Duality methods for stochastic optimal control problems in finance

DOI:
10.26481/dis.20230112tk

Summary

In this addendum, we provide a summary of this dissertation. We first address the main overarching theme. Thereafter, we elaborate in a recapitulating manner on the separate chapters central to the research output of this dissertation. Consistent with its title, the theme of this doctoral thesis concerns *Duality Methods for Stochastic Optimal Control Problems in Finance*. In particular, all three core chapters touch upon convex duality against the economic/financial background of portfolio optimisation. The literature on duality for investment-consumption problems can be classified in accordance with the following three categories: (i) applied studies, (ii) theoretical studies, and (iii) mixtures of the previous two. This dissertation covered all three categories. Concretely, in Chapter 2, we dealt with item (iii), and developed a dual-control mechanism suitable for acquiring analytical near-optimal solutions to constrained investment-consumption problems. In Chapter 3, covering item (ii), we derived a dual formulation corresponding to an optimal consumption problem involving multiplicative habit formation. In Chapter 4, which addressed item (i), we made use of duality techniques to derive optimal policy rules for a pension fund that offers a DC scheme. These three studies jointly constitute the research-based nucleus of this thesis, i.e. its "core". As an introduction to this core, Chapter 1 expanded on the duality-linked theme in a rather general sense. The introductory chapter thereby aimed to highlight what duality theoretically entails and why it is practically useful. In addition to this, it supplied brief synopses of the academic content addressed by the preceding three chapters. Ultimately, in Chapter 5, we concluded this dissertation. Therein, we specifically focused on the contribution of the academically relevant output to the literature on economic/financial duality. As Chapters 1 and 5 serve complementary roles with regard to the remainder of this thesis, we subsequently

summarise Chapters 2, 3 and 4 at greater length. For smaller variants of the following summaries, one can consult the abstracts at the beginning of each core chapter. Correspondingly, for more extensive overviews, the distinct paragraphs provided in the third section of Chapter 1 may be useful.

Chapter 2

In Chapter 2, we developed a dual-control method applicable to a broad class of constrained utility-maximisation problems. Its mechanism rests on a generalisation of the approximating routine proposed by Bick et al. (2013). In order to adequately outline the details underpinning this scheme, we are obliged to look more closely at duality in the context of portfolio optimisation. Due to the inclusion of trading restrictions, constrained investment-consumption problems are difficult to solve. The mathematical complexity associated with a derivation of optimal solutions is almost entirely attributable to the non-uniqueness of equivalent martingale measures. Namely, since the market is constrained and therefore incomplete, there exist infinitely many martingale measures. The most optimal or "least-favourable" martingale measure can be determined by an appropriate minimisation procedure of the dual formulation. However, in most cases, the ensuing first-order conditions cannot be solved in closed-form. In fact, for particular specifications of the conic constraint set, it is not even possible to derive such conditions. As the optimal dual controls give rise to optimal primal rules, the latter phenomena directly encumber a derivation of the optimal decision variables. That is, unless one is able to analytically spell out the dual-optimal martingale measure, closed-form expressions for the primal-optimal investment-consumption strategies are not available. For this reason, most constrained utility-maximisation problems lack analytical tractability and are solved by means of computationally demanding numerical machinery. The general absence of analytical solutions and the interrelated need for numerically intense approaches are two major issues in the applied domain of constrained utility-maximisation.

Our dual-control method deals with the aforementioned issues and manages to generate near-optimal closed-form solutions in a highly efficient way. To this end, it makes use of three-fold approximating scheme. First, it recognises that all analytical nuisance stems from the dual. Therefore, it analytically approximates the optimal dual controls by a modified minimisation procedure of the dual formulation. In particular, it restricts the set of dual controls to a tractable analogue and optimises the dual problem accordingly. The resulting dual approximations bring forth closed-form candidates for the optimal primal controls. To make these candidates *admissible* in the primal environment, they have to be slightly adjusted. In other words, the "raw" candidate solutions generally fail to satisfy the trading/liquidity constraints. Consequently, in the second step, our method projects the candidate solutions into the admissibility set to arrive at near-optimal controls that are primal-feasible. In the third and final step, the approximating routine measures the accuracy of these approximate solutions by a financial evaluation of the corresponding duality gap. As a result, the dual-control method concretely renders analytical approximate policy rules that are accompanied by a "hard" guarantee concerning their accuracy. In the numerical illustrations, the method proved to work well. For the examples under scrutiny, the approximating method resulted in annual welfare losses smaller than 0.051% of the agent's initial endowment. Our conclusive statement on the possible accuracy of the dual-control method is supported by the variety of examined trading constraints and the technical complexity of both the financial environment and the preference qualifications.

Chapter 3

In Chapter 3, we studied the optimal consumption problem with multiplicative habit formation. The habit-linked literature is dominated by studies on additive models. In these models, an agent is assumed to derive utility from the *difference* between consumption and the habit level. While additive configurations are easy to handle in a mathematical sense, they lack economic relevance. As most utility functions only admit strictly positive arguments, consumption is namely required to exceed the habit level at all times. For this reason, in additive frameworks, the habit component is typically interpreted as a subsistence level. Such identifications are plausible from macro-related perspectives, wherein one examines the optimal consumption patterns of nations or large-scaled populations. However, in the confines of a micro-linked setup suitable for e.g. individuals and/or households, similar interpretations cannot be upheld. More specifically, the endogeneity of the habit level complicates such economic/financial identifications of the habit level. Therefore, when studying individuals and/or households, the habit component is generally characterised as a person-specific standard of living. The artificial lower bound imposed upon consumption correspondingly implies that the utility-maximising agent is obliged to consume at least as much as his/her standard of living. Even though this corresponds to a fairly ideal situation, it is not realistic. Adverse shifts in the financial circumstances can always urge a person to scale down consumption below the level to which he/she has become accustomed. Hence, despite the mathematical elegance involved with solving additive problems, they are not economically/financially relevant for all environments.

To arrive at a setup that manages to relax the unnatural lower bound imposed upon consumption, one can make use of multiplicative habit models. In these models, an agent is assumed to derive utility from the *ratio* of consumption to the habit level. This ratio is strictly positive for any budge-feasible consumption strategy. Due to the latter property, the ratio can be incorporated into most conventional preference qualifications. As a consequence, in multiplicative frameworks, consumption is not required to exceed a peculiar lower bound. On account of the relaxation of this bound, the aforementioned frameworks gain a significant amount of economic relevance. The habit-linked configuration is now amenable to micro-related situations consistent with small households and/or individuals. In spite of the ensuing economic advantages, the multiplicative habit models come at a high technical cost. By virtue of the mathematically complicated objective function, consumption problems involving multiplicative habit formation cannot be solved in closed-form. Most studies on multiplicative models therefore resort to numerical applications or the design of approximate solutions. Ordinarily, parts of the mathematical complexity related to the optimal consumption problem are addressed or facilitated by the dual formulation. However, for this problem, there is no dual problem known. In Chapter 3, we filled this gap in the literature, and made an entire branch of dual-related applications accessible, by deriving a corresponding dual formulation. We did so by means of a “concavification” procedure and the less well-known notion of

Fenchel duality. This strong duality result gave rise to a myriad of interesting implications. In our study, we exclusively focused on the duality relations and an evaluation mechanism commonly associated with dual-control methods.

Chapter 4

In Chapter 4, we analysed an optimal terminal wealth problem from an applied point of view. More concretely, we examined a utility-maximising pension fund that operates in conformity with a DC scheme. Terminal wealth setups can easily be identified with the individual-specific nature of conventional DC plans. On the grounds of the personally oriented specification of most preference functions, a participant's attitude towards risk can be included in a very precise manner. This possibility is of significant importance to DC providers, as the participating agents are generally required to carry all retirement-linked risk. The target of optimisation consequently outlines a person-specific function adapted to the preference/risk profile of a unique individual. Another important attribute inherent in a great majority of DC setups is the notion of underfunded starting positions. Individuals typically enter DC schemes with relatively high expectations regarding their retirement wealth. These practically unrealistic outlooks on pension goals result in an initial "mismatch" between the participant's contributions and his/her expectations. Translated into financial jargon, this mismatch can be characterised as an underfunding situation. That is, the pension fund is not in possession of sufficient funds to risk-neutrally cover the pension liabilities/goals. Even if the participants' prospects are adapted to reality, pension funds are still confronted with challenges concerning these retirement goals. These challenges may arise due to a.o. detrimental changes in the economic circumstances. In the context of utility-maximisation, the underfunding positions can easily be accommodated. Given the correspondingly realistic model setup, in Chapter 4, we aimed to answer the following practically relevant question: *Is it possible to increase the likelihood of achieving one's pension goals using target-oriented preferences?*

To be able to answer this question, we considered the LPM operator as a goal-based preference function. This operator essentially specifies a mathe-

matical criterion suitable for problems in the domain of partial hedging. At the same time, its technical definition includes a parameter that accounts for one's personal risk tolerance. By reason of its specification as a hedging criterion, the LPM operator is strongly target-oriented. In conjunction with the person-specific nature implied by the preference parameter, the LPM function is exceptionally appropriate for modelling DC frameworks. Consistent with its specification as a hedging criterion, the LPM operator incorporates a so-called reference level. This reference level is unique to individuals and can be modelled as an explicit retirement goal. For this reason, in our study, we identified the reference level as a person-specific life annuity. We investigated the corresponding utility-maximisation problem in the financial environment proposed by Koijen et al. (2009). Their market model is employed by a.o. the Dutch central bank (DNB) and therefore constitutes a financially meaningful framework. In this environment, the market prices of risk are assumed to be affine in a mean-reverting stochastic process. As an immediate consequence, it is not possible to derive the exact distributional features of the stochastic deflator process. Due to its significant impact on the general optima in the area of continuous-time portfolio optimisation, the latter complicates an analytical retrieval of closed-form solutions. Nevertheless, using inverse Fourier techniques, we were able to derive analytical expressions for the optimal policy rules. Furthermore, we managed to disentangle the distributional properties of retirement wealth in closed-form. Our numerical results demonstrated that the LPM operator is able to significantly improve the likelihood of achieving one's pension goals. Despite this potentially great performance, we also showed that the optimal policy rules are highly sensitive to the estimates for the market prices of risk and may be difficult to implement in reality.