financing. A different approach was introduced by Kabanov [Ka99]. Aimed towards multi-currency-markets, the author does not assign a designated cash account and allows for direct transactions between all assets by introducing a costs matrix $(\lambda^{ij})_{1 \leq i,j \leq d}$. The accumulated transfers from asset i to asset j, measured in cash, is accounted in the increasing process L^{ij}, such that the capital invested in asset i held at time t equals

$$V^i_t = \int_0^t \frac{V^i_{s-}}{S^i_{s-}} dS^i_s + \sum_{i=1}^d L^{ij}_t - \sum_{j=1}^d (1 + \lambda^{ij}) L^{ij}_t.$$

The author observes that the portfolio has to evolve inside the negative of the solvency region at all times, in other words, each change in the portfolio must be attainable at price zero. This idea has led to a new characterization of self-financing portfolios in the successing literature ([KS02], [KRS02], [KL02], [CS06], [Scha04]) both in discrete and continuous time. With this it is possible to consider the portfolio process V directly and abandon the transfer processes L^{ij}.

In this tradition, Schachermayer [Scha04] carries the idea to put the solvency region (rather than the price process) into the center of consideration even further by abolishing the asset price process altogether and directly modelling exchange ratios $(\pi^{ij})_{1 \leq i,j \leq d}$: To obtain one unit of asset j at time t, one has to give up π^{ij}_t

units of asset i. These ratios uniquely imply the solvency region and the bis-ask region.

Some of the results proved in the Kabanov framework are crucial to our work, which is why we discuss the setting in greater detail in Section 2.2.

Over the years, strategies have sometimes been assumed right continuous (e. g. [DN90], [LS97], [Ka99]) and sometimes left continuous (e. g. [CSS95], [CK96]), without one of them becoming standard. Campi and Schachermayer [CS06] and Guasoni et al. [GRS08b], [GRS08a] dismiss both and allow for left and right jumps. We follow this tradition by using a multidimensional counterpart of their framework, which is presented in Chapter 2.

1.3 Notation

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbf{P})$ be a filtered probability space satisfying the usual conditions. Assume that $\mathcal{F}_0$ is trivial. Throughout this work, we we will most of the time consider $\mathbb{R}_+^d$-valued random variables and processes, where $\mathbb{R}_+ = (0,\infty)$. To avoid confusions, we will denote the half line containing zero by $\mathbb{R}_{0+}$=$[0,\infty)$. Random variables and processes are written in capital Roman letters, and there the subscript index is always reserved for the time component. An exception is made for stopping times, which are denoted by small Greek letters. For multidimensional random variables and processes, single components are indicated by superscript indices with brackets, e. g.

$$S = (S_t^{(1)}, S_t^{(2)}, \dots, S_t^{(d)})_{t\in[0,T]}.$$

Deterministic vectors are written with small bold letters, e.g.

$$\mathbf{x} = (x^{(1)}, x^{(2)}, \dots, x^{(d)}) \in \mathbb{R}^d.$$

Especially we will have $\mathbf{0} = (0, 0, \dots, 0)$ and $\mathbf{1} = (1, 1, \dots, 1)$. We will often use component-wise operations on vectors. Therefore we define component-wise multiplication and division of $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^d$ by

$$\mathbf{xy} = (x^{(1)}y^{(1)}, x^{(2)}y^{(2)}, \dots, x^{(d)}y^{(d)})$$
$$\frac{\mathbf{x}}{\mathbf{y}} = \left(\frac{x^{(1)}}{y^{(1)}}, \frac{x^{(2)}}{y^{(2)}}, \dots, \frac{x^{(d)}}{y^{(d)}}\right).$$

The product $\mathbf{xy} : \mathbb{R}_+^d \times \mathbb{R}_+^d \to \mathbb{R}_+^d$ is not to be confused with the standard scalar product on $\mathbb{R}^d$, which we will denote by

$$\mathbf{x}^\top \mathbf{y} = \sum_{i=1}^d x^{(i)} y^{(i)}.$$

Inequalities like $\mathbf{x} \leq \mathbf{y}$ or $\mathbf{x} \leq \epsilon$ with $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $\epsilon \in \mathbb{R}$ are also to be understood component-wise. This will spare us much work on the indices.

For a Borel measure μ on $\mathbb{R}^d$, the *support* of μ is a closed set $\text{supp}(\mu) \subseteq \mathbb{R}^d$ defined by

$$\mathbf{x} \in \text{supp}(\mu) \Leftrightarrow \mu(A) > 0 \text{ for all open } A \subset \mathbb{R}^d \text{ with } \mathbf{x} \in A.$$

For an $\mathbb{R}^d$-valued random variable X on $(\Omega, \mathcal{F}, \mathbf{P})$, the *push-forward measure* $\mathbf{P}^X$ is a measure on $\mathbb{R}^d$ defined by

$$\mathbf{P}^X(A) = \mathbf{P}(X \in A).$$

The support of X w. r. t. $\mathbf{P}$ is defined as $\text{supp}_{\mathbf{P}}(X) = \text{supp}(\mathbf{P}^X)$. Note that the support of X is invariant under equivalent change of measures. We will therefore often omit $\mathbf{P}$ and just write $\text{supp}(X)$.

Definition 1.3.1 *For a d-dimensional process $(X_t)_{t \in [0,T]}$, we distinguish* right-continuous jumps

$$\Delta_- X_t = X_t - X_{t-} = X_t - \lim_{s \nearrow t} X_s$$

and left-continuous jumps

$$\Delta_+ X_t = X_{t+} - X_t = \lim_{s \searrow t} X_s - X_t.$$

Recall the usual definition of total variation *of X:*

$$\text{Var}_t^X = \sup_{0=\tau_0 \leq \tau_1 \leq \ldots \leq \tau_n = t} \sum_{i=1}^{n} \|X_{\tau_i} - X_{\tau_{i-1}}\|,$$

where $\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^d$. If Var^X is finite, then we can decompose X into left and right jumps and a continuous part

$$X = X^c + X^r + X^l,$$

where $X_t^l = \sum_{s<t} \Delta_+ X_s$ and $X_t^r = \sum_{s \leq t} \Delta_- X_s$.

If Var_t^X is a. s. finite for all t, then X^c can be represented as

$$X_t^c = \int_0^t \mathring{X}_s^c d\text{Var}_s^{X^c},$$

with a process $\mathring{X}^c$ taking values only in the d-dimensional unit sphere. Note that usually (e.g. [KS02], [GRS08b]) the Radon-Nikodym derivative with respect to Var^{X^c} is written as $\dot{X}^c$ rather than $\mathring{X}^c$, but we chose this notation to improve readability. $\mathring{X}^c$ is uniquely defined almost everywhere with respect to the random measure induced by the incrasing function $t \mapsto \mathrm{Var}_t^{X^c}$, so all statements on $\mathring{X}^c$ are to be understood Var^{X^c}-a. e.

Chapter 2

Proportional transaction costs

We are now ready to present a model for trading in multiasset models under proportional transaction costs in continuous time. While the first section of this chapter is devoted to the basic definitions and properties of our model, the second one higights the parallels to an alternative model developed by Kabanov, which was briefly discussed in Section 1.2. The main result of this chapter, apart from the presentation of the model, is Theorem 2.2.9.

2.1 The model

As underlying, consider a càdlàg, $\mathbb{R}^d_+$-valued price process

$$S = (S_t^{(1)}, S_t^{(2)}, \dots, S_t^{(d)})_{t\in[0,T]}$$

adapted to the probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbf{P})$. For buying or selling α units of stock i at time $t \in [0,T]$, we have to pay the amount of $\epsilon\alpha S_t^{(i)}$ as transaction costs. Here the constant $\epsilon \in [0,1)$ is the level of costs. This leads to the following definition:

Definition 2.1.1 *A predictable process $H = (H_t^{(1)}, H_t^{(2)}, \dots, H_t^{(d)})_{t\in[0,T]}$ is called a* trading strategy*, if*

1. *the total variation process* $\mathrm{Var}^{H^{(i)}}$ *of each component $H^{(i)}$ is a. s. finite*

2. *$H_0 = H_T = \mathbf{0}$ a. s.*

For $0 \leq \epsilon < 1$*, the* value *of a trading strategy* H *with respect to* S *and* ϵ *is defined as*

$$V_{\epsilon,S}(H) = \sum_{i=1}^{d} \int_0^T H_t^{(i)} dS_t^{(i)} - \epsilon \int_0^T S_t^{(i)} d\mathrm{Var}_t^{H^{(i)}}.$$

Note that the existence of both integrals above in a pathwise Lebesgue-Stieltjes sense is secured by H being of finite variation and S being càdlàg, which in the second case can be seen by pathwise partial integration, see e. g. [JS87]. While the first integral resembles the usual gains and losses accumulated by trading on S, the second integral corresponds to the transaction costs paid during trading. The condition $H_0 = H_T = \mathbf{0}$ means we must both establish our initial position and liquidate our final position at maturity and take into account the corresponding transaction costs. This affects superreplication prices, as we will see. Note furthermore that we are not allowed to swap assets, i. e. if we sell one asset and buy another one at the same time, we pay the transaction costs twice. This restriction will be weakened in Chapter 6
For all statements about trading in risky markets, it is crucial to choose a suitable notion of admissibility. Our notion hinges critically on the choice of the underlying numeraire.

Definition 2.1.2 *Consider an* $\mathbb{R}_+$*-valued process* $(N_t)_{t\in[0,T]}$*. For* $a \in \mathbb{R}$*, we call the strategy* H (ϵ, a)-admissible with respect to the numeraire N*, if for all* $t \in (0, T]$ *we have*

$$V_{\epsilon,S}(H\mathbb{1}_{(0,t)}) \geq -aN_t.$$

H *is called* ϵ-admissible with respect to N *if it is* (ϵ, a)*-admissible w. r. t.* N *for some* a*.*

In this work, we will concentrate on two different numeraires. The most intuitive choice is the constant process $N_t = 1$. Here, the downside risk of every admissible strategy, measured in cash, is bounded from below, or, in other words, for every admissible strategy there is a certain amount of cash such that the trader, endowed with this amount, is guaranteed to be able to liquidate his or her portfolio to a nonnegative amount of cash at any time point. However, short static strategies, i. e. strategies that sell stocks at time 0 and rebuy them at time T are not admissible if the underlying asset is unbounded. This will play a major role in Chapter 3.
An obvious choice that renders all short static strategies admissible is

$$N_t = 1 + \sum_{i=1}^{d} S_t^{(i)}.$$

Instead of a fixed initial endowment of cash, the trader is in this case equipped with an initial portfolio consisting of a units of cash and each asset. This somewhat weaker notion of admissibility proves favourable for the proof of the Fundamental Theorem of Asset Pricing in Chapter 5, as it fits to the definition of an admissible portfolio process in [CS06], see Definition 2.2.4 below. This enables us to translate crucial results from this work to our setting.

It was first shown by [JK95] that when transaction costs are introduced, the fundamental role that equivalent martingale measures play in duality theory is inherited by what is now known as consistent price systems. Therefore the following definition is crucial:

Definition 2.1.3 *For* $\epsilon \in [0,1)$*, an* ϵ-consistent price system (ϵ-CPS) *to a process* S *is a pair* $(M,\mathbf{Q})$ *consisting of a probability measure* $\mathbf{Q} \sim \mathbf{P}$ *and an* $\mathbb{R}^d_+$*-valued* $\mathbf{Q}$*-martingale* $(M_t)_{t\in[0,T]}$ *such that for all* $t \in [0,T]$*:*

$$S_t(1-\epsilon) \leq M_t \leq S_t(1+\epsilon) \;\; \mathbf{Q}-a.s.$$

We will use the shorthand notation $(M,\mathbf{Q}) \sim_\epsilon (S,\mathbf{P})$*.*
If, instead of $\mathbf{Q} \sim \mathbf{P}$*, we have* $\mathbf{Q} \ll \mathbf{P}$*, we call* $(M,\mathbf{Q})$ *an* absolutely continuous ϵ-consistent price system (ϵ-ACCPS) *and use the shorthand notation* $(M,\mathbf{Q}) \ll_\epsilon (S,\mathbf{P})$*.*

Remark 2.1.4

1. The economic interpretation behind this definition is the following: What really matters to the trader is not the process S, which itself does not play any role in actual trading, but rather the bid-ask-spread $[(1-\epsilon)S,(1+\epsilon)S]$. So instead of looking for an equivalent martingale measure for S, one can replace S by any process M inside the bid-ask-spread and look for an EMM for M.

2. Despite this obvious interpretation, CPS are sometimes defined using the condition
$$\frac{1}{1+\epsilon}S_t \leq M_t \leq S_t(1+\epsilon),$$
see for instance [GRS08a]. Obviously this condition implies ours, moreover it makes some calculations easier. Note that both conditions become equivalent if $\epsilon \to 0$, i. e. ϵ-CPS in our sense exist for all $\epsilon > 0$ if and only if they exist for all $\epsilon > 0$ in this sense.

The following lemma delivers some insight in the nature of our transaction costs setting and will be useful several times throughout this work.

Lemma 2.1.5 *Let two càdlàg price processes S and $\tilde{S}$ and constants $\epsilon, \tilde{\epsilon} \geq 0$ be given such that for all $t \in [0, T]$*

$$S_t(1+\epsilon) \geq \tilde{S}_t(1+\tilde{\epsilon}) \geq \tilde{S}_t(1-\tilde{\epsilon}) \geq S_t(1-\epsilon)$$

holds $\mathbf{P}$-a. s. Then we have for all trading strategies H:

$$V_{\epsilon,S}(H) \leq V_{\tilde{\epsilon},\tilde{S}}(H)$$

$\mathbf{P}$*-almost surely.*

Proof: We consider only the one-dimensional case, the general case follows by simply summing up the dimensions. Consider the standard decomposition

$$H = H^c + H^r + H^l$$

form Definition 1.3.1. Furthermore we decompose H^c into two increasing continuous processes $H^{c+} = \frac{1}{2}(\mathrm{Var}^{H^c} + H^c)$ and $H^{c-} = \frac{1}{2}(\mathrm{Var}^{H^c} - H^c)$, such that $H^{c+} - H^{c-} = H^c$ and $H^{c+} + H^{c-} = \mathrm{Var}^{H^c}$. Then we have by partial integration

$$I_c(\epsilon, S) := \int_0^T H_t^c dS_t - \epsilon \int_0^T S_t d\mathrm{Var}_t^{H^c}$$

$$= \int_0^T (H_t^{c+} - H_t^{c-}) dS_t - \epsilon \int_0^T S_t d(H^{c+} + H^{c-})_t$$

$$= \int_0^T (1-\epsilon) S_t dH_t^{c-} - \int_0^T (1+\epsilon) S_t dH_t^{c+} + H_T^c S_T.$$

For the right-continuous jump part we have

$$I_r(\epsilon, S) := \int_0^T H_t^r dS_t - \epsilon \int_0^T S_t d\mathrm{Var}_t^{H^r}$$

$$= \sum_{s \leq T} \Delta_- H_s \left(S_s(1-\epsilon) \mathbb{1}_{\Delta_- H_s < 0} - S_s(1+\epsilon) \mathbb{1}_{\Delta_- H_s > 0} \right) + H_T^r S_T$$

and the same is true for the left jumps, with I_r replaced by I_l, S_s by S_{s-} and Δ_- by Δ_+. If we take into account that

$$\mathrm{Var}^H = \mathrm{Var}^{H^c} + \mathrm{Var}^{H^r} + \mathrm{Var}^{H^l},$$

we have

$$V_{\epsilon,S}(H) = I_c(\epsilon, S) + I_r(\epsilon, S) + I_l(\epsilon, S),$$

and using

$$H_T^c S_T + H_T^r S_T + H_T^l S_T = H_T S_T = 0,$$

we see that these values can be omitted. Now we just have to use the inequalities $S_t(1+\epsilon) \geq \tilde{S}_t(1+\tilde{\epsilon})$ and $\tilde{S}_t(1-\tilde{\epsilon}) \geq S_t(1-\epsilon)$ in all the above expressions to see

$$V_{\epsilon,S}(H) = I_c(\epsilon, S) + I_r(\epsilon, S) + I_l(\epsilon, S) \leq I_c(\tilde{\epsilon}, \tilde{S}) + I_r(\tilde{\epsilon}, \tilde{S}) + I_l(\tilde{\epsilon}, \tilde{S}) = V_{\tilde{\epsilon},\tilde{S}}(H).$$

$\square$

The economic interpretation of this lemma is that the narrower the bid-ask-spread is, the better the trader is off, since we always buy lower and sell higher.

The following statements follow easily from Lemma 2.1.5. The first one will play a role in our comparison of different concepts of arbitrage to come. The second one treats the special case $(\tilde{S}, \mathbf{Q}) \ll_\epsilon (S, \mathbf{P})$ and will be important throughout this work, as it serves as link to many well-known duality results from frictionless finance, as every statement on the frictionless process $\tilde{S}$ implies a statement on S with ϵ transaction costs. The third statement is used to construct arbitrage out of free lunch with bounded risk, see Theorem 5.4.5.

Corollary 2.1.6 *1. By replacing H with $H\mathbb{1}_{(0,t)}$ we get*

$$V_{\epsilon,S}(H\mathbb{1}_{(0,t)}) \leq V_{\tilde{\epsilon},\tilde{S}}(H\mathbb{1}_{(0,t)})$$

for all $0 \leq t \leq T$. So if a strategy is ϵ-admissible on S, it is also $\tilde{\epsilon}$-admissible on $\tilde{S}$.

2. *If $(M, \mathbf{Q}) \ll_\epsilon (S, \mathbf{P})$ for a pair $(M, \mathbf{Q})$, then*

$$V_{\epsilon,S}(H) \leq V_{0,M}(H)$$

holds $\mathbf{Q}$-a.s for each trading strategy H.

3. *Convex combinations of strategies yield a higher value than just the convex combination of the single values: If the strategy G is a convex combination of two strategies, say, $G = \lambda_1 H^1 + \lambda_2 H^2$, then*

$$d\mathrm{Var}^G \leq \lambda_1 d\mathrm{Var}^{H^1} + \lambda_2 d\mathrm{Var}^{H^2},$$

meaning that $\mathrm{Var}^G - \lambda_1 \mathrm{Var}^{H^1} - \lambda_2 \mathrm{Var}^{H^2}$ is a. s. increasing. Hence the value of G satisfies

$$V_{\epsilon,S}(G) \geq \lambda_1 V_{\epsilon,S}(H^1) + \lambda_2 V_{\epsilon,S}(H^2) \ \ a.s.$$

The same is true for convex combinations of countably many strategies.

Lemma 2.1.7 *Let an ACCPS* $(M, \mathbf{Q}) \ll_\epsilon (S, \mathbf{P})$ *be given. Let the numeraire process be either* $N_t = 1$ *or* $N_t = 1 + \sum_{i=1}^d S_t^{(i)}$. *Then for all* ϵ*-admissible strategies* H *the process* $t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)})$ *is a* $\mathbf{Q}$*-supermartingale.*

Proof: Using Corollary 2.1.6.2, we have $\mathbf{Q}$-almost surely

$$V_{0,M}(H\mathbb{1}_{(0,t)}) \geq V_{\epsilon,S}(H\mathbb{1}_{(0,t)}) \geq -aN_t$$

for some $a \in \mathbb{R}$. On the other hand

$$V_{0,M}(H\mathbb{1}_{(0,t)}) = \int_0^t \sum_{i=1}^d H_s^{(i)} dM_s^{(i)}.$$

So if $N_t = 1$, the process $t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)})$ is a stochastic integral with respect to a $\mathbf{Q}$-martingale and bounded from below by $-a$. Therefore by Fatou's Lemma and [JS87] it is a $\mathbf{Q}$-supermartingale.
For the case $N_t = 1 + \sum_{i=1}^d S_t^{(i)}$, note that $\mathbf{Q}$-almost surely

$$V_{0,M}(H\mathbb{1}_{(0,t)}) \geq -a(1 + \sum_{i=1}^d S_t^{(i)}) \geq -a\big(1 + \sum_{i=1}^d (1+\epsilon) M_t^{(i)}\big) \Rightarrow$$

$$V_{0,M}(H\mathbb{1}_{(0,t)}) + a(1+\epsilon)\sum_{i=1}^d S_t^{(i)} = \int_0^t \big(H_s^{(i)} - a(1+\epsilon)\big) dM_s^{(i)} \geq -a,$$

so by the same argument as above

$$t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)}) + a(1+\epsilon)\sum_{i=1}^d M_t^{(i)}$$

is a $\mathbf{Q}$-supermartingale, and so $t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)})$ is. $\square$

2.2 Kabanov's cone-setting and the superreplication theorem

In order to exploit results from [CS06] vital for our theorem, we recall an alternative, more general framework different from the one established in Section 2.1. This framework, which is especially well-suited for multiasset models, was first introduced by Y. Kabanov in [Ka99] and subsequently refined and generalized in the following years [KL02] [KS02]. It is based on a partial order induced by a cone in $\mathbb{R}^{d+1}$ which marks all solvent holdings. However, as we will see, both frameworks are closely connected, allowing us to switch between both systems in our proof of the Fundamental Theorem.

Definition 2.2.1 *For $1 \leq i \leq d+1$, denote by $\mathbf{e}_i$ the i-th unit vector of $\mathbb{R}^{d+1}$. For $\epsilon \geq 0$, we define the ϵ-*solvency cone *$\mathcal{K}^S_{\epsilon,t}$ at time $t \in [0,T]$ as the cone spanned by $\mathbf{e}_i,\ 1 \leq i \leq d+1$ and the $2d$ vectors $\{L^{i+}_t, L^{i-}_t, 1 \leq i \leq d\}$, where*

$$L^{i-}_t = (1+\epsilon)S^{(i)}_t \mathbf{e}_1 - \mathbf{e}_{i+1}$$

$$L^{i+}_t = -(1-\epsilon)S^{(i)}_t \mathbf{e}_1 + \mathbf{e}_{i+1}.$$

We furthermore define the dual ϵ-solvency cone *$\mathcal{K}^{S*}_{\epsilon,t}$ as its dual cone, i. e.*

$$\mathcal{K}^{S*}_{\epsilon,t} = \{\mathbf{y} \in \mathbb{R}^{d+1} : \mathbf{x}^\top \mathbf{y} \geq 0 \text{ for all } \mathbf{x} \in \mathcal{K}^S_{\epsilon,t}\}.$$

We will omit the process S and write $\mathcal{K}_{\epsilon,t}$ or $\mathcal{K}^*_{\epsilon,t}$ if the danger of confusion is not given.

The interpretation of these objects is straightforward. Here a vector in $\mathbb{R}^{d+1}$ represents a trader's holdings in cash (first component) as well as all assets (components 2 to $d+1$). L^{i+} represents a holding of one unit of asset i and exactly that amount of debt in the cash account such that it can be balanced by selling the asset. L^{i-} represents the opposite situation of a short position in stock i and a positive amount of cash. So the holdings in $\mathcal{K}_{\epsilon,t}$ are exactly those which can, at time t, be liquidated to a nonnegative amount of cash, i. e. the solvent ones. On the other hand, $-\mathcal{K}_{\epsilon,t}$ contains all holdings that can be established without initial endowment, possibly by throwing away money or stocks.

Note that if $\epsilon > 0$, we have

$$\mathbf{e}_1 = \frac{L^{1-}_t + L^{1+}_t}{2\epsilon S^{(1)}_t},$$

thus $\mathcal{K}^S_{\epsilon,t}$ is equivalently generated by the vectors $L^{i\pm}_t$ without the $\mathbf{e}_i$. If $\epsilon = 0$, then this is not the case, then $L^{i\pm}_t$ span a d-dimensional linear subspace of $\mathbb{R}^{d+1}$ while $\mathcal{K}^S_{\epsilon,t}$ is the corresponding half space containing $\mathbf{1}$. The polar cone, however, does not change by omitting the unit vectors, even if $\epsilon = 0$.

It is immediately clear both from the mathematical representation and from the interpretation that $\mathbb{R}^{d+1}_+$ is a strict subset of $\mathcal{K}_{\epsilon,t}$, and conversely $\mathcal{K}^*_{\epsilon,t}$ is a subset of $\mathbb{R}^{d+1}_+ \cup \{\mathbf{0}\}$. The following lemma, which is in the spirit of Lemma 2.1.5, is equally plausible: If the pair $(\tilde{S}, \tilde{\epsilon})$ offers more favourable trading than (S, ϵ), then its set of solvent holdings is greater.

Lemma 2.2.2 *Let ϵ, S and $\tilde{\epsilon}$, $\tilde{S}$ be as in Lemma 2.1.5. Then*

$$\mathcal{K}^S_{\epsilon,t} \subseteq \mathcal{K}^{\tilde{S}}_{\tilde{\epsilon},t}$$

holds a. s. for all $t \in [0, T]$.

Proof: For $1 \leq i \leq d$, define the vectors $\tilde{L}^{i\pm}_t$ just like $L^{i\pm}_t$, but with ϵ and S replaced by $\tilde{\epsilon}$ and $\tilde{S}$. Then we have

$$L^{i-}_t = \tilde{L}^{i-}_t + \big((1+\epsilon)S^{(i)}_t - (1+\tilde{\epsilon})\tilde{S}^{(i)}_t\big)\mathbf{e}_1.$$

Since $(1+\epsilon)S^{(i)}_t - (1+\tilde{\epsilon})\tilde{S}^{(i)}_t \geq 0$ and $\mathbf{e}_1 \in \mathcal{K}^{\tilde{S}}_{\tilde{\epsilon},t}$, this proves $L^{i-}_t \in \mathcal{K}^{\tilde{S}}_{\tilde{\epsilon},t}$. By the same argument we get $L^{i+}_t \in \mathcal{K}^{\tilde{S}}_{\tilde{\epsilon},t}$, which completes the proof. □

The dual cone, on the contrary, is linked to our concept of consistent price systems, as the following lemma will show:

Lemma 2.2.3

1. *Let $Z = (Z^{(1)}_t, Z^{(2)}_t, \ldots, Z^{(d+1)}_t)_{t\in[0,T]}$ be a $\mathbf{P}$-martingale satisfying $Z_t \in \mathcal{K}^{S*}_{\epsilon,t} \backslash \{\mathbf{0}\}$ $\mathbf{P}$-a. s. for all t. Then $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$, where $\mathbf{Q}$ is defined by $\frac{d\mathbf{Q}}{d\mathbf{P}} = \frac{Z^{(1)}_T}{\mathbb{E}_{\mathbf{P}}(Z^{(1)}_T)}$ and*

$$M = \left(\frac{Z^{(2)}}{Z^{(1)}}, \frac{Z^{(3)}}{Z^{(1)}}, \ldots, \frac{Z^{(d+1)}}{Z^{(1)}}\right).$$

2. *If an ϵ-consistent price system $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$ is given, then Z is a $\mathbf{P}$-martingale taking values in $\mathcal{K}^{S*}_{\epsilon,t} \backslash \{\mathbf{0}\}$ a. s. for all t, where*

$$Z^{(1)}_t = \mathbb{E}_{\mathbf{P}}(\frac{d\mathbf{Q}}{d\mathbf{P}}|\mathcal{F}_t),$$

$$Z^{(i)} = Z^{(1)}M^{(i-1)},\ 2 \leq i \leq d+1.$$

Proof: As we have seen, $Z_t \in \mathcal{K}_{\epsilon,t}^{S*} \backslash \{\mathbf{0}\}$ implies $Z_t \in \mathbb{R}_+^{d+1}$, so dividing by $Z_t^{(1)}$ is not a problem. We have $Z_t \in \mathcal{K}_{\epsilon,t}^{S*}$ if and only if both $Z_t^\top L_t^{i-} \geq 0$ and $Z_t^\top L_t^{i+} \geq 0$ hold for $1 \leq i \leq d$. On the other hand, if Z is strictly positive, we have both

$$Z_t^\top L_t^{i-} \geq 0 \Leftrightarrow (1+\epsilon) S_t^{(i)} \geq M_t^{(i)}$$

$$Z_t^\top L_t^{i+} \geq 0 \Leftrightarrow (1-\epsilon) S_t^{(i)} \leq M_t^{(i)},$$

hence we end up with

$$\frac{M_t}{S_t} \in [1-\epsilon, 1+\epsilon] \Leftrightarrow Z_t \in \mathcal{K}_{\epsilon,t}^{S*} \backslash \{\mathbf{0}\}.$$

Concerning the martingale property, it suffices to notice that $\mathbb{E}_{\mathbf{P}}(\frac{d\mathbf{Q}}{d\mathbf{P}}|\mathcal{F}_t)$ is by its nature a $\mathbf{P}$-martingale and that the usual change of measure formula for conditional expectation yields

$$\mathbb{E}_{\mathbf{Q}}(M_T^{(i)}|\mathcal{F}_t) = \frac{\mathbb{E}_{\mathbf{P}}(\frac{d\mathbf{Q}}{d\mathbf{P}} M_T^{(i)}|\mathcal{F}_t)}{\mathbb{E}_{\mathbf{P}}(d\mathbf{Q}/d\mathbf{P}|\mathcal{F}_t)}$$

$$= \frac{\mathbb{E}_{\mathbf{P}}(\frac{Z_T^{(1)}}{\mathbb{E}_{\mathbf{P}}(Z_T^{(1)})} \frac{Z_T^{(i+1)}}{Z_T^{(1)}}|\mathcal{F}_t)}{Z_t^{(1)}/\mathbb{E}_{\mathbf{P}}(Z_T^{(1)})} = \frac{\mathbb{E}_{\mathbf{P}}(Z_T^{(i+1)}|\mathcal{F}_t)}{Z_t^{(1)}}.$$

Hence $M^{(i)}$ is a $\mathbf{Q}$-martingale if and only if $Z^{(i+1)}$ is a $\mathbf{P}$-martingale. □

For the remainder of this chapter assume S admits an ϵ-consistent price system, hence a $\mathcal{K}_{\epsilon,\cdot}^{S*} \backslash \{\mathbf{0}\}$-valued $\mathbf{P}$-martingale exists.

Definition 2.2.4 *An $\mathbb{R}^{d+1}$-valued process $(V_t)_{t\in[0,T]}$ is called an* ϵ-self-financing portfolio process*, if*

1. *$\mathbf{P}$-almost every path starts at $V_0 = \mathbf{0}$ and has finite total variation*

2. *V is predictable*

3. *for all stopping times $0 \leq \sigma \leq \tau \leq T$ we have*

$$V_\tau(\omega) - V_\sigma(\omega) \in -\overline{\text{conv}} \left(\bigcup_{t\in[\sigma(\omega),\tau(\omega)]} \mathcal{K}_{\epsilon,t}(\omega) \right).$$

Here $\overline{\text{conv}}$ *denotes the closure of the convex hull with respect to the Euclidean topology of* $\mathbb{R}^{d+1}$*. Such a self-financing portfolio process is called* (ϵ, a)-admissible*, if*

1. $V_T + a\mathbf{1} \in \mathcal{K}_{\epsilon,T}$
2. *for all stopping times* $\tau \in [0,T]$ *and all* $\mathcal{K}^*_{\epsilon,\cdot}\backslash\{\mathbf{0}\}$*-valued* $\mathbf{P}$*-martingales* Z *as in Lemma 2.2.3 we have* $Z_\tau^\top(V_\tau + a\mathbf{1}) \geq 0$.

Again, V *is called* ϵ-admissible *if it is* (ϵ, a)*-admissible for some* a.

As the following lemma will show, $V_T + a\mathbf{1} \in \mathcal{K}_{\epsilon,T}$ already implies the admissibility constant. This, however, does not mean that the second consition in the definition of admissibility can be omitted.

Lemma 2.2.5 *Let an* ϵ*-admissible portfolio process* V *and a constant* b *be given such that* $V_T + b\mathbf{1} \in \mathcal{K}_{\epsilon,T}$ *almost surely. Then* V *is* (ϵ, b)*-admissible.*

Proof: Assume that b is not a valid admissibility constant, i. e. there exists a stopping time $\tau \in [0,T]$ and a $\mathcal{K}^*_{\epsilon,\cdot}\backslash\{\mathbf{0}\}$-valued $\mathbf{P}$-martingale Z such that

$$Z_\tau^\top V_\tau < -bZ_\tau^\top \mathbf{1}$$

on some set $A \in \mathcal{F}_\tau$ with positive probability. $Z^\top\mathbf{1}$ is a $\mathbf{P}$-martingale while [CS06, Lemma 8] tells us that $Z^\top V$ is a $\mathbf{P}$-supermartingale. Thus on some subset of A with positive probability, we have that

$$Z_T^\top V_T < -bZ_T^\top \mathbf{1}.$$

But since $Z_T \in \mathcal{K}^*_{\epsilon,T}$ a. s. we must have that $V_T + b\mathbf{1} \notin \mathcal{K}_{\epsilon,T}$ with positive probability, in contradiction to our assumption. □

Lemma 2.2.6

1. *For each trading strategy* H *there exists an* ϵ*-self-financing portfolio process* V *such that* V_T *is a multiple of* $\mathbf{e}_1$ *(i. e.* V *holds no risky assets in the end) and that* $V_T^{(1)} = V_{\epsilon,S}(H)$ *a. s.*
2. *Let* V *be an* ϵ*-self-financing portfolio process such that* V_T *is a multiple of* $\mathbf{e}_1$*. Then there exists a trading strategy* H *such that* $V_T^{(1)} \leq V_{\epsilon,S}(H)$ *a. s.*

Proof: (1): For $1 \leq i \leq d$, consider the decomposition of the finite-variation process $H^{(i)}$ into two a. s. increasing processes $H^{(i+)}$ and $H^{(i-)}$ like in the proof of Lemma 2.1.5, such that $H^{(i)} = H^{(i+)} - H^{(i-)}$ and $\mathrm{Var}^{H^{(i)}} = H^{(i+)} + H^{(i-)}$. We define the processes

$$X_t^{i+} = \int_0^t -L_s^{i-} dH_s^{(i+)}$$

$$X_t^{i-} = \int_0^t -L_s^{i+} dH_s^{(i-)}$$

as well as the processes $X^i = X^{i+} + X^{i-}$ and

$$V = \sum_{i=1}^d X^i = \sum_{i=1}^d (X^{i+} + X^{i-}).$$

Note that $-L^{i-}$ represents the holding established by buying one unit of asset i at time t, thus X^{i+} counts all the changes in cash and assets caused by buying asset i over time. The analogous is true for $-L^{i+}$ and X^{i-}, respectively. Observe that

$$\Delta X_t^{i\pm} = -L^{i\mp} \Delta H_t^{i\pm}$$

holds true for both left and right jumps, as well as

$$\mathring{X}^{i\pm^c} = -\frac{L_t^{i\mp}}{\|L_t^{i\mp}\|}$$

$\mathrm{Var}^{H^{(i\pm)^c}}$-almost everywhere.

Now $\Delta_+ X^{i+}$ and $\Delta_+ X^{i-}$ are never nonzero at the same time, and the same holds true for $\Delta_- X^{i\pm}$ and $\mathring{X}^{i\pm^c}$ (in other words: We never buy and sell the same asset at the same time, neither in a right-jump, nor in a left-jump or continuous way). So we have that $\Delta_+ X_t^i$, $\Delta_- X_t^i$ and $\mathring{X}_t^i$ lie in $\partial\mathrm{cone}(\{-L_t^{i+}, -L_t^{i-}\})$ a. s. and $\mathrm{Var}^{H^{(i)}}$-a.e. respectively. Now

$$\{-L_t^{i+}, -L_t^{i-}, 1 \leq i \leq d\}$$

are a minimal spanning set and the cone they span is either identical to $-\mathcal{K}_{\epsilon,t}$ or to its boundary, this means that $\Delta_+ V_t$, $\Delta_- V_t$ and $\mathring{V}_t^c$ lie in $-\partial\mathcal{K}_{\epsilon,t}$, a. s. By construction V is also a predictable process starting at $\mathbf{0}$, and thus an ϵ-self financing portfolio process. One easily computes

$$dV_t^{(i+1)} = dH_t^{(i+)} - dH_t^{(i-)} \Rightarrow V_T^{(i+1)} = H_T^{(i)} = 0$$

$$dV_t^{(1)} = \sum_{i=1}^d \Big(-(1+\epsilon) S_t^{(i)} dH_t^{(i+)} + (1-\epsilon) S_t^{(i)} dH_t^{(i-)} \Big)$$

$$= \sum_{i=1}^{d} S_t^{(i)}(-dH_t^{(i)} - \epsilon d\mathrm{Var}_t^{H^{(i)}}) = \sum_{i=1}^{d}\Big(H_t^{(i)} dS_t^{(i)} - \epsilon S_t^{(i)} d\mathrm{Var}_t^{H^{(i)}}\Big),$$

where partial integration was applied in the last equation. This means that we have $V_T = V_{\epsilon,S}(H)\mathbf{e}_1$, as desired.

(2): Given the portfolio process V, define the strategy H by

$$H^{(i)} = V^{(i+1)},\ 1 \leq i \leq d.$$

Now consider the processes $X^{(i\pm)}$ and $X^{(i)}$ corresponding to H like in the first part of the proof, and set

$$W = \sum_{i=1}^{d} X^{(i)} = \sum_{i=1}^{d}(X^{(i+)} + X^{(i-)}).$$

Note that V and W coincide in dimensions 2 to $d+1$, so $V - W$ is at every time point a multiple of $\mathbf{e}_1$.
As we have seen in the first part of the proof, $\Delta_+ W_t$, $\Delta_- W_t$ and $\mathring{W}_t$ lie in $-\partial\mathcal{K}_{\epsilon,t}$, a. s. So for V to be self-financing, we must have that $\Delta_\pm(V - W)_t$ and $\left(V^c \mathring{-} W^c\right)_t$ lie in $-\mathcal{K}_{\epsilon,t}$ a. s. all the time. But since $\mathbf{e}_1 \notin -\mathcal{K}_{\epsilon,t}$, the process $V - W$ must be a decreasing process times $\mathbf{e}_1$, especially

$$V_{\epsilon,S}(H) = W_T^{(1)} \geq V_T^{(1)}.$$

This proves that H has the desired properties.

□

Remark 2.2.7 1. Having learned that the portfolio process $V = V(H)$ constructed in 2.2.6.1 is self-financing, we may interchange the V-notation and the H-notation deliberately by identifying the dimensions 2 to d of V by the process H and extending a strategy H to a portfolio process by adding

$$V_t^{(1)} := \sum_{i=1}^{d} \int_0^t H_s^{(i)} dS_s^{(i)} - \epsilon \int_0^t S_s^{(i)} d\mathrm{Var}_s^{H^{(i)}}$$

as first component. The next lemma will show that the two concepts of admissibility are also consistent with each other.

2. In transaction cost models with continuous processes strategies that allow right- as well as left-continuous jumps, it is possible to “throw away” an

arbitrary amount of cash by buying stock immediately before some time point t and selling it again immediately after t. Using this technique, we would be able to satisfy 2.2.6.2 with equality if we were allowed one final trade immediately after T. In the cone-framework, cash and stocks can be thrown away without such trades, since $-\mathbf{e}_i \in -\mathcal{K}_{\epsilon,t}$ for $1 \leq i \leq d+1$ (This corresponds to the fact that if $\epsilon = 0$, the unit vectors have to be added to the solvency cone and are not spanned by the vectors $L^{i\pm}$). This difference expresses itself by the discrepancy between the portfolio processes V and W in the proof of Lemma 2.2.6.2. In the case of W, every change in the cash account $W^{(1)}$ is explained exactly by some trades, i. e. changes in $W^{(2)}$ to $W^{(d+1)}$, whereas in the case of V, the cash can also be lower. The property of W that $W^{(1)}$ is always as high as possible considering the trades made was called *being on the boundary* in [GRS08b, Definition 4.1], since its mathematical description is that $\mathring{W}_t$ and $\Delta_\pm W_t$ lie in the boundary of $-\mathcal{K}_{\epsilon,t}$. Speaking in terms of financial reporting, a trader whose portfolio is on the boundary can always explain where all the money went.

Lemma 2.2.8 *Let $\epsilon > 0$ be given such that S admits a δ-consistent price system, where*

$$\delta = 1 - \frac{1}{1+\epsilon^{(1/3)}},$$
$$\epsilon^{(1/3)} = (1+\epsilon)^{1/3} - 1.$$

Let V be an ϵ-self-financing portfolio process (ϵ, a)-admissible in the sense of Definition 2.2.4. Then the trading strategy H constructed in Lemma 2.2.6.2 is $\left(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})\right)$-admissible in the sense of Definition 2.1.2, with respect to the sum numeraire $N_t = 1 + \sum_{i=1}^{d} S_t^{(i)}$.

Proof: Consider $(M, \mathbf{Q}) \sim_\delta (S, \mathbf{P})$ and the constant

$$\epsilon^{(2/3)} = (1+\epsilon)^{2/3} - 1,$$

thus $(1+\epsilon^{(1/3)})^2 = 1+\epsilon^{(2/3)}$ and $(1+\epsilon^{(1/3)})^3 = 1+\epsilon$. Note that $1-\delta = \frac{1}{1+\epsilon^{1/3}}$, thus by straightforward computations the inequalities

$$S_t(1+\epsilon^{(1/3)}) \geq M_t \geq \frac{S_t}{1+\epsilon^{(1/3)}} \Leftrightarrow$$

$$S_t(1+\epsilon) \geq M_t(1+\epsilon^{(2/3)}) \geq \frac{M_t}{1+\epsilon^{(2/3)}} \geq \frac{S_t}{1+\epsilon} \Rightarrow$$

$$S_t(1+\epsilon) \geq M_t(1+\epsilon^{(2/3)}) \geq M_t(1-\epsilon^{(2/3)}) \geq S_t(1-\epsilon)$$

hold. Now Lemma 2.2.2 tells us that

$$\mathcal{K}^S_{\epsilon,t} \subseteq \mathcal{K}^M{}_{\epsilon^{(2/3)},t} \Rightarrow \mathcal{K}^{S*}_{\epsilon,t} \supseteq \mathcal{K}^{M*}{}_{\epsilon^{(2/3)},t}.$$

This implies that V is in fact $(\epsilon^{(2/3)}, a)$-admissible with respect to M. Now M admits a martingale measure equivalent to $\mathbf{P}$, namely $\mathbf{Q}$. Then [CS06, Theorem 9] tells us that V satisfies

$$V_\tau + a\mathbf{1} \in \mathcal{K}^M{}_{\epsilon^{(2/3)},\tau}$$

for all stopping times $0 \leq \tau \leq T$.
Just like before we have the inequalities

$$M_t(1+\epsilon^{(2/3)}) \geq S_t(1+\epsilon^{(1/3)}) \geq S_t(1-\epsilon^{(1/3)}) \geq M_t(1-\epsilon^{(2/3)}),$$

hence we get by Lemma 2.1.5,

$$V_{\epsilon^{(1/3)},S}(H\mathbb{1}_{(0,t)}) \geq V_{\epsilon^{(2/3)},M}(H\mathbb{1}_{(0,t)}) \geq V_{\epsilon,S}(H\mathbb{1}_{(0,t)})$$

so the portfolio process $\hat{V}$, induced by H using $\epsilon^{(1/3)}$ and S, dominates V a. s. for all t. Hence we get

$$\hat{V}_t + a\mathbf{1} \in \mathcal{K}^M{}_{\epsilon^{(2/3)},t} \subseteq \mathcal{K}^S_{\epsilon^{(1/3)},t}$$

for all t. So the liquidation value of $\hat{V}_t + a\mathbf{1}$ w. r. t. $\epsilon^{(1/3)}$ is nonnegative, but smaller or equal to

$$V_{\epsilon^{(1/3)},S}(H\mathbb{1}_{(0,t)}) + a + \sum_{i=1}^{d} a(1+\epsilon^{(1/3)})S_t^{(i)},$$

which is thus nonnegative itself, and therefore H is $\left(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})\right)$-admissible with respect to S. $\square$

We are now ready to translate the super-replication theorem from [CS06] into our setting. It will play a crucial role in our proof of the Fundamental Theorem. There, necessary and sufficient conditions were given for a contingent claim to be superreplicable by an admissible strategy without initial endowment. However, we will limit ourselves to sufficient conditions.

Theorem 2.2.9 *Let $\epsilon > 0$ and $\epsilon^{(1/3)}$, δ be given as in Lemma 2.2.8 and assume S admits a δ-consistent price system. Consider a random variable $X \in L^0(\mathcal{F}_T, \mathbf{P})$ satisfying*

$$X \geq -a(1 + \sum_{i=1}^{d} S_T^{(i)})$$

for some constant $a > 0$. If for every $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$ we have

$$\mathbb{E}_{\mathbf{Q}}(X) \leq 0,$$

then there exists an $\left(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})(1+\epsilon)\right)$-admissible trading strategy H with respect to the sum numeraire satisfying $V_{\epsilon^{(1/3)},S}(H) \geq X$ a. s.

Proof: First note that by definition we have $\delta \geq \epsilon$, so an ϵ-consistent price system $(S, \mathbf{P})$ and by Lemma 2.2.3 a $\mathcal{K}^*_{\epsilon,\cdot}\backslash\{\mathbf{0}\}$-valued $\mathbf{P}$-martingale Z do exist. For these we have

$$0 \geq \mathbb{E}_{\mathbf{Q}}(X) = \mathbb{E}_{\mathbf{P}}\left(\frac{Z_T^{(1)}}{\mathbb{E}_{\mathbf{P}}(Z_T^{(1)})} X\right) = \frac{1}{\mathbb{E}_{\mathbf{P}}(Z_T^{(1)})}\mathbb{E}_{\mathbf{P}}(Z_T^\top X \mathbf{e}_1).$$

Since we have chosen Z arbitrarily, we have

$$\mathbb{E}_{\mathbf{P}}(Z_T^\top X \mathbf{e}_1) \leq 0$$

for all $\mathcal{K}^*_{\epsilon,\cdot}\backslash\{\mathbf{0}\}$-valued $\mathbf{P}$-martingales Z.
If we fix $b = \frac{a}{1-\epsilon} > a$, the portfolio $X\mathbf{e}_1 + b\mathbf{1}$ has a terminal liquidation value of

$$X + b + b\sum_{i=1}^{d}(1-\epsilon)S_T^{(i)} \geq -a(1 + \sum_{i=1}^{d} S_T^{(i)}) + b + a\sum_{i=1}^{d} S_T^{(i)} > 0,$$

thus $X\mathbf{e}_1 + b\mathbf{1} \in \mathcal{K}_{\epsilon,T}$ and we can finally apply the superreplication theorem [CS06, Theorem 15] to $X\mathbf{e}_1$. This theorem states that there exists an ϵ-self-financing ϵ-admissible portfolio process V satisfying $V_T \geq X\mathbf{e}_1$. Note that

$$X \geq -a(1 + \sum_{i=1}^{d} S_T^{(i)})$$

implies

$$V_T + a(1+\epsilon)\mathbf{1} \in \mathcal{K}_{\epsilon,T},$$

thus by Lemma 2.2.5 we know that V is $\big(\epsilon, a(1+\epsilon)\big)$-admissible. Now Lemma 2.2.6 and 2.2.8 tell us that there exists an $\big(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})(1+\epsilon)\big)$-admissible trading strategy H with

$$V_{\epsilon^{(1/3)},S}(H) \geq V_{\epsilon,S}(H) \geq X.$$

$\square$

Chapter 3

The face-lifting theorem

3.1 Introduction

In this chapter, we study the impact of proportional transaction costs on the superreplication price of path independent options. In discrete time, option hedging was proven to be possible under transaction costs in different frameworks ([BV92], [BLPS92], [ENU93]), but with the rather surprising result that superreplication can in general be performed cheaper than exact hedging - in discrete time, one cannot throw away money like in continuous time, see Remark 2.2.7. In continuous time, it was first conjectured in [DC94] that the European call cannot be superreplicated at a lower price than that of the stock, a phenomenon which may even appear in frictionless incomplete markets, see e. g. [EJ97]. As mentioned in the introduction, this was generalized to a superreplication principle now known as *face-lifting* pricing. Notable versions of the face-lifting theorem have been proved by Bouchard and Touzi [BT00] in a Kabanov-style framework with a d-dimensional diffusion and by Guasoni et al. [GRS08a] for general one-dimensional continuous processes. We present a generalization of the latter to d-dimensional processes which are not necessarily semimartingales nor continuous, but share the so-called *conditional full support* (CFS) property, a property which is shared by virtually all models used in practice.

In what follows, assume that S takes values in a domain $\mathcal{D} \subseteq \mathbb{R}^d_+$ of the form

$$\mathcal{D} = \prod_{i=1}^{d} (a_i, b_i)$$

with $0 \le a_i < S_0^{(i)} < b_i \le \infty$ for all $1 \le i \le d$. Without loss of generality we assume that there exists some $b \in \{0, 1, \ldots, d\}$ such that $S^{(i)}$ is unbounded for

$1 \leq i \leq b$ and bounded for $b < i \leq d$ (if not so, just swap the dimensions). This includes the special cases where all underlyings are unbounded (i. e. $b = d$) or bounded ($b = 0$).
The (CFS) condition introduced in [GRS08a] will be crucial for our proof. However, we will alter the original notion slightly since we consider other domains than just $\mathcal{D} = \mathbb{R}_+^d$, possibly discontinuous processes and a greater class of payoff functions.

Definition 3.1.1 *A $\mathcal{D}$-valued adapted process $(S_t)_{t\in[0,T]}$ satisfies the* conditional full support (CFS) *property, if for all $\epsilon > 0$, $t \in (0, T)$, and all continuous functions $f : [t, T] \to \mathcal{D}$ we have*

$$\mathbf{P}\big(\|S_t - f(t)\| < \epsilon\big) > 0$$

and on the set $\{\|S_t - f(t)\| < \epsilon\}$ we have almost surely

$$\mathbf{P}\big(\sup_{u\in[t,T]} \|S_u - f(u)\| < \epsilon|\mathcal{F}_t\big) > 0.$$

S satisfies the extended conditional full support (ECFS) *property, if for ϵ, f as above and all Borel sets A contained in an ϵ-ball around $f(T)$ and satisfying $\mathbf{P}(S_T \in A) > 0$, we have*

$$\mathbf{P}\big(\|S_t - f(t)\| < \epsilon\big) > 0$$

and on the set $\{\|S_t - f(t)\| < \epsilon\}$ we have almost surely

$$\mathbf{P}\big(\sup_{u\in[t,T]} \|S_u - f(u)\| < \epsilon \text{ and } S_T \in A|\mathcal{F}_t\big) > 0.$$

The property (CFS) reads as follows: From any time point $t \in [0, T)$ on, given any continuous path, S will run arbitrarily close to this path with positive probability. This does, however, not imply S to be continuous itself. In fact, even pure jump processes may have this property. The difference between (CFS) and (ECFS) becomes crucial, at least if we want to drop the semicontinuity assumption on the payoff function necessary in [GRS08a], as Example 3.3.8 will show.

Also see Example 3.3.6 for another process which fails to satisfy a single aspect of the CPS property, which subsequently prevents the face-lifting from holding.

The main result of this chapter, the extension of the face-lifting theorem proved in [GRS08a] to multiple dimensions, will be split into two separate results. As suggested there, we prove a formula for the superreplication price of a given claim using a detour via the static superreplication price. The term "face-lifting" is derived from the fact that the price formula usually involves some notion of *concave envelope* of the payoff function, which visually lifts the function to a smoother one. The notion we require is the following:

Definition 3.1.2 *Let $g : \mathcal{D} \to \mathbb{R}$ be a Borel-measurable function and μ be a Borel measure on $\mathcal{D}$. Then the* concave μ-envelope *of g is defined as*

$$\mathcal{C}(g, \mu) = \inf\{h : \mathcal{D} \to \mathbb{R},\ h \geq g\ (\mu{-}a.e.), h\ is\ concave.\}$$

Note that the classical concave envelope of a function $\mathcal{D} \to \mathbb{R}$ (see also Definition 3.2.5) is equivalently generated when we replace the set of all concave functions by the countable set of all functions of the form $\mathbf{x} \mapsto a + \mathbf{c}^T\mathbf{x}$ with $a \in \mathbb{Q}$ and $\mathbf{c} \in \mathbb{Q}^d$. So the same is true for the concave μ-envelope. So since $\mathcal{C}(g, \mu)$ is the infimum of only countably many functions which are all greater or equal to g, μ-a. e., we have $\mathcal{C}(g, \mu) \geq g$, μ-almost everywhere.

We discuss the initial endowment needed to superreplicate a given option while paying ϵ-transaction costs as defined in Definition 2.1.1. Unless not noted otherwise, we will only consider the constant numeraire $N_t = 1$ throughout this chapter, with the exception of Theorem 3.2.4, which treats the sum numeraire. The problem is split up into two parts. First, we only consider *static* trading strategies, which are restricted to holding a constant position between 0 and T. The resulting *static superreplication price* p_ϵ^s is then compared to the (dynamic) superreplication price p_ϵ, which takes into account all admissible strategies.

Definition 3.1.3 *Denote the set of all ϵ-admissible strategies by $\mathcal{A}$. For a contingent claim $X \in L^0(F_T, \mathbf{P})$, the* superreplication price *given ϵ is defined as*

$$p_\epsilon(X) := \inf\Big(x : x + V_{\epsilon,S}(H) \geq X \textit{ for some } H \in \mathcal{A}\Big).$$

The static superreplication price given ϵ is defined as

$$p_\epsilon^s(X) := \inf\Big(x : x + V_{\epsilon,S}(H) \geq X \textit{ for some } H = \mathbf{c}\mathbb{1}_{(0,T)} \in \mathcal{A}, \mathbf{c} \in \mathbb{R}^d\Big).$$

In both cases we set $\inf \emptyset = +\infty$.

The first consequence of this definition is the inequality

$$p_\epsilon(X) \leq p_\epsilon^s(X),$$

for each claim X, since the set of admissible static strategies is a subset $\mathcal{A}$. Even though the static superreplication price might be of independent interest, we rather use it as a tool for computing the proper superreplication price as $\epsilon \to 0$. In the first part of the face-lifting theorem we find an explicit formula for p_ϵ^s, which is of course a much easier task than computing p_ϵ straight away. In the second part we prove the equivalence of both superreplication prices as $\epsilon \to 0$ and thereby solve the dynamic superreplication problem for small proportional transaction costs. The main results, that sum up to the face-lifting theorem, are as follows:

Theorem 3.1.4 (Face-lifting theorem, part one) *Consider a contingent claim of the form $X = g(S_T)$ with a function $g : \mathcal{D} \to \mathbb{R}$. Then the static superreplication price of X with respect to the transaction cost level ϵ is given by*

$$p_\epsilon^s(X) = \sup\left(\mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{x}), \mathbf{x} \in \mathcal{R}_\epsilon\right),$$

where $\mathcal{C}(g, \mathbf{P}^{S_T})$ is the concave envelope of g with respect to the push-forward measure $\mathbf{P}^{S_T}$ and the rectangle $\mathcal{R}_\epsilon$ is defined as

$$\mathcal{R}_\epsilon = \prod_{i=1}^{b}\left(0, S_0^{(i)}\frac{1+\epsilon}{1-\epsilon}\right] \times \prod_{i=b+1}^{d}\left[S_0^{(i)}\frac{1-\epsilon}{1+\epsilon}, S_0^{(i)}\frac{1+\epsilon}{1-\epsilon}\right].$$

Theorem 3.1.5 (Face-lifting theorem, part two) *Let S satisfy (ECFS) and let the claim X be of the form $X = g(S_T)$. Then for the superreplication price $p_\epsilon(X)$ and the static supereplication price $p_\epsilon^s(X)$ we have*

$$p_{0+}(X) := \lim_{\epsilon \searrow 0} p_\epsilon(X) = p_0^s(X).$$

While 3.1.5 is essentially the same as in [GRS08a], 3.1.4 is hard to recognize as analogue of the one-dimensional version. Note however that if S_T has the whole space $\mathbb{R}_+^d$ as support and if furthermore g is bounded from below and lower semicontinuous, which was all assumed there, then $p_0^s(X)$ boils down to the concave envelope of g evaluated at S_0, as was the case in [GRS08a].

Remark 3.1.6 For an underlying S and a strategy H, define $\tilde{S}$ and $\tilde{H}$ by $\tilde{S}_t = \frac{S_t}{S_0}$ and $\tilde{H}_t = S_0 H_t$. Then straightforward calculation yields $V_{\epsilon,S}(H) = V_{\epsilon,\tilde{S}}(\tilde{H})$. Moreover, S and $\tilde{S}$ have the same set of admissible strategies. This simple transformation can be used to reduce our problem to strategies starting at $\mathbf{1}$.

3.2 Proof of the face-lifting theorem

Before we prove the first part of the face-lifting theorem, we analyze static strategies in greater detail. Assume $S_0 = \mathbf{1}$ and $\mathbf{c} \in \mathbb{R}^d$. The trading strategy $\mathbf{c}\mathbb{1}_{(0,T)}$ yields the final value

$$\begin{aligned} V_{\epsilon,S}(\mathbf{c}\mathbb{1}_{(0,T)}) = &\Big(\sum_{i:c^{(i)}<0} -c^{(i)}(1-\epsilon) + c^{(i)}(1+\epsilon)S_T^{(i)} \Big) \\ &+ \sum_{i:c^{(i)}>0} \Big(-c^{(i)}(1+\epsilon) + c^{(i)}(1-\epsilon)S_T^{(i)} \Big), \end{aligned}$$

which is bounded from below if and only if $c^{(i)} \geq 0$ for all $i \leq b$. In other words: Static admissible strategies are exactly those without short positions in unbounded underlyings.

Note that for an initial endowment $x \in \mathbb{R}$, the final wealth $x + V_{\epsilon,S}(\mathbf{c}\mathbb{1}_{(0,T)})$ is a deterministic function of S_T, which we will denote by $h_{x,\mathbf{c}}$. In fact,

$$h_{x,\mathbf{c}}(\mathbf{s}) = a^{(0)} + \sum_{i=1}^{d} a^{(i)} s^{(i)},$$

where

$$\begin{aligned} a^{(0)} &= x - \sum_{i:c^{(i)}<0} c^{(i)}(1-\epsilon) - \sum_{i:c^{(i)}>0} c^{(i)}(1+\epsilon), \\ a^{(i)} &= c^{(i)}\Big((1-\epsilon)\mathbb{1}_{c^{(i)}<0} + (1+\epsilon)\mathbb{1}_{c^{(i)}>0} \Big), 1 \leq i \leq d. \end{aligned}$$

Since these equations can be inverted to recover x and $\mathbf{c}$ from $a^{(0)}, a^{(1)}, \ldots, a^{(d)}$, there exists a bijection between all static strategies with initial endowment and all payoffs $h(S_T)$, where h is an affine linear function $\mathbb{R}^d \to \mathbb{R}$ (i. e. a hyperplane).

Proof of Theorem 3.1.4: Following 3.1.6, let us first assume that $S_0 = \mathbf{1}$. For given holdings $\mathbf{c}$ and an initial endowment x, note that $x = h_{x,\mathbf{c}}(\mathbf{z})$, where $\mathbf{z} = (z^{(1)}, z^{(2)}, \ldots, z^{(d)})$ and

$$z^{(i)} = \frac{1+\epsilon}{1-\epsilon}\mathbb{1}_{c^{(i)}\geq 0} + \frac{1-\epsilon}{1+\epsilon}\mathbb{1}_{c^{(i)}<0}.$$

Since $a^{(i)}$ and $c^{(i)}$ have the same sign, $h_{x,\mathbf{c}}$ is strictly increasing in those dimensions where $c^{(i)} > 0$ and strictly decreasing where $c^{(i)} < 0$. This means that that $h_{x,\mathbf{c}}(\mathbf{z}) \geq h_{x,\mathbf{c}}(\mathbf{z}')$ for all $\mathbf{z}' \in \left\{\frac{1+\epsilon}{1-\epsilon}, \frac{1-\epsilon}{1+\epsilon}\right\}^d$.
In other words, for a given hyperplane h, we may compute its price (i. e. the initial endowment $x \in \mathbb{R}$ needed to generate the payoff $h(S_T)$) by

$$x(h) = \max\left(h(\mathbf{z}), \mathbf{z} \in \left\{\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right\}^d\right)$$

$$= \max\left(h(\mathbf{z}), \mathbf{z} \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right]^d\right),$$

since the second maximum is attained at one of the corners.
Recall that the μ-concave envelope is equivalently generated by

$$\mathcal{C}(g,\mu) = \inf\Big(h \geq g(\mu{-}a.e.), h \text{ is affine linear}\Big),$$

thus superhedging $g(S_T)$ or $\mathcal{C}(g, \mathbf{P}^{S_T})(S_T)$ by static strategies are equivalent. Assume for a moment that $b = 0$, i. e. all underlyings are bounded and all static strategies are admissible. Then the static superhedging problem boils down to solving

$$p^s_\epsilon(X) = \inf\Big(x(h) : h \geq \mathcal{C}(g, \mathbf{P}^{S_T}) \text{ is affine linear}\Big).$$

Denote by $\hat{x}$ the maximum of $\mathcal{C}(g, \mathbf{P}^{S_T})$ over $\mathcal{R}_\epsilon = \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right]^d$ and by $\hat{\mathbf{y}} \in \mathcal{R}_\epsilon$ one point where this maximum is attained. Obviously

$$p^s_\epsilon(X) \geq \hat{x},$$

since $h \geq \mathcal{C}(g, \mathbf{P}^{S_T})$ implies $x(h) \geq \hat{x}$. On the other hand, since $\mathcal{C}(g, \mathbf{P}^{S_T})$ is concave, there exists a tangent hyperplane $\hat{h}$ in $\hat{\mathbf{y}}$ that satisfies $x(\hat{h}) = \hat{x}$. Thus $p^s_\epsilon(X) = \hat{x}$.
In the general case (i. e. $b > 0$), we are limited to hyperplanes that represent admissible strategies, i. e. that are are non-decreasing in the first b dimensions.

Thus we may, instead of limiting the set of hyperplanes in consideration, find the optimal tangent hyperplane to the function

$$\inf\Big(h \geq g\ (\mathbf{P}^{S_T}\text{-}a.e.), h(\mathbf{y}) = a^{(0)} + \mathbf{a}^\top \mathbf{y};\ \mathbf{a} \in [0,\infty)^b \times \mathbb{R}^{d-b}\Big),$$

which is itself concave and non-decreasing in the first b dimensions. Once again the minimal initial endowment is the maximum over the rectangle $\left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right]^d$, which coincides with the maximum of $\mathcal{C}(g, \mathbf{P}^{S_T})$ over $\mathcal{R}_\epsilon$. This proves the theorem for $S_0 = \mathbf{1}$.

In case $S_0 \neq \mathbf{1}$, consider the transformed underlying $\tilde{S}$ introduced in 3.1.6, as well as an equally scaled payoff function

$$\tilde{g} : \mathbf{x} \mapsto g(S_0 \mathbf{x}).$$

Using $V_{\epsilon,S}(H) = V_{\epsilon,\tilde{S}}(\tilde{H})$ and $g(S_T) = \tilde{g}(\tilde{S}_T)$ and the fact that $\tilde{S}_0 = \mathbf{1}$, we have

$$\begin{aligned} p^s_\epsilon(X) &= \inf\Big(x : x + V_{\epsilon,\tilde{S}}(\tilde{H}) \geq \tilde{g}(\tilde{S}_T) \text{ for some static } \tilde{H} \in \mathcal{A}\Big) \\ &= \sup\left(\mathcal{C}(\tilde{g}, \mathbf{P}^{\tilde{S}_T})(\mathbf{x}), \mathbf{x} \in \tilde{\mathcal{R}}_\epsilon\right), \end{aligned}$$

where $\tilde{\mathcal{R}}_\epsilon = \left(0, \frac{1+\epsilon}{1-\epsilon}\right]^b \times \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon}\right]^{d-b}$ is the scaled version of $\mathcal{R}_\epsilon$. This maximum clearly coincides with the maximum of $\mathcal{C}(g, \mathbf{P}^{S_T})$ over $\mathcal{R}_\epsilon$, which proves the general case. □

Example 3.2.1 (Two-sided Poisson process) The following simple model satisfies neither (CFS) nor (ECFS), but it demonstrates our idea behind the quite technical proof of Theorem 3.1.5: Consider $T = 1$ and $S = e^{X^{(1)} - X^{(2)}}$, where $X^{(1)}$ and $X^{(2)}$ are independent Poisson processes with intensities $\lambda_1, \lambda_2 > 0$ under $\mathbf{P}$. S is a $\mathbf{P}$-martingale if and only if

$$\lambda_1(e-1) = \lambda_2(1-e^{-1}).$$

Note that the intensities can be changed to arbitrary $\mu_1, \mu_2 > 0$ by the equivalent measure change with density

$$\frac{d\mathbf{Q}}{d\mathbf{P}} = \exp(-\mu_1 - \mu_2 + \lambda_1 + \lambda_2) \left(\frac{\mu_1}{\lambda_1}\right)^{X_1^{(1)}} \left(\frac{\mu_2}{\lambda_2}\right)^{X_1^{(2)}}.$$

It suffices to define the payoff function g on the grid

$$G = \operatorname{supp}(S_1) = \{e^k, k \in \mathbb{Z}\}.$$

Assume the simplest case, where g is bounded from below and therefore

$$p_0^s\big(g(S_T)\big) = \mathcal{C}(g, \mathbf{P}^{S_T})(1),$$

and that there is a tangent to g touching g at some points $a, b \in G$, $a < 1 < b$. Hence $1 = \theta a + (1-\theta)b$ and

$$\mathcal{C}(g, \mathbf{P}^{S_T})(1) = \theta g(a) + (1-\theta)g(b).$$

Define τ as the first time S hits either a or b. By increasing the intensities as described above, we may find $\mathbf{Q}_\delta \sim \mathbf{P}$ such that S is still a $\mathbf{Q}_\delta$-martingale and $\mathbf{Q}_\delta(\tau < 1) > 1 - \delta$, for arbitrary $\delta > 0$. Now consider the event A that S stays constant between τ and T, which has positive probability. So under the measure $\mathbf{R}_\delta \ll \mathbf{P}$, defined by $\frac{d\mathbf{R}_\delta}{d\mathbf{Q}_\delta} = \frac{\mathbb{1}_A}{\mathbf{Q}_\delta(A)}$, S is a martingale satisfying

$$\mathbf{R}_\delta(S_T \in \{a, b\}) > 1 - \delta.$$

Now assume H superreplicates $g(S_T)$ with initial endowment x, hence

$$g(S_T) \leq x + V_{0,S}(H).$$

Taking expectations w. r. t. $\mathbf{R}_\delta$ and minding Lemma 2.1.7, we get $x \geq \mathbb{E}_{\mathbf{R}_\delta}\big(g(S_T)\big)$. Now we have

$$\lim_{\delta \to 0} \mathbb{E}_{\mathbf{R}_\delta}\big(g(S_T)\big) = g(a)\mathbf{R}_\delta(S_T = a) + g(b)\mathbf{R}_\delta(S_T = b),$$

and the probabilities of S_T ending up in a or b tend to θ and $1-\theta$ as S is an $\mathbf{R}_\delta$-martingale and therefore $\mathbb{E}_{\mathbf{R}_\delta}(S_{\tau \wedge T}) = 1$ for all δ, hence

$$x \geq \theta g(a) + (1-\theta)g(b) = p_0^s\big(g(S_T)\big).$$

Two cases were not covered yet: The first one is $g(1) = \mathcal{C}(g, \mathbf{P}^{S_T})(1)$. There we can argue with a measure $\mathbf{R} \ll \mathbf{P}$ such that S is $\mathbf{R}$-a. s. constant. The second one arises when the tangent to $\mathcal{C}(g, \mathbf{P}^{S_T})$ in 1 touches g only once in some $a < 1$. There we argue as in the general case and letting $b \to \infty$. Overall, this proves the face-lifting theorem for S.

So the idea is to argue on tangents supporting $\mathcal{C}(g, \mathbf{P}^{S_T})$ in S_0. In dimension d and for arbitrary S, this of course becomes much more involved, as the tangents there become hyperplanes and the two points a, b have to be replaced by $d+1$ sets of positive $\mathbf{P}^{S_T}$-measure. Moreover, S can in general not be turned into a martingale with a suitable change of measure, which is why we need to argue on a suitable CPS M instead. Most of the technicalities involved in this approach are covered by the following lemma.

Lemma 3.2.2 *Let S satisfy (ECFS), with $S_0 = \mathbf{1}$. Let $\epsilon, \delta > 0$ and vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1} \in \mathcal{D}$ be given such that $\mathbf{1}$ lies in the interior of* $\operatorname{conv}(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1})$. *Let in addition Borel sets*

$$B_m \subseteq \mathbf{x}_m \left[\frac{1}{1+\epsilon}, 1+\epsilon\right]^d, \;\; 1 \leq m \leq d+1$$

be given such that $\mathbf{P}(S_T \in B_m) > 0$ for all m. Then there exists a consistent price system $(M, \mathbf{Q}) \ll_\epsilon (S, \mathbf{P})$ such that

1. $\mathbf{Q}\big(M_T \in \operatorname{conv}(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1})\big) = 1$
2. $\mathbf{Q}\big(M_T \in \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1}\}\big) \geq 1 - \delta$
3. $\mathbf{Q}(S_T \in B_m | M_T = \mathbf{x}_m) = 1$ *for all* $1 \leq m \leq d+1$

Proof: In order to simplify calculations, we will construct a martingale that satisfies the geometric condition

$$\frac{S_t}{1+\epsilon} \leq M_t \leq S_t(1+\epsilon)$$

rather than the original CPS condition, see Remark 2.1.4.2. As $\frac{1}{1+\epsilon} > 1-\epsilon$, this process will indeed be an ϵ-consistent price system.
Fix $K, N \in \mathbb{N}$, the time grid $t_0, t_1, \ldots, t_K$ by $t_k = \frac{kT}{K}$ and vectors $\mathbf{x}_m^n \in \mathcal{D}$ by

$$\mathbf{x}_m^n = \mathbf{1}\frac{N-n}{N} + \mathbf{x}_m \frac{n}{N}, 1 \leq n \leq N, \; 1 \leq m \leq d+1,$$

so we have $\mathbf{x}_m^0 = \mathbf{1}$ and $\mathbf{x}_m^N = \mathbf{x}_m$. The idea is to first construct M as discrete-time Markov chain moving on the grid just defined and being absorbed at either of the end points. To do so, fix $\mu, \eta > 0, \mathbf{x} \in \mathbb{R}_+^d$ and define the set

$$A_{k,\mathbf{x},\eta,\mu} = \left\{\frac{\mathbf{x}}{1+\eta} < S_{t_k} - S_{t_{k-1}} < (1+\eta)\mathbf{x}\right\} \bigcap$$
$$\left\{\frac{S_{t_{k-1}}}{1+\mu} < S_t < S_{t_{k-1}}(1+\mu) \;\forall\, t \in [t_{k-1}, t_k]\right\}.$$

Now define the set

$$\Omega_0 = \bigcup_{m=1}^{d+1} \bigcup_{\mathbf{y}_1 \ldots, \mathbf{y}_K \in \{0, \pm \mathbf{x}_m^1 - \mathbf{1}\}} \bigcap_{k=1}^{K} A_{k,\mathbf{y}_k,\eta^k,\mu}.$$

Note that the sets $\bigcap_{k=1}^{K} A_{k,\mathbf{y}_k,\eta^k,\mu}$ are mutually disjoint if we choose η small enough, which we will assume. On Ω_0, we define a discrete time process M by $M_0 = \mathbf{1}$ and

$$M_{t_{k+1}} = \sum_{m=1}^{d+1} \sum_{\mathbf{y} \in \{0, \pm \mathbf{x}_m^1 - \mathbf{1}\}} (M_{t_k} + \mathbf{y}) \mathbb{1}_{A_{k,\mathbf{y},\eta^k,\mu}}.$$

Right now, M is a random walk starting at $\mathbf{1}$ and running on an infinite grid spanned by the vectors $\pm \mathbf{x}_m^1$. Out of all these trajectories, we choose all paths that satisfy for each $k < K$

(i) $M_{t_k} = \mathbf{x}_m^n$ for some $1 \leq n < N \Rightarrow M_{t_{k+1}} \in \{\mathbf{x}_m^{n-1}, \mathbf{x}_m^{n+1}\}$

(ii) $M_{t_k} = \mathbf{x}_m^N \Rightarrow M_{t_{k+1}} = \mathbf{x}_m^N$

(iii) $M_T = \mathbf{x}_k \Rightarrow S_T \in B_k$

and denote the corresponding subset of Ω_0 by Ω_1. So on Ω_1, M runs only on the set $\{\mathbf{x}_m^n, 1 \leq n \leq N, 1 \leq m \leq d+1\}$ and is stopped whenever it reaches one of the corners. Note that all the paths featured in Ω_1 have positive probability due to the ECFS property. The conditions 1 and 3 of our lemma are now already satisfied.
By construction of Ω_1, the conditional probability

$$\mathbf{P}(A_{k,\mathbf{x}_m^1 - \mathbf{1},\eta^k,\mu} | \mathcal{F}_{t_{k-1}})$$

is a. s. positive on $\Omega_1 \cap \left\{M_{t_{k-1}} = \mathbf{1}\right\}$ and $\Omega_1 \cap \left\{M_{t_{k-1}} \in \{\mathbf{x}_m^1 \ldots, \mathbf{x}_m^{N-1}\}\right\}$ while being zero on all other parts of Ω_1. There exists therefore a probability measure $\mathbf{Q} \sim \mathbf{P}|_{\Omega_1}$ satisfying

$$\mathbf{Q}(A_{k,\mathbf{x}_m^1,\eta^k,\mu} | \mathcal{F}_{t_{k-1}}) = \lambda_m \mathbb{1}_{\left\{M_{t_{k-1}} = \mathbf{1}\right\}} + \frac{1}{2} \mathbb{1}_{\left\{M_{t_{k-1}} \in \{\mathbf{x}_m^1, \ldots, \mathbf{x}_m^{N-1}\}\right\}}$$

$$\mathbf{Q}(A_{k,-\mathbf{x}_m^1,\eta^k,\mu} | \mathcal{F}_{t_{k-1}}) = \frac{1}{2} \mathbb{1}_{\left\{M_{t_{k-1}} \in \{\mathbf{x}_m^1, \ldots, \mathbf{x}_m^{N-1}\}\right\}},$$

where $\mathbf{x} = \sum_{m=1}^{d+1} \lambda_m \mathbf{x}_m^1$ is a representation of $\mathbf{x}$ as strict convex combination of $\mathbf{x}_1^1, \mathbf{x}_2^1, \ldots, \mathbf{x}_{d+1}^1$. This makes M a time-homogeneous discrete-time Markov chain as well as a martingale. Note that $\lambda_1, \lambda_2, \ldots, \lambda_{d+1}$ and therefore the transition properties of the Markov chain do not depend on our choice of K or N.
We now extend M to a $\mathbf{Q}$-martingale on the whole of $[0, T]$ by letting

$$M_t = \mathbb{E}_{\mathbf{Q}} \left(M_{t_k} \mid \mathcal{F}_t \right)$$

for $t_{k-1} < t < t_k$ and $1 \leq i \leq d$. What is left to show is that $\frac{1}{1+\epsilon} \leq \frac{S_t^{(i)}}{M_t^{(i)}} \leq 1+\epsilon$ for all $t \in [0,T]$. For $t \in [t_k, t_{k+1}]$ we have

$$\frac{S_t^{(i)}}{M_t^{(i)}} = \frac{S_t^{(i)}}{S_{t_k}^{(i)}} \frac{S_{t_k}^{(i)}}{M_{t_k}^{(i)}} \frac{M_{t_k}^{(i)}}{M_t^{(i)}}.$$

The first factor $\frac{S_t^{(i)}}{S_{t_k}^{(i)}}$ is bounded by $1+\mu$, whereas the second factor satisfies

$$\frac{S_{t_k}^{(i)}}{M_{t_k}^{(i)}} \leq \prod_{j=1}^{k} (1+\eta^k) \leq e^{\eta/(1-\eta)} \leq 1+2\eta$$

and for the third one, we have $\frac{M_{t_k}^{(i)}}{M_t^{(i)}} \leq \frac{M_{t_k}^{(i)}}{M_{t_{k+1}}^{(i)}}$, which can be bounded arbitrarily close to 1 by choosing N large enough (and therefore all $\mathbf{x}_m^n$ and $\mathbf{x}_m^{n\pm1}$ close enough to each other). Summing up, since we are still free to choose μ, ν and N freely, we can make sure that both $\frac{S_t^{(i)}}{M_t^{(i)}} \leq 1+\epsilon$ and by the same strategy, $\frac{M_t^{(i)}}{S_t^{(i)}} \leq 1+\epsilon$. Thus for sufficiently small μ and η and large enough N, we have $(M, \mathbf{Q}) \ll_\epsilon (S, \mathbf{P})$.

Finally, consider once more the discrete-time $\mathbf{Q}$-Markov chain M, whose transition probabilities do not depend on our choice of K and where the absorbing states are exactly $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1}\}$ while all other states are transient. For such a Markov chain with infinite time horizon, almost every path gets absorbed in finite time, (see e. g. [Be00, Proposition 5.1]). Since the number K of time steps was chosen arbitrarily, we may choose it large enough to ensure that the $\mathbf{Q}$-probability of being absorbed before $t_K = T$ is greater than $1-\delta$. This completes the proof. □

Remark 3.2.3 Property *(iii)* in the above selection of the "good" paths justifies the inclusion of the set A in the definition of extended conditional full support in 3.1.1. Here the sets $B_1, B_2, \ldots, B_{d+1}$ play the role of A. It will become clear in the following proof why these sets are needed. Consider a tangent hyperplane h to $\mathcal{C}(g, \mathbf{P}^{S_T})$ at $\mathbf{1}$. This hyperplane is supported by g (i. e. touches g at) such sets $B_1, B_2, \ldots, B_{d+1}$. If we, instead of trading on S, trade without transaction costs on a CPS M ending up in one of those sets most likely, we can actually forget about g and just try to hedge h. Since M is a martingale, this can't be done at a cheaper price than the value of h at S_0, which results in the missing inequality. This idea is carried out in the following proof. However, the notion of h supporting g in some sets needs to be made rigorous, since in general a tangent h need not be equal to g on a set of positive $\mathbf{P}^{S_T}$-measure.

Proof of Theorem 3.1.5: Once again, assume $S_0 = \mathbf{1}$. Since $p_\epsilon(X) \leq p_\epsilon^s(X)$ was clear and since $p_{0+}^s(X) = p_0^s(X)$ by the continuity of $\mathcal{C}(g, \mathbf{P}^{S_T})$, what is left to show is that $p_{0+}(X) \geq p_0^s(X)$. Let $x \in \mathbb{R}$ and some $H \in \mathcal{A}$ be given such that

$$x + V_{\epsilon,S}(H) \geq g(S_T) \ \mathbf{P}\text{-a.s}$$

for some $\epsilon > 0$. Recall that $H \in \mathcal{A}$ implies

$$C = \operatorname{essinf}_{\mathbf{P}}\big(V_{\epsilon,S}(H)\big)$$

to be finite. Thus the pair (x, H) also superreplicates $g_1(S_T)$, where

$$g_1 = \max(g, x + C).$$

Assume first that $p_0^s(X) = \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1})$. Fix some positive $\delta < \epsilon$. By definition of the concave $\mathbf{P}^{S_T}$-envelope, there exist some $\eta \in (0, \epsilon)$, vectors $\mathbf{x}_m \in \mathcal{D}$ and sets $B_m \subseteq \mathcal{D}, 1 \leq m \leq d+1$, such that $\mathbf{P}(S_T \in B_m) > 0$ for all $1 \leq m \leq d+1$ and such that for all $\mathbf{b}_1 \in B_1, \mathbf{b}_2 \in B_2, \ldots, \mathbf{b}_{d+1} \in B_{d+1}$ we have

1. $1 - \eta \leq \frac{\mathbf{x}_m}{\mathbf{b}_m} \leq 1 + \eta$ for all m
2. $\mathbf{1} \in \operatorname{conv}(\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_{d+1})$
3. if $\mathbf{1} = \sum_{m=1}^{d+1} \lambda_m \mathbf{b}_m$ is the unique convex combination of $\mathbf{1}$ by $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_{d+1}$, then
$$\sum_{m=1}^{d+1} \lambda_m g(\mathbf{b}_m) \geq \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1}) - \delta.$$

Now Lemma 3.2.2 secures the existence of $(M, \mathbf{Q}) \ll_\eta (S, \mathbf{P})$ such that

$$\begin{aligned} \mathbf{Q}\big(M_T \in \operatorname{conv}(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1})\big) &= 1 \\ \mathbf{Q}\big(M_T \in \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{d+1}\}\big) &\geq 1 - \delta \\ \mathbf{Q}(S_T \in B_m | M_T = \mathbf{x}_m) &= 1 \end{aligned}$$

for all m. Thus by Lemma 2.1.5

$$g_1(S_T) \leq x + V_{\eta,S}(H) \leq x + V_{0,M}(H).$$

Now $t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)})$ is a $\mathbf{Q}$-supermartingale by Lemma 2.1.7, so

$$\mathbb{E}_{\mathbf{Q}}\left(\sum_{i=1}^{d}\int_0^T H_t^{(i)} dM_t^{(i)}\right) \leq 0$$

so taking expectations with respect to $\mathbf{Q}$ yields

$$\mathbb{E}_{\mathbf{Q}}\big(g_1(S_T)\big) \leq x.$$

Using the relations between M_T and S_T, we can estimate this by

$$\mathbb{E}_{\mathbf{Q}}\big(g_1(S_T)\big) \;\geq (1-\delta)\,\big(\mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1}) - \delta\big) + \delta(C + x).$$

Since δ was chosen arbitrarily, we have

$$x \geq \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1}) = p_0^s(X),$$

which proves the theorem in this case.
For the case $p_0^s(X) \neq \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1})$, which can only occur when $b > 0$, note that since g_1 is bounded from below, $\mathcal{C}(g_1, \mathbf{P}^{S_T})$ is increasing in the first b components and thus admits its maximum over $\mathcal{R}_0$ at $\mathbf{1}$. Hence we are back in the first case and

$$x \geq p_0^s\big(g_1(S_T)\big) \geq p_0^s\big(g(S_T)\big).$$

Now for general $S_0 \neq \mathbf{1}$, we already know that trading on S and $\tilde{S}$ yields the same set of possible final values, and that $g(S_T) = \tilde{g}(\tilde{S}_T)$. Thus replacing S and g by $\tilde{S}$ and $\tilde{g}$ leaves p_ϵ as well as p_ϵ^s constant, while $\mathcal{R}_\epsilon$ is transformed likewise. This reduces the general case to $S_0 = \mathbf{1}$ which is already solved. $\square$

Having proved the face-lifting theorem for the numeraire process $N_t = 1$, we can easily derive an analogous result for the sum numeraire

$$N_t = 1 + \sum_{i=1}^{d} S_t^{(i)}.$$

Assume again that $S_0 = \mathbf{1}$. Revisiting the value of a static strategy $H = \mathbf{c}\mathbb{1}_{(0,T)}$, we see that

$$V_{\epsilon,S}(H\mathbb{1}_{(0,t)}) \geq -\sum_{i=1}^{d} |c^{(i)}|(1+\epsilon) - \sum_{i=1}^{d} |c^{(i)}|(1+\epsilon)S_t^{(i)}$$

$$\geq -(1+\epsilon)\|\mathbf{c}\|_\infty \sum_{i=1}^{d}(1+S_t^{(i)}) \geq -d(1+\epsilon)\|\mathbf{c}\|_\infty N_t,$$

where $\|\mathbf{c}\|_\infty = \max(|c^{(1)}|, |c^{(2)}|, \ldots, |c^{(d)}|)$. So for this choice of numeraire H is $\big(\epsilon, d(1+\epsilon)\|\mathbf{c}\|_\infty\big)$-admissible. Clearly, admissibility does not change if we let S start at an arbitrary point other than $\mathbf{1}$.

Theorem 3.2.4 (Face-lifting theorem for the sum numeraire) *Let $X = g(S_T)$ as in Theorem 3.1.4. Consider the admissibility condition imposed by the sum numeraire $N_t = 1 + \sum_{i=1}^{d} S_t^{(i)}$. The the following hold:*

1. *the static superreplication price of X with respect to ϵ is given by*

$$p_\epsilon^s(X) = \sup\left(\mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{x}), \mathbf{x} \in \mathcal{R}_\epsilon\right),$$

   where the rectangle $\mathcal{R}_\epsilon$ is defined as

$$\mathcal{R}_\epsilon = \prod_{n=1}^{d}\left[S_0^{(n)}\frac{1-\epsilon}{1+\epsilon}, S_0^{(n)}\frac{1+\epsilon}{1-\epsilon}\right].$$

2. *If S satisfies (ECFS), then we have*

$$p_{0+}(X) := \lim_{\epsilon \searrow 0} p_\epsilon(X) = p_0^s(X).$$

Proof: As we have seen, all static strategies are ϵ-admissible now. Thus the first part of the proof of Theorem 3.1.4, where $b = 0$ was assumed, applies to this case, for arbitrary b. This proves (1).
For (2), the first part of the proof of Theorem 3.1.5 applies, where we assumed $p_0^s(X) = \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1})$. Note that the proof uses Lemma 2.1.7, which works for the sum numeraire, too. □

There are two things about our version of the face-lifting theorem that make it somewhat inconvenient and hard to recognize as multidimensional version of the original theorem from [GRS08a]. The measure μ in Definition 3.1.2 and the distinction between (CFS) and (ECFS) in Definition 3.1.1. These were put in place to allow Theorem 3.1.4 to hold true without (CFS) or (ECFS) satisfied, and to allow a greater variety of payoff functions. In most cases relevant in practice, however, our results can be simplified.

Definition 3.2.5 *For a function $g : \mathcal{D}$, we define the* concave envelope *of g as*

$$\mathcal{C}(g) = \inf\{h : \mathcal{D} \to \mathbb{R},\ h \geq g,\ h\ is\ concave.\}$$

Corollary 3.2.6

1. *If g is in every point either lower semicontinuous, left continuous or right continuous and* $\operatorname{supp}(S_T) = \overline{\mathcal{D}}$*, then*

   $$p^s_\epsilon(X) = \sup\big(\mathcal{C}(g)(\mathbf{x}), \mathbf{x} \in \mathcal{R}_\epsilon\big),$$

   where $\mathcal{R}_\epsilon$ is as in Theorem 3.1.4 or Theorem 3.2.4, depending on the numeraire considered.

2. *If S is unbounded in every dimension and g is bounded from below $\mathbf{P}^{S_T}$-almost everywhere, then for both numeraires*

   $$p^s_\epsilon(X) = \mathcal{C}(g, \mathbf{P}^{S_T})\left(\frac{1+\epsilon}{1-\epsilon} S_0\right).$$

3. *If both (1) and (2) are satisfied, then*

   $$p^s_\epsilon(X) = \mathcal{C}(g)\left(\frac{1+\epsilon}{1-\epsilon} S_0\right).$$

Proof: (1): It suffices to note that in this case $\mathcal{C}(g, \mathbf{P}^{S_T})$ and $\mathcal{C}(g)$ coincide.
(2): In this case $\mathcal{C}(g, \mathbf{P}^{S_T})$ is increasing in every component, hence the maximum of $\mathcal{C}(g, \mathbf{P}^{S_T})$ is admitted in the outermost corner of $\mathcal{R}_\epsilon$.
(3) follows from (1) and (2). □

Corollary 3.2.7 *Assume g is in every point either lower semicontinuous, left continuous or right continuous. Then Theorem 3.1.5 and 3.2.4 hold true with (CFS) instead of (ECFS).*

Proof: It suffices in this case to redo the proofs of Lemma 3.2.2 and Theorem 3.1.5 with the sets B_m replaced by $\mathbf{x}_m[\frac{1}{1+\epsilon}, 1+\epsilon]^d$ and $\mathbf{x}_m[1-\eta, 1+\eta]^d$, respectively. □

Remark 3.2.8 (Open problems) The evident explanation for the fact that under small transaction costs, dynamic strategies are not superior to static ones is that the friction discourages all trades, so strategies that trade less are favoured. If this explanation is right, then the inequality should hold even more so the higher the costs are, i.e. we should have $p_\epsilon(X) = p^s_\epsilon(X)$ for all $\epsilon > 0$ instead of just $p_{0+}(X) = p^s_{0+}(X)$. In the mentioned predecessing works on the European call and similar options, it was usually proved that $p_\epsilon(X) \geq p^s_{0+}(X)$, which trivially follows from our face-lifting theorem and the fact that p_ϵ is increasing in ϵ. Our proof critically depends on the assumption that we may choose ϵ arbitrarily small, so it cannot easily be extended to the case of fixed ϵ.

In a related issue it is still to be shown whether $p_0(X) = p^s_0(X)$. The identity $p^s_{0+}(X) = p^s_0(X)$ is clear by Theorem 3.1.4 and the continuity of $\mathcal{C}(g, \mathbf{P}^{S_T})$. Recall that for $\epsilon = 0$, we formally still impose H to be of finite variation. But the question if $p_{0+}(X) = p_0(X)$ or not still needs to be addressed. If this is the case, then one may argue that transaction costs affect the superreplication price only indirectly in that they lock out strategies of infinite variation. So the real source of market imperfection would be the variation constraint, not the transaction costs.

3.3 Examples

This section presents some examples to clarify the implications, but also the limitations of the face-lifting theorem. In all examples except for 3.3.8, it suffices to assume (CFS) instead of (ECFS), since all considered payoff functions qualify for Theorem 3.2.7. Example 3.3.8, however, is designed to highlight the difference between (CFS) and (ECFS).

As Guasoni et al. point out [GRS08a, Section 4], (CFS) is satisfied for a variety of Markov processes, while (ECFS) was not studied there. As the proof of Theorem 3.1.5 has shown, the distinction between (CFS) and (ECFS) is crucial. However, under fairly mild assumptions they are in fact equivalent:

Lemma 3.3.1 *Let S be a Markov process that satisfies (CFS). If S_T has a density, then (ECFS) holds.*

Proof: For $\mathbf{x} \in \mathbb{R}^d$ and $\delta > 0$, denote by $B_\delta(\mathbf{x})$ the open δ-Ball around $\mathbf{x}$ with radius δ. Let A, t, f and ϵ be given as in Definition 3.1.1. Then $\mathbf{P}(S_T \in A) > 0$ means that A has positive Lebesgue measure. For arbitrary $\mathbf{s} \in B_\epsilon\big(f(t)\big)$ and $\mathbf{x} \in B_\epsilon\big(f(T)\big)$ and $\delta > 0$ such that $B_\delta(\mathbf{x}) \subset B_\epsilon\big(f(T)\big)$, find a continuous function $g : [t, T] \to \mathcal{D}$ and $\eta > 0$ such that $g(t) = \mathbf{s}$, and both

$$B_\eta\big(g(u)\big) \subset B_\epsilon\big(f(u)\big) \ \ \forall\, u \in (0, T)$$

$$B_\eta\big(g(T)\big) \subset B_\delta\big(\mathbf{x}\big)$$

hold true. Then (CFS), applied to g and η yields that

$$\mathbf{P}\big(\sup_{u\in[t,T]}\|S_u - f(u)\| \le \epsilon \text{ and } S_T \in B_\delta(\mathbf{x})|S_t = \mathbf{s}\big) > 0,$$

at least for $\mathbf{P}^{S_t}$-almost all $\mathbf{s} \in B_\epsilon\big(f(t)\big)$. The Markov property now yields

$$\mathbf{P}\big(\sup_{u\in[t,T]}\|S_u - f(u)\| \le \epsilon \text{ and } S_T \in B_\delta(\mathbf{x})|F_t\big) > 0$$

a. s. on the set $\{S_t \in B_\epsilon\big(f(t)\big)\}$. The same is then true if we replace $B_\delta(\mathbf{x})$ by any Borel subset of $B_\epsilon\big(f(T)\big)$ of positive Lebesgue measure, e. g. A. This proves (ECFS). $\square$

Example 3.3.2 (European standard options) The most-quoted results similar to our face-lifting theorem are arguably European put and call options with strike $K > 0$, represented by the payoff functions

$$X_c = g_c(S_T) = \max(S_T - K, 0)$$
$$X_p = g_p(S_T) = \max(K - S_T, 0).$$

It can easily be seen that the concave envelope of g_c is the identity, independent of K, whereas the concave envelope of g_p is constant at K. So for a price process S starting at 1 with conditional full support on $\mathbb{R}_+$, such as in the Black-Scholes model (it will be shown in Theorem 4.3.3 that (ECFS) holds true there), we get $p_{0+}(X_c) = 1$ for the call and $p_{0+}(X_p) = K$ for the put. The call is best superhedged by buying the asset, whereas the put is best superreplicated by not trading and keeping K on the cash account for the worst case. This holds true for both numeraires considered.

Example 3.3.3 (European standard options in the Bachelier model) ·
Although our model only treats strictly positive price processes, the groundbreaking Bachelier model [Bac00] also deserves consideration. For simplicity let S_t be a standard Brownian motion starting at 1. Our definitions of trading strategies, values, admissibility, extended conditional full support and consistent price systems can naturally be extended to this case, and S obviously satisfies (ECFS), as the Black-Scholes model does. It is easily seen that for the numeraire $N_t = 1$, no static strategy is ϵ-admissible but the trivial $H = 0$, even for $\epsilon = 0$. For $N_t = 1 + S_t$ instead, every static strategy is ϵ-admissible, for every $\epsilon < 1$. But on the whole real line, both the concave envelope of g_p and of g_c are equal to $+\infty$, since there exists no concave function above any of the two. In other words, puts and calls cannot be superreplicated in this model. The same holds true for every unbounded payoff function.

Example 3.3.4 (Power options) Similar to the last example, power call options are in some cases not superreplicable under transaction costs, if the underlying is unbounded. Consider a process S starting at 1 with conditional full support on $\mathbb{R}_+$ and the contingent claim

$$X_\alpha = g_\alpha(S_T) = \max(S_T - K, 0)^\alpha$$

for some strike $K > 0$ and an exponent $\alpha > 0$. For $\alpha = 1$ this equals the standard European call. For $\alpha < 1$ we get

$$\mathcal{C}(g_\alpha, \mathbf{P}^{S_T})(x) = qx\mathbb{1}_{(0,p]}(x) + g_\alpha(x)\mathbb{1}_{(p,\infty)}(x),$$

where $p = \frac{K}{1-\alpha} > K$ and $q = \alpha(\frac{\alpha K}{1-\alpha})^{\alpha-1} > 0$. Hence we have

$$p_{0+}(X_\alpha) = q\mathbb{1}_{1\leq p} + (1-K)^\alpha \mathbb{1}_{1>p}.$$

If on the other hand $\alpha > 1$, then g_α outgrows every concave function for $S_T \to \infty$, hence $\mathcal{C}(g_\alpha, \mathbf{P}^{S_T}) = \infty$. So in this case, the option can not be superreplicated for any $\epsilon > 0$.

Example 3.3.5 (Rainbow forwards) Options depending on the maximum or minimum of several stocks are frequently termed *rainbow options* by practitioners, see e. g. [OW06]. Assume a three-dimensional underlying $S = (S^{(1)}, S^{(2)}, S^{(3)})$ with conditional full support on $\mathcal{D} = \mathbb{R}^3_+$, starting at $(1,1,1)$. Consider the following four claims:

$$X_{\text{max}} = \max(S_T^{(1)}, S_T^{(2)}, S_T^{(3)})$$

$$X_{\text{min}} = \min(S_T^{(1)}, S_T^{(2)}, S_T^{(3)})$$

$$X_{\text{med}} = \sum_{i=1}^{3} S_T^{(i)} - X_{\text{max}} - X_{\text{min}}$$

$$X_{\text{mean}} = \frac{1}{3}\sum_{i=1}^{3} S_T^{(i)}.$$

While X_{mean} can be perfectly replicated by buying $\frac{1}{3}$ of each stock and thus $p_{0+}(X_{\text{mean}}) = 1$, The face-lifting theorem yields $p_{0+}(X_{\text{max}}) = 3$ and $p_{0+}(X_{\text{med}}) = \frac{3}{2}$ and finally $p_{0+}(X_{\text{min}}) = 1$. This means: The maximum cannot be superhedged any better way than buying all three assets, X_{min} is best hedged by buying a third of each asset (just like X_{mean}, even though $X_{\text{mean}} < X_{\text{min}}$ a. s.), and X_{med} is best hedged by buying one half each.

Example 3.3.6 (Piecewise, but no joint CFS) To illustrate how crucial the joint CFS in all dimensions (and not only CFS for each single asset) is for our theorem, consider the following two-dimensional example: Let

$$(S^{(1)}, S^{(2)}) = (e^{L^{(1)}}, e^{L^{(2)}}),$$

where $(L^{(1)}, L^{(2)})$ is a compound Poisson process whose jumps are i.i.d. uniformly distributed on the set

$$\big([0,1] \times \{1\}\big) \cup \big(\{1\} \times [0,1]\big) \cup \big([-1,0] \times \{-1\}\big) \cup \big(\{-1\} \times [-1,0]\big).$$

First of all, it is easily seen that both $L^{(1)}$ and $L^{(2)}$ are themselves compound Poisson processes which jump measures consisting of a uniform distribution on $(-1,1)$ plus two Dirac measures on each $\{-1\}$ and $\{1\}$, hence by Theorem 4.3.3 both $S^{(1)}$ and $S^{(2)}$ satisfy the one-dimensional CFS property. Note furthermore that $\mathrm{supp}(L_T) = \mathbb{R}^2$, since e. g. every point in $[-1,1]^2$ can be reached in minimum three jumps. Hence $\mathrm{supp}(S_T) = \mathbb{R}^2_+$.
Now consider the claim $X = g(S_T)$, where $g = \mathbb{1}_{(\frac{3}{2},\infty)^2}$, which pays 1 if both assets gain more than 50%. Using Theorem 3.1.4 (which did not require ECFS), the static superreplication price is easily computed to be

$$p_0^s(X) = \mathcal{C}(g, \mathbf{P}^{S_T})(\mathbf{1}) = \frac{2}{3},$$

where the optimal static strategy consists of holding $\frac{1}{3}$ unit of each stock. Now observe the following strategy for $\epsilon = 0$: Start with a holding of $a = \frac{3e-2}{3(e^2-e-2)}$ in each asset and additionally $b = 1-(1+e)a$ of cash. The initial endowment $2a + b \approx 0.565$ is smaller than $p_0^s(X)$.
The constants a and b are designed to ensure $a(1+e)+b = 1$, which is our minimal wealth in case the first jump goes upwards. So in this case we immediately sell and have therefore superreplicated X. If the first jump goes down, we still own at least $\frac{2a}{e} + b = \frac{2}{3e}$, which is enough to buy $\frac{2}{3}$ of the one stock that has fallen to $\frac{1}{e}$. If this stock should ever reach $\frac{3}{2}$ (i. e. get in the money), we end up with at least 1 and have once again superreplicated X. Thus for $\epsilon = 0$, this strategy beats the optimal static one.
For $\epsilon > 0$, note that the total amount of transaction costs produced by our strategy is bounded by ϵC, with some constant C. This means that for some $\epsilon > 0$ small enough, the strategy remains superior to the static strategy, hence $p_{0+}(X) < p_0^s(X)$. This argument also holds true in the next examples.

Example 3.3.7 Example 3.3.6 did not satisfy (ECFS), for instance if $A = \{S_0\}$ and f forces the process to move at least once. There is an easy example to illustrate that the phenomenon highlighted there can also occur with an (ECFS) process, and that jumps are not necessary, either: Consider a two-dimensional geometric Brownian motion starting at $(1,1)$, conditioned on the event

$$\big\{S_t \notin \{2\} \times [0,4]\ \forall\, t \in [0,T]\big\}.$$

Here both $S^{(1)}$ and $S^{(2)}$ satisfy (ECFS), while S does not. The payoff function $g = \mathbb{1}_{[3,\infty)\times\mathbb{R}_+}$ pays if the first asset gains at least 200%. For this to happen, $S^{(1)}$ must first climb above 4 (denote the first time this happens by τ), thus $H = (0, \frac{1}{4})\mathbb{1}_{(0,\tau\wedge T)}$ superreplicates $g(S_T)$ cheaper than the optimal static strategy $(\frac{1}{3}, 0)\mathbb{1}_{(0,T)}$.

Example 3.3.8 (The set A is essential in (ECFS)) The next is a pathological example to to illustrate that the difference between (CFS) and (ECFS) is crucial when considering non-semicontinuous payoff functions. Let a standard Brownian motion W and countably many Brownian bridges $(W^{(q)})_{q\in\mathbb{Q}}$ be given, each satisfying $W_T^{(q)} = q$ almost surely and all being independent of each other and of W. Further consider a decomposition of $\Omega = \bigcup_{q\in\mathbb{Q}} A_q$ into disjoint sets A_q, all having positive probability and being independent of W and all $W^{(q)}$. The process

$$B = \sum_{q\in\mathbb{Q}} W^{(q)}\mathbb{1}_{A_q}$$

satisfies $\mathbf{P}(B_T \in \mathbb{Q}) = 1$. Finally, define the stopping time

$$\tau = \inf(t \in [0,T] : W_t = -1)$$

and the underlying price process

$$S_t = \exp\left(W_t\Big(1 - \frac{t-\tau}{T-\tau}\mathbb{1}_{[\tau,T]}(t)\Big) + B_t\frac{t-\tau}{T-\tau}\mathbb{1}_{[\tau,T]}(t)\right).$$

We will consider the filtration generated by S. By construction, S runs arbitrarily close to any continuous function f with positive probability, i. e. it satisfies CFS, since both W and B do. However, the part concerning the set A is not fulfilled, take for example the constant function $f = 1$ and $A = \{1\}$: There we have

$$\mathbf{P}\big(|S_t - f(t)| < \epsilon \textit{ for all } t \in [0,T]\big) > 0$$

and $\mathbf{P}(S_T = 1) > 0$, but the intersection of both events has zero probability. Since $\mathbf{P}(S_T = 1) > 0$, the claim generated by $g = \mathbb{1}_{\{1\}}$ yields $p_0^s\big(g(S_T)\big) = 1$. But S almost never ends up at 1 without moving down to e^{-1} before, thus we can superhedge $g(S_T)$ with an initial endowment of $\frac{1}{e}$ and the strategy $H = \mathbb{1}_{[\tau,T)}$. This beats the static strategy for $\epsilon = 0$, thus also for some $\epsilon > 0$ small enough.

Chapter 4

Consistent price systems and CFS for exponential Lévy processes

In this chapter we study the special case of one-dimensional exponential Lévy processes. Over the last decades, exponential Lévy models have frequently been proposed as alternative to the Black-Scholes model (which itself is an exponential Lévy model), either to introduce jumps to the market, or to increase freedom in calibrating returns or implied volatility smiles. See [Scho03] or [R00] for an overview. Lévy processes studied in finance include Brownian motion [BS73], *jump-diffusion* models proposed by Merton [Me76] or Kou [Ko02], pure-jump models like *Normal Inverse Gaussian*- the *Variance Gamma-*, the *Meixner* and the *CGMY*-models, but also α*-stable* and the wide class of *Generalized Hyperbolic* processes ([B-N98], [MS90], [Scho01], [CGMY02], [Ma63], [B-N77]) were considered.
After recalling some basic properties of Lévy processes from standard literature, we give a complete characterization of which processes admit ϵ-consistent price systems for which ϵ. This can be seen as preliminary work to prove a version of the Fundamental Theorem of Asset Pricing for Lévy models under small proportional transaction costs, see Theorem 5.4.9, but may also be of independent interest. The last part of the chapter tries to prove (ECFS) and therefore the face-lifting Theorem 3.1.5 for as many Lévy models used in practice as possible. As will turn out, this can be done for all of the above mentioned models.

4.1 Basic properties of Lévy processes

Let $L = (L_t)_{t\in[0,T]}$ be a one-dimensional Lévy process with respect to $\mathbf{P}$ and let the asset price process be given by $S = e^L$. To avoid complications concerning modifications, *Lévy processes in law*, Lévy vs. $(\mathcal{F}_t)$-Lévy processes etc. (see e. g. [Sat00], [Ch01]), we clarify that we impose L to be càdlàg and have increments $L_{t+h} - L_t$ which are stationary (i. e. the distribution only depends of h) and independent of $\mathcal{F}_t$, both with respect to $\mathbf{P}$. We begin our survey by quoting some well-known results about Lévy processes, a standard source being [Sat00].

Theorem 4.1.1 (Lévy-Khinchine, see [Sat00, Thm. 8.1]) *If L is a Lévy process, then the characteristic function of L satisfies $\mathbb{E}_{\mathbf{P}}(e^{iuL_t}) = e^{t\psi(u)}$ with the characteristic exponent*

$$\psi(u) = -\frac{\sigma^2}{2}u^2 + iu\gamma + \int_{\mathbb{R}} \Big(e^{iux} - 1 - iux\mathbb{1}_{[-1,1]}(x)\Big)\nu(dx)$$

with $\sigma^2 \geq 0$, $\gamma \in \mathbb{R}$ and a Borel measure ν on $\mathbb{R}$ satisfying $\nu(\{0\}) = 0$ and $\int_{\mathbb{R}}(x^2 \wedge 1)\nu(dx) < \infty$.

We call (σ^2, γ, ν) *characteristic triplet* or *Lévy-Khinchine-triplet* of L.

Lemma 4.1.2 (see [Sat00, Thm. 21.9]) *L (and therefore also S) is a. s. of finite variation if and only if $\sigma^2 = 0$ and $\int_{[-1,1]} |x|\nu(dx) < \infty$. If not so, then L (and therefore S) is a. s. of infinite variation.*

If L is of finite variation, then we call

$$\gamma_0 = \gamma - \int_{[-1,1]} x\nu(dx)$$

the *drift* of L. The following statements can be easily verified by comparing the characteristic functions.

Lemma 4.1.3 *Let L be a Lévy process with characteristic triplet (σ^2, γ, ν).*

1. *If $L^{(2)}$ is a Lévy process independent of L with characteristic triplet $(\sigma_2^2, \gamma_2, \nu_2)$, then $L + L^{(2)}$ is a Lévy process with characteristic triplet*

$$(\sigma^2 + \sigma_2^2, \gamma + \gamma_2, \nu + \nu_2).$$

2. *For any $\lambda \neq 0$ the process λL is a Lévy process with characteristic triplet $(\lambda^2\sigma^2, \lambda\gamma, \tilde{\nu})$, where $\tilde{\nu}$ is given by*

$$\tilde{\nu}([a,b]) = \nu\left(\left[\frac{a}{\lambda}, \frac{b}{\lambda}\right] \cup \left[\frac{b}{\lambda}, \frac{a}{\lambda}\right]\right),\ a \leq b.$$

3. *For $\lambda \in \mathbb{R}$, the process $t \mapsto L_t + \lambda t$ is a Lévy process with characteristic triplet $(\sigma^2, \gamma + \lambda, \nu)$.*

Lemma 4.1.4 (see [Sat00, Thms. 24.10 and 27.4 ff.]) *Denote by $\mathcal{L}$ the Lebesgue measure on $\mathbb{R}$.*

1. *If L is of infinite variation, then $\operatorname{supp}(L_t) = \mathbb{R}$ and $\mathbf{P}^{L_t} \sim \mathcal{L}$ for all $t > 0$.*

2. *If $\nu(\mathbb{R}) = \infty$, then $\mathbf{P}^{L_t} \ll \mathcal{L}$ for all $t > 0$*

3. *If $\nu\big((-\delta, 0)\big) > 0$ and $\nu\big((0, \delta)\big) > 0$ for all $\delta > 0$, then $\operatorname{supp}(L_t) = \mathbb{R}$ for all $t > 0$.*

4. *If both (2) and (3) are satisfied, then $\mathbf{P}^{S_T} \sim \mathcal{L}$.*

4.2 Consistent price systems

Not exclusively, but also because of their role in finance, the martingale property (or lack thereof) of Lévy and exponential Lévy processes is a well-researched topic. Since every martingale measure also implies a 0-consistent price system, it is an easy task to completely characterize all exponential Lévy processes that admit ϵ-consistent price systems. As it turns out, there is only a very subtle difference between S admitting an equivalent martingale measure and S admitting ϵ-consistent price systems for all $\epsilon > 0$. The difference between these two classes lies entirely in the class of monotone processes.

Lemma 4.2.1 *Let L be a Lévy process with characteristic triplet (σ^2, γ, ν).*

1. *L is a. s. increasing if and only if it is of finite variation and both $\gamma_0 \geq 0$ and $\nu\big((-\infty, 0)\big) = 0$ hold.*

2. *L is a. s. decreasing if and only if it is of finite variation and both $\gamma_0 \leq 0$ and $\nu\big((0, \infty)\big) = 0$ hold.*

Proof: The first part is proved in [Sat00, Theorem 21.5], the second one follows the first one and Lemma 4.1.3.2 with $\lambda = -1$. □

The following theorem states exponential Lévy processes can be divided up into two classes: monotone processes and martingales, with the only intersection between the two being the constant process.

Proposition 4.2.2 *Let $S = e^L$ be an exponential Lévy process. Assume S is not a. s. constant. Then exactly one of the following holds*

1. *S is almost surely increasing*
2. *S is almost surely decreasing*
3. *S and L both admit martingale measures equivalent to $\mathbf{P}$.*

Proof: It was shown in [Ch01, Theorem 3.2], that S admits an equivalent martingale measure if it is neither increasing nor decreasing. This shows that at least one of the three holds. Finally, note that if two out of the three properties are met, then S is constant. □

Remark 4.2.3 Cherny ([Ch01]) showed that we can even choose the martingale measures for L and S in such a way that L stays a Lévy process. This, however, plays no role in our survey.

Theorem 4.2.4 *Let $S = e^L$ be a non-deterministic exponential Lévy process.*

1. *If S is increasing, then S admits ϵ-CPS for all $\epsilon > \epsilon_1 = \frac{e^{\gamma_0 T}-1}{e^{\gamma_0 T}+1}$, but no ϵ_1-CPS*
2. *If S is decreasing, then S admits ϵ-CPS for all $\epsilon > \epsilon_2 = \frac{1-e^{\gamma_0 T}}{1+e^{\gamma_0 T}}$, but no ϵ_2-CPS*

Proof: (1): Assume S is increasing and non-deterministic, hence

$$\nu\big((0,\infty)\big) > 0.$$

We first show that there is no ϵ_1-CPS. Therefore assume $(M, \mathbf{Q}) \sim_{\epsilon_1} (S, \mathbf{P})$ exists. From Lemma 4.1.3.3 and Lemma 4.2.1 follows that $L_t - \gamma_0 t$ has drift 0 and is still increasing, thus $L_T - \gamma_0 T \geq 0$. Since L_T is not deterministic, this means $S_T \geq e^{\gamma_0 T}$ a. s. and $\mathbf{P}(S_T > e^{\gamma_0 T}) > 0$. Now $\frac{M_t}{S_t} \in [1-\epsilon, 1+\epsilon]$ implies $M_0 \leq (1+\epsilon_1)$ and $M_T \geq (1-\epsilon_1) S_T$ so combined we have

$$M_T \geq (1-\epsilon_1) S_T \geq (1-\epsilon_1) e^{\gamma_0 T} = (1+\epsilon_1) \geq M_0$$

almost surely, while the second inequality is met strictly with positive probability. Hence $\mathbb{E}_{\mathbf{Q}}(M_T) > \mathbb{E}_{\mathbf{Q}}(M_0)$, in contradiction to M being a $\mathbf{Q}$-martingale. So there is no ϵ_1-CPS for $(S, \mathbf{P})$.

If on the other hand $\epsilon > \epsilon_1$, then define $c = \log(1+\epsilon)$ and $\delta = -\frac{1}{T}\log\left(\frac{1-\epsilon}{1+\epsilon}\right)$. It is easily verified that $\epsilon > \epsilon_1$ implies $\delta > \gamma_0$. This means that $L_t - \delta t$ has negative drift and is therefore neither increasing nor decreasing (it still has upwards jumps). So by Proposition 4.2.2 $e^{L_t - \delta t}$ is a martingale with respect to some $\mathbf{Q} \sim \mathbf{P}$, and thus $M_t = e^{c + L_t - \delta t}$ is also a $\mathbf{Q}$-martingale. The way we have chosen c and δ we find that

$$\frac{M_t}{S_t} = e^{c-\delta t} \in [e^{c-\delta T}, e^c] = [1-\epsilon, 1+\epsilon]$$

for all $t \in [0,T]$, which shows $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$. This completes (1).

(2) is proved in the complete same fashion. This time, every ϵ_2-CPS $(M, \mathbf{Q})$ has to meet the inequality

$$M_T \leq (1+\epsilon_2)S_T \leq (1+\epsilon_2)e^{\gamma_0 T} = 1 - \epsilon_2 \leq M_0$$

almost surely and strictly with positive probability, so M cannot be a $\mathbf{Q}$-martingale. If on the other hand $\epsilon > \epsilon_2$, we define $c = \log(1-\epsilon)$ and $\delta = \frac{1}{T}\log\left(\frac{1+\epsilon}{1-\epsilon}\right) > -\gamma_0$, such that $M_t = e^{c+L_t+\delta t}$ admits an equivalent martingale measure and satisfies

$$\frac{M_t}{S_t} \in [e^c, e^{c+\delta T}] = [1-\epsilon, 1+\epsilon]$$

for all t. □

Remark 4.2.5

1. Note especially the case $\gamma_0 = 0$, where $\epsilon_1 = \epsilon_2 = 0$. These are the only cases where S admits ϵ-CPS for all $\epsilon > 0$, but no equivalent martingale measure. These are the processes that are monotone, but only barely. In terms of finance, an increasing process with $\gamma_0 = 0$ would be a product with no downside risk but also no guaranteed lower bound on the return, other than 0. If buying into this product causes transaction costs, then there is the risk of losing money, like in Example 5.2.8 later on. This will play a role in the Fundamental Theorem of Asset Pricing, see Theorem 5.4.9.

2. The only case we have not covered is the case of deterministic, non-constant S, i. e. the case where $\sigma^2 = 0 = \nu$ and $\gamma_0 \neq 0$. Reviewing the proof of Theorem 4.2.4, it is easily seen that in this case the theorem still holds, except that an ϵ_1- resp. ϵ_2-CPS also exists, the solution being

the constant process $M_t = e^c$ in both cases. But for smaller ϵ consistent price systems still do not exist. So with this remark we have finished a complete characterization of which one-dimensional exponential Lévy processes admit which consistent price systems.

4.3 Conditional full support and the face-lifting theorem

We continue our survey of one-dimensional exponential Lévy models by deriving sufficient conditions for (CFS) and (ECFS) to hold. Unlike in the previous section, we do not intend to give a full characterization, i. e. necessary and sufficient conditions, but rather concentrate on properties that most models used in practice share.
Introduce the process

$$L^{\eta}{}_t = L_t - \sum_{0 \le s \le t} \Delta_- L_s \mathbb{1}_{|\Delta_- L_s| \ge \eta},$$

that arises from L by cancelling all jumps of absolute value greater or equal to η. By Lemma 4.1.3 L^η is also a Lévy process with triplet

$$\left(\sigma^2, \gamma - \int_{[-1,1]\setminus(-\eta,\eta)} x\nu(dx), \nu\mathbb{1}_{(-\eta,\eta)}\right).$$

Lemma 4.3.1 *Define the process $\overline{L}$ by*

$$\overline{L}_t = \sup(|L_s|, s \in [0,t]).$$

If L and L^η admit equivalent martingale measures for all $\eta > 0$, then $0 \in \operatorname{supp}(\overline{L}_T)$.

Proof: Although technically L is only defined on $[0,T]$, we may assume without loss of generality that L has infinite time horizon. For $\epsilon > 0$, define τ_ϵ as the first time L leaves $(-\epsilon, \epsilon)$. Then $0 \in \operatorname{supp}(\overline{L}_T)$ for all $T > 0$ is equivalent to τ_ϵ being unbounded for all $\epsilon > 0$. So assume that there exists some $\epsilon > 0$ such that

$$T_\epsilon = \operatorname{esssup}_{\mathbf{P}}(\tau_\epsilon)$$

is finite. Fix some small $\delta \in (0, T_{\epsilon/3})$, where $T_{\epsilon/3} = \operatorname{esssup}_{\mathbf{P}}(\tau_{\epsilon/3}) < \infty$. Further define $N = \lceil \frac{T_\epsilon}{T_{\epsilon/3} - \delta} \rceil + 1$ and $\eta = \frac{\epsilon}{3\lceil N/2 \rceil}$.

Denote the martingale measures of L and L^η equivalent to $\mathbf{P}$ by $\mathbf{Q}$ and $\mathbf{Q}_\eta$. Then the process

$$\hat{L}_t = L_{t\wedge(T_{\epsilon/3}-\delta)} + ({L^\eta}_t - {L^\eta}_{T_{\epsilon/3}-\delta})\mathbb{1}_{(T_{\epsilon/3}-\delta,\infty)},$$

That arises from L by cancelling out all jumps greater or equal to η after $T_{\epsilon/3}-\delta$. $\hat{L}$ is a martingale with respect to $\hat{\mathbf{Q}} \sim \mathbf{P}$, defined by

$$\frac{d\hat{\mathbf{Q}}}{d\mathbf{Q}_\eta} = \mathbb{E}_{\mathbf{Q}_\eta}\left(\frac{d\mathbf{Q}}{d\mathbf{Q}_\eta}\middle|\mathcal{F}_{T_{\epsilon/3}-\delta}\right).$$

Now consider the set

$$A = \Big\{T_{\epsilon/3} - \delta < \tau_{\epsilon/3} \le T_{\epsilon/3}\Big\} \cap \Big\{|\Delta_- L_t| < \eta\ \forall\ t \in (T_{\epsilon/3} - \delta, T_{\epsilon/3}]\Big\}.$$

By definition of $T_{\epsilon/3}$, we have $\mathbf{P}(T_{\epsilon/3} - \delta < \tau_{\epsilon/3} \le T_{\epsilon/3}) > 0$, and, since the number of jumps of L in $(T_{\epsilon/3}-\delta, T_{\epsilon/3}]$ greater or equal to η is is Poisson$\Big(\delta\nu\big(\mathbb{R}\backslash(-\eta,\eta)\big)\Big)$-distributed independent of $F_{T_{\epsilon/3}-\delta}$, we have $\mathbf{P}(A) > 0$. Since $\hat{L}$ is a $\hat{\mathbf{Q}}$-martingale, we have both

$$\hat{\mathbf{Q}}(\hat{L}_{\tau_{\epsilon/3}} > 0|\mathcal{F}_{T_{\epsilon/3}-\delta}) > 0$$
$$\hat{\mathbf{Q}}(\hat{L}_{\tau_{\epsilon/3}} < 0|\mathcal{F}_{T_{\epsilon/3}-\delta}) > 0$$

on A. But since on A, L and $\hat{L}$ coincide for all $t \in [0, T_{\epsilon/3}]$ a. s., this means that both

$$A_+ = A \cap \{L_{\tau_{\epsilon/3}} > 0\}$$
$$A_- = A \cap \{L_{\tau_{\epsilon/3}} < 0\}$$

have positive probability.
Define the sequence $(\tau^{(n)})_{n\in\mathbb{N}}$ by $\tau^{(0)} = 0$ and

$$\tau^{(n+1)} = \inf\Big(t > \tau^{(n)} : |L_t - L_{\tau^{(n)}}| \ge \frac{\epsilon}{3}\Big).$$

Now consider the set of paths of L satisfying all of the following:

1. $|\Delta_- L_{\tau^{(n)}}| < \eta$ for all $1 \le n \le N$
2. $\tau^{(n+1)} - \tau^{(n)} > T_{\epsilon/3} - \delta$ for all $0 \le n < N$
3. $L_{\tau^{(n+1)}} > L_{\tau^{(n)}}$ for all even $n \le N$
4. $L_{\tau^{(n+1)}} < L_{\tau^{(n)}}$ for all odd $n \le N$

Due to the Lévy property of L, the probability of this set is at least

$$\mathbf{P}(A_+)^{\lceil N/2\rceil}\mathbf{P}(A_-^{\lfloor N/2\rfloor}) > 0.$$

By construction, L does not leave $(-\epsilon, \epsilon)$ on this set before

$$\tau^{(N)} \geq N(T_{\epsilon/3} - \delta) > T_\epsilon,$$

in contradiction to the definition of T_ϵ. This completes the proof. $\square$

Remark 4.3.2 All processes covered by Lemma 4.3.1 share the common property $\mathrm{supp}(L_T) = \mathbb{R}$, but this property is neither necessary nor sufficient. For instance, for a standard Poisson process we have $\mathrm{supp}(L_T) = \mathbb{N}_0$, but still $\mathbf{P}(\overline{L}_T = 0) > 0$. On the other hand, consider

$$L_t = t + N_t^{(1)} - \pi N_t^{(2)}$$

with two independent Poisson processes $N^{(1)}, N^{(2)}$. While this process satisfies $\mathrm{supp}(L_1) = \mathbb{R}$, we have $\overline{L}_1 \geq 1$ almost surely.

Theorem 4.3.3 *Let L satisfy properties 1 or 4 from Lemma 4.1.4 Then $S = e^L$ satisfies (ECFS). If only property 3 is satisfied, then (CFS) holds.*

Proof: Assume property 1 or 3 are satisfied. Since $(\frac{S_T}{S_t}|\mathcal{F}_t)$ has the same distribution as S_{T-t}, we may assume $t = 0$ and $f(0) = 1$ and leave the conditioning on $\mathcal{F}_t$ away. Define

$$\eta = \log \inf_{s\in[0,T]} \Big(\frac{f(s)+\epsilon}{f(s)}\Big) > 0.$$

Since $\log(f)$ is uniformly continuous on $[0, T]$, we may find a grid

$$0 = t_0 < t_1 < \ldots < t_N = T$$

and a continuous polygonal function

$$p(x) = \sum_{n=1}^{N} (a_n x + b_n)\mathbb{1}_{(t_{n-1},t_n]}(x)$$

satisfying $p(0) = 0$ and

$$sup_{s\in[0,T]}\Big(|\log\big(f(x)\big) - p(x)|\Big) \leq \frac{\eta}{2}.$$

Not only does L satisfy the conditions of Lemma 4.3.1, but so does $L_s + as$ for all $a \in \mathbb{R}$, by Lemma 4.1.3.3. By this and the independent stationary increment property of L we know that the sets

$$A_n = \left\{ \sup_{s \in [t_{n-1}, t_n]} \big(|L_s - L_{t_{n-1}} - a_n(s - t_{n-1})| \big) < \frac{\eta}{2^{n+2}} \right\}, \; 1 \leq n \leq N$$

are independent of each other and have positive probability, and so has

$$\left\{ \sup_{s \in [0,T]} |L_s - \log(f(s))| < \eta \right\} \supseteq \left\{ \sup_{s \in [0,T]} |L_s - p(s)| \leq \frac{\eta}{4} \right\} \supseteq \bigcap_{n=1}^{N} A_n.$$

So with positive probability we have for all $s \in [0, T]$:

$$\log\big(f(s)\big) - \eta < L_s < \log\big(f(s)\big) + \eta \Rightarrow$$

$$f(s) - \epsilon \leq f(s)e^{-\eta} < S_s < f(s)e^{\eta} \leq f(s) + \epsilon,$$

which secures (CFS). Now if property 1 or 4 is satisfied, then S_T has a density and Lemma 3.3.1 secures (ECFS).

$\square$

Example 4.3.4 To see that only property 3 of Lemma 4.1.4 is not enough to secure (EFCS), assume L is a compound Poisson process with normally distributed jumps. There ν is the normal distribution, so property 3 is satisfied and (CFS) holds. But $\mathbf{P}(S_1 = 1) > 0$, so with $T = 1$,

$$f(x) = (1 + x)\mathbb{1}_{[0,\frac{1}{2})}(x) + (1 - x)\mathbb{1}_{(\frac{1}{2},1]}(x),$$

some small enough ϵ and $A = \{1\}$ it can be seen that (ECFS) is not satisfied: To stick close enough to f, S has to jump, but $\{S_1 = 1\}$ has zero probability conditioned on the event that at least one jump occurs before 1.

Example 4.3.5 (Common models in finance) As we have seen, (ECFS) and therefore Theorem 3.1.5 holds if L has an active Brownian part, i. e. $\sigma^2 > 0$, or if ν is infinite and two-sided. The first is true for the Black-Scholes-Model and for the Jump-Diffusion models. The latter is true for the NIG- the Variance-Gamma- the Meixner- the CGMY- (with parameter $Y < 0$) the α-stable (with $\alpha \in (1, 2)$) and the Generalized-Hyperbolic models. In all those cases ν is infinite and has a strictly positive density on the whole real line, see [Scho03] and [ST94] (α-stable processes) for explicit densities.

Example 4.3.6 (3.2.1 revisited) Recall the model $L = X^{(1)} - X^{(2)}$ from Example 3.2.1. There we could manually prove the face-lifting theorem to be true, even though (CFS) was not given. Now consider the similar case $L = X^{(1)} - 2X^{(2)}$, with $X^{(1)}$ and $X^{(2)}$ still being independent Poisson processes. We will construct a strategy that beats the static one for the digital option $g = \mathbb{1}_{[\frac{1}{4},\frac{1}{2}]}$, which is in the money if and only if $L_T = -1$. Start with $a = \frac{1-e^{-1}}{e-e^{-2}}$ units of the stock and $b = 1 - \frac{e-1}{e-e^{-2}}$ units of cash. This position costs

$$a + b \approx 0.58 < 1 = \mathcal{C}(g, \mathbf{P}^{S_T})(1) = p_0^s\big(g(S_T)\big).$$

If no jump occurs, then $g(S_T) = 0$ and we are done. If the first jump goes up to e^1, then we sell our stocks and have $ae+b = 1$ of cash, and have therefore superreplicated g. If the first jump goes down to e^{-2}, we still have $ae^{-2} + b = e^{-1}$, which is enough to buy e units of the stock. If then $L_T = -1$, our terminal wealth is 1 and we have once again superreplicated g. As in Example 3.3.6, it is now clear that this strategy beats the static one for sufficiently small $\epsilon > 0$.

Chapter 5

The Fundamental Theorem of Asset Pricing

5.1 Introduction

The idea of "no-arbitrage"-conditions has developed to be the central principle in modern mathematical finance since its beginnings in the seventies. Introduced by F. Black, M. Scholes and R. Merton in their groundbreaking 1973 papers ([BS73],[Me73]) as a means to determine unique option prices in one specific market model, it has since then drastically increased its importance within the whole mathematical theory of financial markets: In the seminal papers of M. Harrison, D. Kreps and S. Pliska ([HK79], [HP81], [Kr81]) it was shown that the the now-called *Black-Scholes* model is just a special case of a more general framework and that the concept of no-arbitrage is closely connected to martingales. This relation, often referred to as *duality theory*, has opened the gates to various other applications of the powerful martingale theory in finance, for instance superreplication in incomplete markets, portfolio optimization or optimal consumption. The central piece of duality theory is arguably what is known as the *Fundamental Theorem of Asset Pricing*. It states that a model is arbitrage-free if and only if there exists some kind of riskless pricing functional, in most cases an equivalent martingale measure. This theorem comes in countless versions, differing in the general framework of trading, the notion of arbitrage and the notion of admissibility, see [DS06] for a more involved overview.

The theory of arbitrage-free markets under transaction costs is considerably younger, with the most influential pioneering work by E. Jouini and H. Kallal dating back to 1995 [JK95]. First it has been the main challenge to find the right replacement for the concept of equivalent martingale measures, the weapon of choice now being the consistent price systems introduced in Section 2.1. In this chapter this theory is used to derive a multidimensional analogue of a result by Guasoni et al. [GRS08b] for one-dimensional continuous price processes. There it was shown that a model is arbitrage-free with respect to arbitrarily small proportional transaction costs if and only if it admits consistent price systems for every level of costs. As was pointed out there, the concept of asymptotically small transaction costs is of particular interest because of certain models, most notably fractional Brownian motion, which admit arbitrage in the frictionless case but become arbitrage-free as soon as transaction costs are introduced, however small they may be.

Theorem 5.1.1 (Fundamental Theorem of Asset Pricing) *Let $(S_t)_{t\in[0,T]}$ be an $\mathbb{R}^d_+$-valued, continuous process. Then there exists no ϵ-arbitrage for any $\epsilon > 0$ if and only if $(S, \mathbf{P})$ admits ϵ-consistent price systems for all $\epsilon > 0$.*

The proof is organized in the spirit of the original paper. Some crucial results from [CS06] were translated into our setting in Section 2.2. In Section 5.3, local consistent price systems are constructed with respect to some adequate localizing sequence. These local CPS are shown to converge to a true CPS, or otherwise there exists free lunch with bounded risk, which implies arbitrage. This part of the proof is carried out in Section 5.4. Throughout this and the next chapter, we limit ourselves to the sum numeraire $N_t = 1 + \sum_{i=1}^d S_t^{(i)}$, but refer to [GRS08b, Section 5] for a discussion of the subtle difference between the numeraires.

The following notation will help us to highlight the connection between supports and arbitrage. It is used int the rather technical construction of local consistent price systems in Lemma 5.3.1 and the appendix, in which an important result for this construction is derived.

Definition 5.1.2 *For a bounded $\mathbb{R}^d$-valued random variable X on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, we define the set $\mathcal{I}(X) \subseteq \mathbb{R}^d$ by*

$$\mathcal{I}(X) = \mathrm{relint}\big(\mathrm{conv}(\mathrm{supp}_{\mathbf{P}} X)\big),$$

where relint *denotes the relative interior. If $\mathcal{G}$ is a sub-σ-algebra of $\mathcal{F}$, then the corresponding conditional version of $\mathcal{I}$ is defined as the set-valued $\mathcal{G}$-measurable random variable*

$$\mathcal{I}(X|\mathcal{G}) = \mathrm{relint}\Big(\mathrm{conv}\big(\mathrm{supp}_{\mathbf{P}}(X|\mathcal{G})\big)\Big),$$

where $\mathrm{supp}_{\mathbf{P}}(X|\mathcal{G})$ *stands for the conditional support of X given $\mathcal{G}$.*

It is easily seen that $\mathcal{I}(X)$ equals the set of all possible expected values $\mathbb{E}_{\mathbf{Q}}(X)$ with respect to the set of all $\mathbf{Q} \sim \mathbf{P}$.

5.2 Notions of arbitrage

Definition 5.2.1 1. *A trading strategy H is called* ϵ-arbitrage *for some $\epsilon \geq 0$ if $V_{\epsilon,S}(H) \geq 0$ a. s. and $\mathbf{P}\big(V_{\epsilon,S}(H) > 0\big) > 0$.*

2. *If there is no ϵ-admissible ϵ-arbitrage, we say the property NA(ϵ) is satisfied.*

3. *The property NA(ϵ+) is satisfied if and only if NA(δ) is satisfied for all $\delta > \epsilon$.*

4. *A trading strategy H is called* simple strategy*, if it is of the form*

$$H_t = C\mathbb{1}_{(\sigma,\tau]\cap(0,T)}(t)$$

with two stopping times $0 \leq \sigma \leq \tau \leq T$ and $C \in L^\infty(\mathcal{F}_\sigma, \mathbb{R}^d)$.

5. *The property NSA(ϵ) holds if there exists no ϵ-admissible simple ϵ-arbitrage strategy. NSA(ϵ+) is defined respectively.*

Remark 5.2.2 1. In classical stochastic analysis, sometimes processes of the form 5.2.1.4 are called *simple processes* (e. g. [JS87]), sometimes the term is used for predictable piecewise constant processes, i. e. linear combination of our simple strategies (e. g. [P04]). In [GRS08b] our concept of simple arbitrage as arbitrage that can be achieved with only two trades was coined *obvious arbitrage*, but since we discuss the concept in greater detail here it is therefore convenient to also define the class of underlying strategies, and the name *obvious process* would not have made sense.

2. The indicator $\mathbb{1}_{(\sigma,\tau]\cap(0,T)}$ is designed that way because on the one hand, $\mathbb{1}_{(\sigma,\tau)}$ need not be predictable if τ is not, and on the other hand $\mathbb{1}_{(\sigma,\tau]}$ is predictable but not a valid strategy if $\mathbf{P}(\tau = T) > 0$. These technicalities, however, make no difference as we only argue on continuous processes whenever assuming NSA.

Lemma 5.2.3 *Assume that S is bounded. Then NSA(0+) implies that for each two stopping times $0 \leq \sigma \leq \tau \leq T$ we have*

$$S_\sigma \in \overline{\mathcal{I}(S_\tau|\mathcal{F}_\sigma)}$$

almost surely.

Proof: We apply a separating hyperplane argument: Assume $S_\sigma \notin \overline{\mathcal{I}(S_\tau|\mathcal{F}_\sigma)}$ on some $\mathcal{F}_\sigma$-measurable set of positive probability. Then there exists a set $A \in \mathcal{F}_\sigma$ of positive probability, an $\mathcal{F}_\sigma$-measurable random variable C with $\|C\| = 1$ a. s. and some $\delta > 0$ such that

$$C^\top S_\sigma \leq C^\top S_\tau - \delta$$

almost surely on A. The transaction costs caused by the strategy

$$H = (C\mathbb{1}_A)\mathbb{1}_{(\sigma,\tau]\cap(0,T)}$$

are no greater than $\epsilon\|C\|_\infty^\top(S_\sigma + S_\tau)$ on A, which is bounded and therefore smaller than δ for ϵ small enough. Hence H is simple ϵ-arbitrage for ϵ small enough, in contradiction to NSA(0+). □

Remark 5.2.4 In dimension one, it suffices to consider a subset of the set of simple strategies: In fact, if $H = C\mathbb{1}_{(\sigma,\tau]\cap(0,T)}$ is an ϵ-arbitrage, then C can either be replaced by $\mathbb{1}_{C>0}$ or $-\mathbb{1}_{C<0}$. Thus we could restrict us to processes of the form

$$c\mathbb{1}_A\mathbb{1}_{(\sigma,\tau]\cap(0,T)}(t)$$

with deterministic $c \in \{-1, 1\}$ and $A \in \mathcal{F}_\sigma$. In other words, once σ, τ and A are fixed, there are only two strategies to be considered: long and short.
In multiple dimensions, the case is obviously not that simple, since there is a $d-1$-dimensional continuum of holdings C which are not scalar multiples of each other.
The natural analogue of the one-dimensional case are processes of the form $\mathbf{c}\mathbb{1}_A\mathbb{1}_{[\sigma,\tau]}(t)$ with $\mathbf{c} \in \mathbb{R}^d, \|\mathbf{c}\| = 1$ and $A \in \mathcal{F}_\sigma$, which we will call *simple processes with fixed holdings*. The following can be shown:

Proposition 5.2.5 *Let a simple ϵ-arbitrage H be given for some $\epsilon > 0$. Then, for each $\tilde{\epsilon} < \epsilon$, there exists a simple $\tilde{\epsilon}$-arbitrage with fixed holdings.*

Proof: Denote the simple ϵ-arbitrage by $H_t = C\mathbb{1}_{(\sigma,\tau]\cap(0,T)}(t)$ with some $\mathcal{F}_\sigma$-measurable C. Since, if $C \neq \mathbf{0}$, the absolute value of C does not matter, we may and will assume that $\|C\| \in \{0, 1\}$ almost surely. We furthermore assume, without loss of generality, that there exist some $\delta_1 > \delta_0 > 0$ with $S_\sigma \in (\delta_0, \delta_1)^d$ almost surely on $\{C \neq \mathbf{0}\}$: Just fix δ_0 and δ_1 to ensure

$$\mathbf{P}\big(\{S_\sigma \in (\delta_0, \delta_1)^d\} \cap \{C \neq \mathbf{0}\}\big) > 0$$

and set C to zero outside of this $\mathcal{F}_\sigma$-measurable set. As a consequence, the transaction costs to be paid for H are greater than $\epsilon\delta_0$ almost surely on $\{C \neq \mathbf{0}\}$. Now Fix a positive constant

$$\eta < \frac{(\epsilon - \tilde{\epsilon})\delta_0}{d\delta_1(1+\epsilon)}.$$

Then, for any $\mathcal{F}_\sigma$-measurable $[0,\eta]^d$-valued random variable D we have

$$V_{\epsilon,S}(D\mathbb{1}_{(\sigma,\tau]\cap(0,T)}) = D^\top\big((1-\epsilon)S_\tau - (1+\epsilon)S_\sigma\big) > -d\eta(1+\epsilon)\delta_1 > -(\epsilon-\tilde{\epsilon})\delta_0$$

almost surely on $\{C \neq 0\}$. Fix some $\mathbf{c} \in \mathbb{R}^d$ satisfying

$$\mathbf{P}(\mathbf{c} - C \in [0,\eta]^d | C \neq \mathbf{0}) > 0$$

and define the $\mathcal{F}_\sigma$-measurable set

$$A = \{c - C \in [0,\eta]^d\} \cap \{C \neq \mathbf{0}\}.$$

Then the $\tilde{\epsilon}$-arbitrage is given by $\tilde{H} = \mathbf{c}\mathbb{1}_A\mathbb{1}_{(\sigma,\tau]\cap(0,T)}$:

$$V_{\tilde{\epsilon},S}(\tilde{H}) \geq \underbrace{V_{\tilde{\epsilon},S}(C\mathbb{1}_A\mathbb{1}_{(\sigma,\tau]\cap(0,T)})}_{>(\epsilon-\tilde{\epsilon})\delta_0} + \underbrace{V_{\tilde{\epsilon},S}\big((\mathbf{c}-C)\mathbb{1}_A\mathbb{1}_{(\sigma,\tau]\cap(0,T)}\big)}_{>-(\epsilon-\tilde{\epsilon})\delta_0} > 0$$

on A. This completes the proof. □

Remark 5.2.6 As a consequence of Proposition 5.2.5, the absence of simple ϵ-arbitrage with fixed holdings for every $\epsilon > 0$ implies NSA(0+), so both properties are equivalent. Obviously the theorem does not apply to the case $\epsilon = 0$, and, as the following example will show, simple 0-arbitrage does not imply any simple arbitrage with fixed holdings

Example 5.2.7 (simple 0-arbitrage, but no simple arbitrage with fixed holdings): Consider the function

$$f : x \mapsto \sum_{k\in\mathbb{Z}}(x-2k)\mathbb{1}_{[2k-\frac{1}{2},2k+\frac{1}{2})}(x) + \sum_{k\in\mathbb{Z}}(2k-1-x)\mathbb{1}_{[2k-\frac{3}{2},2k-\frac{1}{2})}(x),$$

which is essentially a sawtooth function mapping $\mathbb{R}$ to $[-\frac{1}{2},-\frac{1}{2}]$. Now consider a two-dimensional Brownian motion $W = (W^{(1)}, W^{(2)})$. The process

$$X_t = \frac{f(\|W_t\|)}{\|W_t\|}W_t$$

behaves like W in the beginning, but is reflected off the circle around $\mathbf{0}$ with radius $\frac{1}{2}$. Here we set $\frac{0}{0} = 0$. Finally consider the $\mathbb{R}^2_+$-valued process

$$(S_t^{(1)}, S_t^{(2)}) = (1 + X_t^{(1)}, 1 + X_t^{(2)}).$$

By construction S has conditional full support inside a circle around $(1,1)$ with radius $\frac{1}{2}$, never leaves the circle and touches the edge of the circle at least once before $T = 1$ with positive probability. Denote by σ the first time this happens. Obviously this process admits a simple 0-arbitrage, the arbitrage strategy being given by

$$H = (S_0 - S_\sigma)\mathbb{1}_{(\sigma,T)}\mathbb{1}_{\sigma<T}.$$

By the nature of the Brownian motion, σ (or any other time S touches the edge of the circle) is the only candidate the first trade. But since $S_0 - S_\sigma$ is the only position that allows for arbitrage and since no value of $S_0 - S_\sigma$ is admitted with positive probability, there exists no 0-arbitrage with fixed holding.
With a similar construction the same can be shown for arbitrary $\epsilon > 0$.

Example 5.2.8 The following easy example shows that NSA(0+) does in fact not imply NSA(0), i. e. these properties are not equivalent. Let two independent exponentially distributed random variables X and Y be given and let the filtration $\mathcal{F}$ be generated by the processes $\mathbb{1}_{[0,X]}(t)$ and $\mathbb{1}_{[0,X+Y]}(t)$. Finally consider the underlying

$$S_t = 1 + \int_0^t \mathbb{1}_{[X,X+Y]}(s)ds.$$

Obviously $H = \mathbb{1}_{(0,T)}$ is an admissible 0-arbitrage, hence NSA(0) is not given. In fact, since S is increasing, there is no other sensible candidate for arbitrage than H. However, for no positive ϵ, it forms an ϵ-arbitrage, as

$$V_{\epsilon,S}(H) = -(1+\epsilon) + (1-\epsilon)\big(1 + (Y \wedge T) - (X \wedge T)\big)$$

becomes negative with positive probability, for all $\epsilon, 0$. Hence NSA(0+) holds.

5.3 Local consistent price systems

Following the example of [GRS08b], we start by localizing the problem, i. e. by constructing an increasing sequence of stopping times, such that S, stopped at each of the times, admits CPS. We therefore recall the notation

$$S^\tau : t \mapsto S_{t\wedge\tau}$$

for a process S and a stopping time τ. As will turn out, NSA(0+) instead of NA(0+) is already enough for such a localizing sequence to exist, and, even better, the sequence can be chosen in such a way that it converges to T in a very strong sense, namely stationary.

Proposition 5.3.1 *Assume S is continuous. If S satisfies NSA($0+$), then there exists, for each $\epsilon > 0$, a sequence $(\tau_n)_{n\in\mathbb{N}_0}$ of stopping times increasing stationarily to T, i. e. τ_n is almost surely increasing and*

$$\lim_{n\to\infty} \mathbf{P}(\tau_n = T) = 1,$$

such that for each $n \in \mathbb{N}$, the stopped process S^{τ_n} admits an ϵ-consistent price system.

Proof: As in the proof of Lemma 3.2.2, we replace the original CPS condition by the stronger one introduced in Remark 2.1.4.2. Therefore it is convenient to define the geometric distance function $\Delta : \mathbb{R}^d_+ \times \mathbb{R}^d_+ \to \mathbb{R}_{0+}$ by

$$\Delta(\mathbf{x}, \mathbf{y}) = \max\left(\|\frac{\mathbf{x}}{\mathbf{y}}\|_\infty, \|\frac{\mathbf{y}}{\mathbf{x}}\|_\infty\right) - 1.$$

Then we have the following triangular inequality:

$$\Delta(\mathbf{x}, \mathbf{z}) \le \Delta(\mathbf{x}, \mathbf{y}) + \Delta(\mathbf{y}, \mathbf{z}) + \Delta(\mathbf{x}, \mathbf{y})\Delta(\mathbf{y}, \mathbf{z}).$$

Define the stopping times τ_n recursively by setting $\tau_0 = 0$ and

$$\tau_{n+1} = \inf\Big(t > \tau_n, \Delta(S_{\tau_n}, S_t) = \frac{\epsilon}{4}\Big).$$

By Lemma 5.2.3, applied to the bounded process S^{τ_1} rather than S, we have that $S_0 \in \overline{\mathcal{I}(S_{\tau_1})}$. Thus for every $\delta > 0$ we can find some vector $\boldsymbol{\lambda} \in (\frac{1}{1+\delta}, 1+\delta)^d$ such that $\frac{S_0}{\boldsymbol{\lambda}} \in \mathcal{I}(S_{\tau_1})$, or equivalently, $S_0 \in \mathcal{I}(M_{\tau_1})$, where

$$M_{\tau_1} = S_{\tau_1}\boldsymbol{\lambda}.$$

We choose δ smaller than $\frac{\epsilon}{4}$.
By Theorem A.5 we can find some $\mathbf{Q}_1 \sim \mathbf{P}$ such that $\mathbb{E}_{\mathbf{Q}_1}(M_{\tau_1}) = S_0$. On $[0, \tau_1]$, we define the process M by

$$M_t = \mathbb{E}_{\mathbf{Q}_1}(M_{\tau_1}|\mathcal{F}_t),$$

so M is a $\mathbf{Q}_1$-martingale up to τ_1. From $\Delta(S_0, S_{\tau_1}) = \frac{\epsilon}{4}$ and $\Delta(S_{\tau_1}, M_{\tau_1}) \leq \delta$ follows

$$\Delta(S_0, M_{\tau_1}) \leq \delta + \frac{\epsilon}{4} + \frac{\epsilon\delta}{4} \leq \frac{\epsilon}{2} + \frac{\epsilon^2}{16},$$

and since $M_0 = S_0$ and M is a martingale, the same bound holds for $\Delta(S_0, M_t)$ for $t \leq \tau_1$. Finally, from this inequality and $\Delta(S_0, S_t) \leq \frac{\epsilon}{4}$ follows

$$\Delta(M_t, S_t) \leq \frac{3\epsilon}{4} + \frac{3\epsilon^2}{16} + \frac{\epsilon^3}{64},$$

which is smaller than ϵ for sufficiently small ϵ. This means

$$(M^{\tau_1}, \mathbf{Q}_1) \sim_\epsilon (S^{\tau_1}, \mathbf{P}).$$

Note that we can still choose δ arbitrarily small.

We now proceed inductively, assuming there exists some consistent price system $(M, \mathbf{Q}_n) \sim_\epsilon (S^{\tau_n}, \mathbf{P})$ satisfying $\Delta(M_{\tau_n}, S_{\tau_n}) \leq \eta$ almost surely for some small $\eta > 0$. Since, once again by Lemma 5.2.3,

$$S_{\tau_n} \in \overline{\mathcal{I}(S_{\tau_{n+1}}|\mathcal{F}_{\tau_n})}$$

almost surely, we can find an $\mathcal{F}_{\tau_n}$-measurable $\left(\frac{1}{1+2^{-n}}, 1+2^{-n}\right)^d$-valued random variable Λ satisfying

$$\frac{S_{\tau_n}}{\Lambda} \in \mathcal{I}(S_{\tau_{n+1}}|\mathcal{F}_{\tau_n}).$$

If we define

$$M_{\tau_{n+1}} = S_{\tau_{n+1}}\ \Lambda\ \frac{M_{\tau_n}}{S_{\tau_n}},$$

we have $M_{\tau_n} \in \mathcal{I}(M_{\tau_{n+1}})$ and by Corollary A.7 we can find a probability measure $\mathbf{Q}_{n+1} \sim \mathbf{Q}_n$ coinciding with $\mathbf{Q}_n$ on $\mathcal{F}_{\tau_n}$ such that

$$\mathbb{E}_{\mathbf{Q}_{n+1}}(M_{\tau_{n+1}}|\mathcal{F}_{\tau_n}) = M_{\tau_n}$$

almost surely. We can now extend M to a martingale on $[0, \tau_{n+1}]$ by setting for all $t \in [\tau_n, \tau_{n+1}]$:

$$M_t = \mathbb{E}_{\mathbf{Q}_{n+1}}(M_{\tau_{n+1}}|\mathcal{F}_t)$$

By construction we have $\Delta(S_{\tau_n}, S_t) \leq \frac{\epsilon}{4}$ and $\Delta(S_{\tau_n}, M_{\tau_n}) \leq \eta$, thus for all $t \in [\tau_n, \tau_{n+1}]$

$$\Delta(S_t, M_{\tau_n}) \leq \frac{\epsilon}{4} + \eta + \frac{\eta\epsilon}{4} \leq \frac{\epsilon}{2} + \frac{\epsilon^2}{16},$$

for $\eta \leq \frac{\epsilon}{4}$. Additionally

$$\Delta(M_{\tau_n}, M_{\tau_{n+1}}) = \Delta(S_{\tau_n}, S_{\tau_{n+1}}\Lambda) \leq \frac{\epsilon}{4}(1 + 2^{-n}),$$

and since M is a martingale, the same bound holds for $\Delta(M_{\tau_n}, M_t)$ for all $t \in [\tau_n, \tau_{n+1}]$. The triangular inequality applied once more yields

$$\Delta(M_t, S_t) \leq \epsilon\frac{3 + 2^{-n}}{4} + \epsilon^2\frac{3 + 2^{-n+1}}{16} + \epsilon^3\frac{1 + 2^{-n}}{64},$$

which is smaller than ϵ for sufficiently small ϵ independent of n. Since $\Delta(M_t, S_t) \leq \epsilon$ implies $1 - \epsilon \leq \frac{M_t}{S_t} \leq 1 + \epsilon$, this completes the proof. □

The sequence $(\tau_n)_{n\in\mathbb{N}}$ just constructed depended on ϵ. But as the following theorem will show, We can also choose it independent of ϵ.

Theorem 5.3.2 *Assume S is continuous and let NSA($0+$) be given. Then there exists a sequence $(\tau_n)_{n\in\mathbb{N}_0}$ of stopping times (where $\tau_0 = 0$) increasing stationarily to T such that for all $n \in \mathbb{N}$ and for each $\epsilon > 0$, the process S^{τ_n} admits an ϵ-consistent price system.*

Proof: Let some $\alpha \in (0, 1)$ be given. For fixed $\epsilon > 0$, denote by $(\tau_n^{(\epsilon)})_{n\in\mathbb{N}}$ the sequence of stopping times constructed in the proof of 5.3.1. Since for each $\epsilon > 0$ this sequence converges stationarily to T, we may choose, for each $k \in \mathbb{N}$, a stopping time $\tau^{(\alpha,k)} \in \{\tau_n^{(\alpha)}, n \in \mathbb{N}\}$ satisfying

$$\mathbf{P}(\tau^{(\alpha,k)} < T) < \alpha 2^{-k}.$$

By construction, the stopping time $\sigma^{(\alpha)} = \bigwedge_{k=1}^{\infty} \tau^{(\alpha,k)}$ admits ϵ-consistent price systems for each $\epsilon > 0$ while satisfying

$$\mathbf{P}(\sigma^{(\alpha)} = T) \geq 1 - \sum_{k=1}^{\infty} \mathbf{P}(\tau^{(\alpha,k)} < T) > 1 - \alpha.$$

Since α was chosen arbitrarily, we may, by letting $\tau_n = \sigma^{(2^{-n})}$, construct a whole sequence $(\tau_n)_{n\in\mathbb{N}}$ of such stopping times such that each S^{σ_n} admits ϵ-CPS for all $\epsilon > 0$ and such that

$$\lim_{n\to\infty} \mathbf{P}(\tau_n = T) = 1.$$

Note furthermore that by choosing $\tau^{(2^{-(n+1)},k)} \geq \tau^{(2^{-(n)},k)}$ for all k and n, we may ensure that τ_n is actually increasing. So this sequence has all the desired properties. □

Remark 5.3.3 Given the similarity of arbitrage and superreplication (arbitrage is a non-perfect superreplication strategy for the trivial claim $X = 0$ with initial endowment $x = 0$) and the lesson learned from the face-lifting theorem (that is: under whatever small transaction costs, there is nothing dynamic strategies can do that can't be already done with static ones), one might wonder whether NSA(0+) already implies NA(0+), that means, whether there even exists a model that satisfies ϵ-arbitrage, but no simple ϵ-arbitrage. The answer is: Already in the one-dimensional case NSA(0+) and NA(0+) are not the same, as was shown in Proposition A.2 of [GRS08b]. The reason why we cannot limit ourselves to static strategies (in this case Lemma 5.2.3 would deliver a satisfactory answer to the question of arbitrage) is that this time, conditional full support is not given.

5.4 Proof of the Fundamental Theorem

For the rest of this chapter, we will assume S to be continuous and NSA(0+) to hold true, such that Theorem 5.3.2 can be applied. We see the sequence $(\tau_n)_{n\in\mathbb{N}}$ constructed in Theorem 5.3.2 as fixed from now on. Starting from this sequence, we furthermore introduce the set of finite compositions of these stopping times: Denote by $\mathcal{S}$ the set of all stopping times of the form

$$\sigma = \sum_{i=1}^{N} \tau_{n_i} \mathbb{1}_{A_i}$$

with $N, n_1, n_2, \ldots, n_N \in \mathbb{N}$ and a decomposition of Ω into mutually disjoint sets $A_i \in \mathcal{F}_{\tau_{n_i}}$. For such a $\sigma \in \mathcal{S}$, denote by

$$c(\sigma, \epsilon) = \{(M, \mathbf{Q}) : (M, \mathbf{Q}) \sim_\epsilon (S^\sigma, \mathbf{P})\}$$

the set of all ϵ-consistent price systems up to σ. Furthermore, we denote the projection of $c(\sigma, \epsilon)$ to its second component by $q(\sigma, \epsilon)$, i. e.

$$q(\sigma, \epsilon) = \{\mathbf{Q} \sim \mathbf{P} : (M, \mathbf{Q}) \in c(\sigma, \epsilon) \text{ for some } M\}.$$

Note that NSA(0+) implies that $c(\sigma, \epsilon) \neq \emptyset$ for all $\sigma \in \mathcal{S}$ and $\epsilon > 0$: Assume $\sigma = \sum_{i=1}^{N} \tau_{n_i} \mathbb{1}_{A_i}$ and $n_i \leq \bar{n}$ for all i, then Theorem 5.3.2 implies that there exists $(M, \mathbf{Q}) \in c(\tau_{\bar{n}}, \epsilon)$, hence $(M^\sigma, \mathbf{Q}) \in c(\sigma, \epsilon)$.

Definition 5.4.1 *Let $(M^1, \mathbf{Q}_1) \in c(\sigma_1, \epsilon_1)$ be given. For $\sigma_2 \in \mathcal{S}, \sigma_2 \geq \sigma_1$, we call a CPS $(M^2, \mathbf{Q}_2) \in c(\sigma_2, \epsilon_2)$ a* sequel *of $(M^1, \mathbf{Q}_1)$, if all the following are given:*

1. *$M^1_t = M^2_t$ a. s. on $\{t \leq \sigma_1\}$, for all t*
2. *$\mathbf{Q}_1|_{\mathcal{F}_{\sigma_1}} = \mathbf{Q}_2|_{\mathcal{F}_{\sigma_1}}$*
3. *$\frac{d\mathbf{Q}_2}{d\mathbf{Q}_1} = 1$ a. s. on $\{\sigma_1 = \sigma_2\}$*

We will try to construct a consistent price system for $(S, \mathbf{P})$ as the limit of a sequence of sequels. The following lemma will tell us that, given NSA(0+), such sequels always exist for every choice of $\sigma_2 \geq \sigma_1$, and $(M^1, \mathbf{Q}_1) \in c(\sigma_1, \epsilon_1)$, as long as $\epsilon_2 > \epsilon_1$.

Lemma 5.4.2 *Let, for some $\sigma_1 \in \mathcal{S}$ and $\epsilon > 0$, a CPS $(M, \mathbf{Q}) \in c(\sigma_1, \epsilon)$ be given. Consider, for some $N \in \mathbb{N}$, another stopping time $\sigma_2 \in \mathcal{S}, \sigma_2 \geq \sigma_1$ of the form*

$$\sigma_2 = \sum_{i=1}^{N} \tau_{n_i} \mathbb{1}_{B_i},$$

where $B_i \in \mathcal{F}_{\sigma_1}$ for $i = 1, 2, \ldots, N$. For all i, let consistent price systems $(\mathbf{R}_i, M^i) \in c(\tau_{n_i}, \delta)$, with some arbitrary $\delta > 0$, be given. Then there exists a sequel

$$(\tilde{\mathbf{Q}}, \tilde{M}) \in c\big(\sigma_2, f(\epsilon, \delta)\big)$$

of $(M^1, \mathbf{Q}_1)$, where

$$f(\epsilon, \delta) = \max\left(\frac{1+\delta}{1-\delta}(1+\epsilon) - 1, 1 - \frac{1-\delta}{1+\delta}(1-\epsilon)\right),$$

that satisfies

$$\tilde{\mathbf{Q}}(B|\mathcal{F}_{\sigma_1}) = \mathbf{R}_i(B|\mathcal{F}_{\sigma_1})$$

almost surely for all $1 \leq i \leq N$ and $B \in \mathcal{F}, B \subseteq B_i$.

Proof: Let

$$\tilde{M} = M \mathbb{1}_{[\![0,\sigma_1]\!]} + \sum_{i=1}^{N} \mathbb{1}_{B_i} \frac{M_{\sigma_1}}{M^i_{\sigma_1}} M^i \mathbb{1}_{]\!]\sigma_1,\tau_{n_i}]\!]}.$$

Note that from

$$1 - \epsilon \leq \frac{M_{\sigma_1}}{S_{\sigma_1}} \leq 1 + \epsilon$$

$$1 - \delta \le \frac{M_t^i}{S_t} \le 1 + \delta$$

on $\{t \le \sigma_2\}$ follows

$$\frac{1-\delta}{1+\delta}(1-\epsilon) \le \frac{\tilde{M}_t}{S_t} \le \frac{1+\delta}{1-\delta}(1+\epsilon).$$

It can then be seen that if we then let

$$\frac{d\tilde{\mathbf{Q}}}{d\mathbf{Q}} = \mathbb{1}_{\sigma_1=\sigma_2} + \sum_{k=1}^{N} \mathbb{1}_{B_i} \frac{d\mathbf{R}_i/d\mathbf{Q}}{\mathbb{E}_{\mathbf{Q}}(d\mathbf{R}_i/d\mathbf{Q}|\mathcal{F}_{\sigma_1})} \mathbb{1}_{\sigma_1<\tau_{n_i}},$$

then $(\tilde{M}, \tilde{\mathbf{Q}})$ has the desired properties. □

The only thing that matters about the function f is that

$$\lim_{\delta \to 0} f(\epsilon, \delta) = \epsilon.$$

This allows us to define the sequence $(\delta_k)_{k\in\mathbb{N}}$ as follows: For fixed $\epsilon > 0$, start with $\epsilon_1 = \frac{\epsilon}{2}$ and then recursively define $\delta_k > 0$ small enough to ensure

$$\epsilon_{k+1} = f(\epsilon_k, \delta_k) \le (1 - 2^{-k-1})\epsilon.$$

We want to construct, for some suitable sequence $(\sigma_k)_{k\in\mathbb{N}}$ in $\mathcal{S}$ increasing stationarily to T, an infinite sequence $(M^k, \mathbf{Q}_k) \in c(\sigma_k, \epsilon_k)$, each CPS being a sequel of its predecessor. Lemma 5.4.2 alongside with the fact that under NSA(0+), $c(\tau_n, \delta_k)$ is never empty for any $k, n \in \mathbb{N}$, tells us that such a sequence can always be extended ad infinitum independent of our choice of σ_k. The question is, however, whether or not the sequence converges to a true CPS $(M, \mathbf{Q}) \in c(T, \epsilon)$, let alone how to define convergence in this context. Obviously, on sets of the form $\bigcup_{k=1}^{N}\{\sigma_k = T\}$, convergence is not a problem, since $\frac{d\mathbf{Q}_k}{d\mathbf{P}}$ and M^k change at most N times here. But outside, we may run into singularities, so we have to choose the sequels carefully. We therefore introduce the process

$$X_t = \left(1 + \sum_{i=1}^{d} S_t^{(i)}\right) \mathbb{1}_{t<T}$$

And look for sequels that satisfy

$$\mathbb{E}_{\mathbf{Q}_k}(X_{\sigma_k}) \le 2^{-k}$$

for all $k \in \mathbb{N}$. As we will see later, the existence of such a sequence implies the existence of a real ϵ-consistent price system for $(S, \mathbf{P})$, while the absence of such a sequence leads to arbitrage. To separate these two cases, the following lemma is crucial.

Lemma 5.4.3 *Consider some $k \in \mathbb{N}$, $\epsilon > 0$ and a stopping time $\sigma_k \in \mathcal{S}$ of the form*

$$\sigma_k = \sum_{i=1}^{N} \mathbb{1}_{A_i} \tau_{n_i}$$

and a consistent price system $(M^k, \mathbf{Q}_k) \in c(\sigma_k, \epsilon_k)$. Then one of the following holds true:

1. *There exists some $\sigma_{k+1} \in \mathcal{S}$ a. s. greater or equal than $\sum_{i=1}^{N} \mathbb{1}_{A_i} \tau_{n_{i+1}}$ and a sequel $(M^{k+1}, \mathbf{Q}_{k+1}) \in c(\sigma_{k+1}, \epsilon_{k+1})$ to $(M^k, \mathbf{Q}_k)$ satisfying*

$$\mathbb{E}_{\mathbf{Q}_{k+1}}(X_{\sigma_{k+1}}) \leq 2^{-(k+1)}$$

2. *There exists some $\bar{n} \in \mathbb{N}$ and a set $A \in \mathcal{F}_{\tau_{\bar{n}}}$ with $\mathbf{P}(A) > 0$ as well a constant $\alpha > 0$, such that*

$$\mathbb{E}_{\mathbf{Q}}(X_{\tau_n} | \mathcal{F}_{\tau_{\bar{n}}}) > \alpha$$

almost surely on A for all $n > \bar{n}$ and all $\mathbf{Q} \in q(\tau_n, \delta_{k+1})$.

Proof: We set $\bar{n} = \max(n_1, n_2, \ldots, n_N) + 1$. Consider the $\mathcal{F}_{\sigma_k}$-measurable random variable

$$\Xi = \text{essinf}\Big(\mathbb{E}_{\mathbf{Q}}(X_{\tau_n} | \mathcal{F}_{\tau_{\sigma_k}}); \mathbf{Q} \in q(\tau_n, \delta_{k+1}), n > \bar{n}\Big)$$

i. e. Ξ is the largest random variable smaller or equal $\mathbb{E}_{\mathbf{Q}}(X_{\tau_n} | \mathcal{F}_{\tau_{\sigma_k}})$ a. s. for all $n > \bar{n}$ and $\mathbf{Q} \in q(\tau_n, \delta_{k+1})$.
If $\mathbb{E}_{\mathbf{Q}_k}(\Xi) < 2^{-(k+1)}$, then there exists a decomposition of Ω into mutually disjoint $\mathcal{F}_{\sigma_k}$-measurable sets $B_1, B_2, \ldots, B_K$ and some integers $m_1, m_2, \ldots, m_K > \bar{n}$ as well as consistent price systems

$$(\mathbf{R}_i, \tilde{M}^{(i)}) \in c(\tau_{m_i}, \delta_{k+1}), \quad 1 \leq i \leq K,$$

such that

$$\mathbb{E}_{\mathbf{Q}_k}\left(\min_{1 \leq i \leq K} \mathbb{E}_{\mathbf{R}_i}(X_{\tau_{m_i}} | \mathcal{F}_{\sigma_k})\right) = \mathbb{E}_{\mathbf{Q}_k}\left(\sum_{i=1}^{K} \mathbb{E}_{\mathbf{R}_i}(X_{\tau_{m_i}} | \mathcal{F}_{\sigma_k}) \mathbb{1}_{B_i}\right) < 2^{-(k+1)}.$$

If we concatenate these K consistent price systems with $(M^k, \mathbf{Q}_k)$ as done in Lemma 5.4.2, we end up with a sequel $(M^{k+1}, \mathbf{Q}_{k+1}) \in c(\sigma_{k+1}, \epsilon_{k+1})$, where

$$\sigma_{k+1} = \sum_{i=1}^{K} \mathbb{1}_{B_i} \tau_{m_i} \geq \sum_{i=1}^{N} \mathbb{1}_{A_i} \tau_{n_{i+1}},$$

almost surely such that

$$\begin{aligned} \mathbb{E}_{\mathbf{Q}_{k+1}}(X_{\sigma_{k+1}}) &= \mathbb{E}_{\mathbf{Q}_{k+1}} \left(\mathbb{E}_{\mathbf{Q}_{k+1}}(X_{\sigma_{k+1}} | \mathcal{F}_{\sigma_k}) \right) \\ &= \mathbb{E}_{\mathbf{Q}_{k+1}} \left(\sum_{i=1}^{K} \mathbb{E}_{\mathbf{Q}_{k+1}}(X_{\tau_{m_i}} | \mathcal{F}_{\sigma_k}) \mathbb{1}_{B_i} \right) \\ &= \mathbb{E}_{\mathbf{Q}_k} \left(\sum_{i=1}^{K} \mathbb{E}_{\mathbf{R}_i}(X_{\tau_{m_i}} | \mathcal{F}_{\sigma_k}) \mathbb{1}_{B_i} \right) < 2^{-(k+1)}, \end{aligned}$$

thus we have found our sequel satisfying (1).
If, on the other hand, $\mathbb{E}_{\mathbf{Q}_k}(\Xi) \geq 2^{-(k+1)}$, we may choose $\alpha = \frac{1}{2}\mathbb{E}_{\mathbf{Q}_k}(\Xi)$ and $A = \{\Xi > \alpha\}$ to satisfy (2). $\square$

Proposition 5.4.4 *Assume there exists a sequence $(\sigma_k)_{k \in \mathbb{N}}$ in $\mathcal{S}$, stationarily increasing to T, and a sequence of consistent price systems $(M^k, \mathbf{Q}_k) \in c(\sigma_k, \epsilon_k)$, each being a sequel of its predecessor, such that*

$$\lim_{k \to \infty} \mathbb{E}_{\mathbf{Q}_k}(X_{\sigma_k}) = 0$$

holds true. Then $(S, \mathbf{P})$ admits an ϵ-consistent price system.

Proof: Since $(\sigma_k)_{k \in \mathbb{N}}$ becomes stationary $\mathbf{P}$-almost surely, the sequence $(\frac{d\mathbf{Q}_k}{d\mathbf{P}})_{k \in \mathbb{N}}$ also becomes stationary a. s. Hence the limit $\frac{d\mathbf{Q}}{d\mathbf{P}} = \lim_{k \to \infty} \frac{d\mathbf{Q}_k}{d\mathbf{P}}$ exists and is strictly positive $\mathbf{P}$-a. s. So we have defined a measure $\mathbf{Q} \sim \mathbf{P}$ on $(\Omega, \mathcal{F})$. Now $\lim_{k \to \infty} \mathbb{E}_{\mathbf{Q}_k}(X_{\sigma_k}) = 0$ implies that especially

$$\lim_{k \to \infty} \mathbb{E}_{\mathbf{Q}_k}(\mathbb{1}_{\sigma_k < T}) = 0,$$

Thus monotone convergence, thogether with the fact that $\{\sigma_k = T\} \in \mathcal{F}_{\sigma_k}$ yields

$$\mathbf{Q}(\Omega) = \mathbf{Q}\Big(\bigcup_{k=1}^{\infty} \{\sigma_k = T\}\Big)$$

$$= \lim_{k\to\infty} \mathbf{Q}(\sigma_k = T) = \lim_{k\to\infty} \mathbf{Q}_k(\sigma_k = T) = 1,$$

hence $\mathbf{Q}$ is in fact a probability measure.
For $t \in [0, T]$, the sequence $(M_t^k)_{k\in\mathbb{N}}$ also has a $\mathbf{P}$-a. s. stationary limit, which we will denote by M_t. Since $1 - \epsilon \leq \frac{M_t^k}{S_t} \leq 1 + \epsilon$ $\mathbf{Q}_k$-a. s. on $\{t \leq \sigma_k\}$ for all k, we also have $1 - \epsilon \leq \frac{M_t}{S_t} \leq 1 + \epsilon$ $\mathbf{Q}$-a. s.
This, together with $\mathbb{E}_{\mathbf{Q}_k}(S_{\sigma_k}\mathbb{1}_{\sigma_k<T}) \to 0$ component-wise also implies

$$\mathbb{E}_{\mathbf{Q}_k}(M_{\sigma_k})\mathbb{1}_{\sigma_k<T} \leq (1+\epsilon)\mathbb{E}_{\mathbf{Q}_k}(S_{\sigma_k}\mathbb{1}_{\sigma_k<T}) \to 0$$

component-wise. Thus we have, again using monotone convergence,

$$\mathbb{E}_{\mathbf{Q}}(M_T|\mathcal{F}_t) = \mathbb{E}_{\mathbf{Q}}(\lim_{k\to\infty} M_T\mathbb{1}_{\sigma_k=T}|\mathcal{F}_t)$$

$$= \lim_{k\to\infty} \mathbb{E}_{\mathbf{Q}}(M_T\mathbb{1}_{\sigma_k=T}|\mathcal{F}_t) = \lim_{k\to\infty} \mathbb{E}_{\mathbf{Q}_k}(M_{\sigma_k}^k\mathbb{1}_{\sigma_k=T}|\mathcal{F}_t)$$

$$= \lim_{k\to\infty} \Big(\mathbb{E}_{\mathbf{Q}_k}(M_{\sigma_k}^k|\mathcal{F}_t) - \mathbb{E}_{\mathbf{Q}_k}(M_{\sigma_k}^k\mathbb{1}_{\sigma_k<T}|\mathcal{F}_t)\Big)$$

$$= \lim_{k\to\infty} M_{t\wedge\sigma_{n_k}}^k - 0 = M_t$$

Thus M is a $\mathbf{Q}$-martingale and hence we have $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$. This completes the proof.

□

In order to prove the Fundamental Theorem, we need to show the absence of consistent price systems leads to arbitrage. We already know that if there is no CPS, then 5.4.3.2 holds true for some $\alpha > 0$. This leads to the situation that

$$\mathbb{E}_{\mathbf{Q}}(X_{\tau_n}|\mathcal{F}_{\tau_{\bar{n}}}) > \alpha$$

on some set A for all consistent price systems up to any stopping time τ_n, while $X_{\tau_n} \to 0$ as $n \to \infty$. This discrepancy on the set A can be used to construct what is called *free lunch with bounded risk* in the literature, i. e. a series of payoffs with bounded downside risk, which can all be superreplicated with admissible strategies and which tend to some strictly nonnegative yet nontrivial value. To bridge the gap between this notion and true arbitrage, we need the following result, which is in spirit of the proof of [GRS08b, Theorem 1.11] and [CS06, Proposition 13]. What this result shows is that the set of claims that can be superreplicated by admissible strategies stopping no later than some τ_n is Fatou-closed. This was shown to be true in [CS06, Proposition 14] for the case that S admits ϵ-consistent price systems for all $\epsilon > 0$ which is not the case here.

But as long as NSA(0+) holds true, it is true for S^{τ_n} i. e. for strategies that end at some fixed τ_n. Using a diagonal procedure, we show the same for the union of all these sets.

Theorem 5.4.5 *Assume $\epsilon > 0$. Consider a sequence $(Y_n)_{n\in\mathbb{N}}$ of random variables such that*

1. *Y_n is $\mathcal{F}_{\tau_n}$-measurable for all $n \in \mathbb{N}$*
2. *there exists some $a \in \mathbb{R}$ such that $Y_n \geq -a(1+\sum_{i=1}^d S^{(i)}_{\tau_n})$ a. s. for all $n \in \mathbb{N}$*
3. *$\mathbb{E}_{\mathbf{Q}}(Y_n) \leq 0$ for all $n \in \mathbb{N}$ and $\mathbf{Q} \in q(\epsilon, \tau_n)$*
4. *$\lim_{n\to\infty} Y_n = Y$ $\mathbf{P}$-a. s. for some random variable $Y \in L^0(\mathcal{F}_T)$*

Then there exists some $\epsilon^{(1/3)}$-admissible strategy H satisfying

$$V_{\epsilon^{(1/3)},S}(H) \geq Y,$$

where $\epsilon^{(1/3)}$ is as in Lemma 2.2.8.

Proof: Following the superreplication theorem 2.2.9, applied to Y_n and S^{τ_n}, we find, for each n, some $\big(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})(1+\epsilon)\big)$-admissible strategy H^n satisfying

$$V_{\epsilon^{(1/3)},S^{\tau_n}}(H^n) \geq Y_n$$

almost surely. Since S^{τ_n} does not move after τ_n, we may choose H^n equal to zero on $[\tau_n, T]$. It follows from Lemma 11 and Proposition 13 of [CS06], applied to the sequence $(H^n)_{n\in\mathbb{N}}$ (or rather their portfolio process counterparts $V(H^n)$, see Remark 2.2.7.1) and S^{τ_1}, that there exists a sequence $(G^{1,n})_{n\in\mathbb{N}}$ of processes, each one satisfying

$$G^{1,n} \in \text{conv}(H^m, m \geq n),$$

and a predictable, finite-variation process $(\tilde{G}^1_t)_{t\in[0,\tau_1]}$ such that

$$\mathbf{P}\Big((G^{1,n})^{\tau_1}_t \to (\tilde{G}^1)^{\tau_1}_t \;\forall\, t \in [0,T]\Big) = 1.$$

We proceed inductively, now with $(G^{1,n})_{n\in\mathbb{N}}$ instead of $(H^n)_{n\in\mathbb{N}}$ and S^{τ_2} instead of S^{τ_1} (note that the properties required to apply Lemma 11 and Proposition 13 of [CS06] are inherited from H^n to $G^{1,n}$), to find a sequence $(G^{2,n})_{n\in\mathbb{N}}$, each one satisfying

$$G^{2,n} \in \text{conv}(G^{1,m}, m \geq n),$$

and a predictable, finite-variation process $(\tilde{G}^2_t)_{t\in[0,\tau_2]}$ such that

$$\mathbf{P}\Big((G^{2,n})^{\tau_2}_t \to (\tilde{G}^2)^{\tau_2}_t \;\forall\, t \in [0,T]\Big) = 1.$$

By construction, $\tilde{G}^1$ and $\tilde{G}^2$ coincide on $[0,\tau_1]$ a. s. If we repeat this procedure, we end up with a sequence of predictable, finite-variation processes $(\tilde{G}^n)_{n\in\mathbb{N}}$, each $\tilde{G}^{n+1}$ coinciding with its predecessor on $[0,\tau_n]$. Since $\tau_n \to T$ stationarily, the limiting process $H_t = \lim_{n\to\infty} \tilde{G}^n_t$ is still predictable and of a. s. finite variation. Since all strategies H^n are $\big(\epsilon^{(1/3)}, a(1+\epsilon^{(1/3)})(1+\epsilon)\big)$-admissible, the same is true for all convex combinations, and hence also for H.
Following Corollary 2.1.6.3 we have that

$$V_{\epsilon^{(1/3)},S}(H) \in \overline{\text{conv}}\,(Y_m; m \geq n) + L^0_+(\mathcal{F}_T)$$

for all $n \in \mathbb{N}$, where $L^0_+(\mathcal{F}_T)$ denotes the set of all a. s. nonnegative $\mathcal{F}_T$-measurable random variables. Hence $V_{\epsilon^{(1/3)},S}(H) \geq Y$ almost surely. □

We are now ready to complete the proof of the Fundamental Theorem of Asset Pricing. So from now on, we do not assume NSA(0+) up front.

Proof of Theorem 5.1.1: Suppose NA(0+) holds true, then so does NSA(0+), and so the sequence $(\tau_n)_{n\in\mathbb{N}}$ from Theorem 5.3.2 exists. If S does not admit an ϵ-consistent price system for some fixed $\epsilon > 0$, then we know from Proposition 5.4.4 that property 5.4.3.2 holds true for some $\alpha > 0$, $\bar{n} \in \mathbb{N}$ and $A \in \mathcal{F}_{\tau_{\bar{n}}}$ with $\mathbf{P}(A) > 0$. It follows that, for all $n > \bar{n}$ and $\mathbf{Q} \in q(\tau_n, \delta_{\bar{n}+1})$:

$$\frac{\mathbb{E}_{\mathbf{Q}}(X_{\tau_n})}{\mathbf{Q}(A)} > \alpha,$$

or equivalently,

$$\mathbb{E}_{\mathbf{Q}}(\alpha \mathbb{1}_A - X_{\tau_n}) < 0.$$

Now the sequence $Y_n = \alpha\mathbb{1}_A - X_{\tau_n}$ satisfies the conditions of Theorem 5.4.5 with $a = 1$. Hence there exists an η-admissible strategy H such that

$$V_{\eta,S}(H) \geq Y = \alpha \mathbb{1}_A$$

almost surely, where

$$\eta = (1+\delta_{\bar{n}+1})^{1/3} - 1 > 0.$$

So H is an η-arbitrage, in contradiction to our assumption NA(0+). hence there exists an ϵ-consistent price system.
On the contrary, note that if we have a CPS $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$, Lemma 2.1.5 tells us that for every ϵ-admissible strategy H we have

$$V_{\epsilon,S}(H) \leq V_{0,M}(H),$$

while on the other hand Lemma 2.1.7 tells us that the process $t \mapsto V_{0,M}(H\mathbb{1}_{(0,t)})$ is a $\mathbf{Q}$-supermartingale. Therefore

$$\mathbb{E}_{\mathbf{Q}}\big(V_{\epsilon,S}(H)\big) \leq \mathbb{E}_{\mathbf{Q}}\big(V_{0,M}(H)\big) \leq V_{0,M}(0) = 0,$$

hence H is not an ϵ-arbitrage. This completes the proof. □

Example 5.4.6 (Example 5.2.8 revisited) Consider once more the one-dimensional Example 5.2.8. There S was shown to satisfy NA(0+), but since S is increasing and not a. s. constant, no equivalent martingale measure exists. Here we can construct ϵ-CPS explicitly for all $\epsilon > 0$: As described in Example 3.2.1, we may change the intensities of X and Y deliberately. So assume X has intensity $\frac{1}{T}$, while Y given $t \geq X$ has intensity $\frac{1}{\epsilon}$ under some measure $\mathbf{Q} \sim \mathbf{P}$. Define the process M by

$$M_t = (1-\frac{\epsilon t}{T})\mathbb{1}_{[0,X)}+(1-\frac{\epsilon X}{T}+\epsilon+t-X)\mathbb{1}_{[X,X+Y)}+(1-\frac{\epsilon X}{T}+Y-X)\mathbb{1}_{(X+Y,T]}.$$

So M decreases with drift $-\frac{\epsilon}{T}$ and jumps up by ϵ at X. Then it moves parallel to S until it jumps down by ϵ at $X+Y$. After that it stays constant, like S. The way the $\mathbf{Q}$-intensities were chosen, M is a $\mathbf{Q}$-martingale, while satisfying $(1-\epsilon)S_t \leq M_t \leq (1+\epsilon)S_t$ in all three intervals. Therefore $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$.

In [GRS08a] it was shown that the CFS property for d-dimensional continuous processes implies the existence of ϵ-CPS for all $\epsilon > 0$. Obviously (CFS) implies NSA(0+), so Proposition 5.3.1 holds in this case. The proof there follows the same strategy as our proof of Proposition 5.3.1, but, since (CFS) implies $\mathbf{P}(\tau_{n+1} = T|\mathcal{F}_{\tau_n}) > 0$ almost surely for each $n \in \mathbb{N}$, the measure canges could be performed in a way there that secures the local CPS to converge. Using this result, the following important corollary is clear:

Corollary 5.4.7 *Let S be continuous and satisfy (CFS). Then NA(0+) holds.*

So all (CFS) processes are arbitrage free under small transaction costs. However, (CFS) does not imply absence of 0-arbitrage, as the following example shows

Example 5.4.8 (CFS, but 0-arbitrage) It was shown in [GRS08a, Lemma 4.5] that if a process X satisfies (CFS) on $\mathbb{R}$ instead of $\mathbb{R}_+$, then so does $Y_t = \int_0^t X_s ds$. So if we choose X to be a standard Brownian motion, then

$$S_t = \exp(\int_0^t X_s ds)$$

satisfies (CFS) on $\mathbb{R}_+$, hence by Corollary 5.4.7 NSA(0+) holds. Now S is continuously differentiable. If we define σ as the first time that $\frac{dS_t}{dt} = 1$ and τ to be the first time after σ that $\frac{dS_t}{dt} = 0$, then

$$H = \mathbb{1}_{[\sigma,\tau\wedge T)}\mathbb{1}_{\sigma<T}$$

is simple 0-arbitrage.

Using the preparatory work done in Section 4.2, it is now an easy task to derive an analogous result for one-dimensional Lévy processes.

Theorem 5.4.9 (FTAP for one-dimensional Lévy processes) *Assume $S = e^L$, where L is a one-dimensional Lévy process with characteristic triplet (σ^2, γ, ν). Then S satisfies NA(0+) if and only if S admits ϵ-consistent price systems for all $\epsilon > 0$.*

Proof: The proof that the existence of an ϵ-CPS implies NA(ϵ) from Theorem 5.1.1 can be used in this case as well, since there continuity of S was not used. So if S admits ϵ-CPS for all $\epsilon > 0$, then NA(0+) holds true.
Now assume that there exists $\epsilon > 0$ such that S does not admit an ϵ-CPS. With Theorems 4.2.2 and 4.2.4 it follows that S is either increasing with drift $\gamma_0 > 0$ or decreasing with $\gamma_0 < 0$. In the first case we have that $S_T \geq e^{\gamma_0 T}$ almost surely, so the strategy $H = \mathbb{1}_{(0,T)}$ yields $V_{\delta,S}(T) > 0$ a. s. for all $\delta \in [0, \frac{e^{\gamma_0 T}-1}{e^{\gamma_0 T}+1})$. In the second case, we have that S is bounded and $S_T \leq e^{\gamma_0}$ a. s. Then the strategy $H = -\mathbb{1}_{(0,T)}$ is δ-admissible and yields $V_{\delta,S}(T) > 0$ a. s. for all $\delta \in [0, \frac{1-e^{\gamma_0 T}}{1+e^{\gamma_0 T}})$. So in both cases NA(0+) is not satisfied. □

Remark 5.4.10 (Open problem) In the setting of Theorem 5.4.9, it proved possible to omit the *for-all*-quantifiers in our version of the FTAP. In both the continuous and the one-dimensional Lévy case, we were able to show not only that the existence of ϵ-CPS *for all* $\epsilon > 0$ implies NA(0+), but also that the existence of an ϵ-CPS for *some fixed* $\epsilon > 0$ implies NA(ϵ), a considerably stronger statement.
In the Lévy case, we were also able to show the opposite direction: If no ϵ-CPS

exists, then ϵ-arbitrage is possible, or, equivalently, NA(ϵ) implies the existence of an ϵ-CPS. In the continuous case, we are far away from such a statement. In fact, even the existence of a function $f : (0, 1) \to (0, 1)$ such that every continuous process that satisfies NA(ϵ) admits an $f(\epsilon)$-CPS, is not clear. But at least, no counterexample disproving the existence of such a function (that would be a family $\mathcal{S}$ of continuous processes such that for every pair $(\epsilon, \delta) \in (0, 1)^2$ there exists some $S \in \mathcal{S}$ satisfying NA(ϵ) but not admitting a δ-CPS) is known to the author, either.

Chapter 6

The FTAP for generalized transaction costs

In the present framework, we assume the level of costs ϵ as equal for all d assets, contrary to the observation that actual costs may vary, i. e. because of differences in liquidity. Furthermore we do not allow swapping of assets: All transactions have to go through the cash account, on which also the transaction costs are registered. Therefore, transferring capital from asset i to asset j results in a proportion of 2ϵ in transaction costs, rather than ϵ. This is a natural assumption for example in equity markets, but seems unnatural for example in multi-currency markets: While transaction costs may be higher if we exchange exotic, rarely traded currencies into one another, direct trading should still be cheaper than always funneling our capital through the same predetermined anchor currency, which is in our case the cash account.
In this chapter, we propose a generalized model for trading in a d-dimensional market with proportional transaction costs, along with the corresponding cone-framework in the spirit of Section 2.2. We show that the model used in the previous chapters is embedded as a special case and that most of the basic theorems derived so far still hold true. The Fundamental Theorem of Asset Pricing follows as corollary of Theorem 5.1.1 if the transaction costs are assumed to be strictly positive. This assumption, however, is then weakened by introducing the notion of efficient friction. With this property, also one-sided transaction csts are allowed, for instance for products that one can buy without costs but that charge a fee for selling.

6.1 Transaction costs matrices

The aforementioned two aspects, varying costs for different assets and the possibility of direct asset swapping, are tackled by the following generalization of our model: Instead of one single constant ϵ, consider a transaction costs matrix

$$C = (c_{ij})_{i,j \in \{0,1,\dots,d\}} \in [0,1)^{(d+1)\times(d+1)},$$

where c_{ij} denotes the proportion to be paid when the i-th asset is exchanged into the j-th one. Here the cash account is written as 0-th asset $(S_t^{(0)})_{t\in[0,T]}$, which is assumed to be equal to 1 for all t. For this new framework we need to renew our definitions of trading strategies, values and admissibility. The d-dimensional strategy H is now replaced by a $(d+1)^2$-dimensional process $H = (H_t^{(ij)})_{t\in[0,T],0\le i,j\le d}$, where the process $H^{(ij)}$ counts the accumulated transactions from asset i to asset j, but only in this direction.

Definition 6.1.1 *For a price process S and a given transaction costs matrix $C \in [0,1)^{(d+1)\times(d+1)}$, a predictable $(d+1)^2$-dimensional process H is called* trading strategy *with respect to S and C, if for all $i, j \in \{0, 1, \dots, d\}$*

1. *$H_0^{(ij)} = 0$ a. s.*
2. *$H_t^{(ii)} = 0$ a. s. for all $t \in [0,T]$*
3. *$H^{(ij)}$ is almost surely increasing*
4. *$\tilde{H}_T^{(i)} = 0$ a. s. for $i \ge 1$, where the process $\tilde{H}^{(i)}$ is defined by*

$$\tilde{H}_t^{(i)} = \sum_{j=0}^{d} H_t^{(ji)} - \sum_{j=0}^{d} \int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} dH_s^{(ij)}$$

The value of a trading strategy H is defined as

$$V_{C,S}(H) = \sum_{i=1}^{d} \int_0^T \tilde{H}_t^{(i)} dS_t^{(i)} - \sum_{i,j=0}^{d} c_{ij} \int_0^T S_t^{(j)} dH_t^{(ij)}.$$

*For $a \in \mathbb{R}$, we call H (C,a)-*admissible *with respect to a numeraire process N, if*

$$V_{C,S}\Big(H^t + \sum_{i=1}^{d} \tilde{H}_t^{(i)} S_t^{(i)} M_{i0} \mathbb{1}_{[t,T]}\Big) \ge -aN_t$$

almost surely for all $t \in [0,T]$. Here the matrix

$$M_{i0} = (m_{kj})_{k,j \in \{0,1,\dots,d\}} \in \mathbb{R}^{(d+1)\times(d+1)}$$

is defined by $m_{i0} = 1$ and $m_{kj} = 0$ otherwise.
*H is called C-*admissible *with respect to N if it is (C,a)-admissible w. r. t. N for some a.*

Remark 6.1.2 1. In this definition, the process $\tilde{H}^{(i)}$ for $i \geq 1$ indicates how many units of asset i the trader holds at every time point. So property 4 states that all positions have to be liquidated by maturity, just like property 2 in Definition 2.1.1. In fact, the following lemma shows that both definitions are consistent with each other.

2. Note that, in contrast to Definition 2.1.1, the definition of a trading strategy depends on S, i. e. a process which is a trading strategy w.r.t. some process S need not be one w.r.t. some other process. The reason is that in the current setting, it is not possible to determine the number of assets held at a specific moment just by looking at the history of H, without knowing the history of S, so we cannot know if H meets the liquidation property 6.1.1.4 without knowing S. The transaction costs matrix, however, plays no role here.

3. Also, the notion of admissibility is more involved, compared to Definition 2.1.1. The process

$$H^t + \sum_{i=1}^{d} \tilde{H}_t^{(i)} S_t^{(i)} M_{i0} \mathbb{1}_{[t,T]}$$

is the same as H until t and then liquidated by selling all assets (i. e. increasing $H^{(i0)}$), analogous to the liquidated process $H\mathbb{1}_{(0,t)}$.

Lemma 6.1.3 *For $\epsilon \in [0,1)$, define the transaction costs matrix C by*

$$C_\epsilon = \begin{pmatrix} 0 & \epsilon & \cdots & \cdots & \epsilon \\ \epsilon & 0 & 2\epsilon & \cdots & 2\epsilon \\ \vdots & 2\epsilon & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 & 2\epsilon \\ \epsilon & 2\epsilon & \cdots & 2\epsilon & 0 \end{pmatrix}$$

1. *For a trading strategy $H = (H^{(ij)})_{0\le i,j\le d}$ in the sense of Definition 6.1.1, define the d-dimensional process $\tilde{H} = (\tilde{H}^{(i)})_{1\le i\le d}$ as in 6.1.1.4. Then $\tilde{H}$ is a trading strategy in the sense of Definition 2.1.1 and we have*

$$V_{C_\epsilon,S}(H) = V_{\epsilon,S}(\tilde{H})$$

almost surely. Furthermore H is (C_ϵ, a)-admissible with respect to some numeraire N in the sense of Definition 6.1.1 if and only if $\tilde{H}$ is (ϵ, a)-admissible with respect to N for in the sense of Definition 2.1.2.

2. *Let $\hat{H} = (\hat{H}^{(i)})_{1\le i\le d}$ be a trading strategy in the sense of Definition 2.1.1 and define the increasing processes $\hat{H}^{(i+)}$ and $\hat{H}^{(i-)}$ as in Lemma 2.1.5. Define the $(d+1)^2$-dimensional process $H = (H^{(ij)})_{0\le i,j\le d}$ by $H^{(00)} = 0$ and*

$$H^{(0i)} = \hat{H}^{(i+)}, \;\; 1 \le i \le d$$

$$H_t^{(i0)} = \int_0^t S_s^{(i)} d\hat{H}_s^{(i-)}, \;\; 1 \le i \le d$$

$$H^{(ij)} = 0, \;\; 1 \le i, j \le d.$$

Then H is a trading strategy in the sense of 6.1.1 and the definitions of value and admissibility coincide, as in the first part.

Proof: (1): Since H is increasing in each component, it is of finite variation. Thus $\tilde{H}$ is by construction predictable and of finite variation. Since $\tilde{H}_0 = 0$ by construction and $\tilde{H}_T = 0$ was imposed, $\tilde{H}$ is a trading strategy. Note that since $H^{(ji)}$ and $\int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} dH_s^{(ij)}$ are all increasing, we have

$$\mathrm{Var}_t^{\tilde{H}^{(i)}} = \sum_{j=0}^d H_t^{(ji)} + \sum_{j=0}^d \int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} dH_s^{(ij)}.$$

We therefore have

$$\sum_{i=1}^d \epsilon \int_0^T S_t^{(i)} d\mathrm{Var}_t^{\tilde{H}^{(i)}} = \sum_{i=1}^d \epsilon \int_0^T S_t^{(i)} \sum_{j=0}^d (dH_t^{(ji)} + \frac{S_t^{(j)}}{S_t^{(i)}} dH_t^{ij})$$

$$= \sum_{i=1}^d \sum_{j=0}^d \epsilon \int_0^T (S_t^{(i)} dH_t^{(ji)} + S_t^{(j)} dH_t^{(ij)})$$

$$= \sum_{i,j=0}^{d} c_{ij} \int_0^T S_t^{(j)} dH_t^{(ij)},$$

Which proves $V_{C_\epsilon,S}(H) = V_{\epsilon,S}(\tilde{H})$. If we repeat this computation with the strategy H liquidated at t instead of H, we find that the corresponding d-dimensional strategy equals $\tilde{H}\mathbb{1}_{(0,T)}$. This proves the equivalence of the notions of admissibility.

(2): H trivially satisfies the first three properties of a $(d+1)^2$-dimensional trading strategy by construction. For the last one, consider the process $\tilde{H}$ constructed out of H in 2.1.1.4. There we have

$$\tilde{H}_t^{(i)} = \sum_{j=0}^{d} H_t^{(ji)} - \sum_{j=0}^{d} \int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} dH_s^{(ij)}$$

$$= H_t^{(0i)} - \int_0^t \frac{S_t^{(0)}}{S_s^{(i)}} dH_s^{(i0)} = \hat{H}_t^{(i+)} - \int_0^t d\hat{H}_s{}^{(i-)} = \hat{H}_t^{(i)},$$

thus $\hat{H} = \tilde{H}$ almost surely. This proves the fourth an last property for H, and, by the first part of this lemma, the equivalence of values and admissibility. $\square$

Remark 6.1.4 A similar model was introduced in [Ka99] and refined in subsequent papers ([KL02], [KS02], [CS06]). There the transaction costs matrix $\Lambda = (\lambda_{ij})_{1\le i,j\le d}$ also indicates a proportion of wealth lost by transaction costs: When transferring capital from asset i to asset j, the amount by which the holding in asset i decreases is $(1+\lambda_{ij})$ times what it would be in the frictionless case. Therefore it is natural to impose the triangular inequality

$$(1+\lambda_{ij}) \ge (1+\lambda_{ik})(1+\lambda_{kj})$$

for $1 \le i,j,k \le d$, which means that direct trade is always equal or more favourable than indirect trade via a third asset. The main difference between this and our model is that there the transaction costs are always accounted for in the selling asset, whereas we always book the costs to the cash account and treat the exchange ratios in the risky assets like in the frictionless case. Therefore our triangular inequality is different, namely

$$c_{ij} \le c_{ik} + c_{kj}$$

for $0 \le i,j,k \le d$.

In order to define an adequate notion of a C-consistent price system for a given transaction costs matrix C, we recall the relation between ϵ-consistent price systems and the solvency cone derived in Lemma 2.2.3. Generalizing the notion of the solvency cone to our new setting is actually pretty straightforward: The analogous version of the vectors $L_t^{i\pm}$ spanning $\mathcal{K}_{\epsilon,t}$ are now vectors $L_t^{ij} \in \mathbb{R}^{d+1}$ representing a long position in some asset i, a short position of equal value in position j and the right amount of cash to pay the transaction costs when these two positions are cancelled out against each other. This leads to the following definition:

Definition 6.1.5 *For a transaction costs matrix $C = (c_{ij})_{i,j\in\{0,1,\ldots,d\}}$, we define the C-*solvency cone *$\mathcal{K}^S_{C,t}$ at time t as the $d+1$-dimensional cone spanned by $\mathbf{e}_i,\ 1 \leq i \leq d+1$ and the $d(d+1)$ vectors $\{L_t^{ij}, 0 \leq i \neq j \leq d\}$, where*

$$L_t^{ij} = \mathbf{e}_{i+1} - \frac{S_t^{(i)}}{S_t^{(j)}}\mathbf{e}_{j+1} + c_{ij}S_t^{(i)}\mathbf{e}_1.$$

The dual C-solvency cone *$\mathcal{K}^{S*}_{C,t}$ is defined as in Definition 2.2.1, i. e.*

$$\mathcal{K}^{S*}_{C,t} = \{y \in \mathbb{R}^{d+1} : x^\top y \geq 0 \text{ for all } x \in \mathcal{K}^S{}_{C,t}\}.$$

As in the case of uniform ϵ, the inclusion of the unit vectors $\mathbf{e}_i,\ 1 \leq i \leq d+1$ is only necessary if the transaction costs are not strictly positive, and does not affect $\mathcal{K}^{S*}_{C,t}$ at all.

As in Lemma 6.1.3, we show that this definition is consistent with the case of fixed ϵ in Definition 2.2.1:

Lemma 6.1.6 *For $\epsilon \in [0,1)$, define the transaction costs matrix C_ϵ as in Lemma 6.1.3. Then the C_ϵ-solvency cone $\mathcal{K}^S_{C_\epsilon,t}$ and the ϵ-solvency cone $\mathcal{K}^S_{\epsilon,t}$ from Definition 2.2.1 coincide for all $t \in [0,T]$.*

Proof: Recall the vectors $L_t^{i\pm}$ from Definition 2.2.1. To see $\mathcal{K}^S_{\epsilon,t} \subseteq \mathcal{K}^S_{C_\epsilon,t}$ it suffices to note that

$$L_t^{i+} = L_t^{i0}$$

$$L_t^{j-} = S_t^{(j)} L_t^{0j}.$$

On the other hand we have for $i, j, \geq 1$

$$L_t^{ij} = L_t^{i+} + \frac{S_t^{(i)}}{S_t^{(j-)}} L_t^{j-},$$

which proves $\mathcal{K}_{C_\epsilon,t} \subseteq \mathcal{K}_{\epsilon,t}$. □

The following lemma is proven by straightforward calculation.

Lemma 6.1.7 *For given S, C and a strategy H, define the $d+1$-dimensional process*

$$V_t^H = -\sum_{i,j=0}^{d} \int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} L_s^{ij} dH_s^{(ij)}.$$

Then for all $1 \leq i \leq d$ the process $\tilde{H}^{(i)}$ equals the $(i+1)$-st component of V^H, almost surely. Moreover, the first component of V_T^H equals $V_{C,S}(H)$ almost surely.

The process V^H can be seen as analogue to the self-financing portfolio process $V(H)$ introduced in Remark 2.2.7. So we have found an alternative representation of $V_{C,S}(H)$ that does not require $\tilde{H}$ explicitly.

Proposition 6.1.8 *Consider two price processes S, M and two transaction costs matrices C, D. Let H be a C-admissible trading strategy with respect to S. If for all $t \in [0,T]$ we have*

$$\mathcal{K}_{C,t}^S \subseteq \mathcal{K}_{D,t}^M,$$

then there exists a D-admissible trading strategy G with respect to M such that

$$V_{C,S}(H) \leq V_{D,M}(G) \; a.s.$$

Proof: The $\mathbb{R}^{(d+1)\times(d+1)}$-valued process G is given by

$$dG_t^{(ij)} = \min\left(1, \frac{S_t^{(j)} M_t^{(i)}}{S_t^{(i)} M_t^{(j)}}\right) dH_t^{(ij)}, \;\; 1 \leq i,j \leq d$$

$$dG_t^{(0j)} = dH_t^{(0j)} + \sum_{i=1}^{d} \max\left(0, 1 - \frac{S_t^{(j)} M_t^{(i)}}{S_t^{(i)} M_t^{(j)}}\right) dH_t^{(ij)}, \;\; 1 \leq j \leq d$$

$$dG_t^{(i0)} = \frac{M_t^{(i)}}{S_t^{(i)}} dH_t^{(i0)} + \sum_{j=1}^{d} \max\left(0, (\frac{S_t^{(j)}}{S_t^{(i)}} - \frac{M_t^{(j)}}{M_t^{(i)}}) M_t^{(i)}\right) dH_t^{(ij)}, \;\; 1 \leq i \leq d.$$

G is obviously a predictable, finite variation process starting at 0. Straightforward calculation yields for $1 \leq i \leq d$

$$d\tilde{G}_t^{(i)} = \sum_{n=0}^{d} dG_t^{(ni)} - \sum_{n=0}^{d} \frac{M_t^{(n)}}{M_t^{(i)}} dG_t^{(in)}$$

$$= \sum_{n=0}^{d} dH_t^{(ni)} - \sum_{n=0}^{d} \frac{S_t^{(n)}}{S_t^{(i)}} dH_t^{(in)} = d\tilde{H}_t^{(i)},$$

thus $\tilde{G} = \tilde{H}$, i. e. H and G hold the same amount of each asset at each time. So $\tilde{G}_T = 0$ and G is in fact an M-trading strategy.
Now recall the alternative representation of $V_{C,S}(H)$ from Lemma 6.1.7. Therefore consider the portfolio processes induced by H and G given by

$$V_t^H = -\sum_{i,j=1}^{d} \int_0^t \frac{S_s^{(j)}}{S_s^{(i)}} L_s^{ij} dH_s^{(ij)}$$

$$V_t^G = -\sum_{i,j=1}^{d} \int_0^t \frac{M_s^{(j)}}{M_s^{(i)}} \tilde{L}_s^{ij} dG_s^{(ij)},$$

where $\tilde{L}_t^{ij}$ denote the spanning vectors of $\mathcal{K}_{D,t}^M$ and L_t^{ij} those of $\mathcal{K}_{C,t}^S$. By the same argument as in Lemma 2.2.6 we see that $\Delta_\pm V_t^G$ and $\mathring{V}_t^G$ live on the boundary of $-\mathcal{K}_{D,t}^M$, while $\Delta_\pm V_t^H$ and $\mathring{V}_t^H$ live on the boundary of $-\mathcal{K}_{C,t}^S \subseteq -\mathcal{K}_{D,t}^M$. But since V^G and V^H can only differ in the first component and since $\mathbf{e}_1 \notin -\mathcal{K}_{D,t}^M$, this means that the first component of $V^G - V^H$ must be increasing, hence

$$V_{D,M}(G) = {V_T^G}^\top \mathbf{e}_1 \geq {V_T^H}^\top \mathbf{e}_1 = V_{C,S}(H).$$

Finally, note that the same argument also holds if we replace H by H liquidated at $t < T$, and the corresponding M-strategy coincides with G liquidated at t. This shows that the (C, a)-admissibility of H implies the (D, a)-admissibility of G. $\square$

Remark 6.1.9 The idea behind the strategy G is the following: If $\frac{S_t^{(j)}}{S_t^{(i)}} > \frac{M_t^{(j)}}{M_t^{(i)}}$, then exchanging asset i into asset j is more lucrative on M than on S. If $H^{(ij)}$ increases in this case, then $G^{(ij)}$ increases at the same rate, i. e. the same number of stocks j are purchased, but by giving up fewer units of stock i. The difference is sold to the cash account, hence $dG_t^{(ij)} = dH_t^{(ij)}$ and

$$dG_t^{(i0)} = dH_t^{(i0)} + \frac{S_t^{(j)} M_t^{(i)}}{S_t^{(i)} M_t^{(j)}} dH_t^{(ij)}.$$

If on the other hand $\frac{S_t^{(j)}}{S_t^{(i)}} < \frac{M_t^{(j)}}{M_t^{(i)}}$, then G gives up the same amount of stock i to get less of stock j. The difference has to be bought out of the cash account. Note especially that if $S = M$, then $G = H$. in this case $\mathcal{K}_{C,t}^S \subseteq \mathcal{K}_{D,t}^M$ is equivalent to $D \leq C$, component-wise.

The definition of a C-consistent price system now runs along the idea of Lemma 2.2.3. A $\mathbf{Q}$-martingale M is called C-CPS if the corresponding $\mathbf{P}$-martingale Z is $\mathcal{K}_{C,\cdot}^{S*}\backslash\{\mathbf{0}\}$-valued. This makes it easy to see the consistence with the original Definition 2.1.3, as the next lemma will show.

Definition 6.1.10 *For a transaction costs matrix C, a C-*consistent price system *to S is a pair $(M, \mathbf{Q})$ consisting of a probability measure $\mathbf{Q} \sim \mathbf{P}$ and an $\mathbb{R}_+^d$-valued $\mathbf{Q}$-martingale $(M_t)_{t\in[0,T]}$ such that*

$$\frac{M_t^{(j)}}{S_t^{(j)}} \leq \frac{M_t^{(i)}}{S_t^{(i)}} + c_{ij}$$

$\mathbf{Q}$-almost surely for all $0 \leq i, j \leq d$ and $t \in [0, T]$. Here both $M^{(0)}$ and $S^{(0)}$ are assumed to be equal to 1. We will use the shorthand notation $(M, \mathbf{Q}) \sim_C (S, \mathbf{P})$. For a given $S_t \in \mathbb{R}_+^d$, the bid-ask region *$\mathcal{B}(C, S_t)$ is the set of all $M_t \in \mathbb{R}_+^d$ that satisfy all of the above inequalities.*

Lemma 6.1.11

1. *Let a probability measure $\mathbf{Q} \sim \mathbf{P}$ and an $\mathbb{R}_+^d$-valued adapted process M be given. Then $(M, \mathbf{Q}) \sim_C (S, \mathbf{P})$ if and only if the process Z defined in Lemma 2.2.3.2 is a $\mathbf{P}$-martingale and satisfies $Z_t \in \mathcal{K}_{C,t}^{S*}\backslash\{\mathbf{0}\}$ a. s. for all t.*
2. *Consider $\epsilon \in [0, 1)$ and the corresponding transaction costs matrix C_ϵ from Lemma 6.1.3. Then we have $(M, \mathbf{Q}) \sim_\epsilon (S, \mathbf{P})$ for a pair $(M, \mathbf{Q})$ if and only if $(M, \mathbf{Q}) \sim_{C_\epsilon} (S, \mathbf{P})$.*

Proof: Like in the case of uniform ϵ, we have $\mathcal{K}_{C,t}^{S*} \in \mathbb{R}_+^{d+1} \cup \{\mathbf{0}\}$ a.s., so dividing by $Z^{(1)}$ is not a problem.
(1): Note that $Z_t \in \mathcal{K}_{C,t}^{S*}$ if and only if $Z_t^\top L_t^{ij} \geq 0$ for all $0 \leq i, j \leq d$, and that we have

$$Z_t^\top L_t^{ij} \geq 0 \Leftrightarrow M_t^{(i)} Z_t^{(1)} + \frac{S_t^{(i)}}{S_t^{(j)}} M_t^{(j)} Z_t^{(1)} + c_{ij} S_t^{(i)} Z_t^{(1)} \geq 0,$$

which itself is equivalent to

$$\frac{M_t^{(i)}}{S_t^{(i)}} - \frac{M_t^{(j)}}{S_t^{(j)}} + c_{ij} \geq 0.$$

The rest of the proof follows analogous to the proof of Lemma 2.2.3.

(2): By Lemma 6.1.6 we know that Z is a $\mathcal{K}^{S*}_{\epsilon,\cdot}\backslash\{\mathbf{0}\}$-valued $\mathbf{P}$-martingale if and only if it is a $\mathcal{K}^{S*}_{C_\epsilon,\cdot}\backslash\{\mathbf{0}\}$-valued $\mathbf{P}$-martingale. This together with the first part of the lemma completes the proof. $\square$

We are now ready to prove a first version of the Fundamental Theorem of Asset Pricing for generalized transaction costs matrices. As we will see, it follows rather simple from the original Theorem 5.1.1. For the sake of completeness we re-define arbitrage for general C.

Definition 6.1.12 *Let C be a transaction costs matrix. A C-admissible strategy is called C-*arbitrage *if it satisfies $V_{C,S}(H) \geq 0$ almost surely and*

$$\mathbf{P}\big(V_{C,S}(H) > 0\big) > 0.$$

Theorem 6.1.13 (FTAP for strictly positive C) *Assume S is continuous. Consider a transaction costs matrix C satisfying $c_{ij} > 0$ for all $i \neq j$. Then there is no ϵC-arbitrage for any $\epsilon > 0$ if and only if S admits ϵC-consistent price systems, for all $\epsilon > 0$.*

Proof: Since C is strictly positive off the diagonal, both $\delta_1 = \frac{1}{2}\min\big(c_{ij}, i \neq j\big)$ and $\delta_2 = \max\big(c_{ij}, i \neq j\big)$ are positive. By construction we have

$$C_{\delta_1} \leq C \leq C_{\delta_2}$$

component wise, where C_{δ_1} and C_{δ_2} are defined as in Lemma 6.1.3. This implies

$$\mathcal{K}^S_{C_{\delta_2},t} \subseteq \mathcal{K}^S_{C,t} \subseteq \mathcal{K}^S_{C_{\delta_1},t}$$

for all $t \in [0, T]$, and hence

$$V_{C_{\delta_2},S}(H) \leq V_{C,S}(H) \leq V_{C_{\delta_1},S}(H)$$

a. s. for all S-strategies H. Hence S being ϵC-arbitrage free for all $\epsilon > 0$ is equivalent to NA(0+).
On the other hand, we have for every pair $M, \mathbf{Q}$:

$$(M, \mathbf{Q}) \sim_{C_{\delta_1}} (S, \mathbf{P}) \Rightarrow (M, \mathbf{Q}) \sim_C (S, \mathbf{P}) \Rightarrow (M, \mathbf{Q}) \sim_{C_{\delta_2}} (S, \mathbf{P}).$$

This means that S admitting ϵC-CPS for all $\epsilon > 0$ is equivalent to S admitting ϵ-CPS for all ϵ. Now Theorem 5.1.1 completes the proof. $\square$

6.2 Efficient friction and one-sided transaction costs

There is one particularly interesting situation that Theorem 6.1.13 does not cover, and that is one-sided transaction costs, i. e. the case $c_{ij} = 0$ and $c_{ji} > 0$. In the one-dimensional case it was argued in [GRS08b, Remark 3.19] that the Fundamental Theorem still holds true for this case. The approach suggested there will be generalized to the multidimensional case and made more rigorous in what follows. For this the concept of efficient friction will turn out helpful, which was introduced by Kabanov et al. [KRS02] and has since then proven to be crucial in the theory of markets with transaction costs. Revisiting the analogy of transaction costs and physical friction once again, it means that friction must be present in every direction at all times, or otherwise the particle would be able to move without friction if choosing the right path. In this case we would have a frictionless model (possibly of smaller dimension) embedded in our model, which would obviously undermine our whole framework. This somewhat blurred motivation is made rigorous with the following definition and results.

Definition 6.2.1 *A transaction costs matrix* $C \in [0,1)^{(d+1)\times(d+1)}$ *is said to satisfy* efficient friction (EF)*, if for each subset* $\{n_1, n_2, \dots, n_k\} \subseteq \{0, 1, 2, \dots, d\}$ *containing at least two elements we have*

$$c_{n_1 n_2} + c_{n_2 n_3} + \ldots + c_{n_{k-1} n_k} + c_{n_k n_1} > 0.$$

Remark 6.2.2

1. The theory of transport problems on finite graphs, such as shortest path problems, offers another analogy to efficient friction. If C is understood as distance matrix of a weighted directed graph (see e. g. [Bu65]), then efficient friction is equivalent to the absence of circles of zero length. A circle of length zero would correspond to a subgraph on which one can travel without costs, i. e. frictionless.

2. Note that if C satisfies the triangular inequality $c_{ij} \leq c_{ik} + c_{kj}$ mentioned in Remark 6.1.4, then failing (EF) implies that there exist $0 \leq i < j \leq d$ such that $c_{ij} = c_{ji} = 0$. So these two assets taken for themselves form a frictionless model.

In the transaction costs literature in the tradition of Kabanov's cone setting (e. g. [CS06]) a different definition of EF is used. There a model is said to satisfy (EF) if the interior of $K^*_{C,t}$ is a. s. nonempty for all t. As the following lemma will show, the two definitions are equivalent.

Proposition 6.2.3 *Consider a transaction costs matrix C. the following are equivalent:*

1. *C satisfies (EF)*
2. *The set $\mathcal{B}(C, \mathbf{1})$ has a nonempty interior.*
3. *the interior of $\mathcal{K}^*_{C,t}$ is a. s. nonempty for all $t \in [0, T]$*

Proof: To show ($1 \Rightarrow 2$), we need to find a strict solution M_t to all inequalities from Definition 2.1.3 for $S_t = \mathbf{1}$. Now $\mathbf{x} = \mathbf{1}$ solves all inequalities, but not necessarily strictly. Denote by I the set of all pairs $(i, j) \in \{0, 1, \ldots, d\}^2$ where $i \neq j$ and $c_{ij} = 0$. If $I = \emptyset$, then $\mathbf{1}$ is a strict solution, so assume this is not the case.
Define the sets $A_n, n \geq 0$ recursively by

$$A_0 = \{j : (i, j) \notin I \text{ for all } 0 \leq i \leq d\}$$

$$A_n = \{j : j \notin \bigcup_{k=0}^{n-1} A_k \text{ and } (i, j) \notin I \text{ for all } i \notin \bigcup_{k=0}^{n-1} A_k\}.$$

If (EF) is satisfied, then this gives us a finite partition of $\{0, 1, \ldots, d\}$ into finitely many nonempty sets $A_0, A_1, \ldots A_N$. Now a strict solution to all inequalities in I is given by

$$\mathbf{x} = \mathbf{1} - \sum_{n=1}^{N} \epsilon \frac{2^n - 1}{2^{n-1}} \sum_{k=0}^{d} \mathbf{e}_k \mathbb{1}_{k \in A_n},$$

Where $\epsilon > 0$ is chosen arbitrarily. For ϵ small enough, the inequalities not corresponding to pairs in I are still met strictly and $\mathbf{x} \in \text{int}\mathcal{B}(C, \mathbf{1})$. If on the other hand (EF) is not satisfied, there is a subset $\{n_1, n_2, \ldots, n_k\} \subseteq \{0, 1, 2, \ldots, d\}$ such that

$$c_{n_1 n_2} + c_{n_2 n_3} + \ldots + c_{n_{k-1} n_k} + c_{n_k n_1} = 0.$$

The inequalities corresponding to these indices read

$$x^{(n_1)} \leq x^{(n_2)} \leq \ldots \leq x^{(n_k)} \leq x^{(n_1)},$$

which can obviously not be met strictly. This proves ($1 \Leftrightarrow 2$).
To show ($2 \Leftrightarrow 3$), note that $\mathbf{x} S_t \in \text{int}\mathcal{B}(C, S_t)$ for all $t \in [0, T]$. Now Lemma 6.1.11 tells us that

$$\mathbf{x} S_t \in \text{int}\mathcal{B}(C, S_t) \Leftrightarrow (1, x^{(1)} S_t^{(1)}, x^{(2)} S_t^{(2)}, \ldots x^{(d)} S_t^{(d)}) \in \text{int}\mathcal{K}^{S*}_{C,t}.$$

This proves ($2 \Leftrightarrow 3$). □

Remark 6.2.4 The set I offers another analogy to graph theory. Consider a network consisting of the nodes $0, 1, 2, \ldots, d$ such that flow from node j to node i is possible if and only if $(i, j) \in I$. Here (EF) means that this network does not allow any circles, therefore (see e. g. [Bu65]) it is possible to decompose the set of nodes into different layers $A_0, A_1, \ldots, A_N$ such that flows are only possible from $i \in A_n$ to $j \in A_m$ if $n < m$. This decomposition was demonstrated in the proof.

Theorem 6.2.5 (FTAP for efficient friction matrices) *Assume S is continuous. Consider a transaction costs matrix C satisfying (EF). Then there is no ϵC-arbitrage for any $\epsilon > 0$ if and only if S admits ϵC-consistent price systems, for all $\epsilon > 0$.*

Proof: Proposition 6.2.3 tells us that there exists $\mathbf{x} \in \mathrm{int}\mathcal{B}(C, \mathbf{1})$. Note that for C_δ as in Lemma 6.1.3 we have

$$\mathcal{B}(C_\delta, \mathbf{x}) = \mathbf{x}[1-\delta, 1+\delta]^d.$$

Since $\mathcal{B}(C, \mathbf{1})$ is bounded, there exist $\delta_1, \delta_2 > 0$ such that

$$\mathcal{B}(C_{\delta_1}, \mathbf{x}) \subseteq \mathcal{B}(C, \mathbf{1}) \subseteq \mathcal{B}(C_{\delta_2}, \mathbf{x}),$$

and multiplying by S_t yields

$$\mathcal{B}(C_{\delta_1}, \mathbf{x}S_t) \subseteq \mathcal{B}(C, S_t) \subseteq \mathcal{B}(C_{\delta_2}, \mathbf{x}S_t).$$

this implies on the one hand for every pair $(M, \mathbf{Q})$

$$(M, \mathbf{Q}) \sim_{C_{\delta_1}} (\mathbf{x}S, \mathbf{P}) \Rightarrow (M, \mathbf{Q}) \sim_C (S, \mathbf{P}) \Rightarrow (M, \mathbf{Q}) \sim_{C_{\delta_2}} (\mathbf{x}S, \mathbf{P}),$$

so S admits ϵC-CPS for all $\epsilon > 0$ if and only if $\mathbf{x}S$ admits ϵ-CPS for all ϵ. On the other hand it implies

$$\mathcal{K}^{\mathbf{x}S*}_{C_{\delta_1},t} \subseteq \mathcal{K}^{S*}_{C,t} \subseteq \mathcal{K}^{\mathbf{x}S*}_{C_{\delta_2},t} \Rightarrow$$

$$\mathcal{K}^{\mathbf{x}S}_{C_{\delta_1},t} \supseteq \mathcal{K}^{S}_{C,t} \supseteq \mathcal{K}^{\mathbf{x}S}_{C_{\delta_2},t}.$$

So by Proposition 6.1.8 we know that every C_{δ_2}-arbitrage on $\mathbf{x}S$ implies C-arbitrage on S, which itsef implies C_{δ_1}-arbitrage on $\mathbf{x}S$. This means that S does not admit ϵC-arbitrage for any $\epsilon > 0$ if and only if $\mathbf{x}S$ satisfies NA(0+).
As in the proof of Theorem 6.1.13, applying Theorem 5.1.1, this time to $\mathbf{x}S$, finishes the proof. $\square$

Remark 6.2.6 (Open problem) The obvious next step to generalizing transaction costs would be to consider random costs, i. e. a a strictly positive process $(\epsilon_t)_{t\in[0,T]}$ or an $\mathbb{R}^{(d+1)\times(d+1)}$-valued process $(C_t)_{t\in[0,T]}$ such that C_t satisfies (EF) almost surely for all t. If we assume ϵ to be bounded away from zero, i. e.

$$\epsilon_t \geq \epsilon_{\min} > 0$$

almost surely for all $t \in [0,T]$ for some constant, $\epsilon_{\min}$, one could apply the same argument as in Theorem 6.1.13 and 6.2.5, arguing that ϵ-arbitrage implies $\epsilon_{\min}$-arbitrage and any $\epsilon_{\min}$-CPS is an ϵ-CPS and then applying Theorem 5.1.1 once again. If on the other hand ϵ can become arbitrarily small, then we are not only unable to use this argument, but also unable to construct an ϵ-CPS in an analogous manner to Chapter 5. There we heavily exploited that ϵ is known from the very beginning, especially in the crucial Propositions 5.3.1 and 5.4.4, where we recursively constructed processes M^n increasingly distant from S. If now the bid-ask region unexpectedly shrinks by ϵ_t getting smaller, this procedure can not be continued. This problem cannot be avoided by localizing ϵ_t.
Regarding the case of a random matrix C_t, we would simply have to make the vector $\mathbf{x} \in \text{int}\mathcal{B}(\mathbf{C}, \mathbf{1})$ from Proposition 6.2.3 random, since $\mathcal{B}(C_t, \mathbf{1})$ is now a random set as well. Depending on what kind of processes are considered for C, this seems much less of a problem than the aforementioned one.
So at this point, the Fundamental Theorem for random ϵ, let alone for random C, remains open, at least in our framework.

Appendix A

On the Esscher transform

The Esscher transform, first introduced by Esscher in 1932 [E32], defines a parametrized family of measures equivalent to a given original measure. It has found various applications mainly in actuarial science, but has also found its way into mathematical finance, e. g. [GS94]. The simplicity of the Esscher transform makes it tractable enough for our needs, but on the other hand the resulting family of measures is sufficiently rich, meaning that it allows us to transform the expectation of a given bounded random variable to any possible point.

Definition A.1 *Consider an $\mathbb{R}^d$-valued random variable X on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$. For arbitrary $\mathbf{a} \in \mathbb{R}^d$, the Esscher transform of $\mathbf{P}$ given $\mathbf{a}$ and X is a probability measure $\mathbf{P}_{\mathbf{a},X}$ on $\mathcal{F}$, defined by its Radon-Nikodym derivative*

$$\frac{d\mathbf{P}_{\mathbf{a},X}}{d\mathbf{P}} = \frac{e^{\mathbf{a}^\top X}}{\mathbb{E}_{\mathbf{P}}(e^{\mathbf{a}^\top X})}.$$

$\mathbf{P}_{\mathbf{a},X}$ is well-defined and equivalent to $\mathbf{P}$ if $\mathbb{E}_{\mathbf{P}}(e^{\mathbf{a}^\top X})$ is finite. In what follows, we will always assume X to be bounded, hence all Esscher transforms exist.

We want to investigate the range of possible expectations of X with respect to Esscher transformed measures. We therefore define the function $\phi_X : \mathbb{R}^d \to \mathbb{R}^d$ by

$$\phi_X(\mathbf{a}) = \mathbb{E}_{\mathbf{P}_{\mathbf{a},X}}(X) = \frac{E_{\mathbf{P}}(Xe^{\mathbf{a}^\top X})}{E_{\mathbf{P}}(e^{\mathbf{a}^\top X})}.$$

Recall the operator $\mathcal{I}$ from Definition 5.1.2. Since $\mathbf{P}_{\mathbf{a},X} \sim \mathbf{P}$ we have

$$\phi_X(\mathbb{R}^d) = \{\mathbb{E}_{\mathbf{P}_{\mathbf{a},X}}(X), \mathbf{a} \in \mathbb{R}^d\} \subseteq \{\mathbb{E}_{\mathbf{Q}}(X), \mathbf{Q} \sim \mathbf{P}\} = \mathcal{I}(X).$$

In this section, we will show that this inclusion is in fact satisfied with equality.

Lemma A.2 *Let X be bounded.*

(a) The function ϕ_X is continuously differentiable in $\mathbf{a}$ and its Jacobi matrix $\nabla\phi_X(\mathbf{a})$ with respect to $\mathbf{a}$ equals the covariance matrix of X with respect to $\mathbf{P}_{\mathbf{a},X}$.

(b) If X has full dimension (i. e. $\operatorname{supp}(X)$ is not contained in any $(d-1)$-dimensional affine subspace of $\mathbb{R}^d$), then $\nabla\phi_X(\mathbf{a})$ is positive definite for all $\mathbf{a} \in \mathbb{R}^d$.

Proof: Since X is bounded and the moment generating function

$$\mathbf{a} \mapsto M_X(\mathbf{a}) = \mathbb{E}_{\mathbf{P}}(e^{\mathbf{a}^\top X})$$

is analytic, Fubini's theorem yields for its first and second derivative:

$$\nabla M_X(\mathbf{a}) = \mathbb{E}_{\mathbf{P}}(\nabla e^{\mathbf{a}^\top X}) = \mathbb{E}_{\mathbf{P}}(X e^{\mathbf{a}^\top X})$$
$$\nabla^2 M_X(\mathbf{a}) = \mathbb{E}_{\mathbf{P}}\big(\nabla(X e^{\mathbf{a}^\top X})\big) = \mathbb{E}_{\mathbf{P}}(XX^\top e^{\mathbf{a}^\top X}).$$

Note that $\phi_X = \frac{\nabla M_X}{M_X}$, hence ϕ is continuously differentiable with

$$\nabla\phi_X(\mathbf{a}) = \frac{\mathbb{E}_{\mathbf{P}}(XX^\top e^{\mathbf{a}^\top X})M_X(\mathbf{a}) - \mathbb{E}_{\mathbf{P}}(X e^{\mathbf{a}^\top X})\mathbb{E}_{\mathbf{P}}(X e^{\mathbf{a}^\top X})^\top}{M_X^2(\mathbf{a})},$$

which simplifies to

$$\nabla\phi_X(\mathbf{a}) = \mathbb{E}_{\mathbf{P}_{\mathbf{a},X}}(XX^\top) - \mathbb{E}_{\mathbf{P}_{\mathbf{a},X}}(X)\mathbb{E}_{\mathbf{P}_{\mathbf{a},X}}(X)^\top = \mathrm{Var}_{\mathbf{P}_{\mathbf{a},X}}(X).$$

This proofs (a). To see the second part of the lemma, it suffices to notice that there is no vector $\mathbf{v} \in \mathbb{R}^d \backslash \{\mathbf{0}\}$ such that $\mathbf{v}^\top X$ is $\mathbf{P}$-almost surely constant. Hence for all $\mathbf{v} \neq \mathbf{0}$:

$$\mathbf{v}^\top \nabla\phi_X(\mathbf{a})\mathbf{v} = \mathrm{Var}_{\mathbf{P}_{\mathbf{a},X}}(\mathbf{v}^\top X) > 0.$$

□

Remark A.3 Note that since X is bounded, say, $\|X\| \leq c$ a. s., the Cauchy-Schwarz inequality yields $|\mathbf{v}^\top X| \leq c$ a. s. for all $\mathbf{v} \in \mathbb{R}^d$ with $\|\mathbf{v}\| = 1$. Thus

$$\mathbf{v}^\top\big(\nabla\phi_X(\mathbf{a})\big)\mathbf{v} = \mathrm{Var}_{\mathbf{P}_{\mathbf{a},X}}(\mathbf{v}^\top X) \leq c^2$$

for all such $\mathbf{v}$. In other words, the largest eigenvalue of $\nabla\phi_X(\mathbf{a})$ possesses a global upper bound independent of $\mathbf{a}$, therefore ϕ_X is globally Lipschitz continuous with Lipschitz constant c^2.

Lemma A.4 *Let X be bounded and of full dimension. Let a sequence $(\mathbf{a}_n)_{n\in\mathbb{N}}$ in $\mathbb{R}^d$ be given such that $\{\|\mathbf{a}_n\|, n \in \mathbb{N}\}$ is unbounded. Then there exists a subsequence $(\mathbf{a}_{n_k})_{k\in\mathbb{N}}$ of $(\mathbf{a}_n)$ such that the distance of $\phi_X(\mathbf{a}_{n_k})$ and the boundary of $\mathcal{I}(X)$ converges to zero as $k \to \infty$.*

Proof: For $n \in \mathbb{N}$, define $\lambda_n = \|\mathbf{a}_n\|$ and $\mathbf{v}_n = \frac{\mathbf{a}_n}{\lambda_n}$. Without loss of generality we assume $\lambda_n \geq n$ for all n, otherwise pass to an adequate subsequence. Since $\mathbf{v}_n$ lives on the compact unit sphere of $\mathbb{R}^d$, there exists a subsequence $(\mathbf{v}_{m_k})$ of $(\mathbf{v}_n)$ converging to some $\mathbf{v} \in \mathbb{R}^d$ with $\|\mathbf{v}\| = 1$. Consider the constant

$$\alpha = \mathrm{esssup}_{\mathbf{P}}(\mathbf{v}^\top X).$$

Since X is bounded, α is finite.
For $\epsilon > 0$, consider some $\mathbf{x} \in \mathbb{R}^d$ with $\mathbf{v}^\top \mathbf{x} < \alpha - \epsilon$. Then we have

$$\mathbf{P}\left(\mathbf{v}^\top (X - \mathbf{x}) > \frac{\epsilon}{2}\right) > 0,$$

hence $\mathbf{v}_{m_k}^\top (X - \mathbf{x}) > \frac{\epsilon}{4}$ for all k large enough on a set $A \subset \Omega$ of positive probability. Since $\lambda_{m_k} \to \infty$ as $k \to \infty$, we have

$$\lim_{k\to\infty} \mathbf{a}_{m_k}^\top (X - \mathbf{x}) = \infty$$

on A. This means that

$$\frac{d\mathbf{P}_{\mathbf{a}_{m_k},X}}{d\mathbf{P}}\Big|_{\{X=\mathbf{x}\}} = \frac{\exp(\mathbf{a}_{m_k}^\top \mathbf{x})}{\mathbb{E}_{\mathbf{P}}\big(\exp(\mathbf{a}_{m_k}^\top X)\big)} = \mathbb{E}_{\mathbf{P}}\Big(\exp\big(\mathbf{a}_{m_k}^\top (X - \mathbf{x})\big)\Big)^{-1}$$

tends to 0 for all $\mathbf{x}$ with $\mathbf{v}^\top \mathbf{x} < \alpha - \epsilon$. This means that

$$\lim_{k\to\infty} \mathbf{P}_{\mathbf{a}_{m_k},X}(\mathbf{v}^\top X \leq \alpha - \epsilon) = 0,$$

hence for all k large enough we find

$$\mathbf{v}^\top \phi_X(\mathbf{a}_{m_k}) = \mathbb{E}_{\mathbf{P}_{\mathbf{a}_{m_k}}}(\mathbf{v}^\top X) > \alpha - 2\epsilon.$$

Since ϵ was chosen arbitrarily, we can construct a subsequence $(\mathbf{a}_{n_k})$ satisfying

$$\lim_{k\to\infty} \mathbf{v}^\top \phi_X(\mathbf{a}_{n_k}) = \alpha,$$

which means that $\phi_X(\mathbf{a}_{n_k})$ approaches the boundary of $\mathcal{I}(X)$.

$\square$

Theorem A.5 *Let X be bounded. Then the image of ϕ_X is $\mathcal{I}(X)$.*

Proof: Assume first that X has full dimension. We already know that $\phi_X(\mathbf{a}) \in \mathcal{I}(X)$ for all $\mathbf{a} \in \mathbb{R}^d$ and from Lemma A.2 that $\phi(X)$ is injective. Let $\mathbf{v} \in \mathbb{R}^d$ with $\|\mathbf{v}\| = 1$ be given. Consider the d-dimensional ordinary differential equation

$$\nabla \mathbf{a}(t) = F\big(\mathbf{a}(t)\big)\mathbf{v}$$

with the $\mathbb{R}^{d\times d}$-valued function $F(\mathbf{a}) = \big(\nabla\phi_X(\mathbf{a})\big)^{-1}$ and the initial condition $\mathbf{a}(0) = 0$. From Lemma A.2 we know that F is well-defined and everywhere continuous, thus by Peano's existence theorem [He05, Theorem 11.1], there exists some solution $\mathbf{a}$ up to some maximal time point $T \in (0, \infty]$. Now the curve $t \mapsto \phi_X\big(\mathbf{a}(t)\big)$ satisfies

$$\frac{d}{dt}\phi_X\big(\mathbf{a}(t)\big) = \nabla\phi_X\big(\mathbf{a}(t)\big)\nabla\mathbf{a}(t) = \mathbf{v},$$

which means that its image equals

$$\{\phi_X(0) + t\mathbf{v}, 0 \le t < T\},$$

thus T is finite. This means that $\mathbf{a}$ must explode at T, so $\mathbf{a}$ is unbounded. With Lemma A.4 it follows that there exists an increasing sequence $(t_n)_{n\in\mathbb{N}}$ in $(0, T)$ such that $\phi_X\big(\mathbf{a}(t_n)\big)$ approaches the boundary of $\mathcal{I}(X)$. This means that the image of $t \mapsto \phi_X\big(\mathbf{a}(t)\big)$ must be

$$\{\phi_X(0) + t\mathbf{v};\ t \ge 0\} \cap \mathcal{I}(X).$$

Because $\mathbf{v}$ was chosen arbitrarily on the unit sphere and $\mathcal{I}(X)$ is convex, all of $\mathcal{I}(X)$ is contained in the image of $\phi(X)$, which completes the proof for the full dimensional case.
Assume now that supp(X) is contained in an n-dimensional affine subspace of $\mathbb{R}^d$, where n is assumed to be minimal, i. e. supp(X) is not contained in any affine subspace of smaller dimension. Then by basic linear algebra X admits the representation

$$X = \mathbf{v} + AZ$$

with a $\mathbb{R}^n$-valued random variable Z of full dimension, a full-ranked matrix $A \in \mathbb{R}^{n\times d}$ with $A^\top A$ being the n-dimensional unit matrix and a vector $\mathbf{v} \in \mathbb{R}^d$ with $A^\top\mathbf{v} = 0$. If $\mathbf{x} \in \mathcal{I}(X)$, then $\mathbf{x} = \mathbf{v} + A\mathbf{z}$ with $\mathbf{z} \in \mathcal{I}(Z)$, and since Z has full dimension, $\mathbf{z} = \phi_Z(\mathbf{b})$ for some $\mathbf{b} \in \mathbb{R}^n$. If we define $\mathbf{a} = \mathbf{v} + A\mathbf{b}$, we end up with

$$\phi_X(\mathbf{a}) = \mathbf{v} + A\frac{\mathbb{E}(Ze^{\mathbf{b}^\top Z})}{\mathbb{E}(e^{\mathbf{b}^\top Z})} = \mathbf{x}$$

$\square$

Remark A.6 We have shown that the derivative of ϕ_X is everywhere positive definite if and only if X has full dimension. In this case, ϕ_X is a bijection between $\mathbb{R}^d$ and $\mathcal{I}(X)$ and for all $\mathbf{y} \in \mathcal{I}(X)$ the equation

$$\mathbb{E}_{\mathbf{P}(\mathbf{a})}(X) = \mathbf{y}$$

has a unique solution $\mathbf{a} \in \mathbb{R}^d$. If X lives on an affine subspace $\mathbf{v} + V$ with a vector space $V \subset \mathbb{R}^d$ of minimal dimension, then ϕ_X is obviously constant along any line orthogonal to V. From the projection argument used in Theorem A.5 follows that ϕ_X is a bijection between V and $\mathcal{I}(X)$. In any case the above equation admits a unique solution $\mathbf{a}$ of minimal norm $\|\mathbf{a}\|$. This allows us, via an argument of *measurable selection*, to derive the following corollary:

Corollary A.7 *Let two $\mathbb{R}^d$-valued bounded random variables X and Y be given on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, such that Y is measurable w.r.t a sub-σ-algebra $\mathcal{G} \subseteq \mathcal{F}$. If the conditional support of X given $\mathcal{G}$ satisfies*

$$Y \in \mathcal{I}(X|\mathcal{G})$$

almost surely, then there exists a probability measure $\mathbf{Q} \sim \mathbf{P}$ identical to $\mathbf{P}$ on $\mathcal{G}$, such that, almost surely,

$$\mathbb{E}_{\mathbf{Q}}(X|\mathcal{G}) = Y.$$

Proof: Conditional on $\mathcal{G}$, we may treat Y and the set $\mathcal{I}(X|\mathcal{G})$ as constant. For each pair of Y and $\mathcal{I}(X|\mathcal{G})$, choose the norm-minimal Esscher parameter $A(\omega) = \mathbf{a}$ satisfying $\mathbb{E}_{\mathbf{P}(\mathbf{a})}(X)(\omega) = Y(\omega)$. This defines a $\mathcal{G}$-measurable random variable A. Since X is bounded, say, $\|X\| \leq c$ a.s., we get by applying the ordinary Cauchy-Schwarz inequality and monotone convergence

$$\mathbb{E}_{\mathbf{P}}(e^{A^\top X}|\mathcal{G}) \leq \mathbb{E}_{\mathbf{P}}(e^{\|A\|\cdot\|X\|}|\mathcal{G})$$

$$= \sum_{k=0}^{\infty} \frac{1}{k!}\mathbb{E}_{\mathbf{P}}(\|A\|^k\|X\|^k|\mathcal{G}) \leq \sum_{k=0}^{\infty} \frac{\|A\|^k c^k}{k!} = e^{\|A\|c} < \infty$$

and we may therefore define a measure $\mathbf{Q} \sim \mathbf{P}$ by

$$\frac{d\mathbf{Q}}{d\mathbf{P}} = \frac{\exp(A^{\top}Y)}{\mathbb{E}_{\mathbf{P}}(e^{A^{\top}Y}|\mathcal{G})},$$

which has all the desired properties. □

Remark A.8 A result equivalent to A.5 in the notion of exponential families was derived by Barndorff-Nielsen [B-N78, Theorem 9.2], for the full-dimensional case and without Remark A.6, which plays a crucial role in our corollary. There it was shown that a necessary and sufficient condition for A.5 to hold true is the so-called *steepness* of the family $\{\mathbf{P}_{\mathbf{a},X}, \mathbf{a} \in \mathbb{R}^d\}$. The verification of this property is more or less equivalent to the assertion of Lemma A.4.

Bibliography

[Bac00] Bachelier, L., (1900) *Théorie de la spéculation*. Annales Scientifiques de l'École Normale Supérieure, 3, 17, pp. 21-86.

[Bau52] Baumol, W. J., (1952) *The transactions demand for cash: an inventory theoretic approach*. Quart. J. Econ., 66, pp. 545–556.

[Be00] Behrends, E., (2000) *Introduction to Markov Chains*. Vieweg.

[BLPS92] Bensaid, B., Lesne, J., Pages, H., Scheinkman, J., (1992) *Derivative asset pricing with transaction costs*. Mathematical Finance, 2, pp. 63–86.

[B-N77] Barndorff-Nielsen, O.-E., (1977) *Exponentially decreasing distributions for the logarithm of particle size*. Proceedings of the Royal Society of London, A, 353, pp. 401–419.

[B-N78] Barndorff-Nielsen, O.-E., (1978) *Information and exponential families in statistical theory*. Wiley, New York.

[B-N98] Barndorff-Nielsen, O.-E., (1998) *Processes of normal inverse gaussian type*. Finance and Stochastics, 2, pp. 42–68.

[BS73] Black, F., Scholes, M., (1973) *The pricing of options and corporate liabilities*. Journal of Political Economy, 81, pp. 637–659.

[BT00] Bouchard, B., Touzi, N., (2000) *Explicit solution to the multivariate super-replication problem under transaction costs*. Ann. Appl. Probab., 10, 3, 685–708.

[Bu65] Busacker, R. G., Saaty, T. L., (1965) *Finite Graphs and Networks*. McGraw-Hill.

[BV92] Boyle, P. P., Vorst, T., (1992) *Option replication in discrete time with transaction costs*. J. Finance, 47, pp. 272–293.

[CGMY02] P. Carr, P., Geman, H., Madan, D. H., Yor, M., (2002) *The fine structure of asset returns: an empirical investigation*. Journal of Business 75, pp. 305–332

[Ch01] Cherny, A., (2001) *No-arbitrage and completeness for the linear and exponential models based on Lévy processes*. Research Report 2001-33, MaPhySto, University of Aarhus.

[CK96] Cvitanić, J., Karatzas, (1996) *Hedging and Portfolio Optimization under Transaction Costs: a Martingale Approach*. Mathematical Finance, 6, 2, pp. 133–165.

[Co37] Coase, R. H., (1937) *The Nature of the Firm*. Economica, 4, 16, pp. 386-405.

[Co60] Coase, R. H., (1960) *The Problem of Social Cost*. Journal of Law and Economics, 3, pp. 1-44.

[CPT99] Cvitanić, J., Pham, H. & Touzi, N., (1999) *A closed-form solution to the problem of super-replication under transaction costs*. Finance and Stochastics, 3, 1, pp. 35–54.

[CS06] Campi, L., Schachermayer, W., (2006) *A Super-Replication Theorem in Kabanov's Model of Transaction Costs*. Finance and Stochastics, 10, pp. 579–596.

[CSS95] Cvitanić, J., Shreve, S. E., Soner, H. M., (1995) *There is no nontrivial hedging portfolio for option pricing with transaction costs*. Ann. Appl. Probab., 5, pp. 327–355.

[DC94] Clark, J. M. C., Davis, M. H. A., (1994) *A note on super replicating strategies*. Philos. Trans. Royal Soc., London, A, 347, pp. 485–494.

[DN90] Davis, M. H. A., Norman, A., (1990) *Portfolio Selection with Transaction Costs*. Math. Oper. Res., 15, pp. 676–713.

[DS06] Delbaen, F., Schachermayer, W., (2006), *The Mathematics of Arbitrage*. Springer Finance Series, Springer Verlag.

[E32] Esscher, F., (1932) *On the probability function in the collective theory of risk*. Skandinavisk Aktuarietidskrift 15, pp. 175-195.

[EJ97] Eberlein, E., Jacod, J., (1997) *On the range of option prices*. Finance and Stocastics, 1, pp. 131–140.

[ENU93] Edirisinghe, C., Naik, V., Uppal, R., (1993) *Optimal replication of options with transaction costs and trading restrictions*. Journal of Financial and Quantitative Analysis, 28, pp. 117–138.

[F65] Fama, H., (1965) *The behavior of stock market prices*. Journal of Business, 38, pp. 34–105.

[GRS08a] Guasoni, P., Rásonyi, M., Schachermayer, W., (2008) *Consistent Price Systems and Face-Lifting Pricing under Transaction Costs*. Ann. Appl. Prob., 18, 2, pp. 491–520.

[GRS08b] Guasoni, P., Rásonyi, M., Schachermayer, W., (2008) *The Fundamental Theorem of Asset Pricing for Continuous Processes under Small Transaction Costs*. Online first in Annals of Finance.

[GS94] Gerber, H. U., Shiu, E. S. W., (1994) *Option Pricing by Esscher Transforms*. Transactions of Society of Actuaries, 46, pp. 99–190.

[Gu06] Guasoni, P., (2006) *No arbitrage with transaction costs, with fractional brownian motion and beyond*. Mathematical Finance, 16, 3, pp. 569–582.

[He05] Heuser, H., (2005) *Gewöhnliche Differentialgleichungen. Einführung in Lehre und Gebrauch*. Vieweg + Teubner.

[HK79] Harrison, J. M., Kreps, D. M., (1979) *Martingales and arbitrage in multiperiod securities markets*. J. Econom. Theory, 20, 3, pp. 381–408.

[HP81] Harrison, J. M., Pliska, S. R., (1981) *Martingales and stochastic integrals in the theory of continuous trading*. Stochastic Process. and Appl., 11, pp. 215–260.

[JK95] Jouini, E., Kallal, H., (1995) *Martingales and arbitrage in securities markets with transaction costs*. J. Econom. Theory, 66, pp. 178–197.

[JS87] Jacod, J., Shiryaev, A. N., (1987) *Limit Theorems for Stochastic Processes*. Springer-Verlag, Grundlehren der mathematischen Wissenschaften, 288.

[Ka99] Kabanov, Y., (1999) *Hedging and liquidation under transaction costs in currency markets*. Finance and Stochastics, 3, pp. 237–248.

[Ke36] Keynes, J. M., (1936) *The general theory of employment, interest, and money*. Harcourt Brace, New York.

[KL02] Kabanov, Y., Last, G., (2002) *Hedging under transaction costs in currency markets: a continuous time model*. Mathematical Finance, 12, 1, pp. 63–70.

[Ko02] Kou, S. G., (2002) *A jump diffusion model for option pricing*. Management Science, 48, pp. 1086–1101.

[Kr81] Kreps, D. M., (1981) *Arbitrage and equilibrium in economies with infinitely many commodities*. J. Math. Econom., 8, pp. 15–35.

[KRS02] Kabanov, Y., Rásonyi, M., Stricker, C., (2002) *No-arbitrage criteria for financial markets with efficient friction.* Finance and Stochastics, 6, 3, pp. 371–382.

[KS02] Kabanov, Y., Stricker, C., (2002) *Hedging of contingent claims under transaction costs*. In: Advances in Finance and Stochastics. Essays in honour of Dieter Sondermann. Eds. K. Sandmann and Ph. Schönbucher, Springer, pp. 125–136.

[Le85] Leland, H. E., (1985) *Option pricing and replication with transaction costs*. J. Finance, 40, pp. 1283–1301.

[LS97] Leventhal, S., Skorohod, A. V., (1997) *On the Possibility of Hedging Options in the Presence of Transaction Costs*. Annals of Applied Probability, 7, 2, pp. 410–443.

[Ma63] Mandelbrot, B., (1963) *The variation of certain speculative prices*. The Journal of Business 36, pp. 394–419.

[Ma68] Mandelbrot, B., (1968) *Fractional Brownain motion, fractional noises and applications*. SIAM Review, 10, pp. 422–437.

[MC76] Magill, M. J. P., Constantinides, G. M., (1976) *Portfolio selection with transaction costs*. J. Economic Theory, 13, pp. 264–271.

[Me69] Merton, R. C., (1969) *Lifetime portfolio selection under uncertainty: The continuous-time case*. The Review of Economics and Statistics, 51, pp. 247-257.

[Me73] Merton, R. C., (1973) *Theory of Rational Option Pricing*. Bell Journal of Economics and Management Science, 4, pp. 141–183.

[Me76] Merton, R. C., (1976) *Option pricing when underlying stocks are discontinuous*. Journal of Finance and Economics, 5, pp. 125–144.

[MS90] Madan, D. B., Seneta, E., (1990) *The variance gamma model for share market returns*. Journal of Business, 63, pp. 511–524.

[N98] Frängsmyr, T. (ed.), (1998) *Les Prix Nobel. The Nobel Prizes 1997*. Nobel Foundation, Stockholm.

[OW06] Ouwehand, P., West, G., (2006) *Pricing rainbow options*. WILMOTT Magazine, 3/06, pp. 74–80.

[P04] Protter, P. E., (2004) *Stochastic Integration and Differential Equations*. Applications of Mathematics, 21, Springer Verlag.

[R00] Raible, S., (2000) *Lévy Processes in Finance: Theory, Numerics and Empirical Facts*. Ph. D. thesis, Albert-Ludwigs-Universität Freiburg i. Br.

[Sam65] Samuelson, P. A., (1965) *Proof that properly anticipated prices fluctuate randomly*. Industrial Management Review, 6, pp. 41–50.

[Sat00] Sato, K. I., (2000) *Lévy Processes and Infinitely Divisible Dirstributions*. Cambridge Studies in Advanced Mathematics, 68.

[Scha04] Schachermayer, W., (2004) *The Fundamental Theorem of Asset Pricing under Proportional Transaction Costs in Finite Discrete Time*. Mathematical Finance, 14, 1, pp. 19–48.

[Scho01] Schoutens, W., (2001) *The meixner process in finance*. EURANDOM Report 2001-2002. EURANDOM, Eindhoven.

[Scho03] Schoutens, W., (2003) *Lévy Processes in Finance - Pricing Financial Derivatives.* Wiley Series in Probability and Statistics.

[ST94] Samorodnitsky, G., Taqqu, M., (1994) *Stable Non-Gaussian Random Processes*. Chapman and Hall, New York.

[Tob56] Tobin, J., (1956) *The Interest Elasticity of the Transactions Demand for Cash*. Review of Economics and Statistics, 38, 3, pp. 241–247.

[Tob78] Tobin, J., (1978) *A proposal for international monetary reform.* Eastern Economic Journal, 4, 3-4, pp. 153-159.