Preface

The present lecture notes are based on several advanced courses that I gave at the University of Vienna between 2011 and 2013. In 2015 I gave a similar course ("Nachdiplom-Vorlesung") at ETH Zürich. The purpose of these lectures was to present and organize the recent progress on portfolio optimization under proportional transaction costs $\lambda > 0$. Special emphasis is given to the asymptotic behavior when λ tends to zero.

The theme of portfolio optimization is a classical topic of mathematical finance, going back to the seminal work of Robert Merton in the early seventies (considering the frictionless case without transaction costs). Mathematically speaking, this question leads to a concave optimization problem under linear constraints. A technical challenge arises from the fact that—except for the case of finite probability spaces Ω—the optimization takes place over infinite-dimensional sets.

There are essentially two ways of attacking such an optimization problem.

The *primal* method consists in directly addressing the problem at hand. Following the path initiated by Robert Merton, this leads to a partial differential equation of *Hamilton–Jacobi–Bellman* type. This PDE method can also be successfully extended to the case of proportional transaction costs. Important work in this line was done by Constantinides [10], Dumas and Luciano [25], Taksar, Klass, and Assaf [75], Davis and Norman [18], Shreve and Soner [72], to name just some of the early work on this topic.

An alternative method consists in passing to the *dual version* of the problem. This dual method is also sometimes called the *martingale method* as one now optimizes over the constraint variables, which in the present context turn out to be martingales (or their generalizations such as supermartingales). In the course of the analysis, an important role is played by the *Legendre transform* or *conjugate function* V of the utility function U appearing in the primal version of the problem.

In these notes we focus on the dual method as well as the interplay between the dual and the primal problem. This approach yields to a central concept of our approach, namely the concept of a *shadow price process*. Mathematically speaking, this is an infinite-dimensional generalization of the fundamental concept of a Lagrange multiplier. It has a clear economic interpretation as a price process, which—without transaction costs—yields the same optimal portfolio as the original price process under transaction costs. This gives a direct link between the present optimization theory under transaction costs and the more classical frictionless theory.

This brings us to a major open challenge for mathematical finance that constitutes much of the motivation for the present notes and the underlying research. The question is how to design an economically, as well as mathematically, meaningful framework to deal with financial models that are based on *fractional Brownian motion*.

This variant of the basic concept of Brownian motion was introduced in 1940 by Kolmogorov under the name of the *Wiener spiral*. It was strongly advocated by Mandelbrot more than 50 years ago as a more realistic approach to financial data than models based on classical Brownian motion, such as the Black–Scholes model.

But there are fundamental problems that have made it impossible to reconcile these models with the main stream of mathematical finance, which is based on the paradigmatic assumption of *no arbitrage*. In fact, fractional Brownian motion fails to be a semimartingale. It is well known ([21, Thm. 7.2]) that processes which fail to be semimartingales always allow for arbitrage. Hence it does not make any economic sense to apply *no-arbitrage* arguments, e.g., in the context of option pricing, if the underlying model for the stock price process already violates this paradigm.

One way to get out of this deadlock is the consideration of transaction costs. It was shown in [40] that the consideration of (arbitrarily small) proportional transaction costs $\lambda > 0$ makes the arbitrage possibilities disappear for the presently considered models based on fractional Brownian motion. This allows for a similar duality theory to the frictionless case. While in the frictionless theory the dual objects are the martingale measures and their variants, their role now is taken by the *λ-consistent price systems*. However, many of the classical concepts from the frictionless theory, such as replication and/or superreplication, do not make any economic sense when considering these models under transaction costs. Indeed, one can give rigorous mathematical *proofs* (the "face-lifting theorem" in [39, 56, 73]) that it is not possible to derive *any nontrivial result* from superreplication arguments in the present context of models under transaction costs.

Yet there is still hope to find a proper framework that allows us to obtain nontrivial results for fractional Brownian motion. There is one financial application that does make perfect sense in the presence of transaction costs, from an economic as well as from a mathematical perspective, namely *portfolio optimization*. This is precisely the theme of the present lecture notes and we shall develop this theory quite extensively. But before doing so let us come back to the original motivation. What does portfolio optimization under transaction costs have to do with the original problem of pricing and hedging options in financial models involving fractional Brownian motion? The answer is that we have hope that finally a well-founded theory of portfolio optimization can shed some light on the original problem of pricing derivative securities via utility indifference pricing. The key fact is the existence of a *shadow price process* $\tilde{S}$ that can serve as a link to the traditional frictionless theory. In Theorem 8.4 we shall prove the existence of a shadow price process under general assumptions in the framework of models based on fractional Brownian motion. This theorem was proved only very recently in [13] and is the main and final result of the present lecture notes. In a sense, the lecture notes aim at providing and developing all the material for proving this theorem. At the same time they try to present a comprehensive introduction to the general theme of portfolio optimization under transaction costs. They are structured in the following way.

In the first two chapters we develop the theory of portfolio optimization in the elementary setting of a finite probability space. Under this assumption all relevant spaces are finite-dimensional and therefore the involved functional analysis reduces to linear algebra. These two chapters are analogous to the summer school course [67], as well as the two introductory chapters in [24] where a similar presentation was given for the frictionless case.

In Chapter 3 we focus rather extensively on the most basic example: the Black–Scholes model under logarithmic utility $U(x) = \log(x)$. A classical result by Merton states that, in the frictionless case, the optimal strategy consists of holding a constant fraction (depending in an explicit way on the parameters of the model) of wealth in stock and the rest in bond (the "Merton line"). If we pass to transaction costs, it is also well known that one has to keep the proportion of wealth within a certain interval and much is known on this interval. Our purpose is to *exactly* determine all the quantities of interest, e.g., the width of this corridor and the effect on the indirect utility. The dual method allows us to calculate these quantities either in closed form as functions of the level $\lambda > 0$ of transaction costs, or as a power series of $\lambda^{1/3}$. In the latter case we are able to explicitly compute all the coefficients of the fractional Taylor series. The key concept for this analysis is the notion of a shadow price process. In the case of the Black–Scholes model, this shadow price process can explicitly be determined. This demonstrates the power of the dual method. We have hope that the analysis of the Black–Scholes model can serve as a role model for a similar analysis for other models of financial markets, e.g., stock price processes based on fractional Brownian motion. Here is a wide open field for future research.

In Chapter 4 we go systematically through the duality theory for financial markets under transaction costs. As regards the degree of generality, we do not strive for the maximal one, i.e., the consideration of general càdlàg price process $S = (S_t)_{0 \le t \le T}$ as in [14] and [16]. Rather we confine ourselves to *continuous* processes S and—mainly for convenience—we assume that the underlying filtration is Brownian. We do so as we have fractional Brownian motion as our final application in the back of our mind. On the other hand, for this application it is important that we do *not* assume that the process S is a semimartingale. The central result of Chapter 4 is Theorem 4.22, which establishes a polar relation between two sets of random variables. On the primal side, this is the set of random variables that is attained from initial wealth $x > 0$ by trading in the stock S under transaction costs λ in an admissible way. On the dual side, the set consists of the so-called supermartingale deflators. The polar relation between these two sets is such that the conditions of the general portfolio optimization theorem in [55] are satisfied, which allows us to settle all central issues of portfolio optimization.

Chapter 5 is a kind of side step and develops a local duality theory. It shows that several traditional assumptions in the theory of portfolio optimization can be replaced by their *local versions* without loss of generality with respect to the conclusions. A typical example is the assumption of "no free lunch with vanishing risk" (NFLVR) from [22]. This well-known concept is traditionally assumed to hold true in portfolio optimization problems, e.g., in [55]. It was notably pointed out by Karatzas and Kardaras [51] that this assumption may be replaced by its local version, which is the condition of "no unbounded profit with bounded risk" (NUPBR). We give equivalent formulations of this latter property in the frictionless setting (Theorem 5.6) and an analogous theorem in the setting of (arbitrarily small) transaction costs $\lambda > 0$ (Theorem 5.11).

After all these preparations we turn to the general theme of portfolio optimization under proportional transaction costs in Chapter 6. The basic duality Theorem 6.2 is a consequence of these preparations and the results from [55]. In Theorem 6.5 we show the crucial fact that—under appropriate assumptions—the dual optimizer, which a priori is only a supermartingale, is in fact a local martingale. The crucial assumption underlying this theorem is that the process S satisfies the "two-way crossing property", a notion introduced recently by Bender [5], which generalizes the concept of a continuous martingale.

In Chapter 7 we show that the local martingale property established in Theorem 6.5 is the key to the existence of a shadow price process (Theorem 7.3). This insight goes back to the work of Cvitanic and Karatzas [11].

In Chapter 8 we finally turn to the case of exponential fractional Brownian motion. It culminates in Theorem 8.4 where we prove that for this model there is a shadow price process. The key ingredient is a recent result by Peyre [60], showing that fractional Brownian motion has the "two-way crossing" property.

Many people have participated in the research efforts underlying these lectures. My sincere thanks go to my coauthors over the past 10 years on this topic: L. Campi, S. Gerhold, P. Guasoni, J. Muhle-Karbe, R. Peyre, M. Rásonyi, and J. Yang. Special thanks go to Ch. Czichowsky. Without his endurance the six joint papers [12, 14, 15, 16, 17, 13], which are the basis of the present research, would never have seen the light of day. I also thank the participants at my lectures at the University of Vienna and ETH Zürich.