Closed-form Approximations to Optimal Investment Policies in Markets with Frictions

Thijs Kamma*
Dept. of Quantitative Economics
Maastricht University
NETSPAR

Antoon Pelsser
Dept. of Quantitative Economics
Maastricht University
NETSPAR

October 15, 2020

Abstract

In this paper, we develop a dual-control method for approximating the optimal dynamic asset allocation in multi-dimensional financial markets under transaction costs. This method consists in confining the space of feasible shadow price process, due to which we are able to derive a continuous trading strategy in closed-form. We then discretise the latter decision rule, such that the investor exclusively trades at a finite amount of dates on the trading interval. Using convex duality, we are able to compute lower and upper bounds on the optimal value function, which enables us to assess the quality of the approximating procedure. The numerical examples show that this method may result in negligible welfare losses.

Keywords: Convex duality, incomplete markets, life-cycle investment, Malliavin calculus, state-dependent utility, transaction costs

JEL Classification: D23, D52, D53, G11

1 Introduction

The inclusion of proportional transaction costs in ordinary utility-maximising investment frameworks typically complicates a derivation of closed-form expressions for the optimal portfolio decisions. Namely, either the technical derivation suffers from severe mathematical complexity, cf. Shreve and Soner (1994), Choi et al. (2013), and Rokhlin (2013), or it requires computationally intense techniques that are by construction analytically intractable, cf. Constantinides (1986), Gennotte and Jung (1994), and Balduzzi and Lynch (1999). These difficulties with regard to the acquirement of closed-form expressions are predominantly attributable to the nature of transaction costs, which makes the conventionally employed martingale techniques in continuous time inadequate. In particular, on account of an application of convex duality approaches, one is able to rely on "surrogate" martingale machinery. Such methodology specifically allows one to solve

*Corresponding author. Mailing address: t.kamma@maastrichtuniversity.nl. PO Box 616, 6200 MD Maastricht, The Netherlands. Phone: +31 (0) 43 388 3695.

an analogous investment problem in a market that excludes transaction costs. In spite of this mathematically elegant problem description, this approach ultimately involves the choice for a so-called shadow price process[1], which is only in exceptional cases obtainable in closed-form, see e.g. Cvitanić et al. (1999), Kabanov and Safarian (2009), or Kallsen and Muhle-Karbe (2015). In order to circumvent this analytical inability and the related computational demand, this paper proposes a dual-control method that enables one to approximate the optimal dynamic investment policies in closed-form.

In market models that rule out the presence of transaction costs, the optimal solutions to the corresponding dynamic investment problems have been derived and analysed by a voluminous body of literature. The pioneering articles in that regard are those by Merton (1969) and Merton (1971). In particular since the introduction of the martingale techniques that are developed in Pliska (1986), Karatzas et al. (1987), and Cox and Huang (1989, 1991), an exhaustive amount of other studies have contributed to the literature on this topic.[2] The success of these methods is primarily a consequence of the fact that, by virtue of its application, a static rather than the initial dynamic problem has to be solved. Traditional Lagrangian techniques in combination with straightforward hedging arguments suffice to solve the latter problem. In the absence of market-related peculiarities such as transaction costs, these techniques therefore not only facilitate but ensure the recovery of closed-form identities for the relevant controls. However, if one relaxes the simplifying assumption of non-existent transaction costs, as was done for the first time in Magill and Constantinides (1976) and Davis and Norman (1990), the conventional application of the preceding martingale techniques is infeasible. Concretely, on the basis of its original dynamic problem specification, we are not capable anymore of defining an analogous static problem, and consequently not of deriving optimal solutions.

In order to bypass the impossibility of solving a convex optimisation problem, one typically derives the dual and tries to solve this instead. Accordingly proceeding, and deriving the dual by means of the convex duality method in Klein and Rogers (2007) and Rogers (2013), we arrive at a problem description that implies a frictionless auxiliary financial market. The papers by Jouini and Kallal (1995) and Lamberton et al. (1998) demonstrate that this auxiliary market comprises of the same asset mix as the original market, howbeit with inclusion of a non-unique strictly positive semi-martingale that

[1]These shadow prices are conceptually equivalent to those in financial environments with convex trading constraints, cf. Cvitanić and Karatzas (1992), Karatzas and Shreve (1998), and Hu et al. (2005). That is, these both serve to engender a least-favourable auxiliary market. The former processes, however, coincide with the complete characterisation of an artificial semi-martingale stock process. This contrasts to the latter counterparts, which exclusively coincide with local drift-terms that are inherent in a non-unique risk-neutral pricing measure. Also the conditions that these processes ought to fulfill diverge: cf. Cvitanić and Karatzas (1996). In our approximate method, these shadow prices play a central role.

[2]The list of studies that rely in their analysis and derivations on such martingale machinery is extraordinarily extensive, due to which it is not sensible to cite all of them here. Hence, we refer the reader to Karatzas and Shreve (1998) for an overview of the method itself as well as of studies that employ this technique. Note that in section 3.3 of this paper, we also apply the martingale method at hand.

delineates the stock process. In line with the literature, we refer to this non-unique asset price as the shadow price process, cf. Kallsen et al. (2010), and Kallsen and Muhle-Karbe (2011), and Gerhold et al. (2012). This shadow price is restricted to attain values between two pre-fixed bounds: the bid and ask prices of the traded stock. In accordance with the absence of transaction costs in the auxiliary environment, we are able to utilise the martingale method to derive closed-form expression for the decision rules that optimise the therein defined investment problem. Subsequently, in conformity with the dual problem specification, we must determine the shadow price process in a fashion that is least-advantageous from the investor's point of view. Thereby, the mechanism underscoring the dual aims to "prune" the auxiliary-optimal rules to the admissibility region.[3] However, by reason of its cumbersome restriction to values between two bounds, an analytical expression for the shadow price process is in general not available.

Hence, if we employ the method above, we are in possession of closed-form expressions for the optimal controls, which are entirely identified up to these shadow prices. The inability to equivalently obtain analytical expressions for the shadow prices, however, guarantees the unidentifiable nature and thus the non-analytical definition of these trading rules. So as to sidestep this analytical obstruction and thereby be able to derive approximate controls in closed-form, we broadly follow the reasoning underscoring similar approximations in markets with trading constraints by Keppo et al. (2007), Bick et al. (2013), and Kamma and Pelsser (2019). That is, taking into consideration that the inability to derive expressions for the optimal investment decisions in analytical form originates completely from the in closed-form generally unattainable shadow prices, the first step in our approximating routine consists of confining the (dual) space of shadow price processes to some closed and convex subspace. Not every such subspace guarantees that one is successfully able to collect analytical expressions for the shadow prices, but there exists *at least* one that does.[4] Moreover, if one still wishes to employ numerical methods (which our framework allows for), this restricted space may unquestionably facilitate potential computations. Ergo, assuming that we opted for a subspace that yields analytical shadow prices, we are in possession of closed-form investment decisions.

Nevertheless, these closed-form decisions are inadmissible, since they presume that the investor is able to continuously trade. In order to make these decisions notwithstanding

[3]We owe this terminology to Bick et al. (2013), who analyse a similar investment problem that instead of transaction costs incorporates convex trading constraints. These authors emphatically stress the inclination of the dual to opt for shadow price processes in such a manner that admissibility of the auxiliary-optimal controls in the actual market holds. We refer the reader to Cvitanić and Karatzas (1992), Davis (1997), and Schachermayer (2004) for similar claims. This dual-induced mechanism resembles the one in case of transaction costs: only the shadow prices are differently specified.

[4]The original dual space contains all shadow prices that attain values between the bid and ask prices of the traded stock, see e.g. Herczegh et al. (2015). Therefore, one always available, feasible and completely identified choice would be to set the shadow price process equal to the process for the risky asset in the original market. Confining the dual-space to a subset that exclusively contains the latter shadow price process will thus always result in a closed-form recovery of primal-optimal trading rules.

their continuous specification admissible in the original market, we discretise them over time. This discretisation is based on a predetermined selection of trading dates. In concrete terms, at each prefixed trading date, the investor adjusts the portfolio holdings in accordance with those that are spelled out by the analytically available investment policies. As for the confined dual space, not every such selection of trading leads to admissibility, but there exists *at least* one that does. Additionally, we assume that the dates are state-dependent, such that their capability to safeguard admissibility is substantially improved.[5] Therefore, assuming that we have chosen an "admissible" selection of trading dates, we are in possession of admissible and closed-form investment decisions. To evaluate the quality of the approximation, we follow de Palma and Prigent (2008) and de Palma and Prigent (2009) and quantify the ensuing duality gap in monetary terms. The approximation is thus accompanied by a strong guarantee regarding its accuracy. To the best of our knowledge, approximate methods similar to ours have not been proposed yet.[6]

We evaluate the potential precision of this approximating routine by casting the matter into an ordinary one-dimensional Black-Scholes framework, see e.g. Soner et al. (1995), and Højgaard and Taksar (1998). Moreover, we consider two different utility functions: a conventional CRRA function and a non-trivial dual-CRRA qualification (cf. Kamma et al. (2020) for details on this fucntion). In the former case, e.g. Zakamouline (2002) and Liu (2004) have found analytical but mathematically comparatively complicated solutions; in the latter case, closed-form expressions are not available. The actual optimal rules for the dual-CRRA investors are because of the underscoring utility function considerably more dynamic in nature that those for the CRRA individuals. Combining both observations: these examples can adequately illustrate how the method performs in case of known and unknown solutions, and at the same in case of fairly static and dynamic truly optimal portfolio rules. The results demonstrate that the annual welfare losses that are incurred due to implementation of these approximations vary between roughly 2 and 12 basis points of the individual's initial endowment. These negligible welfare losses therefore allude to near-optimality of the approximated dynamic allocations of assets.

The remainder of the paper is structured as follows. Section 2 introduces the economic environment. Subsequently, section 3 revisits the convex duality results for investment problems in markets with frictions. Section 4 sets out the approximate method, followed by a corresponding numerical evaluation. Finally, section 5 concludes the paper.

[5] The most trivial and straightforward choice for trading dates that assure admissibility at all times would be to not trade at all, i.e. to select zero of such dates. In light of the fact this approach entirely disregards the analytically available though continuous approximate investment rules, which provide a proper indication as to the actual optimal trading behaviour, let us propose another choice that assures admissibility. By making use of the dates' potential state-dependency, we may allow the investor to trade only if the corresponding transfers maintain a non-negative level of total wealth and if the proportion of wealth allocated to stocks does not exceed the unit interval, cf. Remark 4.2.

[6] Approximate methods in markets with frictions have been developed by Gennotte and Jung (1994), Keppo and Peura (1999), and Zakamouline (2006). However, the latter two do not involve a proper analytical foundation; the remaining one exclusively applies to pricing European call options.

2 Model Setup

2.1 Financial Market Model

Define a horizon $T > 0$, and consider a probability space $(\Omega, \mathcal{F}_t, \{\mathcal{F}_t\}_{t\in[0,T]}, \mathbb{P})$, whereon an $\mathbb{R}^N$-valued independent standard Brownian motion, $\{W_t\}_{t\in[0,T]}$, subsists. The $\mathbb{P}$-augmentation of W_t's canonical filtration $\{\mathcal{F}_t^W\}_{t\in[0,T]}$ reads $\{\mathcal{F}_t\}_{t\in[0,T]}$.[7] In the sequel, (in)equalities between stochastic processes are understood $\mathbb{P}$-almost surely.

We outline a market, $\mathcal{M}$, that populates a single investor who continuously trades over $[0, T]$ in a riskless instrument and N risky assets. The risk-free asset complies with

$$\frac{\mathrm{d}B_t}{B_t} = r_t \mathrm{d}t, \;\; B_0 = 1, \tag{2.1}$$

where r_t defines the $\mathbb{R}$-valued progressively $\mathcal{F}_t$-measurable interest rate, such that $r_t \in \mathbb{D}^{1,2}([0, T])$. The price process of the N risky stocks evolves in conformity with

$$\frac{\mathrm{d}S_{i,t}}{S_{i,t}} = \mu_{i,t}\mathrm{d}t + \sigma_{i,t}^\top \mathrm{d}W_t, \;\; S_{i,0} = 1, \tag{2.2}$$

in which $\mu_{i,t}$ denotes the $\mathbb{R}$-valued appreciation rate and $\sigma_{i,t}$ the $\mathbb{R}^N$-valued volatility process. We assume that both processes are $\mathcal{F}_t$-measurable and satisfy $\|\mu_t\|_{\mathbb{R}^N}, \mathrm{diag}\left(\sigma_t\sigma_t^\top\right) \in L^1\left([0, T]; \mathbb{R}^N\right)$.[8] Additionally, we postulate that σ_t fulfills the strong non-degeneracy assumption, $\phi^\top\sigma_t\sigma_t^\top\phi \geq \epsilon \|\phi\|_{\mathbb{R}^N}^2$ for all $\phi \in \mathbb{R}^N$ and some $\epsilon > 0$. Here, σ_t is $\mathbb{R}^{N\times N}$-valued, with rows $\sigma_{i,t}$, $i = 1, \ldots, N$; μ_t is $\mathbb{R}^N$-valued, with entries $\mu_{i,t}$, $i = 1, \ldots, N$.

The finite-horizon investor who resides in $\mathcal{M}$ has in mind a one-dimensional benchmark, concerning the endogenous controls, that evolves in accordance with the next equation

$$\frac{\mathrm{d}\Pi_t}{\Pi_t} = \mu_{\Pi,t}\mathrm{d}t + \sigma_{\Pi,t}^\top \mathrm{d}W_t, \;\; Y_0 = 1, \tag{2.3}$$

where we assume that $\mu_{\Pi,t}, \sigma_{\Pi,t} \in \mathbb{D}^{1,2}([0, T])$ holds.[9] As long as $\{\Pi_t\}_{t\in[0,T]}$ remains exogenous, we could interpret (2.3) as e.g. one's neighbour's labour income or national GDP. Consider Kamma and Pelsser (2019) for an analogous benchmark.

In $\mathcal{M}$, the investor is allowed to continuously trade in $\{B_t\}_{t\in[0,T]}$ and $\{S_t\}_{t\in[0,T]}$, whilst facing constant proportional transaction costs. For that purpose, we introduce:

$$S_t^B = (1_N - \lambda_1) \odot S_t, \quad \text{and} \quad S_t^A = (1_N + \lambda_2) \odot S_t, \quad \lambda > 0, \; \mu < 1, \tag{2.4}$$

[7]Furthermore, note that we may identify the Wiener space on which $\{W_t\}_{t\in[0,T]}$ lives as follows: $\Omega = \mathcal{C}_0\left([0, T]; \mathbb{R}^N\right)$ and $\mathcal{F} = \mathcal{B}_T^N$, with Wiener measure $\mathbb{P}(\mathrm{d}W)$, cf. Detemple and Rindisbacher (2005). The notion of this space constitutes an essential attribute in Malliavin calculus, see Nualart (2006).

[8]We let $L^p\left([0, T]; \mathbb{R}^N\right)$ be the space of all $\mathbb{P}$-a.s. p-integrable $\mathbb{R}^N$-valued progressively $\mathcal{F}_t$-measurable stochastic processes. Likewise, $L^p\left(\Omega \times [0, T]; \mathbb{R}^N\right)$ is the space of all p-integrable analogous process with finite expectations. In the text, we lighten notation and relegate the dependency on $\mathbb{R}^N$ to the superscript.

[9]The Sobolev-Watanabe space $\mathbb{D}^{1,2}([0, T])$ contains all $L^2(\Omega \times [0, T])$ Malliavin differentiable processes.

as, respectively, the bid price and ask price of the stock. Here, $\lambda_1 \in [0,1)^N$ and $\lambda_2 \in \mathbb{R}_+^N$ represent the proportional transaction costs involved with, respectively, selling and buying the stock. Moreover, we let "$\odot$" denote the so-called Hadamard product. This specification within the model of transaction costs constitutes an N-dimensional generalisation of the ones in Davis and Norman (1990) and Cvitanić and Karatzas (1996).[10]

In line with these papers, we then introduce the following two portfolio holdings

$$\begin{aligned} \mathrm{d}X_t &= (1_N - \lambda_1)^\top \mathrm{d}M_t - (1_N + \lambda_2)^\top \mathrm{d}L_t + r_t X_t \mathrm{d}t, \ X_0 \in \mathbb{R}_+, \\ \mathrm{d}Y_t &= \mathrm{d}L_t - \mathrm{d}M_t + \mu_t \odot Y_t \mathrm{d}t + \sigma_t Y_t \mathrm{d}W_t, \ Y_0 \in \mathbb{R}_+^N, \end{aligned} \tag{2.5}$$

where $\{M_t\}_{t\in[0,T]}$ and $\{L_t\}_{t\in[0,T]}$ are left-continuous and non-decreasing processes that specify the cumulative amount (in monetary units) of sales and purchases of $\{S_t\}_{t\in[0,T]}$. Observe that $M_0 = L_0$ ought to hold. By deviating from the conventional setup, we call a pair of funds transferred from the stock to the bank account, and vice versa, *admissible* if $\{L_t, M_t\}_{t\in[0,T]}$ satisfies (2.5) in addition to $X_t + 1_N^\top Y_t \geq 0$ for all $t \in [0,T]$.[11]

At the end of the trading interval, $t = T$, the finite-horizon investor is exclusively interested in the sum of his/her holdings in S_T and B_T, i.e. $X_T + 1_N^\top Y_T$. Therefore,

$$\begin{aligned} \widehat{X}_T = X_0 + 1_N^\top Y_0 &+ \int_0^T r_t \left\{ \widehat{X}_t - [p_t - q_t]^\top S_t \right\} \mathrm{d}t \\ &+ \int_0^T [p_t - q_t]^\top \mathrm{d}\left(\mu_t \odot S_t + \sigma_t S_t \mathrm{d}W_t\right) - \int_0^T S_t^\top \left[\lambda_1 \odot \mathrm{d}q_t + \lambda_2 \odot \mathrm{d}p_t\right], \end{aligned} \tag{2.6}$$

describes the stochastic quantity at $t = T$ that is central to this paper, in which $\widehat{X}_t := X_t + 1_N^\top Y_t$ holds. Moreover, $\{p_t\}_{t\in[0,T]} \in \mathbb{R}_+^N$ and $\{q_t\}_{t\in[0,T]} \in \mathbb{R}_+^N$ represent, respectively, the cumulative purchases and sales in each asset, $\{S_t\}_{t\in[0,T]}$, implying that $\mathrm{d}L_t = S_t \odot \mathrm{d}p_t$ and $\mathrm{d}M_t = S_t \odot \mathrm{d}q_t$. Concretely, the finite-horizon investor aims to maximise an expected utility criterion that exclusively depends on the quantity in (2.6) over all admissible trading strategies. For this purpose, let us introduce the set containing all admissible trading strategies: $\mathcal{A}_{X_0,Y_0}$.[12] The definition of this set directly implies that the investor is "solvent" at $t = T$, i.e. $X_T + 1_N^\top Y_T \geq 0$, assuming that $\{L_t, M_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ is satisfied.

2.2 Problem Specification and Preferences

The finite-horizon agent in $\mathcal{M}$ invests at $t = 0$ a predetermined amount of cash $(X_0, Y_0) \in \mathbb{R}_+ \times \mathbb{R}_+^N$, or similarly $\widehat{X}_0 = X_0 + Y_0$, in the bank account and in the stocks, and retires at

[10]We deviate from the more general formulations in Kabanov (1999), and Deelstra et al. (2001). In addition to ours, their models namely account for transfers within S_t. On account of the typical absence of such trades in reality, we exclude a corresponding characterisation within the model.

[11]The conventional setups focus on: $X_T + (1-\lambda_1)^\top Y_T \mathbb{1}_{\{Y_T>0\}} + (1+\lambda_2)^\top Y_T \mathbb{1}_{\{Y_T\leq 0\}}$, cf. Liu and Loewenstein (2002) or Kallsen et al. (2010). We deviate from this quantity for ease of exposition. For alternative quantities, the approximating principle namely requires only moderate modifications.

[12]In Klein and Rogers (2007)'s Example 2, $\mathcal{A}_{X_0,Y_0}$ is shown to be a closed and convex cone in $\mathbb{R}^{N+1}$.

$t = T$. This agent aims to maximise expected utility from his/her total amount of terminal non-liquidated holdings in both the stocks and the money market account, in relation to a person-specific stochastic target or benchmark. To that end, the former individual is free to select over the course of the trading interval, $[0, T]$, the specifics underlying the sales and purchases of the stocks. Precisely, the agent ought to determine these controls in agreement with (2.5) and (2.6), such that the admissibility conditions are met. That is, the individual faces the following dynamic optimisation problem[13]:

$$\begin{aligned} \sup_{\{p_t, q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}} \quad & \mathbb{E}\left[U\left(\widehat{X}_T, \Pi_T\right)\right] \\ \text{s.t.} \quad \mathrm{d}\widehat{X}_t &= r_t \left\{\widehat{X}_t - [p_t - q_t]^\top S_t\right\} \mathrm{d}t + [p_t - q_t]^\top \mathrm{d}S_t \\ & \quad - S_t^\top \left[\lambda_1 \odot \mathrm{d}q_t + \lambda_2 \odot \mathrm{d}p_t\right], \quad (X_0, Y_0) \in \mathbb{R}_+ \times \mathbb{R}_+^N. \end{aligned} \tag{2.7}$$

The preference qualification in the latter problem (2.7) requires some special attention. The most salient difference from the current paradigm with regard to utility specifications is the inclusion of a second argument. The inclusion of this component enables one to explicitly model the preferences around the benchmark, in an attempt to acquire optimal controls that are more target-oriented. This potentially concrete focus on a pre-specified goal is in light of recent development in the pension industry[14] and for risk management intents a highly preferred feature, cf. Basak (2002) or Basak et al. (2006).

Technically, we postulate that $U : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ satisfies the following conditions:

$$\lim_{x\to\infty} U'_X(x, y) = 0, \quad \lim_{x\to 0} U'_X(x, y) = \infty, \quad \text{and} \quad \limsup_{x\to\infty} x \frac{U'_X(x, y)}{U(x, y)} < 1, \tag{2.8}$$

for all $y \in \mathbb{R}_+$, in which $U'_X : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ specifies the first derivative of U with respect to its first argument. Likewise, we define $U''_X : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_-$ as the second derivative of U with respect to its first argument. Clearly, we assume that $U'_X > 0$ and $U''_X < 0$ hold. The first two limits in (2.8) compose the ordinary Inada conditions. The last limit in (2.8) outlines the reasonable asymptotic elasticity requirement, cf. Kramkov and Schachermayer (1999). We impose this condition to hold in order to assure amongst others feasibility of convex duality applications. Now, let $I : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ be the inverse of marginal utility, U'_X, such that $U'_X(I(x, y), y) = x$ holds for all $x, y \in \mathbb{R}_+$. Then,

$$V(x, y) = \sup_{z\in\mathbb{R}_+} \{U(z, y) - xz\} = U(I(x, y), y) - xI(x, y), \tag{2.9}$$

[13]The admissibility set, $\mathcal{A}_{X_0,Y_0}$, now contains all pairs $\{p_t, q_t\}_{t\in[0,T]}$, instead of $\{L_t, M_t\}_{t\in[0,T]}$, such that $X_t + 1_N^\top Y_t \geq 0$ for all $t \in [0, T]$ and (2.5) hold. For simplicity, we do not alter the notation.

[14]There is a general shift discernible from collective, defined-benefit, schemes to more individual, defined-contribution, analogues, see e.g. Gao (2008). Due to a shift of risk from the employer to the employee (i.e. to the individual under consideration) in the latter case, the individual-specific interests ought to play a more dominant role. In Kamma et al. (2020), the authors show how a utility function as displayed in (2.7) may facilitate the emphasis on indvidiual-specific interests.

defines the convex conjugate of U, for all $x, y \in \mathbb{R}_+$. Ultimately, we let $U'_Y : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_-$ be the first derivative of U in the second argument, for which we obviously assume that $U'_Y < 0$ holds. As to the differentiability of the foregoing functions, we note that U is once continuously differentiable in both arguments, $U \in \mathcal{C}(\mathbb{R}_+, \mathbb{R}_+; \mathbb{R})$. Similarly, we assume that U'_X and I are at least once piecewise continuously differentiable in both arguments, $U'_X, I \in \mathcal{PC}(\mathbb{R}_+, \mathbb{R}_+; \mathbb{R}_+)$. The preceding postulate deviates from the standard continuous differentiability assumption, and significantly enlarges the class of utility functions, without harming a proper recovery of the optimal portfolio rules, cf. Lakner and Nygren (2006).[15] Note that the convex conjugate V must naturally as well be in $\mathcal{PC}(\mathbb{R}_+, \mathbb{R}_+; \mathbb{R})$.

3 Convex Duality and Optimality

We proceed by studying the dynamic optimisation problem (2.7), wherein we particularly focus on the derivation of optimal controls. Due to the complexity of the budget constraint, it is not possible to solve (2.7) for the optimal decision rules, $\{p_t\}_{t\in[0,T]}$ and $\{q_t\}_{t\in[0,T]}$, on the mere basis of its primal formulation. In view of this inability, we first resort to an examination of the corresponding dual problem, for which purpose we rely on convex duality methods. Afterwards, we address the auxiliary market that may be derived from this dual specification. So as to establish a proper foundation for our approximate method, we conclude by considering the related situation wherein one is able to replicate terminal wealth by a continuous trading strategy in the absence of transaction costs.

3.1 Specification of the Convex Dual Problem

In this subsection, we derive and analyse the dual specification of the optimal control problem in (2.7). For that purpose, we utilise standard convex duality methods, consistent with the approach in Rogers (2013). To facilitate the exposition, we divide all relevant monetary variables by the strictly positive money market account, $\{B_t\}_{t\in[0,T]}$, which leads to: $\overline{X}_t := \widehat{X}_t B_t^{-1}$, $\overline{S}_t = S_t B_t^{-1}$, and $\overline{\Pi}_t = \Pi_t B_t^{-1}$ for all $t \in [0,T]$. As a consequence, by adjusting the arguments inside the objective in accordance with the former procedure, the re-scaled counterpart of the problem at hand is equivalent to (2.7) and reads:

$$\begin{aligned} &\sup_{\{p_t,q_t\}_{t\in[0,T]}\in\mathcal{A}_{X_0,Y_0}} \mathbb{E}\left[U\left(\overline{X}_T B_T, \overline{\Pi}_T B_T\right)\right] \\ &\qquad \text{s.t. } \overline{X}_T = \overline{X}_0 + \int_0^T \left[p_t - q_t\right]^\top \mathrm{d}\overline{S}_t - \int_0^T \overline{S}_t^\top \left[\lambda_1 \odot \mathrm{d}q_t + \lambda_2 \odot \mathrm{d}p_t\right]. \end{aligned} \tag{3.1}$$

[15]In concrete terms, Lakner and Nygren (2006) show that $I(X,Y) \in \mathbb{D}^{1,2}$ for all $X, Y \in \mathbb{D}^{1,2}$, $I \in \mathcal{PC}(\mathbb{R}, \mathbb{R}; \mathbb{R})$. We refer to the general utility functions as introduced in paper by Detemple and Zapatero (1991), as well as to the paper by Di Nunno and Øksendal (2009) regarding applications of Malliavin calculus in mathematical finance, from which it becomes clear how the latter result enlarges the class of conventional utility functions and simultaneously constitutes a desirable feature. See Kamma and Pelsser (2019) for a specific application of this principle to the *dual CRRA* utility qualification.

For purposes that pertain to a later stage of the convex duality mechanism, we may rewrite the integrals on the right-hand side (RHS) of (3.1) in terms of two analogues, in which the integrators coincide with $\{p_t\}_{t\in[0,T]}$ and $\{q_t\}_{t\in[0,T]}$. Specifically, it holds that $\int_0^T p_t^\top \mathrm{d}\overline{S}_t = \int_0^T \left[\overline{S}_T - S_t\right]^\top \mathrm{d}p_t$ and $\int_0^T q_t^\top \mathrm{d}\overline{S}_t = \int_0^T \left[\overline{S}_T - S_t\right]^\top \mathrm{d}q_t$.[16] Moreover, let us emphasise that $\{p_t, q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ implies $p_t, q_t \geq 0$ for all $t \in [0,T]$.

Then, as in Rogers (2013), we introduce a strictly positive semi-martingale process, $\{Z_t\}_{t\in[0,T]}$, which will serve as a dynamic Lagrange multiplier that enforces the equality constraint upon $\overline{X}_T$. In particular, this semi-martingale process evolves as follows

$$Z_t = Z_0 \exp\left\{\int_0^t \alpha_s \mathrm{d}s - \frac{1}{2}\int_0^t \|\lambda_s\|^2_{\mathbb{R}^N} \mathrm{d}s + \int_0^t \lambda_s^\top \mathrm{d}W_s\right\},\ Z_0 \in \mathbb{R}_+, \tag{3.2}$$

in which we assume that $\alpha_t \in \mathbb{D}^{1,2}([0,T])$ and $\lambda_t \in \mathbb{D}^{1,2}([0,T])^N$ hold. By reason of the latter assumption, $\{\lambda_t\}_{t\in[0,T]}$ is such that it satisfies: $\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^T \|\lambda_s\|^2_{\mathbb{R}^N} \mathrm{d}s\right)\right] < \infty$, cf. Novikov's condition in Karatzas and Shreve (1991). Now, let H_T be identical to the RHS of the budget constraint (3.1), then $\overline{X}Z_T - H_T Z_T = 0$ holds. We continue in a Lagrangian fashion, and insert the $\mathbb{P}$-expectation of the latter equation's left-hand side (LHS) into the value function of (3.1). Considering the objective, we pose according complementary slackness (CS) conditions on $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$. Minimising over $\{Z_t\}_{t\in[0,T]}$ in turn renders the subsequent specification of the dual problem, consider Theorem 3.1.

Theorem 3.1. *Consider the dynamic optimisation problem in* (3.1)*. Introduce*

$$\mathrm{d}Z_t = Z_t\left[\alpha_t \mathrm{d}t + \lambda_t^\top \mathrm{d}W_t\right], \tag{3.3}$$

for some $Z_0 \in \mathbb{R}_+$, where $\alpha_t \in \mathbb{D}^{1,2}_-([0,T])$ and $\lambda_t \in \mathbb{D}^{1,2}([0,T])^N$. Furthermore, define the set $\mathcal{H}_{\mathcal{A}_{X_0,Y_0}}$ as the barrier cone, which contains all $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$, such that $\mathbb{E}\left[Z_T\overline{S}_T \mid \mathcal{F}_t\right] \geq \overline{S}^B_t \mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]$ and $\mathbb{E}\left[Z_T\overline{S}_T \mid \mathcal{F}_t\right] \leq \overline{S}^A_t \mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]$, or equivalently:

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}_\lambda}\left[B_T^{-1}S_T \mid \mathcal{F}_t\right] &\geq (1_N - \lambda_1) \odot B_t^{-1}S_t, \quad \forall\ t \in [0,T]\\ \mathbb{E}^{\mathbb{Q}_\lambda}\left[B_T^{-1}S_T \mid \mathcal{F}_t\right] &\leq (1_N + \lambda_2) \odot B_t^{-1}S_t, \quad \forall\ t \in [0,T],\end{aligned} \tag{3.4}$$

in which we define $\overline{S}^A_t := B_t^{-1}S^A_t$, $\overline{S}^B_t := B_t^{-1}S^B_t$, and where $\mathbb{Q}_\lambda \sim \mathbb{P}$ specifies the non-unique probability measure induced by the martingale, $\left\{\mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]^{-1} Z_T\right\}_{t\in[0,T]}$. Then,

$$\inf_{\{\alpha_t,\lambda_t\}_{t\in[0,T]}\in\mathcal{H}_{\mathcal{A}_{X_0,Y_0}}, Z_0\in\mathbb{R}_+} \mathbb{E}\left[V\left(B_T^{-1}Z_T, \Pi_T\right)\right] + \widehat{X}_0\mathbb{E}\left[Z_T\right], \tag{3.5}$$

[16] It is rather simple to show that $\int_0^T p_t^\top \mathrm{d}\overline{S}_t = \int_0^T \int_0^t \mathrm{d}p_s^\top \mathrm{d}\overline{S}_t$ holds. Subsequently, a straightforward change in the order of integration gives rise to the following: $\int_0^T \int_t^T \mathrm{d}\overline{S}_s^\top \mathrm{d}p_t = \int_0^T \left[\overline{S}_T - S_t\right]^\top \mathrm{d}p_t$. The same procedure applies to the analogous integral that embeds $\{q_t\}_{t\in[0,T]}$ as the integrand. Inserting the ensuing two expressions into the dynamic budget constraint, for $\overline{X}_T$, in (3.1) then furnishes a more elegant identity: $\overline{X}_T = \overline{X}_0 + \int_0^T \left[\overline{S}_T - (1_N + \lambda_2) \odot \overline{S}_t\right]^\top \mathrm{d}p_t - \int_0^T \left[\overline{S}_T - (1_N - \lambda_1) \odot \overline{S}_t\right]^\top \mathrm{d}q_t$.

characterises for all $X_0, Y_0 \in \mathbb{R}_+$ the dual formulation corresponding to (3.1)*. Optimal terminal wealth therefore must conform to $\widehat{X}_T^{\text{opt}} = I\left(B_T^{-1} Z_T, \Pi_T\right)$. Moreover,*

$$J_P\left(\widehat{X}_0\right) = \inf_{\{\alpha_t, \lambda_t\}_{t\in[0,T]} \in \mathcal{H}_{\mathcal{A}_{X_0,Y_0}}, Z_0 \in \mathbb{R}_+} J_D\left(\widehat{X}_0, Z_0, \{\alpha_t, \lambda_t\}_{t\in[0,T]}\right), \tag{3.6}$$

is true for all $\widehat{X}_0 \in \mathbb{R}_+$, i.e. strong duality holds. Herein, we define $J_P : \mathbb{R}_+ \to \mathbb{R}$ as $J_P\big(\widehat{X}_0\big) := \sup_{\{p_t,q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}} \mathbb{E}\left[U\left(\overline{X}_T B_T, \overline{\Pi}_T B_T\right)\right]$ and the other mapping, $J_D : \mathbb{R}_+ \times \mathbb{R}_+ \times \mathcal{H}_{\mathcal{A}_{X_0,Y_0}} \to \mathbb{R}$, as $J_D\big(\widehat{X}_0, Z_0, \{\alpha_t, \lambda_t\}_{t\in[0,T]}\big) := \mathbb{E}\left[V\left(B_T^{-1} Z_T, \Pi_T\right)\right] + \widehat{X}_0 Z_0$.[17]

Proof. By means of the procedure in Klein and Rogers (2007), we derive the results in this theorem. For a formal proof that this is the dual, we refer the reader to Cvitanić and Wang (2001).[18] Following the strategy as described in the text, we find that

$$\begin{aligned} \mathcal{L} = \sup_{\{p_t,q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}, \overline{X}_T \in L^2(\Omega)} \mathbb{E}\Bigg[& U\left(\overline{X}_T B_T, \overline{\Pi}_T B_T\right) - \overline{X}_T Z_T + \overline{X}_0 Z_T \\ & + \int_0^T \left[\overline{S}_T - (1_N + \lambda_2) \odot \overline{S}_t\right]^\top Z_T \mathrm{d}p_t - \int_0^T \left[\overline{S}_T - (1_N - \lambda_1) \odot \overline{S}_t\right]^\top Z_T \mathrm{d}q_t \Bigg] \end{aligned} \tag{3.7}$$

specifies the so-called "reduced" Lagrangian. In consideration of $\{p_t, q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ implying $p_t, q_t \geq 0$, we find that $\mathbb{E}\left[\overline{S}_T - (1_N + \lambda_2) \odot \overline{S}_t \,\middle|\, \mathcal{F}_t\right] \leq 0$ must hold for all $t \in [0,T]$ as well as $\mathbb{E}\left[\overline{S}_T - (1_N - \lambda_1) \odot \overline{S}_t \,\middle|\, \mathcal{F}_t\right] \geq 0$ for all $t \in [0,T]$. Note that we apply Fubini's Theorem here. These CS conditions result in the restrictions on $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$, given by (3.4). Furthermore, it is clear that $\widehat{X}_T^{\text{opt}} = I\left(B_T^{-1} Z_T, \Pi_T\right)$ optimises terminal wealth, which in turn yields the $V\left(B_T^{-1} Z_T, \Pi_T\right)$ term in (3.5). As a result, we arrive at the following identity for $\mathcal{L}$, which establishes a function of $\{Z_t\}_{t\in[0,T]}$:

$$\mathcal{L} = \mathbb{E}\left[V\left(B_T^{-1} Z_T, \Pi_T\right)\right] + \widehat{X}_0 \mathbb{E}\left[Z_T\right]. \tag{3.8}$$

Consequently, $\mathcal{L}$ coincides with the objective in (3.5), and all the preceding restrictions as to $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$ are embedded within the set $\mathcal{H}_{\mathcal{A}_{X_0,Y_0}}$. Minimising $\mathcal{L}$ over $\{Z_t\}_{t\in[0,T]}$, such that the conditions implied by $\mathcal{H}_{\mathcal{A}_{X_0,Y_0}}$ are met, then spawns (3.5). Observe that $p_t = q_t = 0_N$ qualifies as a basic feasible solution, such that (3.6) holds. □

Remark 3.1. *In order to reconcile the results in Theorem 3.1 with the conventional specifications, let us derive an alternative, howbeit identical, formulation for the dual in*

[17] In order to erase potential ambiguity concerning the notation, let us stress that the the Lagrange multiplier, $Z_0 \in \mathbb{R}_+$, which optimises (3.5), ought to satisfy $\mathcal{H}(Z_0) = \mathbb{E}\left[I\left(B_T^{-1} Z_T, \Pi_T\right) B_T^{-1} Z_0^{-1} Z_T\right] = \widehat{X}_0 \mathbb{E}\left[Z_0^{-1} Z_T\right]$. In view of the fact that $\mathcal{H} : \mathbb{R}_+ \to \mathbb{R}_+$ is a monotonically increasing function, the optimal multiplier is given by $Z_0^{\text{opt}} = \mathcal{H}^{-1}\big(\widehat{X}_0 \mathbb{E}\left[Z_0^{-1} Z_T\right]\big)$, in which $\mathcal{H}^{-1} : \mathbb{R}_+ \to \mathbb{R}_+$ characterises the inverse of $\mathcal{H}$. Observe here that $Z_0^{-1} Z_T$ excludes the $Z_0 \in \mathbb{R}_+$ term entirely in its specification.

[18] The method at hand, as spelled out plainly in Rogers (2003), is of a sole heuristic nature. Concretely, this approach is not a formal proof; it can exclusively be utilised to arrive at the possibly "right" dual. Hence, for a formal proof, cf. Bouchard (2002), Kallsen and Li (2013), and Czichowsky et al. (2016).

(3.5). *To that end, introduce two semi-martingales, $\{Z_t^0\}_{t\in[0,T]}$ and $\{Z_t^1\}_{t\in[0,T]}$, that evolve according to $\mathrm{d}Z_t^0 = Z_t^0\left[\alpha_t \mathrm{d}t + \beta_t \mathrm{d}W_t\right]$ and $\mathrm{d}Z_t^1 = Z_t^1\left[\gamma_t \mathrm{d}t + \delta_t \mathrm{d}W_t\right]$ for some $Z_0^0, Z_0^1 \in \mathbb{R}_=$ $\alpha_t, \gamma_t \in \mathbb{D}^{1,2}\left([0,T]\right)$ and $\beta_t, \delta_t \in \mathbb{D}^{1,2}\left([0,T]\right)^N$. Rather than enforcing the constraint upon $\{\widehat{X}_t\}_{t\in[0,T]}$ in isolation, we employ $\{Z_t^0\}_{t\in[0,T]}$ and $\{Z_t^1\}_{t\in[0,T]}$ to enforce the separate constraints in (2.5) upon, respectively, $\{X_t\}_{t\in[0,T]}$ and $\{1_N^\top Y_t\}_{t\in[0,T]}$. Then, $\mathcal{L}$ reads:*

$$\begin{aligned}\mathcal{L} = \sup_{\{p_t,q_t\}_{t\in[0,T]}\in\mathcal{A}_{X_0,Y_0}} \mathbb{E}\Bigg[& V\left(Z_T^0, \Pi_T\right) + \widehat{X}_0 Z_0^0 + \int_0^T \left(\mu_t + \gamma_t 1_N + \sigma_t\delta_t\right)^\top Y_t Z_t^1 \mathrm{d}t \\ & + \int_0^T \left(\alpha_t - r_t\right) X_{Z,t}\mathrm{d}t + \int_0^T \left[Z_t^1 S_t - Z_t^0 S_t^A\right]^\top \mathrm{d}p_t - \int_0^T \left[Z_t^1 S_t - Z_t^0 S_t^B\right]^\top \mathrm{d}q_t\Bigg],\end{aligned} \tag{3.9}$$

which develops from inserting the solutions to the stochastic differential equations (SDE's) at $t = T$ for the processes $\{X_t Z_t^0\}_{t\in[0,T]}$ and $\{Y_t Z_t^1\}_{t\in[0,T]}$ into the objective in (3.1), as described in the text.[19] *Due to $p_t, q_t \geq 0$, we now find that $S_t Z_t^1 \geq S_t^B Z_t^0$ and $S_t Z_t^1 \leq S_t^A Z_t^0$ must hold for all $t \in [0,T]$. Additionally, since $X_{Z,t} = X_t Z_t^0 \in \mathbb{R}_+$ is true if one aims to optimise over $\{p_t, q_t\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$, we find that $\alpha_t \leq -r_t$ ought to be satisfied $\mathrm{d}t\otimes\mathbb{P}-\mathrm{a.e.}$ By the same token, i.e. since $Y_t Z_t^1 \in \mathbb{R}_+$, we know that $\mu_t + \gamma_t 1_N + \sigma_t\delta_t \leq 0_N$ must hold $\mathrm{d}t\otimes\mathbb{P}-\mathrm{a.e.}$ Note here that $\widehat{X}_T^{\mathrm{opt}} = I\left(Z_T^0, \Pi_T\right)$ and $Z_T^0 = Z_T^1$ are true. Now,*

$$\inf_{Z_T^0\in\mathcal{Z}(Z_0)} \mathbb{E}\left[V\left(Z_T^0, \Pi_T\right)\right] + \widehat{X}_0 Z_0^0, \tag{3.10}$$

represents the dual formulation corresponding to the problem in (3.1). Monotonicity arguments imply that $\alpha_t = -r_t$ and $1_N\gamma_t = -\mu_t - \sigma_t\delta_t$ are true. The set $\mathcal{Z}\left(Z_0\right)$ contains all Z_T^0 such that $B_t Z_t^0 = Z_0^0 \frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}\Big|_{\mathcal{F}_t} = \mathbb{E}\left[Z_T^0 \mid \mathcal{F}_t\right]$, $Z_T^0 = Z_T^1$ and $S_t^B \leq \frac{Z_t^1 S_t}{Z_t^0} \leq S_t^A$ hold for $\mathbb{Q}_\lambda \sim \mathbb{P}$ and a semi-martingale $\{Z_t^1\}_{t\in[0,T]}$ that satisfies $\mathbb{E}\left[Z_T^1 S_T \mid \mathcal{F}_t\right] = Z_t^1 S_t \;\forall\; t\in[0,T]$. More, let $\mathcal{H}\left(Z_0^0\right) = \mathbb{E}\big[I\left(Z_T^0, \Pi_T\right) Z_0^{0^{-1}} Z_T^0\big]$, then $Z_0^{0,\mathrm{opt}} = Z_0^{1,\mathrm{opt}} = \mathcal{H}^{-1}\big(\widehat{X}_0\big)$ holds, where $\mathcal{H}^{-1}$ is the inverse of $\mathcal{H}$. As a result of setting $\alpha_t = 0$ in (3.3), $\{Z_t\}_{t\in[0,T]}$ becomes a martingale, and all results in Theorem 3.1 coincide with those in the remark.[20]

The dual equivalent of the primal problem specification corresponding to the investment problem in markets with frictions, as described in Theorem 3.1, clearly differs from the well-known dual formulations in ordinary constrained markets, cf. for instance Cvitanić and Karatzas (1992) or Bardhan (1994). The central difference consists in the constraints that must be satisfied by the so-called shadow price process $\{Z_t\}_{t\in[0,T]}$, or by the in

[19] The SDE's for $\{X_t Z_t^0\}_{t\in[0,T]}$ and $\{1_N^\top Y_t Z_t^0\}_{t\in[0,T]}$ look as follows: $\mathrm{d}X_t Z_t^0 = Z_t^0 S_t^\top\left[\left(1-\lambda_1\right)\odot\mathrm{d}q_t - \left(1+\lambda_2\right)\odot\mathrm{d}p_t\right] + \left(r_t + \alpha_t\right) X_t Z_t^0 \mathrm{d}t + \beta^\top X_t Z_t^0 \mathrm{d}W_t$, and $\mathrm{d}1_N^\top Y_t Z_t^1 = Z_t^1 S_t^\top\left[\mathrm{d}p_t - q_t\right] + Z_t^1 Y_t^\top\left(S_t^{-1}\odot\mathrm{d}S_t\right) + \left[\gamma_t + \sigma_t\delta_t\right]^\top Y_t Z_t^1 \mathrm{d}t + 1_N^\top Y_t Z_t^1 \delta_t \mathrm{d}W_t$. Note that we do not use $\{B_t\}_{t\in[0,T]}$ *a priori* as numéraire.

[20] Even though it is omitted in the "proof" of Theorem 3.1, it can be shown that $\alpha_t \leq 0$ must be true $dt\otimes\mathbb{P}-\mathrm{a.e.}$, cf. Klein and Rogers (2007). Since $V: \mathbb{R}_+\times\mathbb{R}_+\to\mathbb{R}$ is decreasing in its first argument (and thus in $\{\alpha_t\}_{t\in[0,T]}$), it is optimal to set $\alpha_t = 0$. The equality of $\{B_t^{-1} Z_t\}_{t\in[0,T]}$ with $\{Z_t^0\}_{t\in[0,T]}$ is therefore latent in Theorem 3.1. For the "conventional" specifications of the dual problem, we refer the reader to e.g. Dai et al. (2009), Altarovici et al. (2017), and Bayraktar and Yu (2019).

that regard implicit processes $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$, as displayed in equation (3.4). Namely, contrary to the case of markets with convex trading restrictions, these necessary conditions do not impose direct, concrete restrictions on the relevant $\{\alpha_t\}_{t\in[0,T]}$ and $\{\lambda_t\}_{t\in[0,T]}$ processes. The nature of these restrictions in conjunction with this absence of explicit restraints as to the preceding processes complicate the construction of an auxiliary (fictitious) market, which is ordinarily a rather straightforward procedure, see e.g. Cvitanić and Karatzas (1993), Karatzas and Kou (1996), or Karatzas and Kou (1998).

3.2 Auxiliary Frictionless Market

We continue the line of subsection 3.1, and spell out the artificial, auxiliary financial market model that is implied by the dual problem in Theorem 3.1. This artificial market, say $\widehat{\mathcal{M}}$, eliminates the frictions that exist in its original counterpart, $\mathcal{M}$. As a result of this implied absence of transaction costs, an analytical recovery of optimal solutions for the investment policies (in $\widehat{\mathcal{M}}$) is significantly facilitated. These investment policies and the auxiliary model institute essential grounds for our approximate method.

In line with the literature on convex duality in markets with frictions[21], we introduce

$$\tilde{S}_t = \mathbb{E}\left[\frac{Z_T}{\mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]}\frac{B_t}{B_T}S_T \,\middle|\, \mathcal{F}_t\right] = \frac{Z_t^1 S_t}{Z_t^0}, \quad \forall\, t \in [0,T], \tag{3.11}$$

as the so-called shadow price process. This process springs from rewriting the $\mathbb{Q}_\lambda$-expectation in Theorem 3.1. In line with the CS conditions in (3.4), we therefore must require that $\left\{\tilde{S}_t\right\}_{t\in[0,T]}$ satisfies $S_t^B \leq \tilde{S}_t \leq S_t^A$ for all $t \in [0,T]$. The shadow price process characterises the $\mathbb{R}^N$-valued price evolution of the stock in the artificial market without frictions, $\mathcal{M}$. Due to the indefiniteness of the $\left\{\mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]^{-1} Z_T\right\}_{t\in[0,T]}$ process, it is clear that $\left\{\tilde{S}_t\right\}_{t\in[0,T]}$ is likewise subject to an unspecified input, which is inherent in the fact that $\alpha_t \in \mathbb{D}^{1,2}\left([0,T]\right)$ and $\lambda_t \in \mathbb{D}^{1,2}\left([0,T]\right)^N$ are not identified.

These properties enable us to particularise that the following assets compose $\widehat{\mathcal{M}}$[22]:

$$\frac{\mathrm{d}\tilde{B}_t}{\tilde{B}_t} = r_t \mathrm{d}t, \quad \text{and} \quad \frac{\mathrm{d}\tilde{S}_{i,t}}{\tilde{S}_{i,t}} = \left[\mu_{\tilde{S},i,t}\mathrm{d}t + \sigma_{\tilde{S},i,t}^\top \mathrm{d}W_t\right], \tag{3.12}$$

in which $\tilde{B}_0 = 1$, $\tilde{S}_{i,0} = 1_N \frac{Z_0^1}{Z_0^0}$ holds for all $i = 1,\ldots,N$. We define an $\mathbb{R}^N$-valued process $\left\{\mu_{\tilde{S},t}\right\}_{t\in[0,T]}$ as the vector with entries $\mu_{\tilde{S},i,t}$ for all $i = 1,\ldots,N$; similarly, we define the

[21] In expounding the artificial market $\widehat{\mathcal{M}}$, we will heavily depend on the findings in the studies by Kühn and Stroh (2010), Benedetti et al. (2013), and Kallsen and Muhle-Karbe (2017). Even though our focus is on the accompanying analytical features that accentuate the artificial market, we aim to explicitly couple the dual to $\widehat{\mathcal{M}}$, in an attempt to justify the principal idea of our approximating routine.

[22] We specify the auxiliary market, $\widehat{\mathcal{M}}$, on the same probability triplet, $(\Omega, \mathcal{F}_t, \{\mathcal{F}_t\}_{t\in[0,T]}, \mathbb{P})$, that we employed to detail the original market model, $\mathcal{M}$; also, we rely on the same notation, unless stated differently. As a result, we work with the same N-dimensional Brownian motion $\{W_t\}_{t\in[0,T]}$. Note that the latter is logical, since (3.11) does obviously not accommodate additional drivers of uncertainty.

$\mathbb{R}^{N\times N}$-valued $\{\sigma_{\tilde{S},t}\}_{t\in[0,T]}$ process as a matrix with rows $\sigma_{\tilde{S},i,t}$ for $i=1,\ldots,N$. Like before, we assume that both processes are $\mathcal{F}_t$-measurable and satisfy $\|\mu_{\tilde{S},t}\|_{\mathbb{R}^N}, \operatorname{diag}(\sigma_{\tilde{S},t}\sigma_{\tilde{S},t}^\top) \in L^1([0,T];\mathbb{R}^N)$. Moreover, we postulate that $\phi^\top\sigma_{\tilde{S},t}\sigma_{\tilde{S},t}^\top\phi \geq \epsilon\|\phi\|_{\mathbb{R}^N}$ holds for all $\phi\in\mathbb{R}^N$ and some $\epsilon>0$. We stress here that both $\{\mu_{\tilde{S},t}\}_{t\in[0,T]}$ and $\{\sigma_{\tilde{S},t}\}_{t\in[0,T]}$ may be non-unique, and must obey to the following form: $\mu_{\tilde{S},t}=\left[r_t-\widehat{\alpha}_t^\top\widehat{\lambda}_t\right]1_N-\sigma_t\widehat{\alpha}_t$ and $\sigma_{\tilde{S},t}=\sigma_t+\widehat{\lambda}_t^\top\odot 1_{N\times N}$, for unspecified $\widehat{\alpha}_t,\widehat{\lambda}_t\in\mathbb{D}^{1,2}([0,T])^N$, inherent in equation (3.11).[23]

The artifical environment implied by $\tilde{B}_t$ and $\tilde{S}_t$ induces a stochastic deflator process,

$$\tilde{B}_t^{-1}\frac{\mathbb{E}\left[Z_T\mid\mathcal{F}_t\right]}{\mathbb{E}\left[Z_T\right]}=\exp\left\{-\int_0^t r_s\mathrm{d}s-\frac{1}{2}\int_0^t\|\widehat{\alpha}_s\|_{\mathbb{R}^N}^2\mathrm{d}s+\int_0^t\widehat{\alpha}_s^\top\mathrm{d}W_s\right\},\tag{3.13}$$

in which $\widehat{\alpha}_t\in\mathbb{D}^{1,2}([0,T])^N$ is non-unique, such that $\mathbb{E}\left[\exp\left(\frac{1}{2}\int_0^T\|\widehat{\alpha}_s\|_{\mathbb{R}^N}^2\mathrm{d}s\right)\right]<\infty$ holds, and where $\{\tilde{B}_t\}_{t\in[0,T]}$ serves as numéraire. Obviously, $\tilde{B}_t^{-1}\mathbb{E}\left[Z_T\right]^{-1}\mathbb{E}\left[Z_T\mid\mathcal{F}_t\right]$ identifies with $Z_0^{-1}Z_t^0$, cf. Remark 3.1. For simplicity, we choose to work with $Z_0^{0^{-1}}Z_t^0$. We may employ $Z_0^{0^{-1}}Z_t^0$ to price assets in $\widehat{\mathcal{M}}$: straightforwardly, $Z_0^{0^{-1}}Z_t^0\tilde{S}_t$ identifies a $\mathbb{P}$-martingale. The deflator process is decomposable as $Z_0^{-1}Z_t^0=\tilde{B}_t^{-1}\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}$, where $\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}$ is such that $\mathbb{Q}_\lambda\sim\mathbb{P}$ and $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ defines the market price of risk. Here, $\mathbb{Q}_\lambda$ details the pricing measure in $\widehat{\mathcal{M}}$, under which $\mathrm{d}W_t^{\mathbb{Q}_\lambda}=\mathrm{d}W_t-\widehat{\alpha}_t\mathrm{d}t$ describes a standard Brownian motion.

Correspondingly, the dynamic wealth process of the investor in $\widehat{\mathcal{M}}$ must submit to

$$\mathrm{d}\widehat{X}_{\lambda,t}=\widehat{X}_{\lambda,t}\left[r_t\mathrm{d}t+\pi_{\lambda,t}^\top\left\{\sigma_t+\widehat{\lambda}_t^\top\odot 1_{N\times N}\right\}\left(\mathrm{d}W_t-\widehat{\alpha}_t\mathrm{d}t\right)\right],\tag{3.14}$$

where $\widehat{X}_{\lambda,0}=\widehat{X}_0\in\mathbb{R}_+$, and $\{\pi_{\lambda_t}\}_{t\in[0,T]}$ defines the $\mathcal{F}_t$-measurable, $\mathbb{R}^N$-valued portfolio process, containing the fractions of $\{\widehat{X}_t\}_{t\in[0,T]}$ that the agent allocates to $\{\tilde{S}_t\}_{t\in[0,T]}$. We call a trading strategy admissible in $\widehat{\mathcal{M}}$, if it satisfies $\pi_{\lambda,t}^\top\left(\sigma_t+\widehat{\lambda}_t^\top\odot 1_{N\times N}\right)\widehat{\alpha}_t\in L^1([0,T])$, $\pi_{\lambda,t}^\top\left(\sigma_t+\widehat{\lambda}_t^\top\odot 1_{N\times N}\right)\in L^2([0,T])$, in addition to $\widehat{X}_{\lambda,t}\geq 0$ for all $t\in[0,T]$. We denote the set of all such investment strategies by $\widehat{\mathcal{A}}_{X_0,Y_0}$. In contrast to $\mathcal{M}$, the artificial market $\widehat{\mathcal{M}}$ entirely excludes the presence of transaction costs. Therefore, noting that $\tilde{S}_t$ trades in $\widehat{\mathcal{M}}$ if $S_t^B\leq\tilde{S}_t\leq S_t^B$, maximal utility in $\mathcal{M}$ cannot exceed that in $\widehat{\mathcal{M}}$.

Proposition 3.2. *In conformity with* (3.5), *the investment problem in* $\widehat{\mathcal{M}}$ *reads*

$$\begin{aligned}\sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}}\quad&\mathbb{E}\left[U\left(\widehat{X}_{\lambda,T},\Pi_T\right)\right]\\ \text{s.t.}\quad&\mathrm{d}\widehat{X}_{\lambda,t}=\widehat{X}_{\lambda,t}\left[r_t\mathrm{d}t+\pi_{\lambda,t}^\top\left\{\sigma_t+\widehat{\lambda}_t^\top\odot 1_{N\times N}\right\}\left(\mathrm{d}W_t-\widehat{\alpha}_t\right)\right],\end{aligned}\tag{3.15}$$

[23] To establish $\widehat{\mathcal{M}}$, we rely on the findings in Remark 3.1. In particular, as displayed in equation (3.11), we find that $Z_t^0=B_t^{-1}\mathbb{E}\left[Z_T\mid\mathcal{F}_t\right]$ and $Z_t^1S_t=\mathbb{E}\left[B_T^{-1}Z_TS_T\mid\mathcal{F}_t\right]$ must be true; the latter process is a $\mathbb{P}$-martingale with respect to $\{\mathcal{F}_t\}_{t\in[0,T]}$, the former induces a state price density, $B_t^{-1}\frac{d\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}\big|_{\mathcal{F}_t}$. Hence, so as to unravel the structure of $\{\tilde{S}_t\}_{t\in[0,T]}$, it is sufficient to analyse $\{Z_t^{0^{-1}}Z_t^1S_t\}_{t\in[0,T]}$. The SDE of the preceding term equates to $\mathrm{d}Z_t^{0^{-1}}Z_t^1S_t=Z_t^{0^{-1}}Z_t^1S_t[\odot\left(\{r_t+\beta_t^\top(\beta_t-\delta_t)\}1_N-\sigma_t\beta_t\right)\mathrm{d}t+(\{\delta_t-\beta_t\}^\top\odot 1_{N\times N}+\sigma_t)\mathrm{d}W_t]$. Letting $\widehat{\lambda}_t:=\delta_t-\beta_t$ and $\widehat{\alpha}_t:=\beta_t$, we acquire the process in (3.12).

in which we are able to rewrite the dynamic budget constraint in terms of the shadow price process, $\{\tilde{S}_t\}_{t\in[0,T]}$, *in the following way:* $\mathrm{d}\widehat{X}_{\lambda,t} = \widehat{X}_{\lambda,t}\big[\left(1 - 1_N^\top \pi_{\lambda,t}\right)\frac{\mathrm{d}\tilde{B}_t}{\tilde{B}_t} + \pi_{\lambda,t}^\top\{\tilde{S}_t^{-1}\odot \mathrm{d}\tilde{S}_t\}\big]$. *Consequently, the dual problem specification in* (3.5) *is equivalent to*

$$\inf_{\{\tilde{S}_t\}_{t\in[0,T]}\in\mathcal{S}} \ \sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}} \mathbb{E}\left[U\left(\widehat{X}_{\lambda,T},\Pi_T\right)\right], \tag{3.16}$$

where $\mathcal{S}$ *contains all strictly positive semi-martingales* $\{\tilde{S}_t\}_{t\in[0,T]}$ *such that* $\tilde{S}_T = S_T$, $S_t^B \leq \tilde{S}_t \leq S_t^A\ \forall\, t\in[0,T]$, *and* $\tilde{S}_t Z_t^0$ *is a* $\mathbb{P}$*-martingale. As a an evident result, the dual* (3.16) *implies that* $\sup_{\{p_t,q_t\}_{t\in[0,T]}\in\mathcal{A}_{X_0,Y_0}} \mathbb{E}\big[U\big(\widehat{X}_T,\Pi_T\big)\big] \leq \sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}} \mathbb{E}\big[U\big(\widehat{X}_{\lambda,T},\Pi_T\big)\big]$ *for all* $\{\tilde{S}_t\}_{t\in[0,T]}\in\mathcal{S}$. *By strong duality, minimising the latter objective over all* $\{\tilde{S}_t\}_{t\in[0,T]}\in\mathcal{S}$ *transforms the preceding inequality into an equality. As a result,*

$$J_P\left(\widehat{X}_0\right) = \inf_{\{\widehat{\alpha}_t,\widehat{\lambda}_t\}_{t\in[0,T]}\in\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}} \widehat{J}_D\left(\widehat{X}_0,\left\{\widehat{\alpha}_t,\widehat{\lambda}_t\right\}_{t\in[0,T]}\right), \tag{3.17}$$

must be true, in which $\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$ *contains all the drift and diffusion terms,* $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ *and* $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ *respectively, associated with the shadow price process* $\{\tilde{S}_t\}_{t\in[0,T]}$ *such that* $\{\tilde{S}_t\}_{t\in[0,T]}\in\mathcal{S}$. *We define* $\widehat{J}_D\big(\widehat{X}_0,\{\widehat{\alpha}_t,\widehat{\lambda}_t\}_{t\in[0,T]}\big) := \sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}} \mathbb{E}\big[U\big(\widehat{X}_{\lambda,T},\Pi_T\big)\big]$. *The formulation in* (3.17) *highlights the aforementioned drift and diffusion terms.*[24]

Proof. We give a brief and rough sketch of the proof, wherein we predominantly aim to stress the analogy between the dual specification (3.5) and the dynamic optimisation problem in (3.15). For a formal proof, we refer the reader to any of the earlier cited papers on the relevant subject, and to section 4 of Bayraktar and Yu (2019) in specific. Let us re-consider the dual problem formulation in (3.10). Therein, observe that we are able to rewrite the isolated objective, with exclusion of the "supremum" term, as follows:

$$\begin{aligned}\mathbb{E}\left[V\left(Z_T^0,\Pi_T\right)\right] + \widehat{X}_0 Z_0^0 &= \mathbb{E}\left[U\left(I\left(Z_T^0,\Pi_T\right),\Pi_T\right)\right] \\ &\quad - \left(\mathbb{E}\left[I\left(Z_T^0,\Pi_T\right)Z_T^0\right] - \widehat{X}_0 Z_0^0\right) = \mathbb{E}\left[U\left(\widehat{X}_{\lambda,T}^{\text{opt}},\Pi_T\right)\right].\end{aligned} \tag{3.18}$$

In Remark 3.1, as well as in Theorem 3.1, we have argued that the optimal "Lagrange multiplier", $Z_0^{0,\text{opt}}\in\mathbb{R}_+$, ought to be chosen such that $Z_0^{0,\text{opt}} = \mathcal{H}^{-1}\big(\widehat{X}_0\big)$ is true, where we define $\mathcal{H}:\mathbb{R}_+\to\mathbb{R}_+$ as $\mathcal{H}\left(Z_0^0\right) = \mathbb{E}\left[I\left(Z_T^0,\Pi_T\right)Z_T^0\right]$. This clarifies the second equality in (3.18). Now, Theorem 3.1 shows that the dual objective is equal to the Lagrangian corresponding to the problem in (3.15), explaining the remaining equality. In combination with the former outcome, the dual in (3.16) follows, and logically (3.17) ensues. □

[24] In spite of the fact that in both differently characterised, though identical, dual specifications, (3.5) and (3.10), we are essentially minimising the objective over a single control (the diffusion coefficients of the respective deflator processes, $\{Z_t\}_{t\in[0,T]}$ and $\{Z_t^0\}_{t\in[0,T]}$), the two controls in Proposition 3.2, $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ reconcile through the conditions implied by (3.11), as classified among $\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$.

Remark 3.2. *So as to underscore the automated choice for a least-favourable artificial market* $\widehat{\mathcal{M}}$*, as notably revealed by the delineation of the dual in Proposition 3.2, we wish to shift the focus from the shadow price, i.e. the fictitiously traded frictionless stock process* $\{\tilde{S}_t\}_{t\in[0,T]}$*, to the pricing measure,* $\mathbb{Q}_\lambda \sim \mathbb{P}$*. The alternative dual would read*

$$\inf_{\left(\mathbb{Q}_\lambda, \{\widehat{\lambda}_t\}_{t\in[0,T]}\right)\in\mathcal{M}_{\mathbb{Q}}} \sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}} \mathbb{E}\left[U\left(\widehat{X}_{\lambda,T}, \Pi_T\right)\right], \tag{3.19}$$

wherein $\mathcal{M}_{\mathbb{Q}}$ *contains all pairs* $\mathbb{Q}_\lambda \sim \mathbb{P}$ *and* $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ *such that, for* $\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}\big|_{\mathcal{F}_t} = \mathbb{E}\left[\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}} \mid \mathcal{F}_t\right]$ *and* $\{Z_t^1 S_t\}_{t\in[0,T]}$*, these satisfy* $\tilde{B}_t^{-1}\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}} = Z_T^1 S_T$ *and* $S_T^B \leq \tilde{B}_t \frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}\big|_{\mathcal{F}_t}^{-1} Z_t^1 S_t \leq S_t^A$ *for all* $t \in [0,T]$*. To substantiate* $\{\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}\big|_{\mathcal{F}_t}\}_{t\in[0,T]}$*'s effect on* $\widehat{X}_{\lambda,T}$*, observe that we may rewrite the budget constraint in* (3.15) *as* $\mathrm{d}\widehat{X}_{\lambda,t} = \widehat{X}_{\lambda,t}\left[r_t \mathrm{d}t + \pi_{\lambda,t}^\top\{\sigma_t + \widehat{\lambda}_t^\top \odot 1_{N\times N}\}\mathrm{d}W_t^{\mathbb{Q}_\lambda}\right]$*. Having regard to the fact that each* $\left(\mathbb{Q}_\lambda, \{\widehat{\lambda}_t\}_{t\in[0,T]}\right) \in \mathcal{M}_{\mathbb{Q}}$ *prompts a unique artificial market, it is unmistakable from* (3.19) *that the dual opts for the least-favourable* $\widehat{\mathcal{M}}$*.*[25]

In compliance with the dual in (3.5), Proposition 3.2 sets out the optimal dynamic investment problem in $\widehat{\mathcal{M}}$. Most importantly, the previous proposition shows that the dual objective spells out a Lagrangian functional that harmonises with the aforementioned problem in $\widehat{\mathcal{M}}$. The absence of transaction costs in this problem betokens that we are able to acquire analytical solutions for $\{\pi_{\lambda,t}\}_{t\in[0,T]}$ that are completely identified up to the shadow price processes, specific to $\{\tilde{S}_t\}_{t\in[0,T]}$. In turn, the dual determines the shadow prices in a manner such that these optimal decision rules are "pruned" to optimality and the admissibility region in $\mathcal{M}$. Subsequently, we exemplify this mechanism.

Example 3.1. *(No Transaction Costs) Assume that* $\mathcal{M}$ *omits frictions:* $\lambda_1 = \lambda_2 = 0_N$*. In that case, we acquire from Proposition 3.2 that* $\tilde{S}_t = S_t$ *holds for all* $t \in [0,T]$*. Furthermore, from the equality in* (3.11) *that essentially gives rise to the* $\mathcal{S}$ *set, we find that*

$$\mathbb{E}\left[B_T^{-1} Z_T S_T \mid \mathcal{F}_t\right] = B_t^{-1}\mathbb{E}\left[Z_T \mid \mathcal{F}_t\right] S_t \quad \Leftrightarrow \quad Z_t^0 = Z_t^1 \ \forall\ t \in [0,T], \tag{3.20}$$

must be true. The first equation suggests that $\{Z_t\}_{t\in[0,T]}$ *defines a Radon-Nikodym derivative that induces a pricing measure in* $\mathcal{M}$*, in which* $\alpha_t = 0$*, and where the market price of risk equates to* $\lambda_t = -\sigma_t^{-1}(\mu_t - r_t 1_N)$ *for all* $t \in [0,T]$*. The second equality evinces that* $\widehat{\lambda}_t = \delta_t - \beta_t = 0_N$ *and* $\gamma_t = -r_t$ *ought to hold for all* $t \in [0,T]$*. Moreover, from the requirement that* $\tilde{S}_t\tilde{B}_t^{-1}Z_t^0$ *must define a* $\mathbb{P}$*-martingale, we derive that* $\widehat{\alpha}_t = \beta_t = \delta_t = -\sigma_t^{-1}(\mu_t - r_t 1_N)$ *for all* $t \in [0,T]$*. As a consequence,* $\mathcal{M} = \widehat{\mathcal{M}}$*, given that* $\{\tilde{S}_t\}_{t\in[0,T]} \in \mathcal{S}$ *holds.*

[25] As well as in Proposition 3.2, we would like to stress that, irrespective of the fact that $\{\widehat{X}_{\lambda,t}\}_{t\in[0,T]}$ (as specified in Remark 3.2) incorporates apart from the non-unique standard Brownian motion under $\mathbb{Q}_\lambda$, the $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ term, the latter term is related to the Radon-Nikodym derivative $\frac{\mathrm{d}\mathbb{Q}_\lambda}{\mathrm{d}\mathbb{P}}$ by means of the conditions integrated in $\mathcal{M}_{\mathbb{Q}}$. Consider Czichowsky et al. (2014), Czichowsky et al. (2017) for a technical analysis in that regard; Herczegh et al. (2015) may be consulted for a concrete example.

Example 3.2. *(Non-traded Asset) Suppose that the* i^{th} *element of a vector, say* ϕ_t*, is denoted by* $\phi_{i,t}$*. Then, assume that* $\sigma_{N,i,t} = 0$ *holds for all* $i = 1, \ldots, N-1$*: only* $S_{N,t}$ *is driven by* $W_{N,t}$*, accompanied by* $\sigma_{N,N,t}$ *as the volatility process. Consider the case wherein* $\lambda_{i,1} = \lambda_{i,2} = 0$ *for* $i = 1, \ldots, N-1$*,* $\lambda_{N,1} \to 1$ *and* $\lambda_{N,2} \to \infty$*, i.e. a situation in which* $\mathcal{M}$ *excludes frictions and wherein* $\{S_{N,t}\}_{t\in[0,T]}$ *is not traded. Then,* $\forall\ t \in [0,T]$

$$\mathbb{E}\left[\frac{B_T^{-1} Z_T}{B_t^{-1}\mathbb{E}\left[Z_T \mid \mathcal{F}_t\right]} S_{i,T} \,\middle|\, \mathcal{F}_t\right] = S_{i,t} \quad \Leftrightarrow \quad \frac{Z_t^1 S_{i,t}}{Z_t^0} = S_{i,t}, \tag{3.21}$$

ought to be true, where $i = 1, \ldots, N-1$*, in addition to* $\mathbb{E}^{\mathbb{Q}_\lambda}\left[B_T S_{N,T} \mid \mathcal{F}_t\right] \in \mathbb{R}_+$ *and* $\frac{Z_t^1 S_{N,t}}{Z_t^0} \in \mathbb{R}_+$*. As to* $\{Z_t\}_{t\in[0,T]}$*, like in Example 3.1, it must define a Radon-Nikodym derivative that renders a pricing measure in* $\mathcal{M}$*, wherein* $\alpha_t = 0$ *and where the market price of risk,* $\lambda_t = \left[\lambda_{1:N-1,t}, \lambda_{N,t}\right]^\top$*, equates to* $\lambda_{1:N-1,t} = -\sigma^{-1}_{(1:N-1)\times(1:N-1),t}\left(\mu_{1:N-1,t} - r_t 1_{N-1}\right)$ *for a yet unspecified* $\lambda_{N,t} \in \mathbb{D}^{1,2}\left([0,T]\right)$*, and all* $t \in [0,T]$*. Accordingly, for* $\{Z_t^0\}_{t\in[0,T]}$ *and* $\{Z_t^1\}_{t\in[0,T]}$*,* $\widehat{\lambda}_{i,t} = \delta_{i,t} - \beta_{i,t} = 0$*,* $\gamma_t = -r_t$*, and* $\widehat{\alpha}_{1:N-1,t} = \lambda_{1:N-1,t}$ *hold for* $i = 1, \ldots, N-1$ *and all* $t \in [0,T]$*; i.e.* $\widehat{\alpha}_{N,t}, \widehat{\lambda}_{N,t} \in \mathbb{D}^{1,2}\left([0,T]\right)$ *are unspecified. Therefore,*

$$\begin{gathered}\inf_{\widehat{\alpha}_{N,t},\widehat{\lambda}_{N,t}\in\mathbb{D}^{1,2}([0,T])} \ \sup_{\{\pi_{\lambda,t}\}_{t\in[0,T]}\in\widehat{\mathcal{A}}_{X_0,Y_0}} \mathbb{E}\left[U\left(\widehat{X}_{\lambda,T}, \Pi_T\right)\right], \quad \text{where} \\ \widehat{\lambda}_{1:N-1,t} = 0_{N-1}, \quad \text{and} \quad \widehat{\alpha}_{1:N-1,t} = -\sigma^{-1}_{(1:N-1)\times(1:N-1),t}\left(\mu_{1:N-1,t} - r_t 1_{N-1}\right),\end{gathered} \tag{3.22}$$

from (3.16)*, or equivalently* $\inf_{\lambda_{N,t}\in\mathbb{D}^{1,2}([0,T]),Z_0\in\mathbb{R}_+} \mathbb{E}\left[V\left(B_T^{-1}Z_T, \Pi_T\right)\right] + \widehat{X}_0\mathbb{E}\left[Z_T\right]$ *from* (3.5)*, define the identical dual problems for this example, in which we assume that all the foregoing conditions are met. In agreement with* (3.22)*, the investor in* $\widehat{\mathcal{M}}$ *chooses* $\{\pi_{\lambda,t}\}_{t\in[0,T]}$ *as if* $\{S_{N,t}\}_{t\in[0,T]}$ *is traded, albeit in the form of* $\left\{\tilde{S}_{N,t}\right\}_{t\in[0,T]}$*. Thereafter, it can be shown that the dual chooses* $\widehat{\alpha}_{N,t} \in \mathbb{D}^{1,2}\left([0,T]\right)$ *such that* $\pi_{N,\lambda,t} = 0$ *holds* $\mathrm{d}t \otimes \mathbb{P}-$a.e.[26] *Since* $\left\{\widehat{\alpha}_{N,t}\right\}_{t\in[0,T]}$ *alone is in possession of enough freedom to arrive at* $\pi_{N,\lambda,t} = 0$ $\mathrm{d}t \otimes \mathbb{P}-$a.e., $\widehat{\lambda}_{N,t} \in \mathbb{D}^{1,2}\left([0,T]\right)$ *plays a redundant role in this case.*[27]

Although these two "extreme" examples exclude transaction costs, they adequately demonstrate how the mechanism underlying the dual is able to connect the artificial market, $\widehat{\mathcal{M}}$, with its original counterpart, $\mathcal{M}$. That is, the mere task of the dual is to determine the shadow price process, $\left\{\tilde{S}_t\right\}_{t\in[0,T]} \in \mathcal{S}$ in a least-favourable, such that the analytically available optimal controls in $\widehat{\mathcal{M}}$ are admissible and optimal in $\mathcal{M}$. Nevertheless, contrary

[26]We rely for this example on section 3.3 in Kamma and Pelsser (2019). Therein, the authors show that $\pi^{\text{opt}}_{\lambda,t} = \sigma^{\top^{-1}}_{\tilde{S},t}\left(\widehat{X}_{\lambda,t} Z_0^{0^{-1}} Z_t^0\right)^{-1}\left\{\mathbb{E}\left[\mathcal{D}_t^W \widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0 \mid \mathcal{F}_t\right] - \widehat{\alpha}_t \mathbb{E}\left[\widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0 \mid \mathcal{F}_t\right]\right\}$ is the optimal strategy in $\widehat{\mathcal{M}}$. Moreover, they prove that the $\widehat{\alpha}^{\text{opt}}_{N,t}$ that optimises (3.22) is derivable from $\widehat{\alpha}^{\text{opt}}_{N,t} = \mathbb{E}\left[\widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0 \mid \mathcal{F}_t\right]^{-1}\mathbb{E}\left[\mathcal{D}_t^{W_N} \widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0 \mid \mathcal{F}_t\right]$. As a result, $\pi_{N,\lambda,t} = 0$ holds, irrespective of $\widehat{\lambda}_{N,t}$'s specification, and the latter does not affect $\pi^{\text{opt}}_{1:N-1,\lambda,t}$. Due to the fact that $\widehat{\lambda}_{N,t}$ may be omitted from (3.22), this problem identifies with the dual formulation in Cvitanić and Karatzas (1992).

[27]For some vector, say ϕ_t, $\phi_{1:i,t}$ is the vector containing the first i elements of the original analogue, ϕ_t. For a matrix, say σ_t, $\sigma_{(1:i)\times(1:i),t}$ contains the first i rows, of equal length i, of the original matrix, σ_t.

to these special cases, in the presence of market frictions, the characterisation of the shadow price process by the dual is non-trivial. Namely, despite the possible existence of such a process, cf. Kallsen and Muhle-Karbe (2011) or Gu et al. (2017), these are typically not available in closed-form or require extensive mathematical derivations, see Gerhold et al. (2013) or Gerhold et al. (2014). However, noting that $\{\tilde{S}_t\}_{t\in[0,T]} \in \mathcal{S}$ only trades in $\mathcal{M}$ when $\tilde{S}_t = S_t^B, S_t^A$, and that $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$ suggests what the optimal portfolio looks like, we can make an educated "guess" about the analytical structure of π_t^{opt}.

3.3 Optimality in Auxiliary Market

In order to disentangle the analytical configuration that characterises the optimal portfolio in the auxiliary market, we continue this section by solving the optimal dynamic investment problem in $\widehat{\mathcal{M}}$, as provided in (3.15). To that end, we observe that the problem rules out trading restrictions and frictions, due to which we are able to fall back on the well-known martingale method, cf. Pliska (1986), Karatzas et al. (1987), and Cox and Huang (1989, 1991), to transform the dynamic problem into its static variational counterpart. The consequential specification embeds instead of a dynamic constraint, a static one:

$$\begin{aligned} \sup_{\widehat{X}_{\lambda,T}\in L^2(\Omega)} \quad & \mathbb{E}\left[U\left(\widehat{X}_{\lambda,T}, \Pi_T\right)\right] \\ \text{s.t.} \quad & \mathbb{E}\left[\widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0\right] \leq \widehat{X}_0. \end{aligned} \tag{3.23}$$

Expressly, the individual now faces a static problem, in which the sole control concerns terminal wealth, $\widehat{X}_{\lambda,T}$. Appropriately solving (3.23) occasions a closed-form expression for optimal horizon wealth, $\widehat{X}_{\lambda,T} = I\big(\eta Z_0^{0^{-1}} Z_T^0, \Pi_T\big)$ for some $\eta \in \mathbb{R}_+$, which ensures that the static budget constraint constraints is binding. Hereinafter, replicating arguments inclusive of the dynamic budget constraint in (3.15) enable us to collect the corresponding optimal allocation to assets, $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$.[28] Let us observe that the static constraint guarantees that the acquired optimal policies are attainable and self-financing.[29]

Proposition 3.3. *Consider the optimisation problem* (3.23)*, for a finite-horizon investor with wealth dynamics given by* (3.14)*. Then, optimal horizon wealth reads as*

$$\widehat{X}_{\lambda,T}^{\text{opt}} = I\left(\eta^{\text{opt}} Z_0^{0^{-1}} Z_T^0, \Pi_T\right), \tag{3.24}$$

[28] To be more precise, let $\mathcal{D}_t^W$ denote the so-called Malliavin derivative kernel in the W_t-direction, see e.g. Nualart (2006), and assume that $\widehat{X}_T^{\text{opt}}$ represents optimal terminal wealth. Then, the harmonious optimal dynamic portfolio policies, i.e. those that hedge $\widehat{X}_{\lambda,T}^{\text{opt}}$ in $\widehat{\mathcal{M}}$, emerge from the next identity (the Clark-Ocone formula): $\widehat{X}_{\lambda,T} Z_0^{0^{-1}} Z_T^0 = \widehat{X}_0 + \int_0^T \mathbb{E}\big[\mathcal{D}_t^W X_{\lambda,T}^{\text{opt}} Z_0^{0^{-1}} Z_T^0 \mid \mathcal{F}_t\big]^\top \mathrm{d}W_t$.

[29] We supply the static problem in (3.23) with inclusion of the Z_T^0 process rather the $B_T^{-1} Z_T$ process, with the object of staying as close as possible to $\widehat{\mathcal{M}}$'s outline in section 3.2. If we would utilise the process $B_T^{-1} Z_T$ instead, the static constraint in (3.23) would read $\mathbb{E}\big[\widehat{X}_{\lambda,T} B_T^{-1} \mathbb{E}\left[Z_T\right]^{-1} Z_T\big] \leq \widehat{X}_0$. Note that monotonicity arguments suffice to conclude that the static constraint ought to bind at the optimum.

where $\eta^{\mathrm{opt}} \in \mathbb{R}_+$ *denotes the Lagrange multiplier, such that* $\widehat{X}^{\mathrm{opt}}_{\lambda,T}$ *satisfies the budget constraint in* (3.23)*,* $\eta^{\mathrm{opt}} = \mathcal{H}^{-1}\big(\widehat{X}_0\big)$*. Furthermore, we define:*

$$\mathcal{R}_{1,T} = -\widehat{X}^{\mathrm{opt}}_{\lambda,T}\frac{U''_X\left(\widehat{X}^{\mathrm{opt}}_{\lambda,T},\Pi_T\right)}{U'_X\left(\widehat{X}^{\mathrm{opt}}_{\lambda,T},\Pi_T\right)}, \quad \text{and} \quad \mathcal{R}_{2,T} = \frac{I'_Y\left(\eta Z_0^{0^{-1}}Z_T^0,\Pi_T\right)\Pi_T}{\widehat{X}^{\mathrm{opt}}_{\lambda,T}}, \tag{3.25}$$

as the relevant RRA coefficients, in which $I'_Y : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ *defines the derivative of* $I : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ *in its second argument.*[30] *The market fair value of* $\widehat{X}^{\mathrm{opt}}_{\lambda,T}$ *in* $\widehat{\mathcal{M}}$ *equals* $\widehat{X}^{\mathrm{opt}}_{\lambda,t} = \mathbb{E}\big[\frac{Z_T^0}{Z_t^0}\widehat{X}^{\mathrm{opt}}_{\lambda,T} \mid \mathcal{F}_t\big]$ *for all* $t \in [0,T]$*. The optimal allocation to assets,* $\{\pi^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$*, corresponding to* $\widehat{X}^{\mathrm{opt}}_{\lambda,T}$ *decomposes into* $\pi^m_{\lambda,t} + \pi^r_{\lambda,t} + \pi^{\Pi}_{\lambda,t} + \pi^{\widehat{\alpha}}_{\lambda,t}$*, in which*

$$\begin{aligned}
\pi^m_{\lambda,t} &= -\mathbb{E}\left[\frac{1}{\mathcal{R}_{1,T}}\frac{\widehat{X}^{\mathrm{opt}}_{\lambda,T}Z_T^0}{\widehat{X}^{\mathrm{opt}}_{\lambda,t}Z_t^0}\,\middle|\,\mathcal{F}_t\right]\sigma^{\top^{-1}}_{\tilde{S},t}\widehat{\alpha}_t, \\
\pi^r_{\lambda,t} &= -\sigma^{\top^{-1}}_{\tilde{S},t}\mathbb{E}\left[\frac{1}{\widehat{\mathcal{R}}_{1,T}}\frac{\widehat{X}^{\mathrm{opt}}_{\lambda,T}Z_T^0}{\widehat{X}^{\mathrm{opt}}_{\lambda,t}Z_t^0}\int_t^T \mathcal{D}_t^W r_s \mathrm{d}s\,\middle|\,\mathcal{F}_t\right],
\end{aligned} \tag{3.26}$$

for all $t \in [0,T]$ *and* $\widehat{\mathcal{R}}^{-1}_{1,T} := 1 - \mathcal{R}^{-1}_{1,T}$ *identify, respectively, the tangency mean-variance and interest rate hedge demands. Here,* $\mathcal{D}_t^W : \mathbb{D}^{1,2}([0,T]) \to L^2(\Omega \times [0,T])^N$ *denotes the Malliavin derivative kernel. The remaining two portfolio weights comply with*

$$\begin{aligned}
\pi^{\Pi}_{\lambda,t} &= \sigma^{\top^{-1}}_{\tilde{S},t}\mathbb{E}\left[\mathcal{R}_{2,T}\frac{\widehat{X}^{\mathrm{opt}}_{\lambda,T}Z_T^0}{\widehat{X}^{\mathrm{opt}}_{\lambda,t}Z_t^0}\left\{\int_t^T \mathcal{D}_t^W\left(\mu_{\Pi,s}\mathrm{d}s + \sigma_{\Pi,s}\mathrm{d}W_s^{\sigma_\Pi}\right) + \sigma_{\Pi,t}\right\}\,\middle|\,\mathcal{F}_t\right], \\
\pi^{\widehat{\alpha}}_{\lambda,t} &= -\sigma^{\top^{-1}}_{\tilde{S},t}\mathbb{E}\left[\frac{1}{\widehat{\mathcal{R}}_{1,T}}\frac{\widehat{X}^{\mathrm{opt}}_{\lambda,T}Z_T^0}{\widehat{X}^{\mathrm{opt}}_{\lambda,t}Z_t^0}\left(-\int_t^T\left[\mathcal{D}_t^W\widehat{\alpha}_s\right]\widehat{\alpha}_s\mathrm{d}s + \int_t^T \mathcal{D}_t^W\widehat{\alpha}_s\mathrm{d}W_s\right)\,\middle|\,\mathcal{F}_t\right],
\end{aligned} \tag{3.27}$$

for all $t \in [0,T]$ *and* $W_t^{\sigma_\Pi} = W_t + \int_0^t \sigma_{\Pi,s}\mathrm{d}s$*, and characterise, respectively, a hedge against the benchmark,* $\{\Pi_t\}_{t\in[0,T]}$*, and a hedge against* $\{\widehat{\alpha}_t\}_{t\in[0,T]}$*, the fictitious market price process. Observe that* $\{\pi^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$ *of* $\{\widehat{X}^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$ *is allocated to the artificial stock, and that* $\{1 - 1_N^\top\pi^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$ *of* $\{\widehat{X}^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$ *is invested in the (original) cash account.*

Proof. The equality for $\widehat{X}^{\mathrm{opt}}_{\lambda,T}$ in (3.24) is trivial from the static problem in (3.23), and naturally coincides with the identity for $\widehat{X}^{\mathrm{opt}}_T$ as provided in Remark 3.1; the components that constitute the optimal portfolio process, $\{\pi^{\mathrm{opt}}_{\lambda,t}\}_{t\in[0,T]}$, in (3.26) and (3.26), are procurable from the Clark-Ocone formula, cf. Nualart (2006). For a more formal proof of these results, we refer the reader to section 3 in Kamma and Pelsser (2019). □

[30] According to the assumptions in section 2.2, $I'_Y : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ is well-defined. In order to understand how $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ may be related to ordinary RRA coefficients, suppose that $U(x,y) = (1-\gamma)^{-1}[(\frac{x}{y})^{1-\gamma} - 1]$ for some $\gamma \in \mathbb{R}_+ \setminus [0,1]$ and all $x, y \in \mathbb{R}_+$. In that case, clearly $\mathcal{R}_{1,T} = \gamma$ has to hold. Moreover, it is straightforward to show that $\mathcal{R}_{2,T} = 1 - \frac{1}{\gamma}$ is true. Consider for instance Brennan and Xia (2002) for such a ratio utility function in a less general financial market model.

Considering that Proposition 3.3 contains the optimality conditions imposed on the portfolio rules, $\{\pi^{\text{opt}}_{\lambda,t}\}_{t\in[0,T]}$, and on terminal wealth, $\widehat{X}^{\text{opt}}_{\lambda,T}$, in the auxiliary market $\widehat{\mathcal{M}}$, these artificially optimal rules have without an adequate characterisation of the shadow price process, $\{\tilde{S}_t\}_{t\in[0,T]}$, little to no economic meaning in the original market with frictions $\mathcal{M}$. Confining ourselves to an economic interpretation of Proposition 3.3 in $\widehat{\mathcal{M}}$, it suffices to refer to standard papers on this subject: consider for instance Merton (1969), Karatzas and Shreve (1998), or Detemple and Rindisbacher (2009). To draw attention to the optimality in $\mathcal{M}$ based on that in $\widehat{\mathcal{M}}$ as showcased in Proposition 3.3, ignoring the aforementioned economic interpretation for now, let us fix an $\omega \in \Omega$ and consider

$$p_t = \left(\widehat{\theta}_{\lambda,\tau_i} - \widehat{\theta}_{\lambda,\tau_{i-1}}\right)^+ + p_{\tau_{i-1}} \quad \text{and} \quad q_t = -\left(\widehat{\theta}_{\lambda,\tau_i} - \widehat{\theta}_{\lambda,\tau_{i-1}}\right)^- + q_{\tau_{i-1}}, \tag{3.28}$$

for all $t \in [\tau_i, \tau_{i+1})$, in which $(x)^+ = \max\{0, x\}$ and $(y)^- = \min\{y, 0\}$ for all $x, y \in \mathbb{R}$. Moreover, we define $\widehat{\theta}_{\lambda,\tau_i} := \left[\pi^{\text{opt}}_{\lambda,\tau_i} \odot \tilde{S}^{-1}_{\tau_i}\right]\widehat{X}^{\text{opt}}_{\lambda,\tau_i}$ such that $\widehat{\theta}_{\lambda,\tau_0} = 0$, using Propositions 3.2 and 3.3, where $\tau_i \in [0, T]$ are for $i = 1, \ldots, M$ all dates on the trading interval $[0, T]$ at which either $\tilde{S}_{\tau_i} = S^B_t$ or $\tilde{S}_{\tau_i} = S^A_t$ holds. The identities in (3.28) spell out the optimal transfers from the $\{X_t\}_{t\in[0,T]}$ account to $\{Y_t\}_{t\in[0,T]}$, and vice versa, derived from Proposition 3.3, under the assumption that there exists an optimal $\{\tilde{S}_t\}_{t\in[0,T]}$. Note that, since the portfolio in Proposition 3.3 merely trades in $\mathcal{M}$ if $\{\tilde{S}_t\}_{t\in[0,T]}$ fulfills either of the foregoing conditions, the optimal transfer processes in $\mathcal{M}$, $\{p_t\}_{t\in[0,T]}$ and $\{q_t\}_{t\in[0,T]}$, are continuously specified (over $[0, T]$) but do not imply continuous trading behaviour.[31]

4 Approximate Method

This section advances by the formulation and a corresponding assessment regarding the accuracy of the approximate method. By virtue of the general computational burden and the ordinary analytical trouble on the issue of solving the dual, the recovery of optimal solutions to the transaction problem in (2.7) is rather bothersome. However, Propositions 3.2 and 3.3 analytically articulate what the optimal portfolio must look like, cf. (3.28). The aforementioned specification establishes the core of our approximating procedure. First, we explain the rationale behind and the specifics of the approximate method. Second, we exemplify this routine by means of a few concrete situations. Third, and finally, we append this section by a numerical evaluation of the method's precision.

[31]Regarding the identities in (3.28), let us remark that these, with inclusion of the accompanying conditions, ought to hold for all $\omega \in \Omega$ and all $t \in [0, T]$. Moreover, let us stress that these optimality conditions for $\{p_t\}_{t\in[0,T]}$ and $\{q_t\}_{t\in[0,T]}$ follow from both Propositions 3.2 and 3.3: also, they carry over to the $\{L_t\}_{t\in[0,T]}$ and $\{M_t\}_{t\in[0,T]}$ processes, by stating that L_t and M_t must be flat off the sets $\{t \in [0, T] \mid \tilde{S}_t = S^B_t\}$ and $\{t \in [0, T] \mid \tilde{S}_t = S^A_t\}$, respectively, see Theorem 6.1 in Cvitanić and Karatzas (1996). Ultimately, recalling the end of section 3.2, we emphasise that (3.28) concretises a proper analytical structure for an educated "guess" about the dynamics of $\{p^{\text{opt}}_t\}_{t\in[0,T]}$ and $\{q^{\text{opt}}_t\}_{t\in[0,T]}$.

4.1 Confined Dual Space and Projection

The rationale underscoring our approximating procedure primarily consists in the observation that it is not possible to derive an optimal shadow price process on the mere basis of the dual formulation in (3.17). More particularly, it is not possible or at least complicated to derive a closed-form representation for this shadow price process in basic economic environments, under standard preference qualifications.[32] This inability or complexity is problematic, because it directly obstructs a recovery of the optimal decision rules in $\mathcal{M}$. It would therefore be plausible to confine the dual space, $\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$, to some closed and convex equivalent, say $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$, in an attempt to acquire analytical expressions for the shadow price. Accordingly, we determine the processes as follows:

$$\left\{\widehat{\alpha}_t^*, \widehat{\lambda}_t^*\right\}_{t\in[0,T]} := \underset{\left\{\widehat{\alpha}_t, \widehat{\lambda}_t\right\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}}{\arg\inf} \mathbb{E}\left[V\left(Z_T^0, \Pi_T\right)\right], \tag{4.1}$$

where $\{\widehat{\alpha}_t^*\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t^*\}_{t\in[0,T]}$ represent the approximate processes that jointly define the approximate shadow price, $\{\tilde{S}_t^*\}_{t\in[0,T]}$.[33] By construction, $\{\widehat{\alpha}_t^*, \widehat{\lambda}_t^*\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} \subseteq \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$ holds. As indicated by (4.1), we determine these approximate shadow prices by minimising the dual over this confined dual space. We determine the approximate processes in that way for two purposes. One, this approach provides us with an upper bound on the actual optimal value function. Two, the shadow prices that minimise the dual are inclined to "tilt" the strategies in Proposition 3.3 to their corresponding optima in $\widehat{\mathcal{M}}$. The first reason concerns performance evaluation of the method, and the second reason relates to the subsequent step of the approximating routine.

More to the point, the approximate dual pair implies an approximate strategy, $\{\pi_{\lambda,t}^*\}_{t\in[0,T]}$. We subsequently "project" this strategy into $\mathcal{M}$ in the following way:

$$p_t^* = \left(\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*\right)^+ + p_{\tau_{i-1}}^* \quad \text{and} \quad q_t^* = -\left(\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*\right)^- + q_{\tau_{i-1}}^*, \tag{4.2}$$

for all $t \in [\tau_i, \tau_{i+1})$ and $i = 1, \ldots, M$, where $\tau_i \in [0, T]$ such that $\tau_i \leq \tau_{i+1}$ define pre-determined points on the trading interval at which we wish to make a transaction. Observe that these pre-determined trading dates, as well as the amount of trading dates, M, may be state-dependent: e.g. for all $\omega \in \Omega$ we may *a priori* (before $t = 0$) characterise $\tau_i \in [0, T]$ such that $\left|\tilde{S}_{\tau_i} - S_{\tau_i}^A\right| < \epsilon$ holds, where $\epsilon > 0$ is true. For ease of exposition, we subsequently suppress the dependency of τ_i on $\omega \in \Omega$. We restrict τ_i and M to be

[32] Consider for instance Zakamouline (2002), wherein the author already encounters enormous mathematical difficulties when it comes to solving (2.7) for conventional and simple CRRA preferences in an elementary financial market. Similarly, cf. Liu (2004) for analogous mathematical complications as to deriving the shadow price process within a market model that includes CARA preferences.

[33] Henceforth, in line with the specification of the auxiliary market and the according optimality conditions in Proposition 3.2 and 3.3, we work with the undetermined (shadow price) pair $\{\widehat{\alpha}_t^*, \widehat{\lambda}_t^*\}_{t\in[0,T]}$ instead of the actual shadow price process, $\{\tilde{S}_t\}_{t\in[0,T]}$, for the sake of consistency and clarity.

such that that $\{p_t^*, q_t^*\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ holds; the state-dependency concerning τ_i and M facilitates the fulfillment of the admissibility requirement. Note that these two parameters suffice to establish admissibility at all times, since $M = 0$ results in $\widehat{X}_t \geq 0$ for all $t \in [0,T]$. In (4.2), $\widehat{\theta}^*_{\lambda,\tau_i}$ equates to $\widehat{\theta}_{\lambda,\tau_i}$ in (3.28), wherein $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ are replaced by, respectively, $\{\widehat{\alpha}^*_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}^*_t\}_{t\in[0,T]}$, ensuring that $\{p_t^*, q_t^*\}_{t\in[0,T]}$ is analytically fully identified.[34] We justify the choice for this approximation to $\{p_t^{\text{opt}}\}_{t\in[0,T]}$ and $\{q_t^{\text{opt}}\}_{t\in[0,T]}$, by noting that these are, up to the restricted $\{\widehat{\alpha}^*_t, \widehat{\lambda}^*_t\}_{t\in[0,T]}$ and the exact trading dates, equal to the actual optima in (4.1). Approximate wealth accordingly reads:

$$
\begin{aligned}
\widehat{X}^*_t = X_0 + 1_N^\top Y_0 + \int_0^t r_s \left\{ \widehat{X}^*_s - [p_s^* - q_s^*]^\top S_s \right\} \mathrm{d}s \\
+ \int_0^t [p_s^* - q_s^*]^\top \mathrm{d}\left(\mu_s \odot S_s + \sigma_s S_s \mathrm{d}W_s\right) - \int_0^t S_s^\top \left[\lambda_1 \odot \mathrm{d}q_s^* + \lambda_2 \odot \mathrm{d}p_s^*\right].
\end{aligned}
\tag{4.3}
$$

By means of inserting approximate terminal wealth, i.e. (4.3) evaluated at $t = T$, into the primal objective in (2.7), we acquire a lower bound on the actual optimal value function, say $\widehat{J}_P(\widehat{X}_0)$. Similarly, let us recall that the minimisation of equation (4.1)'s RHS corresponds to minimising the dual, conditional on the fact that $Z_0^0 = \mathcal{H}^{-1}(\widehat{X}_0)$ holds, as in Proposition 3.3. Therefore, inserting $\{\widehat{\alpha}^*_t, \widehat{\lambda}^*_t\}_{t\in[0,T]}$ as well as $Z_0^0 = \mathcal{H}^{-1}(\widehat{X}_0)$ into the dual objective in (3.10) provides us with an upper bound on the actual optimal value function, say $\widehat{J}_D(\widehat{X}_0)$. Due to strong duality in case of the transaction cost problem at hand, all deviations of $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$ from the actual dual space $\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$ and/or all departures of the prefixed trading dates from the truly optimal ones ought to result in a non-negative difference between the approximate dual and primal value functions.

So as to quantify this utilitarian difference in monetary terms, let us consider:

$$
\widehat{J}_P\left(\widehat{X}_0 + \mathcal{CV}\right) = \inf_{\left\{\widehat{\alpha}_t, \widehat{\lambda}_t\right\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}} \widehat{J}_D\left(\widehat{X}_0, \left\{\widehat{\alpha}_t, \widehat{\lambda}_t\right\}_{t\in[0,T]}\right),
\tag{4.4}
$$

in which $\mathcal{CV} \in \mathbb{R}_+$ represents the amount of monetary units that one must add to his/her initial endowment in order to "close" the duality gap. In line with de Palma and Prigent (2008) and de Palma and Prigent (2009), we refer to the latter quantity as the compensating variation. One may interpret the compensating variation as an (annual) management fee that the agent pays to some representative investor in order to be assured of obtaining

[34] In contrast to the approximating procedure in Kamma and Pelsser (2019), we directly insert the specified $\{\widehat{\alpha}^*_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}^*_t\}_{t\in[0,T]}$ into the approximate primal investment rules. Specifically, in that study, the authors use $\{p_t\}_{t\in[0,T]}$ and $\{p_t\}_{t\in[0,T]}$ in which they limit the therein embedded $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ to the confined dual space, after which these are identified in a manner such that these maximise the primal value function. However, this study demonstrates that the values for the primal and dual equivalents of the approximate shadow prices are only marginally different. This small difference is sensible, because the dual naturally aims to determine $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t\}_{t\in[0,T]}$ in a fashion that is least favourable from $\widehat{\mathcal{M}}$'s point of view, but simultaneously optimal from $\mathcal{M}$'s perspective.

the most optimal amount of terminal wealth (under the current circumstances involving market frictions). The approximating procedure consists of two fundamental aspects: (i) confining the dual space, and (ii) the proposal of corresponding approximate optimal controls in the true environment ($\mathcal{M}$). Depending on the quality of the choice for an auxiliary dual space, and on the discrepancy between the approximate investment policy and the optimal analogue, the duality gap shrinks or grows. Particularly, for poor choices concerning both aspects, the gap will grow. Therefore, the size of the gap and thus of $\mathcal{CV}$ render proper indicators for the accuracy of the approximated strategies.

Remark 4.1. *We address three points concerning the approximating procedure as presented above. First, in order to highlight the explicit closed-form feature of the routine with regard to the therefrom derived investment strategies, let us observe that the trading strategies are approximated on the basis of the analytics in Proposition 3.3. Since these approximate policies are calculated, like the actual ones in* (3.28)*, as displayed in* (4.2)*, the analytical nature exclusively consists in the tractability of the amount of stocks one possesses over the trading interval. That is, we cannot infer conretely what the closed-form expressions for* p_t^* *and* q_t^* *look like over the entire trading interval in the form of a single (non-recursive) expression. However, for all* $t \in [\tau_{i-1}, \tau_i)$*, the relevant analytics are avaiable.*[35]

Remark 4.2. *In order to equip the user of the routine with an amount freedom pertaining to the admissibility requirements and corresponding trading dates, we have left the choice for* τ_i*,* $i = 1, \ldots, M$*,* $M \in \mathbb{Z}$ *unspecified in* (4.2)*. We are, however, able to spell out restrictions on* τ_i *and on* $M \in \mathbb{Z}$ *such that* $\{p_t^*, q_t^*\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ *holds at all times. This concurrently exemplifies how* τ_i *could be specified. To this end, let us note that*

$$\frac{\left[p_{\tau_i}^* - q_{\tau_i}^*\right]^\top S_{\tau_i}}{\widehat{X}_{\tau_i}^*} \in [0,1)\,, \quad \text{and} \quad \widehat{X}_{\tau_i}^* - \lambda_1 \Delta q_{\tau_i}^{*\top} S_{\tau_i} - \lambda_2 \Delta p_{\tau_i}^{*\top} S_{\tau_i} > 0 \tag{4.5}$$

characterise for a pre-determined set of trading dates, $\{\tau_i\}_{i=1}^M$*, where* $\Delta q_{\tau_i}^* = q_{\tau_i}^* - q_{\tau_{i-1}}^*$ *and* $\Delta p_{\tau_i}^* = p_{\tau_i}^* - p_{\tau_{i-1}}^*$*, two possible conditions that ensure* $\widehat{X}_t^* \geq 0$ *for all* $t \in [0,T]$*. In line with this observation, we suppose that* $\mathcal{T}_i$ *contains all* $\{p_{\tau_i}^*\}_{j=i-1}^i$*,* $\{q_{\tau_i}^*\}_{j=i-1}^i$*, and* $\widehat{X}_{\tau_i}^*$ *such that the requirements in* (4.5) *are satisfied. Then, if we define* $\tau_i = (\tau_i - \tau_{i-1})\,\mathbb{1}_{\{(\{p_{\tau_i}^*\}_{j=i-1}^i, \{q_{\tau_i}^*\}_{j=i-1}^i, \widehat{X}_{\tau_i}^*) \in \mathcal{T}_i\}} + \tau_{i=1}$ *for* $i = 1, \ldots, M$*,* $X_t^* \geq 0$ *holds for all* $t \in [0,T]$*. Consequently, the investor only trades at some date* $\tau_i \in [0,T]$*, if it is certain that* $\widehat{X}_{\tau_i}^* \geq 0$ *is satisfied, otherwise* τ_i *is set equal to* τ_{i-1} *and no trade takes place.*

[35]Note that these assertions hold true, given that we are in possession of entirely identified shadow price processes, $\{\widehat{\alpha}_t^*\}_{t\in[0,T]}$ and $\{\widehat{\lambda}_t^*\}_{t\in[0,T]}$. From the first step in the approximating routine, (4.1), we indeed acquire such identified processes. However, observe that it still may be necessary to numerically determine these approximate shadow price processes, due to which the qualification of the approximation to the optimal investment strategies as purely analytical is not totally accurate. The need for numerical procedures in the characterisation of these processes strongly depends on the specifics of the confined dual space, $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$. Obviously, this need also affects the analytical availability of $\widehat{J}_D(\widehat{X}_0)$.

Remark 4.3. *The approximating procedure as described in the main text essentially distinguishes five fundamental steps. In order to clarify these steps, and to thereby summarise the approximation in a concise overview, we subsequently provide the routine in a step-wise manner. Without paying much regard to the technical details and the potential requirement for numerical numerical procedures that are involved with each step (for which we refer the reader to the main text), we present the approximate method as follows:*

- *Confine the dual space,* $\widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$*, in which the shadow price processes,* $\{\widehat{\alpha}_t\}_{t\in[0,T]}$ *and* $\{\widehat{\lambda}_t\}_{t\in[0,T]}$*, ought to live, to some closed and convex subset,* $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} \subseteq \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}}$*.*

- *Identify the approximate shadow price processes in a manner such that they jointly minimise the dual objective,* $\{\widehat{\alpha}^*_t, \widehat{\lambda}^*_t\}_{t\in[0,T]} = \arg\inf_{\{\widehat{\alpha}_t,\widehat{\lambda}_t\}_{t\in[0,T]}\in\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}} \mathbb{E}\left[V\left(Z^0_T, \Pi_T\right)\right]$*, and set* $\widehat{J}_D(\widehat{X}_0)$ *equal to the ensuing dual value function.*

- *Employ the approximate shadow price processes to acquire the optimal portfolio strategy in* $\widehat{\mathcal{M}}$*, i.e.* $\{\pi^{\text{opt}}_{\lambda,t}\}_{t\in[0,T]}$ *with inclusion of the preceding approximate shadow prices,* $\{\widehat{\alpha}^*_t, \widehat{\lambda}^*_t\}_{t\in[0,T]}$*, and calculate* $\widehat{\theta}^*_{\lambda,t} := \left[\pi^{\text{opt}}_{\lambda,t} \odot \tilde{S}^{-1}_t\right]\widehat{X}^{\text{opt}}_{\lambda,t}$*.*

- *Determine* $M \in \mathbb{Z}$ *trading dates,* $\{\tau_i\}_{i=1}^M$*, and specify the approximate transaction processes according to the following recursive relations:* $p^*_t = (\widehat{\theta}^*_{\lambda,\tau_i} - \widehat{\theta}^*_{\lambda,\tau_{i-1}})^+ + p^*_{\tau_{i-1}}$ *and* $q^*_t = -(\widehat{\theta}^*_{\lambda,\tau_i} - \widehat{\theta}^*_{\lambda,\tau_{i-1}})^- + q^*_{\tau_{i-1}}$ *for all* $t \in [\tau_i, \tau_{i+1})$*. Furthermore, calculate* $\{X^*_t\}_{t\in[0,T]}$ *and set* $\widehat{J}_P(\widehat{X_0})$ *equal to the ensuing primal value function.*

- *Assess the quality of the approximation by computing the compensating variation*[36]*, i.e.* $\mathcal{CV} \in \mathbb{R}_+$ *such that* $\widehat{J}_P(\widehat{X}_0 + \mathcal{CV}) = \widehat{J}_D(\widehat{X}_0)$*.*

Let us conclude the outline/explication of the approximating procedure by addressing two critical features inherent in the preceding routine that distinguish it from the yet established methods for numerically determining or approximating the optimal strategies in markets with frictions. The first feature concerns the ability to concretise the quality of the approximation to the optimal investment strategies by means of the duality gap. Ordinary methods typically qualify the performance of their method by relying on an external algorithm (e.g. a grid-search one) that calculates the optimal solution, to which

[36]The possibility to acquire the compensating variation in closed-form depends exclusively on the choice with regard to the investor's preference qualification. For the conventional (ratio) CRRA preferences, it is possible to calculate it in closed-form; on the contrary, for e.g. *dual-CRRA* preferences as utilised in Kamma and Pelsser (2019), this is not possible. Since these identities are not of primary interest to this study, in the numerical illustrations we compute these quantities numerically (which requires little computational effort). The emphasis that this study puts on the closed-form nature of the routine should namely be considered predominantly with regard to the approximate strategies themselves.

they compare their approximated analogue. The major downside of such methods is the need for an external procedure to measure the quality of the approximation. Our approach circumvents this need and is in that regard self-contained. The second feature concerns the availability of (semi-)closed-form expressions for the investment strategies. In general, ordinary approximating methods namely solely provide numerical answers that are not re-traceable to any analytical configuration. Hence, our method engenders comparably more transparency regarding the dynamics that underscore the approximated policies.

4.2 Illustrations of Approximate Method

In this subsection, we exemplify the approximating procedure by utilising two concrete specifications of the economic environment in section 2. To evaluate and illustrate the approximate method's potential accuracy (in the sequel), we will make use of these two environments. Both characterisations rely on a standard one-dimensional Black-Scholes model. Consequently, we distinguish two explicit model setups by varying the description of the preference qualification. On the grounds of the inability or severe complexity with regard to acquiring closed-form expressions in even elementary economic frameworks, see e.g. Shreve and Soner (1994), Soner et al. (1995), and Højgaard and Taksar (1998), we opt for this simple model; the different utility functions induce sufficient technical heterogeneity. Subsequently, we explicate $\widehat{\mathcal{M}}$ for the ordinary Black-Scholes model, followed by an illustration of section 4.1's routine for the two settings of interest.

Example 4.1. *(Black-Scholes Model) Consider the market model in section 2.1. Suppose that $N = 1$, such that there is only one risk-driver, $\{W_t\}_{t\in[0,T]}$, and thus one stock that is traded in the market, $\mathcal{M}$. Moreover, assume that the interest rate, the stocks's expected return and corresponding volatility are all constant: $\mathrm{d}B_t = rB_t\mathrm{d}t$ and $\mathrm{d}S_t = S_t\left[\mu\mathrm{d}t + \sigma\mathrm{d}W_t\right]$ spell out the according SDE's. Turning to the auxiliary market, $\widehat{\mathcal{M}}$, in the spirit of the approximate method, we already confine the space in which the shadow prices that are unique to $\{Z_t^0\}_{t\in[0,T]}$ and $\{Z_t^1\}_{t\in[0,T]}$ live to the real line: $\{\widehat{\alpha}_t, \widehat{\lambda}_t\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$. We henceforth set $\widehat{\alpha} = \widehat{\alpha}_t$ and $\widehat{\lambda} = \widehat{\lambda}_t$. Based on the requirements that are incorporated in $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$, we know that $Z_t^0\tilde{S}_t$ must be a $\mathbb{P}$-martingale and that $Z_T^0 = Z_T^1$ ought to hold. Therefore, for all $(\widehat{\alpha}, \widehat{\lambda}) \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$, it must be true that:*

$$\widehat{\lambda} = \delta - \beta = 0, \quad \gamma_t = -\mu - \widehat{\alpha}\sigma, \quad \text{and} \quad Z_0^1 = Z_0^0 e^{(-r+\mu+\widehat{\alpha}\sigma)T}. \tag{4.6}$$

To clarify the latter, let us stress that the first and third equality follows clearly from the $Z_T^0 = Z_T^1$ requirement; the second equality ensures that the necessary $\mathbb{P}$-martingale property of $Z_t^0\tilde{S}_t$ is satisfied; the third equality then logically follows from a combination of both aforementioned conditions. The final requirement that must be met, and which simultaneously poses restrictions on $\widehat{\alpha} \in \mathbb{R}$, is the one which states that $S_t^B \leq \tilde{S}_t \leq S_t^A$ for *all* $t \in [0,T]$*. By virtue of the requirements in* (4.6)*, the latter inequality reduces to* $1-\lambda_1 \leq e^{[\mu - r + \widehat{\alpha}\sigma](T-t)} \leq 1+\lambda_2$*. For this reason, we derive that* $\frac{\log(1-\lambda_1)}{\sigma T} - \frac{\mu - r}{\sigma} \leq \widehat{\alpha} \leq \frac{\log(1+\lambda_2)}{\sigma T} - \frac{\mu - r}{\sigma}$ *must be true. Note that for all* $(\widehat{\alpha}, \widehat{\lambda}) \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$*, the reduction of the two available shadow prices inherent in* $\frac{Z_t^1 S_t}{Z_t^0}$ *to one, i.e.* $\widehat{\alpha}$*, nicely corresponds to the inclusion of merely one shadow price process in* Z_t*, cf.* (3.11)*. Moreover, we choose* $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$ *for purposes related to the numerical evaluation of the approximation later. From here on out, we assume as well that* $\mu_{\Pi,t} = \mu_\Pi$ *and* $\sigma_{\Pi,t} = \sigma_\Pi$ *are constant over* $t \in [0,T]$*.*

We now demonstrate how the approximating procedure technically works under the Black-Scholes model in Example 4.2, for two different preference qualifications: (i) a standard CRRA utility specification, and (ii) Kamma and Pelsser (2019)'s dual-CRRA utility function. For simplicity and economic elegance, in both cases, we assume that the agent derives utility from the ratio of terminal wealth to the individual-specific benchmark, $\widehat{X}_T/\Pi_T$. We opt for the CRRA function due to its frequent application in optimal investment frameworks; for the dual-CRRA function, because of its salient departure from the conventionally employed preferences, i.c. the CRRA-paradigm, primarily discernible from its non-constant (wealth-dependent) RRA. In outlining the routine hereafter, we aim to adhere to the scheme of the procedure as described in Remark 4.3.

Example 4.2. *(CRRA Utility) Consider the dynamic investment problem* (2.7) *in a Black-Scholes market with frictions, cf. Example 4.1. Additionally, assume that the investor's preferences abide by the ratio CRRA utility function:* $U(x,y) = \frac{(\frac{x}{y})^{1-\gamma}}{1-\gamma}$ *for all* $x, y \in \mathbb{R}_+$ *and* $\gamma \in \mathbb{R}_+ \setminus [0,1]$*. In line with the first step of the routine, cf. Remark 4.3, we postulate in the corresponding artificial market,* $\widehat{\mathcal{M}}$*, that* $\{\widehat{\alpha}_t, \widehat{\lambda}_t\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$*; we set* $\widehat{\alpha} = \widehat{\alpha}_t$ *and* $\widehat{\lambda} = \widehat{\lambda}_t$*. Observing that* $V : \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}$ *is decreasing in its first argument, we find that* $\widehat{\alpha}^* = -\frac{\mu - r}{\sigma} + \frac{\log(1+\lambda_2)}{\sigma T}$ *and* $\widehat{\lambda}^* = 0$*. Proposition 3.3 infers that the corresponding analytics for* $\{\widehat{X}_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$ *and* $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$*, respectively, read as follows*

$$\widehat{X}_{\lambda,t}^{\text{opt}} = \widehat{X}_0 C_t \left(Z_0^{0^{-1}} Z_t^0 \Pi_t \right)^{-\frac{1}{\gamma}} \Pi_t, \quad \text{and} \quad \pi_{\lambda,t}^{\text{opt}} = -\frac{1}{\gamma}\frac{\widehat{\alpha}^*}{\sigma} + \left(1 - \frac{1}{\gamma}\right)\frac{\sigma_\Pi}{\sigma}, \tag{4.7}$$

for all $t \in [0,T]$*, in which* $C_t = \exp\left\{ -(1-\frac{1}{\gamma})\left[-r + \mu_\Pi + \widehat{\alpha}^*\sigma_\Pi - \frac{1}{2}\frac{1}{\gamma}\left(\sigma_\Pi + \widehat{\alpha}^*\right)^2 \right] t \right\}$*. Hence,* $\widehat{\theta}_{\lambda,t}^* = \frac{1}{\widehat{S}_t}\pi_{\lambda,t}^{\text{opt}}\widehat{X}_{\lambda,t}^{\text{opt}}$ *specifies for* $(\widehat{\alpha}, \widehat{\lambda}) \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$ *the* $\widehat{\mathcal{M}}$*-optimal number of stocks that the investor must purchase. Note that* $\widehat{\theta}_{\lambda,t}^*$ *is for all* $t \in [0,T]$ *available in closed-form and fully identified. To make the continuous strategy* $\{\widehat{\theta}_{\lambda,t}^*\}_{t\in[0,T]}$ *admissible in* $\mathcal{M}$*, we follow the fourth step in Remark 4.3. In principle, this choice is arbitrary as long as* $\{p_t^*, q_t^*\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ *holds, but in view of our subsequent (numerical) illustration, we set* M *equal to* T*: the investor re-balances his/her portfolio at the start of each year. Then, we define* $p_t^* = (\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^+ + p_{\tau_{i-1}}^*$ *and* $q_t^* = -(\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^- + q_{\tau_{i-1}}^*$*, wherein* τ_i *for* $i = 1, \ldots, T$ *are specified in accordance with those that are given in Remark 4.2.*

In the previous Example 4.2, let us stress that a closed-form expression is available for the approximate dual value function, howbeit not for its primal counterpart: $\widehat{J}_D(\widehat{X}_0) = \frac{1}{1-\gamma}(\widehat{X}_0^{1-\gamma} C_T^{-\gamma} - 1)$. Regarding the readily accessible closed-form expressions that underscore $\{p_t^*\}_{t\in[0,T]}$ and $\{q_t^*\}_{t\in[0,T]}$, in (4.7), we note that these resemble the analytical identities that are provided in Theorems 1 and 3 of Brennan and Xia (2002). In fact, this similarity is straightforward, since the latter study relies on a comparable model setup. The mere diversion from the frictionless setup as to $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$ consists in the incorporation therein of the $\widehat{\alpha}^* = -\frac{\mu-r}{\sigma} + \frac{\log(1+\lambda_2)}{\sigma T}$ term, rather than exclusively the $-\frac{\mu-r}{\sigma}$ term. Evidently, the presence of market frictions forces the malleable shadow price process to tilt itself in a downward direction. In the corresponding mean-variance demand, this downward tendency prompts a more prudent trading strategy, despite the trading dates.

Example 4.3. *(Dual-CRRA Utility) Consider the dynamic investment problem* (2.7) *in a Black-Scholes market with frictions, cf. Example 4.1. Furthermore, suppose that the investor's preferences live by the ratio dual-CRRA utility function:* $U(x, y) = \frac{(\frac{x}{y})^{1-\gamma_d}}{1-\gamma_d} \mathbb{1}_{\{x\le y\}} + \frac{(\frac{x}{y})^{1-\gamma_u}}{1-\gamma_u} \mathbb{1}_{\{x>y\}}$ *for all* $x, y \in \mathbb{R}_+$ *and* $\gamma_d, \gamma_u \in \mathbb{R}_+ \setminus [0, 1]$*. We remark that for this preference qualification* $U \in \mathcal{C}(\mathbb{R}_+, \mathbb{R}_+; \mathbb{R})$ *and* $U_X' \in \mathcal{PC}(\mathbb{R}_+, \mathbb{R}_+; \mathbb{R}_+)$ *holds, where* U_X'*'s (and thus* I*'s) break-point is logically positioned at* $x = y$*. Following the approximating procedure's scheme, see Remark 4.3, we assume in* $\widehat{\mathcal{M}}$ *that* $\{\widehat{\alpha}_t, \widehat{\lambda}_t\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$ *holds; we set* $\widehat{\alpha} = \widehat{\alpha}_t$ *and* $\widehat{\lambda} = \widehat{\lambda}_t$*. As in Example 4.2, it is true that* $\widehat{\alpha}^* = -\frac{\mu-r}{\sigma} + \frac{\log(1+\lambda_2)}{\sigma T}$ *and* $\widehat{\lambda}^* = 0$ *must hold. According to Proposition 3.3, under the condition that* $(\widehat{\alpha}, \widehat{\lambda}) \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}}$*, the expressions for* $\{\widehat{X}_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$ *and* $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$*, respectively, read as follows*

$$
\begin{aligned}
\widehat{X}_{\lambda,t}^{\text{opt}} &= \widehat{X}_{t,\lambda,\gamma_d} \Phi(d_{t,T,d}) + \widehat{X}_{t,\lambda,\gamma_u} \Phi(d_{t,T,u}), \quad \text{and} \\
\pi_{\lambda,t}^{\text{opt}} &= -\frac{\widehat{\alpha}^*}{\sigma} \Psi_{t,T}(\gamma_d, \gamma_u) + \frac{\sigma_\Pi}{\sigma} \{1 - \Psi_{t,T}(\gamma_d, \gamma_u)\},
\end{aligned}
\tag{4.8}
$$

in which $X_{t,\lambda,\gamma_i} = (\eta^{\text{opt}} Z_0^{0^{-1}} Z_t^0 \Pi_t)^{-\frac{1}{\gamma_i}} \Pi_t C_{\gamma_i,t,T}$ *for* $C_{\gamma_i,t,T} = \exp\{(1 - \frac{1}{\gamma_i})[-r + \mu_\Pi + \sigma_\Pi \widehat{\alpha}^* - \frac{1}{2}\frac{1}{\gamma_i}(\sigma_\Pi + \widehat{\alpha}^*)^2]\{T - t\}\}$*, and* $d_{t,T,i} = -\frac{1}{(\sigma_\Pi+\widehat{\alpha}^*)\sqrt{T-t}}(\log(\eta^{\text{opt}} Z_t^0 \Pi_t) - [r - \mu_\Pi - \sigma_\Pi \widehat{\alpha}^* + \frac{1}{2}(\sigma_\Pi + \widehat{\alpha}^*)^2 - (1 - \frac{1}{\gamma_i})(\sigma_\Pi + \widehat{\alpha}^*)^2]\{T - t\}) (1 - 2\mathbb{1}_{\{i=\{u\}\}})$*: the standard normal CDF function is given by* $\Phi(\cdot)$*. Moreover, in the identity for* $\pi_{\lambda,t}^{\text{opt}}$*, we define* $\Psi_{t,T}(\gamma_d, \gamma_u) = \left(\sum_{i\in\{d,u\}} \widehat{X}_{t,\lambda,\gamma_i} \Phi(d_{t,T,i})\right)^{-1} \sum_{i\in\{d,u\}} \frac{1}{\gamma_i} \widehat{X}_{t,\lambda,\gamma_i} \Phi(d_{t,T,i})$*. Observe that* $\eta^{\text{opt}} \in \mathbb{R}_+$ *must be acquired numerically, and can be recovered from the following equality:* $\widehat{X}_{\lambda,0}^{\text{opt}} = \widehat{X}_0$*. Consequently, up to the numerical value for* $\eta^{\text{opt}} \in \mathbb{R}_+$*,* $\widehat{\theta}_{\lambda,t}^* = \frac{1}{\widetilde{S}_t} \pi_{\lambda,t}^{\text{opt}} \widehat{X}_{\lambda,t}^{\text{opt}}$ *is entirely identified and available in closed-form. In conformity with Example 4.2, to "project"* $\{\widehat{\theta}_{\lambda,t}^*\}_{t\in[0,T]}$ *into* $\mathcal{M}$*'s admissibility set* $\mathcal{A}_{X_0,Y_0}$*, we set* $M = T$ *(the investor re-balances his/her portfolio at the start of each year). Additionally, we let* $p_t^* = (\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^+ + p_{\tau_{i-1}}^*$ *and* $q_t^* = -(\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^- + q_{\tau_{i-1}}^*$*. Here,* τ_i *for* $i = 1, \ldots, T$ *is given in Remark 4.2.*

As for Example 4.2, in the latter Example 4.2, there is only an analytical identity available for the dual value function: $\widehat{J}_D(\widehat{X}_0) = \sum_{i\in\{d,u\}} \frac{1}{1-\gamma_i}\eta^{\text{opt}}\widehat{X}_{0,\lambda,\gamma_d}\Phi\left(d_{0,T,i}\right) - \sum_{i\in\{d,u\}} \frac{1}{1-\gamma_i}\Phi\left(d_i\right)$, where d_i is equal to $d_{0,T,i}$ in which $1-\frac{1}{\gamma_i}=0$ holds. For a formal derivation and an appropriate economic analysis of these optimality conditions, cf. Kamma and Pelsser (2019) or Kamma et al. (2020). Note that $\widehat{\alpha}^*$ (logically) tilts the portfolio again to a more prudent counterpart. Unlike the analytical identities in Example 4.2, the present approximation is based on semi-closed-form expressions, because of the mathematically unidentifiable $\eta^{\text{opt}} \in \mathbb{R}_+$. Nevertheless, the general analytical structure is completely available to us, see (4.7). We would like to emphasise that $\{p_t^*\}_{t\in[0,T]}$ and $\{q_t^*\}_{t\in[0,T]}$ are derived from a $\{\pi_{\lambda,t}^{\text{opt}}\}_{t\in[0,T]}$ that is wealth-dependent. Consequently, $\{p_t^*\}_{t\in[0,T]}$ and $\{q_t^*\}_{t\in[0,T]}$ are subject to larger changes in the composition of the portfolio. The magnitude of these changes generally grows for larger differences between γ_d and γ_u.

Remark 4.4. *The log-utility function constitutes a special limiting case of both the CRRA and dual-CRRA functions:* $\lim_{\gamma\to 1}\frac{(\frac{x}{y})^{1-\gamma}}{1-\gamma} = \log\frac{x}{y}$ *and* $\lim_{\gamma_d,\gamma_u\to 1}\{\frac{(\frac{x}{y})^{1-\gamma_d}}{1-\gamma_d}\mathbb{1}_{\{x\le y\}} + \frac{(\frac{x}{y})^{1-\gamma_u}}{1-\gamma_u}\mathbb{1}_{\{x>y\}}\} = \log\frac{x}{y}$ *for all* $x, y \in \mathbb{R}_+$.[37] *As a straightforward consequence, setting* $\gamma = 1$ *and* $\gamma_d = \gamma_u = 1$ *in, respectively, Examples 4.2 and 4.3 renders the* $\widehat{\mathcal{M}}$*-optimal policies for* $U(x,y) = \log\frac{x}{y}$, *stipulated that* $\left\{\widehat{\alpha}_t, \widehat{\lambda}_t\right\}_{t\in[0,T]} \in \widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$ *is true, wherein we set* $\widehat{\alpha} = \widehat{\alpha}_t$ *and* $\widehat{\lambda} = \widehat{\lambda}_t$. *To be precise, the policies read* $\widehat{X}_{\lambda,t}^{\text{opt}} = \frac{\widehat{X}_0}{Z_0^{0^{-1}} Z_t^0}$ *and* $\pi_{\lambda,t}^{\text{opt}} = -\frac{\widehat{\alpha}^*}{\sigma}$ *for all* $t \in [0,T]$. *The corresponding value function does obviously not follow from setting* $\gamma = 1$ *and* $\gamma_d = \gamma_u = 1$, *but lives by* $\widehat{J}_D(\widehat{X}_0) = \log\widehat{X}_0 - (-r - \frac{1}{2}\widehat{\alpha}^{*2})T - (\mu_\Pi - \frac{1}{2}\sigma_\Pi^2)T$. *By analogy with Examples 4.2 and 4.3, we define* $\widehat{\theta}_{\lambda,t}^* = \frac{1}{\tilde{S}_t}\pi_{\lambda,t}^{\text{opt}}\widehat{X}_t^{\text{opt}}$, *and "project" this process to* $\mathcal{A}_{X_0,Y_0}$ *by setting* $M = T$ *and* $p_t^* = (\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^+ + p_{\tau_{i-1}}^*$ *together with* $q_t^* = -(\widehat{\theta}_{\lambda,\tau_i}^* - \widehat{\theta}_{\lambda,\tau_{i-1}}^*)^- + q_{\tau_{i-1}}^*$. *Here, the trading dates,* τ_i *for* $i = 1,\ldots,T$, *are defined as in Remark 4.2. Therefore,* $\{p_t^*, q_t^*\}_{t\in[0,T]} \in \mathcal{A}_{X_0,Y_0}$ *is indeed true.*

Let us make three remarks. First, since we have not accentuated this, we observe from $\{\tilde{S}_t\}_{t\in[0,T]}$'s specification in (3.12) that if both $\{\widehat{\alpha}^*\}_{t\in[0,T]}$ and $\{\widehat{\lambda}^*\}_{t\in[0,T]}$ are identifiable, the shadow price process is likewise entirely identified. In consideration of the identifiability of $\{\widehat{\alpha}^*\}_{t\in[0,T]}$ and $\{\widehat{\lambda}^*\}_{t\in[0,T]}$ in all three cases, $\{\tilde{S}_t\}_{t\in[0,T]}$ is here analytically available and identified. Second, for the examples at hand, we point out that it is impossible to acquire a closed-form expression for the primal value function, $\widehat{J}_P(\widehat{X}_0)$. For that purpose, one could rely on e.g. simple, computationally effortless Monte-Carlo simulations: we will rely on these in the subsequent numerical evaluation. Third, it is clear that the constant RRA property of CRRA functions results in relatively "stable" portfolio compositions, and

[37] In spite of the facts that the dual-CRRA preference qualification implicitly embeds the specifications for both the CRRA utilities – for that purpose, note that requiring $\gamma_d = \gamma_u$ in the dual-CRRA function results in equivalence of its specification with the standard CRRA function – and that the CRRA utility implies the log-utility – as a limiting case $\gamma \to 1$ –, we separate the results for specific clarity when it comes to the subsequent numerical evaluation of our approximating technique. Namely, therein we conform to an analogously unmistakable separation of cases. Therefore, the comprehensibility concerning the underscoring approximation-related details of the aforementioned individual numerically evaluated cases enhances as a consequence of this division of results for the intertwined utility functions.

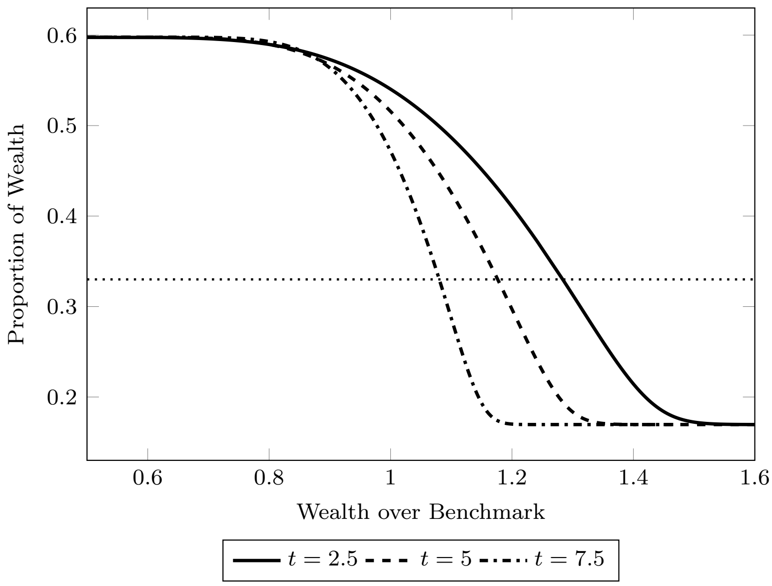

Figure 1. Artificial optimal proportional allocation to stock. The figure depicts the optimal allocation to the stock ($\pi^{\text{opt}}_{\lambda,t}$) in the artificial market ($\widehat{\mathcal{M}}$), expressed in terms of the proportions of wealth, in relation to the corresponding ratio of optimal wealth to the individual-specific benchmark at three different points in time on the trading interval, for two individual CRRA (dotted line) and dual-CRRA (three non-dotted lines) investors. The results are based on the optimality conditions in $\widehat{\mathcal{M}}$ (cf. Proposition 3.3) under a reduction of the therein embedded time-and state-dependent shadow price processes, $\{\widehat{\alpha}_t, \widehat{\lambda}_t\}_{t\in[0,T]}$, to deterministic constants – in line with Example 4.3. The risk-profiles of the CRRA and dual-CRRA agents are characterised by, respectively, $\gamma = 4$ and $\gamma_d = 2$, $\gamma = 10$. In the financial market, we set $T = 10$, $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$.

hence that the losses in wealth due to transactions are marginal for CRRA investors. To provide an idea of the performance for agents whose portfolios are more fluctuating, which thus contrast those of the foregoing CRRA individuals, we have included the case of the dual-CRRA preference qualification. The discrepancy between γ_d and γ_u namely may inflate the wealth-dependency, i.e. variability, of the $\widehat{\mathcal{M}}$-optimal portfolio rule. For this comparatively more dynamic nature of the latter portfolio, consider Figure 1 .

4.3 Evaluation of Approximate Method

We complete this section by an evaluation of the in section 4.1 proposed approximate method's accuracy. For that purpose, we rely on the economic frameworks from Examples 4.1, 4.2 and 4.3. To sum up, a common Black-Scholes model in combination with CRRA and dual-CRRA preference qualifications delineates the environmental configuration within which we gauge the aforementioned precision. In all cases, we confine the space in which the dual controls must attain values to one that exclusively contains deterministic constants: $\widehat{\mathcal{P}}_{\mathcal{A}_{X_0,Y_0}} = \widehat{\mathcal{H}}_{\mathcal{A}_{X_0,Y_0}} \cap \mathbb{R}$. Despite the fact that the examples presume (for illustrative purposes that were at the heart of that subsection) a trading strategy for which $M = T$ holds, i.e. one in which the investor adjusts his/her holdings in both the money market

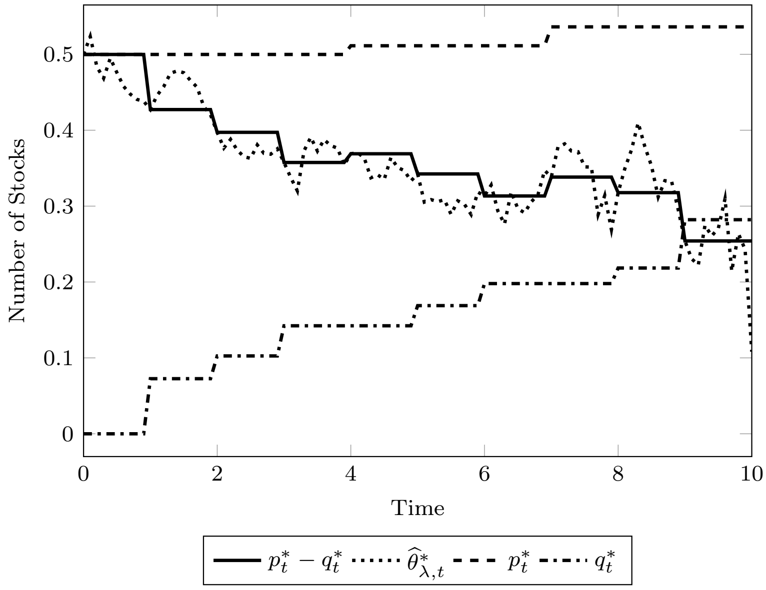

Figure 2. Approximate purchases and sales of stock. The figure depicts the trajectories of the approximations to the optimal purchases (p_t^*) and sales (q_t^*) of the stock, corresponding to a single sample path, for a $(\gamma_d, \gamma_u) = (2, 10)$ dual-CRRA agent (see Example 4.3) – as well as the difference between these trajectories: $p_t^* - q_t^*$. In addition to this, it displays the trajectory of the optimal continuous purchases of the stock, based on the optimality conditions in $\widehat{\mathcal{M}}$ (cf. Proposition 3.3) under a reduction of the therein embedded time-and state-dependent shadow price processes, $\{\widehat{\alpha}_t, \widehat{\lambda}_t\}_{t\in[0,T]}$, to deterministic constants, i.e. $\{\widehat{\theta}_{\lambda,t}^*\}$. The approximation is based on a strategy that is outlined in Example 4.3: the investor adjusts his/her holdings in the money market account and in the risky asset at the start of each year. In the financial market, we set $T = 10$, $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$.

account and the single risky asset at the start of each year on the trading interval, we will vary M in the present evaluation. Depending on the magnitude of these costs, we consider the following strategies: buy-and-hold ($M = 1$), adjustment at the start of every single year ($M = T$), adjustment at the start of every two years ($M = T/2$), adjustment twice at equidistant points in time ($M = 2$), and adjustment five times at equidistant points in time ($M = 5$). Moreover, we assume that $\lambda_1 = \lambda_2$ holds and examine the situation for proportional transaction costs being equal to 1%, 5%, 10%, and 30%. Inspired by the parameter initialisation of Brennan and Xia (2002) and van Bilsen et al. (2019), we fix $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$. In addition to this, we employ two time-horizons, $T = 10$ and $T = 40$. All processes are simulated by Monte-Carlo and a standard Euler scheme: we use $100,000$ sample paths and $T * 10$ time-steps.

In order to make the dynamic evolution of the approximated trading strategies, which straightforwardly underscores the entire subsequent analysis, in relation to their artificial continuous counterparts tangible, we demonstrate in Figure 2 the development of the relevant strategies corresponding to a single sample-path of a $\gamma_d = 2$, $\gamma_u = 10$ dual-CRRA agent. The depicted approximated policies cohere with the specification of the trading dates as outlined in Remark 4.3: $M = T$ for $T = 10$. Indeed, from the plot, we are able to

γ	$T = 10$ 1.5	2	5	10	$T = 40$ 1.5	2	5	10
TC 1%								
B&H	1.900	2.585	3.392	3.185	3.575	7.784	11.178	9.764
2-year	3.042	3.484	2.916	2.276	3.491	4.330	3.878	3.007
1-year	3.777	4.503	3.604	2.671	4.101	5.230	4.503	3.355
TC 5%								
B&H	2.215	2.750	3.170	2.977	3.889	7.946	11.123	9.712
2-year	10.561	11.654	8.700	6.300	12.928	16.522	13.757	9.956
1-year	15.660	17.367	12.838	9.144	17.335	22.621	18.780	13.457
TC 10%								
B&H	2.708	2.950	2.931	2.776	4.759	8.474	11.185	9.734
2 times	10.026	9.916	7.047	5.368	7.504	10.175	10.6275	8.833
5 times	22.168	21.779	14.777	10.653	13.552	16.3307	13.767	10.372
TC 30%								
B&H	0.936	0.974	1.243	1.764	10.810	12.486	12.351	10.475
2 times	11.671	10.509	8.429	7.920	23.861	24.264	18.653	14.303
5 times	28.724	25.784	20.259	18.455	46.126	46.248	33.528	24.424

Table 1. Annual welfare losses for CRRA agents. The table reports the annual welfare losses, i.e. the annualized equivalents of the compensating variations expressed in terms of basis points of the agent's initial endowment, that the CRRA investor suffers from executing the approximated trading strategy. We consider two trading horizons: $T = 10$ and $T = 40$. The different panels display the outcomes for $\lambda_1 = \lambda_2 = 0.01$ (*TC* 1%), $\lambda_1 = \lambda_2 = 0.05$ (*TC* 5%), $\lambda_1 = \lambda_2 = 0.1$ (*TC* 10%) and $\lambda_1 = \lambda_2 = 0.3$ (*TC* 30%). The rows in the first two panels showcase the findings for a buy-and-hold strategy (B&H), and the strategies in which the agent adjusts the portfolio at the start of every 2 years (2-year) and every single year (1-year). The rows in the second two panels demonstrate the findings for a buy-and-hold strategy (B&H), and the strategies in which the portfolio is adjusted twice (2 times) and five times (5 times) at equidistant points in time. The different levels of risk-aversion are given by the "γ" row. In the financial market, we set $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$.

discern that the approximated decision rules are adjusted at the start of each year and accordingly defined in a step-wise fashion such that their difference sticks rather close to the fictitious optimal equivalent. The trajectories for alternative values of M can on the basis of the displayed dynamics be easily imagined. To measure the accuracy of the approximate method in application to the model-specific situations at hand, we focus on the annualised values of the compensating variations, expressed in terms of basis points of the investor's initial endowment: $[(\mathcal{CV}/\widehat{X}_0)^{\frac{1}{T}} - 1] \times 10,000$. In Tables 1 and 2, we report these quantities for a variety of, respectively, four CRRA and four dual-CRRA investors. In Table 3, for descriptive reasons, we present the lower and upper bounds on the optimal value function for a limited selection of the in Tables 1 and 2 uttered cases.

Let us first of all focus on the findings for the CRRA investors in Table 1. Irrespective of the investor's risk-profile, the chosen trading horizon, and the size of proportional transaction costs, we notice that the values under the most optimal choice of trading dates, i.e. the one that is accompanied by the lowest value for the present welfare loss amongst

γ_d, γ_u	$T = 10$				$T = 40$			
	10, 2	6, 4	4, 6	2, 10	10, 2	6, 4	4, 6	2, 10
TC 1%								
B&H	2.998	3.019	3.118	10.492	10.349	11.813	12.722	12.075
2-year	2.092	2.423	2.670	5.660	3.600	4.584	5.565	6.984
1-year	2.487	3.025	3.470	6.641	3.952	5.126	6.304	7.905
TC 5%								
B&H	2.806	2.828	2.875	8.254	10.268	11.703	12.611	11.907
2-year	6.131	7.692	9.186	15.186	10.503	13.525	16.359	19.149
1-year	8.975	11.439	13.830	22.470	13.999	18.100	21.891	25.377
TC 10%								
B&H	2.629	2.636	2.660	5.900	10.255	11.687	12.638	12.073
2 times	5.222	6.280	7.376	11.321	9.351	10.976	12.340	12.938
5 times	10.507	13.253	16.114	23.675	10.883	13.601	16.165	18.748
TC 30%								
B&H	1.747	1.337	1.067	0.843	10.883	12.491	13.870	15.026
2 times	7.900	8.195	8.687	10.408	14.699	17.979	21.243	26.127
5 times	18.443	19.615	21.119	25.685	24.802	31.380	37.957	47.807

Table 2. Annual welfare losses for dual-CRRA agents. The table reports the annual welfare losses, see Table 1 for an according definition of this annual welfare loss, that the dual-CRRA investor suffers from executing the approximated trading strategy. We consider two trading horizons: $T = 10$ and $T = 40$. The different panels display the outcomes for $\lambda_1 = \lambda_2 = 0.01$ (*TC* 1%), $\lambda_1 = \lambda_2 = 0.05$ (*TC* 5%), $\lambda_1 = \lambda_2 = 0.1$ (*TC* 10%) and $\lambda_1 = \lambda_2 = 0.3$ (*TC* 30%). The rows in the first two panels showcase the findings for a buy-and-hold strategy (B&H), and the strategies in which the agent adjusts the portfolio at the start of every 2 years (2-year) and every single year (1-year). The rows in the second two panels demonstrate the findings for a buy-and-hold strategy (B&H), and the strategies in which the portfolio is adjusted twice (2 times) and five times (5 times) at equidistant points in time. The different risk-profiles of the dual-CRRA investors are spelled out by the "γ_d, γ_u" row. In the financial market, we set $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$.

the showcased ones, roughly vary between 2 and 12 basis points. This concretely implies that if the investor relegates the implementation of the strategy to a representative external party, he/she would be obliged to pay on an annual basis between 2 and 12 basis points of his/her initial endowment in order to be assured of acquiring an (expected) amount of terminal wealth that is identical to the amount in case of the execution of the truly optimal investment policy. In view of this interpretation, conditional on the choice for the "optimal" set of trading dates, we may conclude that the losses that are incurred on account of the implementation of these approximated allocations are negligible. Correspondingly, we may state that the approximated policies constitute near-optimal trading strategies. The latter statement is strengthened by the (illustrative howbeit certainly representative) findings that are reported in Table 3. To be more precise, most of the therein reported confidence bands for the estimated primal lower bounds on the actual optimal value functions encapsulate or are at least very close to the analytical dual upper bounds.

γ_d, γ_u	CRRA 1.5, 1.5	CRRA 2, 2	CRRA 5, 5	CRRA 10, 10	Dual-CRRA 10, 2	Dual-CRRA 6, 4	Dual-CRRA 4, 6	Dual-CRRA 2, 10
TC 1%								
LB	−0.339	−0.363	−0.784	−2.968	−2.937	−1.008	−0.613	−0.372
UB	−0.336	−0.360	−0.774	−2.911	−2.916	−0.994	−0.606	−0.365
CI-L	−0.346	−0.371	−0.793	−2.999	−3.036	−1.019	−0.620	−0.378
CI-U	−0.331	−0.356	−0.773	−2.937	−2.839	−0.996	−0.605	−0.365
TC 5%								
LB	−0.306	−0.394	−0.828	−3.167	−3.132	−1.065	−0.648	−0.403
UB	−0.304	−0.390	−0.815	−3.088	−3.092	−1.049	−0.640	−0.393
CI-L	−0.314	−0.400	−0.836	−3.195	−3.221	−1.075	−0.655	−0.409
CI-U	−0.297	−0.387	−0.819	−3.138	−3.043	−1.054	−0.641	−0.397

Table 3. Upper bounds, lower bounds, and confidence intervals. The table reports the lower (LB) and upper (UB) bounds on the optimal value function corresponding to the most accurate approximations, based on the magnitudes for "AL" in Tables 1 and 2 (i.e. the most accurate approximation is the one with the lowest value for AL), for CRRA and dual-CRRA agents that face a time-horizon of 10 years ($T = 10$). Additionally, it comprises the confidence intervals for the estimated lower bounds on the optimal value function. The first and second panels contain the resutls for, respectively, the $\lambda_1 = \lambda_2 = 0.01$ (TC 1%) and $\lambda_1 = \lambda_2 = 0.05$ (TC 5%) cases. The different risk-profiles of the dual-CRRA investors are spelled out by the "γ_d, γ_u" row; we note here that $\gamma_d = \gamma_u$ spawns the ordinary CRRA utility function. In the financial market, we set $\widehat{X}_0 = 1$, $r = 0.01$, $\mu = 0.04$, $\sigma = 0.16$, $\mu_\Pi = 0.05$, $\sigma_\Pi = 0.01$.

Let us now shed some light on the differences in performance and their relation with the risk-profiles, selected trading dates, magnitude of transaction costs, and trading horizons. The differences in performance do not vary much over the choice for the coefficient of risk-aversion. At most, one would be able to state that for lower levels of risk-aversion (i.c. $\gamma = 1.5, 2$), under small-sized proportional transaction costs (i.c. 1% and 5%), it is slightly more optimal to opt for a smaller set of trading dates. Namely, due to these comparatively small coefficients of risk-aversion, the artificial optimal portfolios are more sensitive to economic shocks, which may increase the magnitude of the relevant transactions and thus of the incurred costs involved with these. On the grounds of this moderately more dynamic portfolio, it is for these agents more sensible to trade less frequently. One may apply the opposite reasoning to the $\gamma = 5, 10$ agents in case of small transaction costs, to explain the more optimal performance of strategies that involve frequent trading.

For all agents, we observe that substantial increases in the proportional transaction costs, from 1% and 5% to 10% and 30%, force them naturally to opt for an ordinary buy-and-hold strategy. The more or less stable portfolios inherent in the CRRA preferences ensure that the buy-and-hold strategy performs in general quite great. In fact, the less fluctuating a portfolio is, the less one ought to adjust its composition in an effort to arrive at optimality. Pertaining to the trading horizon, the discernible increase therein gives *ceteris paribus* rise to larger welfare losses. This phenomenon is intuitive, since the difference between the approximation and the optimal analogue is subject to a longer period of time, which therefore magnifies the approximation error. From the preceding, it is clear that the choice for the amount of trading dates and its effect on the approximate method's performance is highly dependent on the magnitude of the transaction costs. That is, for lower levels of transaction costs, it is more optimal to choose a larger amount of trading dates; the opposite holds for larger levels of transaction costs.

Regarding the findings for the dual-CRRA agents in Table 2, we observe that the values for the annual welfare losses under the most optimal choice of trading dates vary roughly between 3 and 14 basis points. On the basis of these values, we are again able to conclude that the approximated trading strategies are near-optimal. This assertion is likewise corroborated by the values that are reported in Table 3. With regard to the relation between the approximate method's performance and the relevant model-parameters, we can be short. Namely, as to the relationship between this performance and the trading horizons, the selected number of trading dates, and the magnitude of the proportional transaction costs, we observe similar patterns as in the CRRA case, to which one consequently may apply similar reasoning.[38] The most salient difference is, however, visible in the fact that the dual-CRRA portfolios are in comparison to the CRRA's substantially more dynamic in nature. This more variable character forces the dual-CRRA agents in situations with low transaction costs to opt for an increased number of trading dates. Naturally, this effect mitigates once one raises the size of the proportional transaction costs.

5 Conclusion

This paper has proposed a dual-control approximate method in multi-dimensional financial markets with frictions for acquiring closed-form expressions to the optimal trading rules. This procedure predominantly rests on the formulation of an auxiliary economic environment in which the presence of transaction costs is relaxed. In this fictitious market, one is therefore able to resort to standard (martingale) techniques with the purpose of acquiring analytical expressions for the corresponding investment problem. These closed-form expressions embed shadow price processes that are in general analytically unavailable, and are due to their continuously variable nature infeasible in markets with frictions. So as to "prune" or project these in principle inadmissible trading policies into admissible counterparts, we (i) confine the space in which the shadow price processes must live and (ii) simultaneously discretise the continuous strategy by mandating the finite-horizon investor to only trade at a predetermined set of dates. Ultimately, we are able to assess the quality of the ensuing decision rule by comparing the (monetised) discrepancy between the related primal lower and dual upper bounds on the actual optimal value function.

[38] In short, Table 2 shows that the performance amongst the risk-profiles is more or less equal. However, for more conservative risk-profiles ($\gamma_d > \gamma_u$), it is moderately more optimal to select a smaller set of trading dates. Moreover, ncreases in the magnitude of transaction costs force the agents to trade less frequently. Last, increases in the trading horizon logically result in less accurate approximations.

To evaluate the potential accuracy of our proposed approximation, we have cast the matter into a technically simple one-dimensional Black-Scholes environment. In addition to that, we have assumed that the preferences of the investor are described by a CRRA function or by a dual-CRRA qualification. In the case of CRRA preferences, closed-form expressions are analytically troublesome to acquire; in the case of dual-CRRA preferences, closed-form expression are not available at all. The results have demonstrated that the approximating routine performs well, and thus that the approximate strategies are near-optimal. In fact, the annual welfare losses that an investor would suffer due to implementation of the (most optimal) approximate policy would roughly vary between a negligible 2 and 12 basis points of the individual's initial endowment. Moreover, the results have shown that there exists a negative relationship between the magnitude of transaction costs and the amount of pre-selected trading dates, in terms of the approximation's accuracy; the same holds for the relationship that involves time-horizons rather than trading dates. Last, since the truly optimal portfolios for dual-CRRA agents are more dynamic in nature than those for the CRRA individuals, the latter ones are more rapidly inclined to a buy-and-hold-like strategy than the former ones for larger levels of proportional transaction costs.

References

Altarovici, A., Reppen, M., and Soner, H. M. (2017). Optimal consumption and investment with fixed and proportional transaction costs. *SIAM Journal on Control and Optimization*, 55(3):1673–1710.

Balduzzi, P. and Lynch, A. W. (1999). Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics*, 52(1):47–78.

Bardhan, I. (1994). Consumption and investment under constraints. *Journal of Economic Dynamics and Control*, 18(5):909–929.

Basak, S. (2002). A comparative study of portfolio insurance. *Journal of Economic Dynamics and Control*, 26(7-8):1217–1241.

Basak, S., Shapiro, A., and Teplá, L. (2006). Risk management with benchmarking. *Management Science*, 52(4):542–557.

Bayraktar, E. and Yu, X. (2019). Optimal investment with random endowments and transaction costs: Duality theory and shadow prices. *Mathematics and Financial Economics*, 13(2):253–286.

Benedetti, G., Campi, L., Kallsen, J., and Muhle-Karbe, J. (2013). On the existence of shadow prices. *Finance and stochastics*, 17(4):801–818.

Bick, B., Kraft, H., and Munk, C. (2013). Solving constrained consumption–investment problems by simulation of artificial market strategies. *Management Science*, 59(2):485–503.

Bouchard, B. (2002). Utility maximization on the real line under proportional transaction costs. *Finance and Stochastics*, 6(4):495–516.

Brennan, M. J. and Xia, Y. (2002). Dynamic asset allocation under inflation. *The Journal of Finance*, 57(3):1201–1238.

Choi, J. H., Sirbu, M., and Zitkovic, G. (2013). Shadow prices and well-posedness in the problem of optimal investment and consumption with transaction costs. *SIAM Journal on Control and Optimization*, 51(6):4414–4449.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *Journal of political Economy*, 94(4):842–862.

Cox, J. C. and Huang, C.-F. (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory*, 49(1):33–83.

Cox, J. C. and Huang, C.-F. (1991). A variational problem arising in financial economics. *Journal of mathematical Economics*, 20(5):465–487.

Cvitanić, J. and Karatzas, I. (1992). Convex duality in constrained portfolio optimization. *Annals of applied probability*, 2(4):767–818.

Cvitanić, J. and Karatzas, I. (1993). Hedging contingent claims with constrained portfolios. *The Annals of Applied Probability*, pages 652–681.

Cvitanić, J. and Karatzas, I. (1996). Hedging and portfolio optimization under transaction costs: a martingale approach. *Mathematical finance*, 6(2):133–165.

Cvitanić, J., Pham, H., and Touzi, N. (1999). A closed-form solution to the problem of super-replication under transaction costs. *Finance and stochastics*, 3(1):35–54.

Cvitanić, J. and Wang, H. (2001). On optimal terminal wealth under transaction costs. *Journal of Mathematical Economics*, 35(2):223–231.

Czichowsky, C., Muhle-Karbe, J., and Schachermayer, W. (2014). Transaction costs, shadow prices, and duality in discrete time. *SIAM Journal on Financial Mathematics*, 5(1):258–277.

Czichowsky, C., Schachermayer, W., et al. (2016). Duality theory for portfolio optimisation under transaction costs. *The Annals of Applied Probability*, 26(3):1888–1941.

Czichowsky, C., Schachermayer, W., and Yang, J. (2017). Shadow prices for continuous processes. *Mathematical Finance*, 27(3):623–658.

Dai, M., Jiang, L., Li, P., and Yi, F. (2009). Finite horizon optimal investment and consumption with transaction costs. *SIAM Journal on Control and Optimization*, 48(2):1134–1154.

Davis, M. H. (1997). Option pricing in incomplete markets. *Mathematics of derivative securities*, 15:216–226.

Davis, M. H. and Norman, A. R. (1990). Portfolio selection with transaction costs. *Mathematics of operations research*, 15(4):676–713.

de Palma, A. and Prigent, J.-L. (2008). Utilitarianism and fairness in portfolio positioning. *Journal of Banking & Finance*, 32(8):1648–1660.

de Palma, A. and Prigent, J.-L. (2009). Standardized versus customized portfolio: a compensating variation approach. *Annals of Operations Research*, 165(1):161.

Deelstra, G., Pham, H., and Touzi, N. (2001). Dual formulation of the utility maximization problem under transaction costs. *Annals of Applied Probability*, pages 1353–1383.

Detemple, J. B. and Rindisbacher, M. (2005). Closed-form solutions for optimal portfolio selection with stochastic interest rate and investment constraints. *Mathematical Finance*, 15(4):539–568.

Detemple, J. B. and Rindisbacher, M. (2009). Dynamic asset allocation: Portfolio decomposition formula and applications. *The Review of Financial Studies*, 23(1):25–100.

Detemple, J. B. and Zapatero, F. (1991). Asset prices in an exchange economy with habit formation. *Econometrica: Journal of the Econometric Society*, pages 1633–1657.

Di Nunno, G. and Øksendal, B. (2009). Optimal portfolio, partial information and malliavin calculus. *Stochastics: An International Journal of Probability and Stochastics Processes*, 81(3-4):303–322.

Gao, J. (2008). Stochastic optimal control of dc pension funds. *Insurance: Mathematics and Economics*, 42(3):1159–1164.

Gennotte, G. and Jung, A. (1994). Investment strategies under transaction costs: the finite horizon case. *Management Science*, 40(3):385–404.

Gerhold, S., Guasoni, P., Muhle-Karbe, J., and Schachermayer, W. (2014). Transaction costs, trading volume, and the liquidity premium. *Finance and Stochastics*, 18(1):1–37.

Gerhold, S., Muhle-Karbe, J., and Schachermayer, W. (2012). Asymptotics and duality for the davis and norman problem. *Stochastics An International Journal of Probability and Stochastic Processes*, 84(5-6):625–641.

Gerhold, S., Muhle-Karbe, J., and Schachermayer, W. (2013). The dual optimizer for the growth-optimal portfolio under transaction costs. *Finance and Stochastics*, 17(2):325–354.

Gu, L., Lin, Y., and Yang, J. (2017). On the existence of shadow prices for optimal investment with random endowment. *Stochastics*, 89(6-7):1082–1103.

Herczegh, A., Prokaj, V., et al. (2015). Shadow price in the power utility case. *The Annals of Applied Probability*, 25(5):2671–2707.

Højgaard, B. and Taksar, M. (1998). Optimal proportional reinsurance policies for diffusion models with transaction costs. *Insurance: Mathematics and Economics*, 22(1):41–51.

Hu, Y., Imkeller, P., and Müller, M. (2005). Utility maximization in incomplete markets. *Annals of Applied Probability*, 15(3):1691–1712.

Jouini, E. and Kallal, H. (1995). Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory*, 66(1):178–197.

Kabanov, Y. and Safarian, M. (2009). *Markets with transaction costs: Mathematical Theory*. Springer Science & Business Media.

Kabanov, Y. M. (1999). Hedging and liquidation under transaction costs in currency markets. *Finance and Stochastics*, 3(2):237–248.

Kallsen, J. and Li, S. (2013). Portfolio optimization under small transaction costs: a convex duality approach. *arXiv preprint arXiv:1309.3479*.

Kallsen, J. and Muhle-Karbe, J. (2011). Existence of shadow prices in finite probability spaces. *Mathematical Methods of Operations Research*, 73(2):251–262.

Kallsen, J. and Muhle-Karbe, J. (2015). Option pricing and hedging with small transaction costs. *Mathematical Finance*, 25(4):702–723.

Kallsen, J. and Muhle-Karbe, J. (2017). The general structure of optimal investment and consumption with small transaction costs. *Mathematical Finance*, 27(3):659–703.

Kallsen, J., Muhle-Karbe, J., et al. (2010). On using shadow prices in portfolio optimization with transaction costs. *The Annals of Applied Probability*, 20(4):1341–1358.

Kamma, T., Balter, A., De Jong, F., and Pelsser, A. (2020). Investing for retirement with an explicit benchmark. *Working Paper*.

Kamma, T. and Pelsser, A. (2019). Near-optimal dynamic asset allocation in financial markets with trading constraints. *Available at SSRN 3403176*.

Karatzas, I. and Kou, S. G. (1996). On the pricing of contingent claims under constraints. *The annals of applied probability*, pages 321–369.

Karatzas, I. and Kou, S. G. (1998). Hedging american contingent claims with constrained portfolios. *Finance and Stochastics*, 2(3):215–258.

Karatzas, I., Lehoczky, J. P., and Shreve, S. E. (1987). Optimal portfolio and consumption decisions for a small investor on a finite horizon. *SIAM Journal on Control and Optimization*, 25(6):1557–1586.

Karatzas, I. and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus.* New York: Springer-Verlag.

Karatzas, I. and Shreve, S. E. (1998). *Methods of mathematical finance*, volume 39. Springer.

Keppo, J., Meng, X., and Sullivan, M. G. (2007). A computational scheme for the optimal strategy in an incomplete market. *Journal of Economic Dynamics and Control*, 31(11):3591–3613.

Keppo, J. and Peura, S. (1999). Optimal portfolio hedging with nonlinear derivatives and transaction costs. *Computational Economics*, 13(2):117–145.

Klein, I. and Rogers, L. C. (2007). Duality in optimal investment and consumption problems with market frictions. *Mathematical Finance*, 17(2):225–247.

Kramkov, D. and Schachermayer, W. (1999). The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Annals of Applied Probability*, 9(3):904–950.

Kühn, C. and Stroh, M. (2010). Optimal portfolios of a small investor in a limit order market: a shadow price approach. *Mathematics and Financial Economics*, 3(2):45–72.

Lakner, P. and Nygren, M. (2006). Portfolio optimization with downside constraints. *Mathematical Finance*, 16(2):283–299.

Lamberton, D., Pham, H., and Schweizer, M. (1998). Local risk-minimization under transaction costs. *Mathematics of Operations Research*, 23(3):585–612.

Liu, H. (2004). Optimal consumption and investment with transaction costs and multiple risky assets. *The Journal of Finance*, 59(1):289–338.

Liu, H. and Loewenstein, M. (2002). Optimal portfolio selection with transaction costs and finite horizons. *The Review of Financial Studies*, 15(3):805–835.

Magill, M. J. and Constantinides, G. M. (1976). Portfolio selection with transactions costs. *Journal of economic theory*, 13(2):245–263.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: the continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.

Nualart, D. (2006). *The Malliavin calculus and related topics*, volume 1995. Springer.

Pliska, S. R. (1986). A stochastic calculus model of continuous trading: optimal portfolios. *Mathematics of Operations Research*, 11(2):371–382.

Rogers, L. C. (2003). Duality in constrained optimal investment and consumption problems: a synthesis. In *Paris-Princeton Lectures on Mathematical Finance 2002*, pages 95–131. Springer.

Rogers, L. C. (2013). *Optimal investment.* Springer.

Rokhlin, D. B. (2013). On the game interpretation of a shadow price process in utility maximization problems under transaction costs. *Finance and Stochastics*, 17(4):819–838.

Schachermayer, W. (2004). Utility maximisation in incomplete markets. In *Stochastic methods in finance*, pages 255–293. Springer.

Shreve, S. E. and Soner, H. M. (1994). Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, pages 609–692.

Soner, H. M., Shreve, S. E., Cvitanic, J., et al. (1995). There is no nontrivial hedging portfolio for option pricing with transaction costs. *The Annals of Applied Probability*, 5(2):327–355.

van Bilsen, S., Bovenberg, A. L., and Laeven, R. J. (2019). Consumption and portfolio choice under internal multiplicative habit formation. *Journal of Financial and Quantitative Analysis*, pages 1–38.

Zakamouline, V. I. (2002). Optimal portfolio selection with transaction costs for a CARA investor with finite horizon. *Norwegian School of Economics and Business Administration.*

Zakamouline, V. I. (2006). Efficient analytic approximation of the optimal hedging strategy for a european call option with transaction costs. *Quantitative Finance*, 6(5):435–445.