Duality methods for stochastic optimal control problems in finance

Valorisation

In this attachment, we address the possible impact of this dissertation on both society and the non-academic industry. Despite the thoroughly technical nature of the overarching theme, each core chapter can be related to a topic of practical importance. In fact, all topics concern optimisation problems with regard to portfolio choice and/or savings decisions. The latter phenomena constitute crucial parts in the lives of individuals and institutional investors alike. We are therefore able to evaluate the societal and/or industry-linked impact of this dissertation along a wide array of practical dimensions. In an attempt to classify the corresponding domains of impact, we narrow the subjects of this valorisation down to three interlinked and substantial fields: (i) the pension industry, (ii) asset-liability management, and (iii) the (re-)insurance industry. We believe that this categorisation is meaningful, as most individuals are either directly or indirectly affected by *at least* one of the aforementioned domains. In addition to this, for a great majority of people, the mere connection to portfolio choice problems is established through one of the preceding items. Against the former background, it is noteworthy that this research has actively contributed to the Dutch debate on pension reforms. As a companion paper to Chapters 2, 3 and 4, we have co-authored an industry-oriented article that was conducive to the new Dutch pension agreement.[1] On the grounds of the pension

[1]This research was financially supported by the Network for Studies on Pensions, Aging and Retirement (NETSPAR). NETSPAR forms a platform on which researchers and practitioners contribute to both academic and pension-related discourse. NETSPAR thereby aims to bridge the gap between academia and the industry. Under the umbrella of the NETSPAR-linked theme "*Design of Pension Contracts in Incomplete Markets and under Uncertainty*", we have written this dissertation. Parts of this research have accordingly been presented at numerous NETSPAR seminars and conferences. The article mentioned in the main text concerns Balter et al. (2020) and appeared in NETSPAR's *Design Series*. Their *Design Series* consists of articles that bear relevance to the Dutch pension debate. Subsequently, we address our article's content in more detail.

industry's close ties to a.o. asset-liability management and (re-)insurance, the foregoing contribution exemplifies this dissertation's impact on the remaining two domains. In the sequel, we visit the separate fields and elaborate on the impact associated with the core chapters.

Pension Industry

All chapters included in this dissertation address problems that are relevant to the pension industry. We first focus on the distinct chapters and then comment on the previously mentioned article. Pension funds are in general concerned with acquiring the best possible replacement ratios. In doing so, they have to deal with multiple sources of unhedgeable risk. The ensuing incompleteness can be attributed to e.g. mortality risk or extremely long-dated/illiquid cash-flows. In addition to this, most pension funds in the EU-area are legally obliged to keep up with clearly defined solvency requirements. These requirements pose direct restrictions upon the funds' feasible set of policy rules. The corresponding situation can therefore be modelled by an ordinary constrained terminal wealth or utility-maximisation framework. In that regard, it is clear that the topic central to Chapter 2 becomes highly relevant. We recall that this chapter introduces an approximate dual-control method suitable for constrained optimal control problems. Its mechanism manages to generate near-optimal approximations to the optimal decision variables in closed-form. The advantages associated with this approximating routine are highly beneficial to the pension industry. It concretely enables pension funds to reduce the real-time computational effort required to implement their investment policies. For similar reasons, the dual-control method furnishes analytically tractable and useful insights regarding the impact of market incompleteness on their executed policy rules. Moreover, the mechanism endows a fund with an explicit framework appropriate for effectively managing non-traded risk-drivers. Due to the cumbersome nature of market incompleteness, specifically in the pension-linked context, such a framework constitutes a valuable asset.

In Chapters 3 and 4, we studied utility-maximisation problems involving different types of reference levels. Chapter 3 concerned a theoretical treatment of multiplicative habit formation. The reference level was accordingly defined as an endogenous or internal process. In accordance with the concept of habit formation, we specified utility with respect to consumption over the entirety of the trading interval. For the corresponding optimal consumption problem, we derived a dual formulation. This primarily theoretical result gave rise to numerous applications/insights of practical importance. From a pension fund's perspective, the most meaningful attributes are given by the semi-analytical specification of optimal consumption and the dual-induced evaluation mechanism. To this end, it is important to note that the habit formation framework can be employed to study the optimal consumption/savings behaviour required to ensure person-specific satisfaction with regard to one's standard of living.[2] In a pension-linked configuration, this setup consequently allows for the exact calculation of individually optimal deposits into e.g. a DC scheme. Due to the semi-analytical specification of the optimal consumption process, the fund is able to infer more precisely how these optimal deposits are affected by changes in the financial circumstances. This closed-form element therefore enables a fund to raise realistic expectations on the subject of the participant's defined contributions. The evaluation mechanism can be utilised to gauge the approximate closed-form contributions' accuracy. Since it is difficult to handle problems of this kind analytically, the preceding evaluation routine saves time and opens doors to the construction of more tractable deposit-related policies. The latter touches upon the creation and implementation of dual-control methods adapted to setups involving multiplicative habit formation.

In Chapter 4, we made use of an exogenous or externally defined reference level. Moreover, instead of consumption over the trading interval, the agent was assumed to derive utility from terminal wealth alone. The exogenous nature of the reference level fits well in the confines of a terminal wealth problem.

[2] We stress that this point of view is unique to the multiplicative branch within the literature on habit formation. In additive setups, an agent is required to keep consumption above the habit level at all times. As a consequence, it is hard to interpret the habit component as a standard of living. Adverse changes in the financial circumstances are very likely to negatively affect an individual's savings behaviour. One may realistically be required to scale down consumption below the level to which he/she has become accustomed. Due to the evident relation to pension contributions, it is clear that the additive framework is too restrictive for individual-specific pension schemes. The multiplicative setup is in that respect useful, as it allows for consumption below the habit component. This attribute makes the setup amenable to interpretations provided in the main text.

Specifically in the context of a DC scheme, the latter configuration admits a variety of pertinent interpretations. In the chapter of interest, we predominantly concentrated on the situation for a pension fund offering a DC scheme. For this purpose, we identified the reference level as an individual-specific life annuity. Moreover, to model the agent's preferences, we relied on the LPM operator. This operator outlines a partial hedging criterion adjusted for the agent's risk-tolerance. The target of optimisation correspondingly accommodates a strong orientation towards the predefined reference level. By means of this unique operator, we studied whether it is possible to improve the likelihood of achieving one's pension goals. In consideration of the new Dutch pension agreement, this study is highly relevant for two concrete reasons. First, in conformity with the new agreement, retirement wealth is adjusted based on the fund's performance over the life-cycle. Furthermore, the participant's contributions are more or less held fixed throughout the entire accumulation phase. The nature of the ensuing scheme closely resembles the DC configuration at the heart of our chapter. Second, in line with the corresponding shift in risk from the employer to the employee, the new agreement puts more emphasis on the participant-specific preferences. Qualifications adapted to an individual's personal risk profile should generate retirement-linked outcomes adequate for his/her specific situation. The LPM framework clearly allows for a person-specific identification of preferences, whilst retaining the possibility of favourable pension outcomes. Apart from some technical downsides[3], Chapter 4 indeed demonstrates that the LPM criterion can significantly improve the likelihood of achieving one's desired pension target. Chapter 4's setup could therefore be interesting to pension funds operating in a.o. the Dutch second pillar.

The industry-oriented paper accompanying Chapters 2, 3 and 4 is closely linked to the previous topic. This article is written by Balter et al. (2020) and carries the title "*Investing for Retirement with an Explicit Benchmark*". It concretely studies the impact of a goal-based utility function on the recovery potential of

[3]These downsides concern the optimal policy rules. The numerical results of Chapter 4 namely revealed that the optimal solutions are highly sensitive to the estimates for the market prices of risk. Moreover, in light of particular solvency requirements, the numerically assessed investment strategies are difficult to implement in practice. Both downsides can easily be handled by slight modifications of the optimisation framework. We have addressed potential modifications in Chapter 5.

a pension fund. For this purpose, the paper relies on a DC setup similar to Chapter 4. The primary differences consist in the preference qualifications and the postulates underpinning the financial environment. Whereas the former chapter resorts to the market proposed by Koijen et al. (2009), the industry-oriented article utilises an ordinary uni-dimensional Black-Scholes model. The corresponding stylised market suffices to appropriately convey the benefits associated with the target-oriented preference function. Unlike the aforementioned chapter, the article makes use of a novel utility qualification: the dual-CRRA function. As the name suggests, this preference qualification constitutes a minor modification of the ordinary CRRA operator. It incorporates two CRRA functions with different coefficients of relative risk-aversion. Depending on whether retirement wealth exceeds a prefixed benchmark, utility is derived from either of the two CRRA operators. This newly defined preference paradigm correspondingly allows one to explicitly characterise an individual's preferences around a reference level or target. For this reason, a pension fund is merely required to estimate/calibrate two preference-linked parameters, i.e. the separate coefficients of risk-aversion. In the context of a DC scheme, we identified the reference level as a person-specific life annuity. Moreover, we studied a participant who becomes notably less risk-averse when wealth falls below the reference level. This risk-related behavior approaches the gamble-for-resurrection phenomenon unique to prospect theory. The numerical results suggested that the dual-CRRA function is capable of substantially improving a pension fund's recovery potential. This improvement was based on relative performance with respect to the ordinary CRRA operator. Even though our study can be embedded in a larger body of similar NETSPAR-linked contributions, it was part of an explicit discussion that led to the new pension agreement. The clause on more individual-specific and preference-linked investment strategies can particularly be related to the article at hand.

Asset-Liability Management

Even though asset-liability management, henceforth ALM, constitutes a crucial practice within the pension industry, we are able to highlight some unique corresponding aspects of impact related to this dissertation. For this reason, it is noteworthy that the practice of ALM is not unique to pension funds. In the domains of, for instance, (re-)insurance, banking or trading, ALM plays a prominent role as well. We are best able to stress the impact of this dissertation on the domain of ALM along the lines of Chapters 2 and 4. In Chapter 2, we developed a dual-control method suitable for utility-maximisation frameworks involving convex trading constraints. Constraints of this kind can be modelled to account for solvency requirements and/or (partially) unhedgeable risk-drivers. For example, by enforcing restrictions upon investments in well-defined volatility derivatives, one arrives at a model with non-traded volatility risk. Similar reasoning applies to e.g. mortality risk or inflation risk. By the same token, the former type of constraints can be employed to keep an investor from taking (extremely) large and/or short positions in any of the traded assets. As practitioners in the field of ALM are generally confronted with such restrictions/requirements, our duality framework bears significant relevance to this domain. The latter also held true for the pension industry. However, ALM as such entails a larger body of (un)hedgeable risk factors. In addition to this, the ALM-specific interpretation of the utility-maximisation framework considerably differs from the pension-related one. Consistent with the configuration of Chapter 4, due to the possibly person-specific nature of a utility function, the most obvious pension-linked interpretation pertains to a DC setup. Therein, terminal wealth should be identified as retirement wealth, and the reference level as a person-specific pension goal. To make the setup amenable to ALM in a broader sense, this interpretation has to be generalised. Terminal wealth ought to be identified as the asset process, and the reference level(s) as the liability process. The preference function can correspondingly be specified in a "risk-neutral" manner by means of e.g. hedging operators or mean-variance criteria. Note here that our dual-control method applies to generally defined state-dependent utility qualifications incorporating exogenous benchmark variables governed by broadly specified semi-martingales.

The relevance of Chapter 2 for the extensive field of ALM is straightforward and resembles the pension-linked importance. More specifically, the dual-control method endows asset-liability managers with a tractable and time-efficient tool for calculating/implementing the optimal investment strategies. Both the efficiency and the tractability are attributable to the closed-form nature of the ensuing near-optimal policy rules. These features save time and allow the managers to clearly communicate the details underpinning their strategy-linked choices. Furthermore, due to the solutions' analytical tractability, the "black box" surrounding the conventional numerical applications disappears. In other words, the exact impact of particular model specifications on the approximate decision variables is clear and explainable. Similarly, comparative analyses amongst a set of dual-induced strategies is fairly uncomplicated and consequently facilitates the construction of unique policy rules. On a slightly less fundamental level, Chapter 4's impact on ALM can be situated in the use of goal-based hedging criteria. We recall that this chapter studied the possibility of improving a pension fund's recovery potential using a strongly target-oriented LPM operator. Akin to the preceding dual-control method, this setup can easily be adapted to a more general ALM problem. Under a modified identification of the terminal wealth process as well as the reference level, the setup can be aligned with general ALM frameworks. As a result, the aforementioned recovery potential coincides with a solvency ratio quantifying the degree up to which an institution is able to meet its liabilities. Given the numerically verified positive impact of the LPM operator on the recovery potential, it is clear that this target-oriented function is also able to positively affect the preceding solvency ratio. By means of different hedging benchmarks, our study consequently suggests that asset-liability managers are capable of improving their results/performance. On account of the endogeneity of the reference level in Chapter 3, the link of this chapter to ALM is not clearly visible. Nevertheless, for situations wherein asset-liability managers are required to withdraw capital from their asset process(es), the multiplicative dual formulation can come in handy. Particularly in the spirit of possible dual-control mechanisms, this framework can be employed to facilitate numerical computations in an analytical-friendly manner. As this closely resembles the conceptual impact of the approximating routine developed in Chapter 2, we do not elaborate on the technical details.

Insurance Industry

Great parts of the (re-)insurance industry are concerned with the design and related pricing of products. Well-known products crucial to the life and non-life sectors are health insurances and car insurances, respectively.[4] More abstract and technical products are handled by re-insurers, which may entail e.g. the evaluation of options on risk carried by an insurer. In pricing these products, the insurance industry as well as the re-insurance industry are obliged to deal with a great body of unhedgeable risk-drivers. For life-linked setups, products straightforwardly hinge upon a.o. mortality/longevity risk. Likewise, in configurations relevant for the non-life branch, claim sizes and corresponding frequencies are typically subject to unhedgeable sources of uncertainty. Due to the involvement of (partially) non-traded risk-drivers, the aforementioned pricing process is highly nontrivial. In agreement with our analysis on optimal investment in the presence of trading constraints, this process generally requires computationally demanding applications that lack analytical tractability. Put differently, the risk-neutral evaluation of insurance-linked products can be time-consuming and may pose mathematical challenges. In an attempt to tackle both issues in a relatively understandable manner, the dual-control method central to Chapter 2 proves useful. Moreover, to improve the performance of replicating strategies associated with particular products, the results in Chapter 4 are helpful. We note that the technical finding outlined in Chapter 3 is not compatible with conventional pricing schemes. By virtue of the endogeneity of internal habit components, the framework cannot be reconciled with most evaluation methods. Therefore, in the sequel, we solely elaborate on the precise impact of Chapters 2 and 4 on the (re-)insurance industry.

On account of the universal mechanism underscoring pricing routines, we subsequently do not distinguish between insurers and re-insurers. For the same reason, the impact of this dissertation on general pricing techniques reaches beyond the insurance industry. Investment banks or private investors can possibly benefit from our research as well. However, in view of the fact that (re-)insurance companies occupy a substantial part of the market for financial/actuarial products, we confine ourselves to an impact-related

[4]In this regard, we deem it noteworthy that almost all EU-citizens are in possession of at least a health insurance. In, for example, the Netherlands and Germany, having a health insurance is required by law. Similarly, although not necessarily enforced by legislation, most individuals in possession of a car have an automobile insurance. In addition to these widespread products, a great amount of people wish to purchase e.g. mortgage-linked insurances, life insurances, or personal liability insurances. This stresses the omnipresence of insurance products and their importance for society.

assessment of this sector alone. The latter choice is corroborated by the large amount of individuals who are actively involved in the purchasing process of insurance-linked products, e.g. health insurances or car insurances.

We have already addressed the possible computational and mathematical issues involved with the pricing procedure of insurance-linked products. The burdensome nature of this process is principally attributable to the partial or full non-tradeability of certain risk-drivers. In the confines of utility-maximisation, the dual-control method developed in Chapter 2 is capable of coping with this unhedgeability in a tractable and efficient way. For the upsides of this dual-control method, one can visit the preceding sections or Chapter 1. Hence, to underline the relevance of this duality mechanism for pricing schemes, we must disclose the link between utility-maximisation and risk-neutral evaluation techniques. As it is debatable whether investors at insurance companies can be classified as risk averse or risk-seeking individuals/agents, preference qualifications do not appear to be the greatest targets of optimisation. Nevertheless, the generality of the utility operators included in our dual-control framework allows for more "risk-neutral" objective functions. Examples of such operators include, but are not limited to, the LPM criterion from Chapter 4 or concavified variants of the celebrated mean-variance function. We recall that the state-dependent preference qualifications in Chapter 2 may incorporate exogenous benchmark processes or reference levels. Therefore, under the additional identification of these reference levels as insurance products, the utility-maximisation problem reduces to a setup suitable for finding the best replicating strategies. These replicating strategies would correspond to a fixed initial endowment. To find the "best" price, the (re-)insurer can determine this endowment in such a manner that the replacement ratio exceeds 100% with a probability of, say, 99%. The latter implies that, in 99% of the cases, the near-optimal analytical replicating strategy generates a proper (partial) hedge against the uncertainty induced by the product. As a result, the corresponding price seems appropriate. The outcomes reported in Chapter 4 demonstrate that a pricing approach of this kind is able to render viable outcomes.

In a similar fashion, one can adapt our dual-control mechanism to utility indifference pricing techniques. These setups do not necessarily depend on "risk-neutral" characterisations of the preference functions. In fact, this pricing technique makes explicit use of Inada-type utility functions in an attempt to compute fair evaluations of partially non-traded insurance products. Utility indifference pricing works along the lines of two interrelated problems as follows. In the first problem, one simply solves for the optimal trading strategy. Utility is here derived from terminal wealth alone. Taking into account this is far from straightforward in the presence of non-traded risk, our dual-control method comes in useful. In the second problem, one also solves for the optimal trading strategy. However, the agent is endowed with his/her initial capital *minus* a constant amount of monetary units. In addition to this, utility is derived from terminal wealth *plus* the insurance product. For reasons addressed around the first problem, our dual-control method may be utilised to arrive at tractable near-optimal solutions for the relevant policy rules. According to the principle of indifference pricing, the fair price for the preceding product is equal to the aforementioned amount of monetary units. This amount must namely be determined such that the objective functions of both problems are equal to each other. Indifference pricing typically relies on utility functions from the exponential family to derive closed-form expressions. Other utility functions in general pose problems with regard to an analytical retrieval of optimal replicating strategies and/or indifference prices. Our approximating routine enlarges this narrow class of applications. The (re-)insurance industry can consequently employ our dual-control routine to find tractable indifference prices in an efficient way for a considerably larger class of (more realistic) utility choices. This shows how our research facilitates and improves the financial/actuarial fair pricing of insurance products. The fair nature benefits both the (re-)insurers and the large number of insured agents.