Continuous Time Mean-Variance Approximations

Thesis Msc Quantitative Finance and Actuarial Science

Tilburg University

Supervisor: Dr. Nikolaus Schweizer

Second Reader: Dr. Anne Balter

Thomas Kohlbacher, 185281

October 8, 2023

The quality of Mean-Variance approximations in continuous time appears relatively understudied. In contrast to the extensively studied single-period or multi-period myopic cases, in a continuous time setting optimal Mean-Variance policies are found to perform very poorly in approximating optimal policies under comparable expected utility criteria. This thesis traces these results to the underlying mechanisms of the Martingale method, which is used to compute optimal pre-commitment policies as the overall optimal policies under the given criteria in continuous time. Based on theoretical assessment and analysis of the underlying wealth processes generated, it is argued that under the optimal pre-commitment solution, the Mean-Variance investor fails to satisfy criteria of Second-Order stochastic dominance, which, in turn, causes considerable loss in value to expected utility maximizers, whose policies are to be approximated. Furthermore, it is shown that while costly to the Mean-Variance investor herself, time-inconsistency of the Mean-Variance criterion is indirectly value-enhancing with respect to Mean-Variance approximations, as the dynamically consistent Mean-Variance solution approximates optimal pre-commitment solutions under the respective expected utility criteria more closely. Analysis of investment patterns and wealth processes generated is provided to investigate the underlying reasons for the above. Furthermore, the quality of approximation is quantified in terms of certainty equivalent returns, in order to provide comparable measures of value.

Contents

1 Introduction

Much has been written on the differences between optimal Mean-Variance and Expected Utility investment. Continuous time Mean-Variance investment, however, challenges us to once again start at first principles and reassess the adherence of the Mean-Variance investor to principles of rational economic decision making. We will see that even under conditions of normality, the overall optimal Mean-Variance policy differs considerably from the strategy, which would be chosen under a comparable expected utility criterion, and trace this difference to the Mean-Variance investor's failure to adhere to axioms of rationality, as they have been developed within the Expected Utility framework. Moreover, while generally seen as implying costs[1], the adaptation of the dynamic Mean-Variance investor's objective function given the time-inconsistency of her criterion[2] will be shown to improve her ability to approximate optimal Expected Utility policies. The picture then emerges of an investor, whose optimal policy relative to her own criterion is starkly rejected by expected utility maximizers and whose ability to serve the latter is improved by constraints, which are seen as costly from her own perspective. To set the stage in this introductory section, we will first briefly review the criticism of Mean-Variance investment from the view of Expected Utility Theory and discuss Mean-Variance investment's failure with respect to Stochastic Dominance. We will then identify the relevant gaps in the current literature with respect to continuous time Mean-Variance investing and argue that normality is not sufficient to constrain the continuous-time Mean-Variance investor with a view to ensuring adherence to principles of economic rationality. Finally, the issue of time-inconsistency will briefly be discussed.

From a theoretical point of view, interest in the differences between Mean-Variance and Expected Utility investing, as well as the ability of strategies derived within the former to approximate optimal strategies within the latter paradigm, derives from the normative appeal of the axioms, on which the Expected Utility framework is built (Machina, 1990). Furthermore, for wide sets of reasonable utility functions, Expected Utility investing is equivalent to the satisfaction of criteria of First-and Second Order Stochastic Dominance (Hadar & Russel, 1969), both of which provide rather broad additional criteria for rational economic decision-making (Hanoch & Levy, 1969). Hence, it seems natural to assess optimal Mean-Variance portfolio choices against Expected Utility criteria, if economic rationality, as embodied in the underlying axioms and requirements, is regarded as important.
Mean-Variance investing has been criticised by proponents of Expected Utility analysis for its failure to reliably yield indifference curves in the (μ, σ^2) plane[3] (Borch, 1969), which, in turn, prevents representability within the Expected Utility framework (Machina, 1990) and casts doubt on the economic rationality of the underlying choices. (Borch, 1969; Machina, 1990; Lindley & Johnstone, 2013) Existence of indifference curves in the (μ, σ^2) is limited to cases of quadratic utility, normally distributed assets or, similarly, axiomatized derivations of indifference curves based on the presumption that all available choices are fully described in the eye of the investor by their means and variances. (Johnstone & Lindley, 2013; Baron, 1977). In the case of normality[4] or quadratic utility, expected utility becomes a function of means

[1] Björk & Murgoci (2010), see blow.
[2] Basak & Chabakauri (2010), see below.
[3] Or, that is, of course the (μ, σ)-plane.
[4] That is, in this case, using the fact that the (single period) Mean-Variance optimization problem is equivalent to the problem of maximizing expected CARA, that is negative exponential, utility (see Munk, 2017).

and variances of terminal prortfolio wealth, allowing for the derivation of linear indifference curves in the (μ, σ^2)-plane (Lindley & Johnstone, 2013; Munk, 2017; Baron, 1977). Moreover, second-order Taylor approximations of expected future utility become exact (Munk, 2017) and expected utility with n-th order polynomial utility functions is fully described by the first n moments of the return distribution and derivatives of the utility function, respectively (Hadar & Russel, 1969). In case of general utility functions, however, the requirement of normality restricts the use of stochastic portfolio policies and essentially confines the investor to using linear portfolio weights (Johnstone & Lindley, 2013).
With reference to criteria of Stochastic Dominance, Mean-Variance investing has, furthermore, been found to violate intuitive criteria of basic economic rationality. On the one hand, variance, which captures the Mean-Variance investor's disutility from the risk related to stochastic prospects, would penalize the investor for risk that is purely to the upside. While no rational economic agent would reasonably choose a deterministic payment of 2 over a risky prospect of 2 or 4 with equal probability, the Mean-Variance investor's choice would depend on her individual risk-aversion parameter (Levy, 2006; Hanoch & Levy, 1969), thus violating First-Order stochastic dominance, as any investor with an increasing utility function would unequivocally have chosen the 'risk' (Hadar & Russel, 1969; Levy, 1990). Moreover, the Mean-Variance investor penalizes downside risk equally as she appreciates upside potential[5] (Machina & Rothschild, 1990). In contrast, risk-averse investors[6] would prefer portfolios, which, as seen from each potential outcome, put a larger probability on smaller downside risk on average, providing better protection against downside risk in this sense and leading to equivalent characterizations in terms of Second-Order stochastic dominance (Hadar & Russel, 1969; Levy, 1990; Rothschild & Stiglitz, 1970; Rothschild & Stiglitz, 1971)[7]. Building on Hanoch & Levy (1969)[8], a Mean-Variance investor may neglect downside risk protection, as a shift in probability mass to higher outcomes may at the same time lead to an increase in variance. Let risk X and probability mass functions $p(X)$ and $q(X)$ be given by:

X	0	50	100
p(X)	0.85	0.02	0.13
q(X)	0.6	0.28	0.12

Risk X under $q(X)$ Second-Order[9] stochastically dominates X under $p(X)$, as it offers greater downside risk protection in the sense discussed above at each possible outcome and is, in this sense, stochastically larger (Hadar & Russel, 1969). The Mean-Variance investor, however, is not able to express unequivocal preference for X under $q(X)$, as both its mean and variance, respectively, are larger under measure q, that is $\mu^p = 14 < \mu^q = 26$ and $\sigma_p^2 = 1154 < \sigma_q^2 = 1224$. While this shows that Second-Order stochastic dominance is not sufficient for preference under Mean-Variance objective functions, vice versa preference under Mean-Variance preferences is neither sufficient for Second-Order Stochastic Dominance (Hanoch & Levy, 1969). Let us consider an investor in times of turbulent stock-markets with a wealth of 10, who faces the

[5] Relatedly, variance has more generally been rejected as a reasonable risk-measure based on an axiomatic approach (see Artzner et al., 1999).

[6] That is, strict risk-aversion in terms of strict concavity yields a strict preference for certain pay-offs (see Hadar & Russel, 1969)

[7] An interesting parallel to these theoretical results are the behavioral results studied through Allais' paradox (see Allais, 1990)

[8] That is, in particular, their 'Examples' 1 and 2'

[9] First-Order Stochastic dominance does not hold in this case, as the probability mass is not always greater to the upside under q (see Hadar & Russel, 1969).

prospect of having his wealth halved to 5 at a probability of 0.6, or having it doubled with a probability of 0.4. The investor considers staying fully invested, or selling the position, investing 8 in a safe bond and 2 in portfolio of deep-out-of-the-money options[10], yielding vast upside potential at very low probability.

X	P(X)	Y	P(Y)
5	0.6	8	0.99
15	0.4	100	0.01

While risky prospect X is preferred to Y under Mean-Variance preferences ($\mu_X = 9 > \mu_Y = 8.92$, while $\sigma^2_X = 24 < \sigma^2_Y = 83.79$), X does not Second-Order stochastically dominate Y, as X performs worse in terms of downside risk for realizations between 5 and 8, whereas for realizations between 5 and 15, Y performs worse. As a result, there are Expected Utility maximizers, who would prefer Y over X (compare Hadar & Russel, 1969; Hanoch & Levy, 1969), such as a CRRA Expected Utility investor with $u(x) = \ln(x)$, as $\mathbb{E}[\ln(X)] = 2.05$ and $\mathbb{E}[\ln(Y)] = 2.1$.
If risky prospects are fully described by means and variance, however, Mean-Variance preference-orderings are equivalent to Second-Order stochastic dominance orderings[11] (Hanoch & Levy, 1969). More generally speaking, both First - and Second Order stochastic dominance provide broad efficiency criteria in that both criteria provide preference orders, ranking portfolios that are 'stochastically larger' higher than their inferior alternatives. (Hadar & Russel, 1969)
Thus, it is precisely at the stated objective of reducing the universe of available portfolios to an 'efficient set' prior to applying individual preferences and selecting optimal portfolios (Markovitz, 1952; Sharpe, 1964), where Mean-Variance analysis may fail in practice (Malavesi et al., 2021).

The latter discussion is also related to concerns from a practical view point concerning the difference between Mean-Variance and Expected Utility investing. If investment choices bear relevance for our consumption choices across time and states of the world[12], then the value of a risky prospect should be evaluated according to the individual utility derived from it, respectively (Bernoulli, 1954). Even if different individuals apply the same utility function to evaluate a given set of risks, their value from a given prospect may differ due to differences in their wealth, or their remaining portfolio, respectively (ibid). Hence, more objectifiable measures of value, such as through Mean-Variance objectives, may after all not be more universally acceptable, due to the inherent dependence of the evaluation of risks on individual circumstances. It is, thus, reasonable to evaluate in particular investment choices that are fundamentally linked to consumption choices, such as for example choices with respect to pension savings, according to Expected Utility criteria (see e.g. Bovenberg et al., 2007; Metselaar et al., 2022)[13]. At the same time, Mean-Variance analysis (Markovitz, 1952; Sharpe 1964) has become a fundamental building block of modern finance, with applications ranging from ex-ante valuation to ex-post

[10]That is, a 'convexity' type of trade (see e.g. Figlewski & Freund, 1994). While the appropriateness of Mean-Variance analysis has particularly been doubted with respect to these trades, however, reality may not be as clear cut (see Markowitz, 2014).

[11]Thus, confirming results in the previous paragraph that indifference curves may be derived and (single period) negative exponential expected utility maximization is equivalent to Mean-Variance optimization (Munk, 2017; Johnstone & Lindley, 2013)

[12]See for example Cochrane (2001) or Eeckhoudt et al. (2005).

[13]While the obvious question as to the choice of utility criterion then arises, a strong case can be made that there are utility functions, which represent preferences with certain generally acceptable characteristics, as for example argued in Rubinstein, M. (1976), and as also follows from the discussion above and in section 2 below.

performance evaluation (Cochrane, 2001; Pedersen, 2015). The appeal of Mean-Variance analysis derives not least from its relative simplicity and ease of application (Markovitz, 1952; Sharpe, 1964; Pulley, 1981), at least its standard-formulation in the single-period case, as well as the smaller burden it places on paramemter estimation both regarding the return distribution and individual risk-aversion (Markowitz, 2014). The question of the goodness of fit of investment decisions based on Mean-Variance criteria with respect to preferences of Expected Utility investors is the question of Mean-Variance approximations. (Levy & Markowitz, 1979)

In an early contribution concerning the assessment of Mean-Variance approximations, Levy & Markowitz (1979) showed that based on empirical returns, Mean-Variance approximations to different utility functions performed well in terms of correlations with a respective Expected Utility criterion. Pulley (1981 & 1983) and Kroll et al. (1984) expand on these results by maximizing Expected Utility criteria and their respective mean-variance approximation separately, to explicitly allow for differences in strategies in response to different optimization criteria (Kroll et al, 1984). Pulley (1981 & 1983) finds that approximations are close to perfect with respect to portfolio weights, returns generated and value derived under the Expected Utility criteria, noting however, that approximations become worse as variance increases. Kroll et al. (1984) likewise find that the optimal protfolios chosen and value achieved, respectively, under the Expected Utility criterion are similar between the two strategies. Kroll et al. (1984) also adduce evidence as to the robustness of results in case of non-normality, using empirical return data, for which normality is rejected. Pulley (1981 & 1983), in turn, suggests that his empirical results are robust, using simulations on the basis of log-normal and t-distributions. Markowitz (1991) provides further reflection. The authors' results are also further discussed in Pulley (1985) and Reid & Tew (1986), who furthermore attest to the superiority of certainty equivalents as a measure for comparison of utility across individuals (see also Kallberg & Ziemba, 1979). The limitations of using standard, single-period Mean-Variance optimization to approximate multi-period expected utility optimization have been noted by Harkansson (1971). While the scope of Mean-Variance analysis has subsequently expanded to analysis in multi-period settings, the assumption of myopic investors was maintained until the arrival of continuous-time Mean-Variance analysis (Richardson, 1989; Basak & Chabakauri, 2010)[14]. In a review of the ensuing debate, Markowitz (2014) further discusses that for reasonably well behaved return distributions, Mean-Variance approximations tend to perform well, for large risk-aversion, they may become problematic, even though the author notes that very large risk-aversion and the absence of a risk-free asset may have to be assumed to generate large differences, and discusses in how far simulations may complement empirical return distributions in order to test Mean-Variance approximations. While Markowitz extensively reviews the empirical literature and discusses the quality of single-period approximations, no reference is made to continuous time Mean-Variance analysis and respective quality of approximation in continuous time. More recently, Mean-Variance approximations have been challenged by the direct computation of sets of Second-Order stochastic dominance (SSD) - efficient portfolios. For example, Malavesi et al. (2021) find that portfolios on the Mean-Variance efficient frontier are partly outperformed by SSD efficient portfolios in terms of common performance measures such as the Sharpe ratio. As the authors discuss, the comparison with SSD-efficient portfolios provides a more encompassing test for the relative efficiency of Mean-Variance optimal portfolios and portfolios optimized under Expected utility criteria, respectively, as rather than picking specific utility functions for the comparison, as in the approaches discussed above, SSD concerns efficiency criteria for the entire class of risk-averse, non-satiable investors (ibid). The authors' results are in line with

[14]For a discussion and review of the literature, see Basak & Chabakauri, 2010.

Hodder et al.(2015), who find SSD-efficient portfolios, which slightly outperform Mean-Variance efficient portfolios according to their Sharpe ratio in out-of-sample analysis.[15]

Following the advent of solutions to portfolio choice problems in continuous time, the solution to continuous-time Mean-Variance problems in a complete market setting under the Martingale approach (Pliska, 1986; Cox & Huang, 1989) was first published by Richardson (1989), with Bajeux-Besnainou & Portait (1998) further deriving and analyzing the 'dynamic efficient frontier', whereas solutions under the dynamic programming approach (Merton, 1969; Merton, 1971) and in a more general market setting were derived by Basak & Chabakauri (2010). Myopic Mean-Variance strategies in a continuous-time setting, in turn, remain an active area of research to this day (Chen & Zhou, 2022; Balter et al., 2021). The dynamic hedging strategy with continuous rebalancing opportunities, which form part of the Martingale approach, lead to an expansion of the set of potential trading strategies, which may lead to efficiency gains relative to portfolio choices under the static Markowitz (1952) approach (Bajeux-Besnainou & Portait, 1988). However, the Mean-Variance criterion is plagued by time-inconsistency (Basak & Chabakauri, 2010; Björk & Murgoci, 2010), so that while theoretically superior, it cannot generally be said that the commitment required under the Martingale approach is realistic (Strotz, 1956; Basak & Chabakauri, 2010)[16]. Hence, analysis of Mean-Variance approximations should include continuous-time modelling and continuous-time Mean-Variance portfolio choice problems should be studied under both the Martingale and dynamic programming approaches, respectively.

The closeness of Mean-Variance approximations in continuous time appears relative understudied, as compared to the static and myopic discrete time case. Bielecki et al. (2005), Zhou & Li (2000) and Li et al. (2000) all include reference to continuous-time expected utility maximization and contrast Mean-Variance portfolio optimization with the latter, however, they do not provide any further discussion on the issue. Wang et al. (2007) study continuous-time CARA utility - and Mean-Variance investing within an insurance setting, whereas Basak & Chabakauri (2010) include a more detailed discussion of the differences between expected CARA utility- and Mean-Variance portfolio strategies, while also including a brief reference to the respective value functions in a continuous time dynamic programming setting, without comparing the two value functions in depth, however. Cvitanic et al. (2008), in turn, briefly touch upon a comparison of expected CRRA utility and Mean-Variance investing on the basis of a Martingale approach. However, the authors do not include any further discussion on the quality of Mean-Variance approximations in terms of value achieved for expected utility investors in continuous time settings. Zhao & Ziemba (2002) essentially ask the reverse question by assessing, which target expected return should be chosen by expected logarithmic utility investors under a Martingale approach so that the latter would outperform Mean-Variance investors in terms of financial wealth achieved. As noted by Chen & Zhou (2022), research has more recently shifted toward the study of dynamic programming solutions to the Mean-Variance investment problem in continuous time, and related problems of time-inconsistency. To the author's best knowledge, there is to this day no research that covers in-depth analysis of continuous-time Mean-Variance approximations on the basis of comparable value-criteria such as certainty equivalents.
Furthermore, due to the constraints it implies with respect to the policies that the Mean-

[15] Simaan (2014) and Lassance (2022) include further reviews of the more recent literature.

[16] The requirement for commitment should, however, not be seen as restrictive relative to the single-period setting, as the opportunity to deviate at intermediate rebalancing opportunities (Strotz, 1956) will challenge the latter solution in the similar way as the solution under the Martingale approach, see discussion below.

Variance investor may actually be able to follow (Strotz, 1956), the effect of time inconsistency on Mean-Variance approximations should be studied. Studies concerning optimal dynamic Mean-Variance portfolio choice in face of the criterion's time-inconsistency include Basak & Chabakauri (2010), Björk & Murgoci (2010), Björk et al. (2014) and Björk et al. (2017). While in particular Basak & Chabakauri (2010) provide in-depth discussion and intuition as to the effect of time-inconsistency on the optimal dynamic investment policy and value function, respectively, the effect of time-inconsistency certainty equivalents achievable for the Mean-Variance investor or the effect on the closeness of Mean-Variance approximations have not yet been studied.

This thesis will study Mean-Variance approximations in continuous time. As its aim is the study of fundamental mechanisms, the analysis will be restricted to the setting of Black-Scholes Markets. From a methodological viewpoint, it borrows heavily from Balter et al. (2021) in its use of certainty equivalent returns to evaluate trading strategies in continuous time and assess potential costs and commitment problems due to time-inconsistency. It contributes to the literature by extending the in-depth analysis of Mean-Variance approximations of CARA and CRRA utilities to the continuous time setting and by assessing the effect of the time inconsistency of Mean-Variance optimization on the closeless of approximations. As will be discussed, the risk-profile reflected in the Mean-Variance investor's pre-commitment (Strotz, 1956) strategy is strongly rejected by Expected Utility maximizers. Their disutility from the latter strategy will be traced to the pre-committed Mean-Variance investor's failure with respect to Second-Order Stochastic Dominance, whereas it will be argued that the roots of the latter violation lie in the mechanics of the Martingale approach. Hence, the Mean-Variance investor's constrained dynamic optimum is preferred by Expected Utility maximizers over the Mean-Variance investor's overall optimal strategy in a continuous time setting.

As to the structure of the following discussion, section 2 will introduce the financial market setting, individual utility - and objective functions, as well as parameterizations to obtain comparable levels of risk-aversion among all investors. Section 3 will briefly discuss the underlying optimization problems and derive optimal policies, respectively. Second 4 will study these policies in greater depth as to their risk-profile over time and the certainty equivalent returns that may be derived. Section 5 will conclude.

2 The Model Setting

2.1 The Financial Market

Investors are assumed to face Black-Scholes financial markets (Black & Scholes, 1973). That is, let $W(t)$, $0 \leq t \leq T$ be a standard Brownian Motion on a probability space $(\Omega, F, \mathbb{P})$ and $(\mathcal{F}_t)$ the filtration generated by the Brownian motion. The financial market consists of a single risky asset, the stock S, and a single risk-free asset, the bond B. The value of the stock is assumed to follow a Geometric Brownian Motion, whereas the value of the bond is assumed to grow exponentially at the deterministic short rate r. Thus, in differential form, asset prices follow the two processes

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{1}$$

and

$$dB_t = rB_t dt, \tag{2}$$

respectively. The assets are traded continuously and in frictionless markets. Unless explicitly stated otherwise, there are no short-selling constraints and investors may hold potentially unlimited long-or short positions in either asset. As a result, investors may also borrow and lend unlimited amounts at the fixed short rate r.

Within the given market, the investors' general optimization problem is to maximize expected utility - or value [17], respectively, from future financial wealth at the fixed terminal time point T by choosing an optimal portfolio based on the financial assets in the market. The strategies available to the agents are given by adapted processes $\pi_t^{S,B}$ and $\theta_t^{S,B}$, respectively, whereas $\pi_t^{S,B}$ denotes the fractions of financial wealth invested in the respective asset at time t, and $\theta_t^{S,B}$ the absolute exposure in terms of amounts invested. We assume that there is no other source of income besides total accumulated financial wealth X_t and there is no intermediate consumption reducing financial wealth at any intermediate time point. Hence, it holds that $\theta_t^S + \theta_t^B = X_t$ and $\pi_t^S + \pi_t^B = 1$, for all $t \in [0, T]$. The resulting wealth dynamics are given by

$$dX_t = X_t(\pi_t^B dB_t + \pi_t^S dS_t) \tag{3}$$

and

$$dX_t = \theta_t^B dB_t + \theta_t^S dS_t, \tag{4}$$

respectively. Imposing $\theta_t^B = X_t - \theta_t^S$ in (3) and $\pi_t^B = 1 - \pi_t^S$ in (4), renaming the relative and absolute exposure to the stock π_t and θ_t, respectively, and plugging in the asset price dynamics given in (1) and (2), we obtain the wealth dynamics

$$dX_t = X_t(r + \pi_t(\mu - r)dt + \pi_t \sigma dW_t) \tag{5}$$

and

$$dX_t = X_t r dt + \theta_t(\mu - r)dt + \theta_t \sigma dW_t. \tag{6}$$

[17]The outcome relative to the Mean-Variance criterion is described as 'value', whereas we refer to the 'utility' of the CARA - and CRRA investor, respectively (see below).

There are three investors in the market with different utility - and objective [18] functions and, relatedly, different risk-attitudes, respectively. They will be referred to as CARA-, CRRA - and Mean-Variance investors below. The CARA investor prefers constant absolute exposure to the risky asset in terms of amounts invested, irrespective of her current wealth level. (Pratt, 1964, Merton, 1969) Her preferences are represented by the negative exponential utility function[19]

$$u^{CA}(x) = -\frac{1}{\alpha}e^{-\alpha x} \quad \text{, with } \alpha > 0. \tag{7}$$

However, as utility functions are only defined up to positive affine transformations and the preferences reflected in utility functions are preserved by such transformations (see e.g. Mas-Colell et al., 1995), the same risk preferences are encoded by

$$u^{CA}(x) = -e^{-\alpha x} \quad \text{, with } \alpha > 0. \tag{8}$$

For computational convenience, we will make use of both forms of CARA-utility below. The CRRA investor, in turn, prefers constant exposure to proportional risk within her portfolio and, thus, constant investment fractions, irrespective of current wealth. (Pratt, 1964, Merton, 1969) Her underlying utility function is given by

$$u^{CR}(x) = \frac{x^{1-\gamma}}{1-\gamma} \quad \text{, with } \gamma > 1.^{20} \tag{9}$$

Finally, Mean-Variance preferences with respect to future - and as yet stochastic - wealth X_T are represented by the following objective function:

$$MV(X_T) = \mathbb{E}[X_T] - \frac{\delta}{2}\text{Var}[X_T]. \tag{10}$$

Our evaluation of the closeness of Mean-Variance approximations of expected utility investment, as well as the cost of time-inconsistency, is based on a comparison of the certainty equivalent, which the three investors would derive from each others' strategies, and analysis of the fundamental drivers behind differences in these certainty equivalents.

[18] The Mean-Variance objective function is not a utility function and does not have Expected Utility property (see e.g. Mas-Colell et al., 1995). Hence, we will refer to the Mean-Variance 'objective' or 'value' function rather than a utility function, as for the CARA and CRRA investors.

[19] We follow the convention to denote 'Bernoulli utility functions' over deterministic wealth x by $u(x)$, whereas $U_t(X_T)$ denote Von-Neumann - Morgenstern utility functions, that is the expectation conditional on information at time t over possible realizations of X_T of the utility of terminal wealth $u(X_T)$ (for terminology see Mas-Colell et al., 1995).

[20] As is commonly known (see e.g. Munk, 2017), it can be shown by use of L'Hôpital's rule that $u^{CR} = ln(x)$ for $\gamma = 1$. For L'Hôpital's rule to be applicable, the above utility function would have to be defined as $u^{CR}(x) = \frac{x^{1-\gamma}-1}{1-\gamma}$, which, in terms of utility functions is equivalent to the formulation chosen here. However, as we will always assume that $\gamma > 1$, the form given above is used for simplicity.

2.2 Parameterizations to obtain comparable levels of risk aversion

Individuals, whose utility functions differ by more than just affine transformations, generally differ as to their degree of risk-aversion (Pratt, 1964), which affects their choice between and evaluation of risky prospects. (Bernoulli, 1954) Thus, to allow for meaningful comparisons of certainty equivalents from a given risky prospect across investors, utility - and objective functions first have to be parameterized such that the implied degree of risk-aversion becomes comparable with respect to a given exogenous risk. This is complicated by the inter-temporal - and at times dynamic - nature of the optimization problem considered, as well as the particular character of the utility - and value functions concerned. This section, thus, motivates the particular choice of parameterization used in the ensuing analysis. It also discusses why and where having comparable degrees of risk aversion does not imply the same preferences among risky prospects or confine the three investors to choosing the same investment strategies.

Risk aversion can be measured in multiple ways. On the one hand, global measures of risk aversion, such as certainty equivalents, concern the investor's attitude towards risks at any scale, given the current level of wealth. Certainty equivalents represent the certain pay-off, CE, which would make the individual indifferent between receiving the given fixed amount and holding a proposed risk $\tilde{Z}$, respectively. (Eeckhoudt et al., 2005) To calculate certainty equivalents, the (future) risky prospect is evaluated by the investor at decision time t (Pratt, 1964), given reference-wealth level at any time $t^R \leq t$, X_{t^R}[21]. With this flexible definition of reference-wealth levels, the certainty equivalent can, thus, be computed via[22]

$$u(X_{t^R} + CE_t) = \mathbb{E}_t \left[u\left(X_{t^R} + \tilde{Z}_{t,T} \right) \right]^{23}, \tag{11}$$

with $\tilde{Z}_{t,T}$ being the risk, which materializes at future time T, as seen conditional on information at decision time t. If $t^R < t$, there appears to be some tension between the choice of reference level and the assumption that decisions are taken based on the information set at time t, as current wealth is part of the information set. At least within the given set of utility functions, this tension cannot be resolved. However, the thought experiment of adding a risky prospect to a given wealth level (Bernoulli, 1954) may equally be performed with reference to previous wealth levels. The effect of the choice of reference wealth level will further be discussed below. On the other hand, Pratt (1964) showed that global measures such as the certainty equivalent may locally be approximated by what became known as the Arrow-Pratt measure of (absolute) risk aversion (Eeckhoudt et al., 2005),

[21]Usually, the reference wealth level is given by current wealth at decision time t. However, as explained below, due to the character of the utility - and value functions discussed and the given purpose of parameterization, this more flexible understanding of 'reference' wealth level is used.

[22]The risks this thesis is concerned with are endogenous risks, whereas the endogeneity results from the fact that the size of the risk within the investment portfolio may be controlled by the exposure to the risky asset, as chosen by the investor. As argued by Briys et al. (1989), for endogenous risks the definition of certainty equivalents should be refined to reflect the fact that optimally chosen strategies may change once the expectation is taken over the stochastic state variables in the course of calculating the certainty equivalent. In a sense, then, the certainty equivalent is the 'optimal' certainty equivalent achievable by the investor. For the present purposes, however, this consideration is not relevant. In order to determine parameterizations of the utility- and value functions, which would lead to comparable levels of risk-aversion among the three investors, it is sufficient to consider exogenous risks.

[23]The neglect for wealth $X_t - X_{t^R}$ in this formula is an inevitable consequence of the reference wealth level chosen and a noted shortcoming of the resulting parameterization. Both will further be discussed below.

$$r_t(X_{t^R}) = -\frac{u''(X_{t^R})}{u'(X_{t^R})}.^{24} \tag{12}$$

Pratt's (1964) derivations, adapted to the present context, are given in appendix A11.

As shown by Pratt (1964), knowledge of the local risk aversion measure for all wealth levels is equivalent to knowledge of the certainty equivalents of any risk at any wealth level, as well as knowledge of the individual's utility function. (Pratt, 1964) This is also evident from another result due to Pratt (1964), which implies that knowledge of local risk aversion at all levels encodes knowledge of the utility function, that is [25]

$$u = \int e^{-\int r}. \tag{13}$$

Thus, if two investors share the same degree of local risk aversion at all wealth levels, then they apply the same utility function to the evaluation of future wealth. It is, thus, also logical that their certainty equivalents from any given risky prospect will be equal in this case.

Ideally, one would, thus, equate local - or global risk-aversion measures for the three investors across all wealth levels to obtain the desired parameterization. However, this is not possible within our setting. Neither is it possible to find closed form solutions for the risk-aversion parameters of all three objective functions to equate global measures for a given initial wealth level. We, thus equate local measures for CARA - and CRRA investors and global measures for CARA - and Mean-Variance investors.
The thought experiment of adding a random prospect to a given reference wealth level at decision time t (Bernoulli, 1954) is based on preference relations reflected in the following expected utility - and value functions, indexed to decision time t,

$$U_t^{CA}(X_T) = \mathbb{E}_t\left[-e^{-\alpha(X_{t^R}+\tilde{Z}_{t,T})}\right], \quad U_t^{CR}(X_T) = \mathbb{E}_t\left[\frac{(X_{t^R}+\tilde{Z}_{t,T})^{1-\gamma}}{1-\gamma}\right]$$

and

$$MV_t(X_T) = \mathbb{E}_t[(X_{t^R}+\tilde{Z}_{t,T})] - \frac{\delta}{2}\mathrm{Var_t}[(X_{t^R}+\tilde{Z}_{t,T})]. \tag{14}$$

Equating the CARA - and CRRA investors' Arrow-Pratt measures of (absolute) risk aversion (Pratt, 1964), we obtain

$$-\frac{u''^{CA}(X_{t^R})}{u'^{CA}(X_{t^R})} = -\frac{-\alpha^2 e^{-\alpha X_{t^R}}}{\alpha e^{-\alpha X_{t^R}}} = -\frac{u''^{CR}(X_{t^R})}{u'^{CR}(X_{t^R})} = -\frac{-\gamma(1-\gamma)X_{t^R}^{-\gamma-1}}{(1-\gamma)X_{t^R}^{-\gamma}}$$
$$\iff \alpha = \frac{\gamma}{X_{t^R}}. \tag{15}$$

[24] Risk aversion is measured at decision time t and for the given reference wealth level. As $t^R \leq t$, Bernoulli utility functions are used for the local approximation around known wealth levels. The sub-script in r_t could, thus, equally be omitted, but is kept here to underline that the evaluation takes place at time t.

[25] This notation, used by Pratt (1964), may not satisfy the mathematical purist, but it is mainly intended as a rough sketch of the underlying idea.

Thus, for the CARA - and CRRA investors to share the same aversion with respect to invested amounts in a risky prospect across reference wealth levels, the CARA investor's risk-aversion parameter α would need to vary inversely with reference wealth level X_{t^R}. However, this would contravene the character of her preferences. Indeed, as can be seen from (13), any agent with absolute risk aversion equal to $\frac{\gamma}{x}$ is automatically a CRRA investor, as

$$u(x) = \int e^{-F(x)} dx \quad \text{, with} \quad F(x) = \int -\frac{\gamma}{x} dx,$$

so that (16)

$$u(x) = \int e^{-\gamma \log(x)} dx = \int x^{-\gamma} dx = \frac{x^{1-\gamma}}{1-\gamma}.$$

In this case, however, the investor would have constant aversion to proportional risk rather than to fixed absolute exposure to the risky asset (Eeckhoudt et al., 2005, Pratt, 1964). Hence, if the CRRA investor's preferences are taken as such, implying that γ is constant, then the CARA investor's absolute risk aversion cannot be constant, if we impose at the same time that the two investors' degrees of risk aversion are to be equal across (reference) wealth levels. Therefore, we parameterize the utility - and value functions so as to obtain comparable levels of risk aversion for the particular wealth level assumed at the beginning of the investment horizon, that is X_0, and assess the effect of choosing a different reference wealth level for robustness below.

As for the Mean-Variance - and CARA investors, we first determine the certainty equivalents of a given single-period, exogenous, normally distributed risk, which we shall also call X_T[26]. We impose again that the certainty equivalents must be equal for both investors and determine the risk-aversion coefficients α and δ accordingly. Thus, following the definition given in (11), we obtain for the Mean-Variance investor, that

$$\begin{aligned} \mathbb{E}_t[X_{t^R} + CE^{MV}] - \frac{\delta}{2}\mathrm{Var_t}[X_{t^R} + CE^{MV}] &= \mathbb{E}_t\left[\mathbb{E}_t[X_{t^R} + \tilde{Z}_{t,T}] - \frac{\delta}{2}\mathrm{Var_t}[X_{t^R} + \tilde{Z}_{t,T}]\right] \\ \iff X_{t^R} + CE^{MV} &= X_{t^R} + \mathbb{E}_t\left[\tilde{Z}_{t,T}\right] - \frac{\delta}{2}\mathrm{Var}_t\left[\tilde{Z}_{t,T}\right] \\ \iff CE^{MV} &= \mathbb{E}_t[\tilde{Z}_{t,T}] - \frac{\delta}{2}\mathrm{Var_t}[\tilde{Z}_{t,T}]. \end{aligned} \tag{17}$$

The certainty equivalent from risk $\tilde{Z}_{t,T}$ - and, consequently from X_T - is, thus, independent of reference wealth level, X_{t^R}, which underlines the strong connection between Mean-Variance preferences and CARA-utility in the single-period context (see e.g. Munk, 2017). For the CARA investor, in turn, the certainty equivalent may be determined via

[26] The choice is motivated by the fact that for additive risk within our market setting, terminal wealth under the resulting wealth process before any particular portfolio strategy is applied is normally distributed, as seen below in appendix A7.

$$
\begin{aligned}
-\exp\left(-\alpha\left(X_{t^R}+CE^{CA}\right)\right) &= \mathbb{E}_t\left[-\exp\left(-\alpha\left(X_{t^R}+\tilde{Z}_{t,T}\right)\right)\right] \\
&= -\exp\left(-\alpha\mathbb{E}_t\left[X_{t^R}+\tilde{Z}_{t,T}\right]+\frac{\alpha^2}{2}\mathrm{Var_t}\left[X_{t^R}+\tilde{Z}_{t,T}\right]\right) \\
&= -\exp\left(-\alpha X_{t^R}-\alpha\mathbb{E}_t\left[\tilde{Z}_{t,T}\right]+\frac{\alpha^2}{2}\mathrm{Var_t}\left[\tilde{Z}_{t,T}\right]\right) \\
\iff CE^{CA} &= \mathbb{E}_t[\tilde{Z}_{t,T}]-\frac{\alpha}{2}\mathrm{Var_t}[\tilde{Z}_{t,T}],
\end{aligned} \tag{18}
$$

whereas the last equality follows from merely taking logarithms on both sides and simplifying. Again, as expected from investors with constant absolute risk aversion, the certainty equivalent of the risky prospect does not depend on reference wealth level X_{t^R}. It, thus, follows immediately from the results in (17) and (18) that

$$
CE^{MV}=CE^{CA} \iff \alpha=\delta. \tag{19}
$$

Moreover, from (15) and (19), $\delta=\frac{\gamma}{X_{t^R}}$, in order to obtain comparable levels of risk-aversion between CRRA - and Mean-Variance investors. Similar to the comparison between CRRA - and CARA investors, we will set $\delta=\frac{\gamma}{X_0}$. While allowing for time-varying risk-aversion coefficient δ would lead to interesting additional source of time-inconsistency (Björk et al., 2014), we abstract from this issue in favor of maintaining constant absolute risk aversion as a characteristic feature of Mean-Variance preferences, as shown above.

Thus, we hold the reference wealth level fixed at the given initial wealth level $X_0=\bar{X}_0$ and equate local measures of risk aversion between CARA-and CRRA investors, and global measures between CARA - and Mean-Variance investors, respectively. There are a number of shortcomings with respect to the proposed procedure. First, given the dynamic and stochastic context assumed, wealth levels are bound to vary over time, so that the risk-aversion measures of CRRA - and CARA-investors (and, hence, also the risk-aversion measures between CRRA and Mean-Variance investors) will diverge over time. In particular, if $X_t>X_0$ ($X_t<X_0$), then given the parameterization $\alpha=\delta=\frac{\gamma}{X_0}$, CARA-and Mean-Variance investors will be more (less) risk-averse than CRRA investors at the current actual wealth level. Second, given that Brownian shocks have unbounded support, the risks under consideration are large risks. The Arrow-Pratt measure at a single given reference wealth level is only a local approximation to risk aversion measures for large risks, however. (Eeckhoudt et al., 2005, Pratt, 1964) Thus, within the current context, the investors' actual risk aversion may differ already at the beginning of the investment horizon. This 'size-effect' may be remedied by numerically determining parameters of the utility functions such that certainty equivalents of single-period exogenous risks are the same across all three investors. However, given the remaining challenges concerning the intended parameterization, it is unclear whether this would lead to a significant improvement with respect to the intended parameterization. Finally, as discussed in footnote (6), the definition of certainty equivalents should be adapted to reflect the endogeneity of risks, in particular given the dynamic context of the present analysis. However, given the intended goal of evaluating one investor's strategy according to another investor's preferences, it is unclear how certainty equivalents should be adapted in this particular case. Moreover, as investors treat each others' proposed strategies as exogenous risks, it is also reasonable to use single-period exogenous risks for the paramterization as discussed.

Besides the three shortcomings discussed in the previous paragraph, there is another challenge to the suggested parameterization, which, however, is inherently linked to the underlying optimization problem. As Bajeux-Besnainou & Portait (1998) point out, the 'Mean-Variance Dynamic Efficient Frontier' (ibid) includes portfolios that are beyond the single-period efficient frontier, essentially because the use of the Martingale Method (Cox & Huang, 1989) together with continuous rebalancing opportunities opens up trades of risk that would not be accessible to standard single-period optimization problems. This is intuitive, because the Martingale Method first determines an optimal terminal wealth profile across states of the world, for which it then provides a perfect hedge within a given complete market setting to determine the optimal investment strategy [27] (ibid), thus dispensing with the restriction to linear portfolio weights, which is regularly imposed in single-period Mean-Variance optimization problems (Johnstone & Lindley, 2013; Baron, 1977). As a result, however, terminal wealth under the Martingale Method may not be normally distributed[28]. However, normality of all possible portfolios - and, hence, wealth processes - is one of the few special cases, where Mean-Variance optimization does generally not violate basic axioms of rational decision making under uncertainty, as formalized by Von-Neumann & Morgenstern within the expected utility framework, or connected criteria of Stochastic Dominance. (Johnstone & Lindley, 2013; Hannoch & Levy, 1969). As is known from different contexts[29], caution should then be used in measuring risk-aversion using certainty equivalents, which relies on Von-Neumann & Morgenstern's axioms being satisfied (Mas-Colell et al., 1995). Furthermore, equations (17), (18) and (19) show that the condition $\alpha = \delta$ only guarantees equality of certainty equivalents between CARA and Mean-Variance investors under the assumption of normally distributed risky prospects. For the wealth process resulting from the Mean-Variance pre-commitment solution[30], this assumption does not hold, even in the given Black-Scholes market setting. While it is beyond the scope of this thesis to fully explore these issues, there may, thus, be certain limitations with respect to the intended parameterization of utility functions, especially when the latter are used as optimization criteria within the context of the Martingale Method.

Finally, the above limitations also imply that there is no circularity in the use of certainty equivalents to obtain a comparable measure of 'value' across investors (see e.g. Balter et al., 2021) after risk-aversion parameters have been determined to obtain comparable levels of risk-aversion based on certainty equivalents in the first place. Parameterization based on given single-period exogenous risks allows for considerable space for the optimal choices across investors and investment strategies to differ, as also evidenced in sections 3 and 4.

[27] To be discussed in greater detail in section 3.

[28] See the results for the Mean-Variance pre-commitment strategy discussed below and derived in appendices A5 and A6

[29] See for example the case of state-dependent preferences (Yaari, 1969; Karni, 1990; Polemarchakis, 1990). Here, marginal rates of substitution between states are state-dependent and, thus, do not rely solely on ratios of probabilities, as would be assumed within the Von-Neumann-Mrogenstern expected utility framework (Mas-Colell et al., 1995; Arrow, 1996), hence, leading to a requirement for an extended expected utility framework and the respective adjustments of certainty equivalents (Mas-Colell et al., 1995)

[30] See below in section 3

3 The Optimization Problem

3.1 Four Optimization Strategies

This section discusses four optimization strategies available to the investors, which differ as to how information in $(\mathcal{F}_t)$ is incorporated into rebalancing decisions, given the continuous rebalancing opportunities in the market. Differences in this respect lead to static, myopic, consistent planning - and precommitment strategies (Strotz, 1956). First, the optimization problems behind each of these strategies will be discussed. Then, the respective strategies will be derived under the assumption of CRRA - and CARA utilities, as well as Mean-Variance preferences, respectively.

3.1.1 Consistent-planning strategies

The term 'consistent planning' strategies was first coined by Strotz and refers to strategies in which the investor takes into account her future self's optimal decisions in determining an optimal decision today (Strotz, 1956). Thus, the term refers to strategies as they would emerge from dynamic programming approaches, in which today's strategy's effect on the optimum achievable at future instances is consciously modelled at each rebalancing time (see e.g. Björk, 2020). Given the Black-Scholes market setting, wealth serves as a Markovian state variable on the basis of which the agent applies backward induction to solve for the optimal dynamic strategy (see e.g. Munk, 2017). With a feedback control law $u_t = g(X_t, t)$ (Björk, 2020), whereas u_t is either given by $\theta_t = \theta(t, X_t)$ or $\pi_t = \pi(t, X_t)$, respectively, the investor solves the following program:

$$\max_{\{u_s\}_{s=0}^{T}} \mathbb{E}\left[V(X_T)\right] \text{[31]} \tag{20}$$

subject to wealth dynamics

$$dX_t^u = \mu^u(t, X_t^u)dt + \sigma^u(t, X_t^u)dW_t \tag{21}$$

and initial condition

$$X_0 = \bar{X}_0. \tag{22}$$

Wealth dynamics in (21) are deliberately chosen to be generic to accommodate wealth dynamics as specified by (5) and (6), respectively, as well as different strategies chosen by the investors.

[31] whereas V stand for 'value' and covers both utiliy, as for CRRA - and CARA investors, and the value derived by the mean-variance investor from terminal wealth. The same notation is used for the consistent planning problem below.

3.1.2 Optimal pre-commitment strategies

Pre-commitment strategies, in turn, are based on the assumption that the investor can be forced to adhere to the optimal strategy, as determined at the beginning of the investment horizon, at any rebalancing opportunity thereafter. (Strotz, 1956) This strategy is routinely[32] computed by use of the Martingale Method (Cox & Huang, 1989), which, in a nutshell, is based on a two-step procedure. First, the investor chooses the terminal terminal wealth level to maximize her objective in each state of the world at the end of the investment horizon subject to her budget constraint, creating a contingent claim to be replicated by financial assets in the market. By maximizing value in each future state of the world, the chosen terminal wealth profile also maximizes the expected future value ex ante. In the second step, a dynamic hedging strategy is determined, which replicates the market consistent value of this optimal terminal wealth profile at each point in time t such that $0 \leq t \leq T$ on the basis of the financial assets traded in the market. The dynamic hedge itself then provides the optimal portfolio strategy for the investor to maximize expected future utilty. (see e.g. Björk, 2020; Cox & Huang, 1989)

The complete financial market assumptions underlying this dynamic hedging approach (see e.g. Munk, 2017) are satisfied within Black-Scholes financial markets, as assumed. As a result, we may define the unique stochastic discount factor M_t, as well as the unique Radon-Nikodym process ξ_t, respectively, underlying the required valuation and replication steps as follows (see e.g. Shreve, 2010; Munk, 2017):

$$\begin{aligned} dM_t &= -rM_t dt - \lambda M_t dW_t, \quad M_0 = 1 \\ \iff M_t &= \exp\left(-\left(r + \frac{1}{2}\lambda^2\right)t - \lambda W_t\right) \end{aligned} \tag{23}$$

and

$$\begin{aligned} d\xi_t &= -\lambda dW_t, \quad \xi_0 = 1 \\ \iff \xi_t &= \exp\left(-\frac{1}{2}\lambda^2 t - \lambda W_t\right), \end{aligned} \tag{24}$$

whereas λ is the uniquely determined market price of risk, as given by $\lambda = \frac{(\mu - r)}{\sigma}$. The agent's optimization problem in the first step described above is then given by the static constrained optimization problem

$$\max_{X_T} \mathbb{E}\left[V(X_T)\right] \tag{25}$$

subject to the budget constraint

$$\bar{X}_0 = \mathbb{E}\left[M_T X_T\right]. \tag{26}$$

The dynamic hedging strategy in the second step, in turn, is dependent on the respective market consistent value process of optimal terminal wealth and is, thus, specific to the particular objective function assumed.

[32] That is, where applicable.

By construction, where applicable, the martingale method delivers the optimal policy to maximize expected future utility - or value, within a given market, as it directly maximizes the objective function and determines the dynamic hedge to deliver the optimal policy. In contrast, consistent-planning strategies are constrained by the requirement that Bellman's principle of optimality needs to be satisfied (Basak & Chabakauri, 2010; Björk & Murgoci, 2010; Björk, 2020). Once the optimal terminal wealth and the dynamic hedging strategy have been determined at the initial time point, the investor may throw away the proverbial key and rest assured that within the given complete market setting, she will arrive at her optimally chosen wealth level whichever state of the world materializes. (Cox & Huang, 1989; Björk 2020)

The strategy's operation on autopilot turns into a behavioral constraint, however, where the underlying optimization problem is time-inconsistent so that the investor may wish to deviate at future time-points. (Strotz, 1956; Björk & Murgoci, 2010) For various reasons, ranging from legal to financial and psychological, it appears unrealistic that investors will always be able to exercise the restraint required to follow pre-commitment strategies. In this case, consistent-planning strategies need to be revisited, with the objective function being adapted to take into account the investor's future self's incentive to deviate from previous plans (Strotz, 1956; Basak & Chabakauri, 2010) The resulting solution may be characterized as a constrained optimum, akin to second-best optima, as are commonly known from incentive theory (see e.g. Laffont & Martimort, 2002). As Chen & Zhou (2022) point out, ascription of 'dynamic optimality' to dynamic programming solutions in face of time-inconsistency is a misnomer to begin with.

3.1.3 Myopic Strategies

The third set of strategies available to the investor is given by myopic strategies. Various interpretations could be given to the concept of myopia in intertemporal (continuous time) optimal investment problems. Their common core is, however, that the effect of today's decision on tomorrow's preferences, and hence, the potential wish to deviate tomorrow from a plan chosen today, is not taken into acccount. Rather, a decision based on today's preferences is taken, irrespective of tomorrow's preferences. (Strotz, 1956)

The precise way, in which this broad definition is operationalized depends on the context in which myopic strategies are applied. Chen & Zhou (2022) suggest that myopic strategies consist of sequences of optimal pre-commitment strategies, each pursued only over infinitesimally short periods of time, however, until the next rebalancing opportunity. The strategies are then derived using a limiting argument. First, a given finite set of rebalancing opportunities are assumed, at each of which the optimal pre-commitment strategy for the entire remaining time-horizon is determined, whereas this strategy is ultimately only pursued until the next rebalancing opportunity. Hence, the myopic strategy takes the form of a simple (in the sense of step[33] -) function applied to continuous-time asset price processes to form a wealth process. Subsequently, the authors let the time between rebalancing opportunities go to zero and derive the limiting wealth process. The myopic strategy, in turn, is the strategy that gives rise to the limiting wealth process. This definition is rather general and applicable within settings, in which the optimal pre-commitment solution is stochastic. (Chen & Zhou, 2022)

This underlying understanding of myopic strategies as pre-and re-commitment strategies agrees

[33]See Shreve, 2010.

with the understanding of the concept in Balter et al. (2021). The latter contribution derives myopic strategies for the case of time-consistent pre-commitment solutions so that future selves' optimal (p)re-commitment strategies agree with former selves' strategies provided for those later time-points. Thus, the resulting myopic strategies are deterministic and at most time-dependent, so that no limiting argument as in Chen & Zhou (2022) is required. Moreover, the authors prove that myopic strategies may be derived via auxiliary constant investment strategies, which facilitates the solution of the problem.

We follow Chen & Zhou (2022) to derive optimal myopic strategies for Mean-Variance investors via the proposed limiting argument. However, we re-derive the result, given that the specification of mean-variance pre-commitment strategies in this thesis differs from that in Chen & Zhou. In turn, we follow Balter et al.'s (2021) results for myopic strategies of CRRA investors within the Black-Scholes setting, while we follow the authors' solution method but provide full derivations for the case of CARA-utility, given the difference in character between the pre-commitment solutions of these two investors.

3.1.4 Static Strategies

Finally, the simplest among the four strategies discussed in this thesis is a static, buy-and-hold strategy. In this case, the investor determines the optimal investment strategy based on information at time zero and refrains from rebalancing along the investment - horizon. The value derived from static investment strategies provides a benchmark, against which the other three strategies may be evaluated. On the one hand, the strategy appears very simple as rebalancing opportunities are being ignored, even though behavioral constraints may seem no less pressing than in the context of optimal pre-commitment solutions. In terms of performance, in turn, the static strategy would be generally be expected to be - potentially highly - sub-optimal, as the investor refrains from reacting to new information, not even in pre-determined ways. As compared to the optimal pre-commitment strategy, one would expect static strategies to deliver significantly lower 'value' to the investor, as buy-and-hold strategies are in a sense akin to static hedges of the present value of optimal future wealth, whereas optimal pre-commitment hedges are based on dynamic hedges of the present value of optimal future wealth. (see the discussion of pre-commitment strategies and the Martingale Method in section 3.1.2 and appendices A1-A6.) Hence, if any given strategy performs worse than a static investment strategy, then this is indicative of a (very) high value to commitment in the particular case.

3.2 Optimal Strategies

Given the four optimization strategies discussed in the previous section, we now introduce the resultant optimal investment policies for each of the three investors.

3.2.1 CRRA investors

The distinguishing characteristic of the CRRA utility is that the investor generally prefers exposure to risk in constant proportion to overall financial wealth (Pratt, 1964). Her consistent-planning strategy is given by the well-known Merton-fraction (Merton, 1969), that is

$$\pi_t^{CR,DC} = \frac{(\mu - r)}{\gamma\sigma^2}. \tag{27}$$

The investor's preference for a constant fraction of total wealth to be invested in the risky asset translates into an optimal value function[34], which is itself CRRA (see Merton, 1969), so that relative risk-aversion at any point in time with respect to the risky prospect resulting from the dynamic trading strategy is constant. The investor solves for the optimal strategy by backward induction, knowing that future risk-preferences will be independent from wealth at that point in time and that investment opportunities within Black-Scholes markets are constant. Hence, there is no need to hedge for the latter, or to take into account the effect of today's decisions on tomorrow's preferences. Thus, given the desired constant fraction invested in the risky asset at each future time point, the investor's optimal choice is to choose the desired constant fraction today as well. Moreover, the lack of inter-temporal linkages between investment choices at intermediate time-points also suggests that the relative underperformance of static strategies may not be very large.

It, thus, also follows immediately that the CRRA investor's optimization problem in the given setting is time-consistent. As a result, the dynamic programming solution attains the overall optimal pre-commitment solution (Merton 1969 & 1971; Cox & Huang, 1989), which is derived in appendix A1 and given by

$$\pi_t^{CR,PC} = \frac{(\mu - r)}{\gamma\sigma^2}. \tag{28}$$

The hedging strategy delivering a constant relative exposure at end of the investment horizon is to hold the given fraction at each intermediate time point as well.

Given that the optimal pre-commitment solution is to hold a constant fraction, irrespective of the realization of wealth at intermediate time points, the investor chooses the same pre-commitment policy when she re-commits at any time $t \in [0, T]$. Hence, as also shown by Balter et al. (2021), her optimal myopic strategy is, likewise, given by the Merton-fraction, so that

$$\pi_t^{CR,MY} = \frac{(\mu - r)}{\gamma\sigma^2}. \tag{29}$$

Finally, if the investor chooses a static strategy, her best choice today is to choose the fraction, which she hopes will prevail at the end of the investment horizon, even though due to intervening stochasticity of returns, the realization will ultimately differ from the intended fraction with probability one. With the time subscript reflecting the fact that the investment fraction will,

[34] By which we mean the value function of the dynamic optimization problem at the optimal strategy (see Bjrk, 2020)

thus, ex-ante only be known for time $t = 0$, the static buy-and-hold strategy is derived in appendix A9 and given by

$$\pi_0^{CR,ST} = \frac{(\mu - r)}{\gamma\sigma^2}. \tag{30}$$

Finally, constant investment fractions, as provided for the pre-commitment, dynamically consistent and myopic strategies do not imply passivity, but require active rebalancing at each intervening opportunity. Moreover, given CRRA preferences and continuous trading, the implied wealth process will not turn negative even in absence of borrowing - or short selling constraints. As wealth approaches zero from above, the desired constant fraction implies that the investment in the stock in terms of amounts invested will automatically approach zero as well, thus reducing the impact of future stochastic shocks on wealth and mitigating the danger of wealth turning negative after further negative shocks to the stock price.

3.2.2 CARA investors

The CARA investor generally desires constant amounts invested in the risky asset, irrespective of overall financial wealth (Merton, 1969; Pratt, 1964). Her consistent-planning strategy is generally known (Merton, 1971; Basak & Chabakauri, 2010) to be given by

$$\theta_t^{CA,DC} = \frac{(\mu - r)}{\alpha\sigma^2} e^{-r(T-t)}. \tag{31}$$

The CARA investor earns a continuously compounded fixed return r on her exposure to the risk-free asset and additionally a stochastic excess return on her exposure to the risky asset. As a result, the risk-free rate may be seen as a baseline- expected growth rate of the exposure to any asset, to which a stochastic excess return is added for the risky asset. Given that all assets, thus, share this given baseline growth rate, it is reasonable to invest less in a risky asset today if a given exposure to the risky asset is desired at the future point in time. This is reflected in the discount term $e^{-r(T-t)}$. Similarly, we could imagine future wealth being discounted to the point in time when a rebalancing decision is taken so that the optimal exposure is decided concerning present-discounted future wealth in order to account for the time value of money (see also footnote 9 in Basak & Chabakauri (2010) and Merton (1971)). The optimal value function of the dynamic optimization problem is of CARA-type as well (see Merton, 1969), so that at each time the investor prefers constant absolute exposure to the risk resulting from the dynamic investment prospect. The investor solves the optimization problem by backward induction, knowing that future optimal decisions will be independent of intermediate wealth levels and that investment opportunities are constant. Thus, there is again little intertemporal linkage between optimal choices at various time-points to exploit to achieve an optimal dynamic strategy, which, thus, takes the given deterministic, time-varying form.

CARA utility maximization is known to be time-consistent (Merton, 1971, Cox & Huang, 1989). As is also derived in appendix A3, the optimal pre-commitment strategy is, hence, given by

$$\theta_t^{CA,PC} = \frac{(\mu - r)}{\alpha\sigma^2} e^{-r(T-t)}. \tag{32}$$

It seems intuitive for the CARA investor, who desires a constant absolute exposure to the risky asset within her terminal portfolio, to invest the given constant amount discounted by the risk-free rate at each time-point, merely rebalancing where and in so far as stochastic returns move the present exposure away from the desired discounted exposure. However, for both the CRRA - and CARA investor, the given strategy may also directly be derived from the hedging strategy in the second step of the Martingale Method. This will be discussed below after outlining the pre-commitment strategy for the Mean-Variance investor, as this also underlines why the Mean-Variance investor's hedging efforts lead to a stochatic pre-commitment strategy, whereas these efforts yield deterministic strategies for the CARA - and CRRA investors.

The investor's optimal myopic strategy, in turn, is derived according to the method proposed by Balter et al. (2021). As shown in appendix A11, the policy is given by

$$\theta_t^{CA,MY} = \frac{(\mu - r)}{\alpha\sigma^2} e^{-r(T-t)}. \tag{33}$$

Given the time-consistency of the problem, it is again intuitive that the myopic-, pre-commitment - and consistent planning solutions coincide.

Finally, the investor's optimal static strategy, as derived in appendix A8 is given by

$$\theta_0^{CA,ST} = \frac{2(\mu - r)}{\alpha\sigma^2(e^{rT} + 1)}, \tag{34}$$

which can be seen as a weighted average over optimal pre-commitment policies. This is intuitive, as the investor now aims at achieving the desired terminal absolute exposure while refraining from rebalancing at intermediate time points. Hence, her best strategy in this case is to choose a weighted average of the pre-commitment strategies, which would otherwise have been choses if rebalancing opportunities had been used. As above, the strategy carries time subscript 0, as the given absolute exposure will only hold at the initial time point, due to the stochasticity of returns.

Unless short-selling and borrowing constraints are put in place, there is a non-zero probability that wealth will turn negative. As wealth approaches zero from above, the CARA - investor still desires the given deterministic, time-varying exposure to the risky asset, implying that she would eventually borrow at the risk-free rate in order to keep the exposure to the risky asset at the desired level. Hence, the likelihood of wealth turning negative increase when current wealth is lower. This is also related to the main criticism against CARA-utility, which is that absolute risk aversion and, hence, the preferences with respect to gambles in absolute terms do not change as the current wealth level varies (Eeckhoudt et al., 2005).

3.2.3 Mean-Variance investors

It follows from equation (17) that with constant δ and the additive risk as represented by the wealth process in (6), the certainty equivalent and, thus, the degree of absolute risk aversion of the Mean-Variance investor, are constant, just as in the case of single-period exogenous risks discussed in section 2. To illustrate, plugging the solution for X_T from (6), as derived in (92) in appendix A7, into the Mean-Variance objective function, as given in (10), yields

$$
\begin{aligned}
MV_t(X_T) &= \mathbb{E}_t\left[e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right] \\
&\quad - \frac{\delta}{2}\mathrm{Var}_t\left[e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right] \\
&= e^{r(T-t)}X_t + \mathbb{E}_t\left[\int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right] \\
&\quad - \frac{\delta}{2}\mathrm{Var}_t\left[\int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right],
\end{aligned} \tag{35}
$$

so that by (17), the certainty equivalent is given by

$$
\begin{aligned}
CE^{MV} &= \mathbb{E}_t\left[\int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right] \\
&\quad - \frac{\delta}{2}\mathrm{Var}_t\left[\int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right],
\end{aligned} \tag{36}
$$

as X_t is a known constant at time t. Thus, the the Mean-Variance investor shares the CARA investor's constant absolute risk aversion. Moreover, they would be expected to exhibit the same preferences with respect to normally distributed, single period, exogenous risks, as their certainty equivalents are the same in this case, subject to $\alpha = \delta$. However, this does not necessarily carry over to the dynamic context, where differences in choice behavior re-emerge due to differences in the investors' utility - and objective functions.

The consistent-planning strategy for the Mean-variance investor is derived in appendix A7 and is given by

$$
\theta_t^{MV,DC} = \frac{(\mu-r)}{\delta\sigma^2}e^{-r(T-t)}. \tag{37}
$$

Thus, the time-consistent Mean-Variance investor, who takes into account her incentive to deviate over time, chooses the same policy as the optimal dynamically consistent or pre-committed CARA-investor. Whereas Basak & Chabakauri (2010) explain the presence of the time-discount factor in this case by the increase in the variance of wealth with the investment horizon, the discount factor may also be explained by reference to CARA-utility. As the Mean-Variance investor solves her dynamic optimization problem by backward induction and knows that her future selves have constant absolute risk aversion, irrespective of future wealth levels, she essentially chooses the dynamic CARA-investor's strategy, as discussed in the previous section, and with the same interpretation given to the discount factor as above.[35]

[35] Another important connection between the dynamically consistent CARA - and Mean-Variance investor is that under normally distributed terminal wealth and non-stochastic (linear) portfolio weights, the optimal mean-variance portfolios satisfy Second-Order stochastic dominance and can, in principle, be computed by use of negative exponential utility (Johnstone & Lindley, 2013). Given the strategy and the resulting terminal

The Mean-Variance investor's pre-commitment strategy differs from her consistent-planning strategy, due to the time-inconsistency of the underlying optimization problem and the resulting need to adapt the dynamic programming problem, which causes divergence between the obtained constrained optimum and the overall optimum as given by the pre-commitment solution (Basak & Chabakauri, 2010, Björk & Murgoci, 2010, Björk et al, 2014). Moreover, there is reason to believe that in contrast to her consistent planning solution, the Mean-Variance investor's pre-commitment solution does not satisfy the requirements of Second-Order stochastic dominance, which causes further divergence between the two solutions, as will be discussed in section 4. The investor's optimal pre-commitment policy, as derived in appendix A5 is given by

$$\theta_t^{MV,PC} = \frac{1}{\delta}\frac{(\mu - r)}{\sigma^2}e^{-r(T-t)+\lambda^2(T-t)+rt}M_t. \tag{38}$$

In contrast to the CARA - and CRRA investors' pre-commitment strategies, the Mean-Variance investor's pre-commitment strategy is stochastic. While the literature appears to draw on the underlying objective function over terminal wealth for intuition, further insights may also be gained from looking at the underlying hedging argument.

The literature includes various representations of the pre-commitment solution for the Mean-Variance investment problem, as well as various choices as to the formulation of the underlying optimization problem (Richardson, 1989; Bajeux-Besnainou & Portait, 1998; Basak & Chabakauri, 2010; Chen & Zhou, 2022). The difference in representation of the result, however, is largely due to the choice as to whether the investor's problem is formulated in terms of discounted wealth or not (Richardson, 1989). The result as presented here in (38) agrees with the pre-commitment solution stated (without derivation) in Basak & Chabakauri (2010). In spite of the diversity in representations, the emerging underlying mechanisms governing the Mean-Variance investor's optimal pre-commitment solutions, as proposed in the literature, are the same: The investor chooses a portfolio on the dynamic Mean-Variance frontier (Bajeux-Besnainou, 1998), balancing the mean and variance of terminal wealth[36], respectively, on the basis of her risk-aversion, as represented by δ. A negative shock to the value of the stock at time $t \in [0, T]$ reduces current wealth and, as seen in (93) in appendix A7, furthermore reduces expected future wealth conditional on current wealth. To offset this drop in expected future wealth, the investor increases exposure to the stock in order to increase the share of current wealth that would benefit from the risk-premium and, thus, aims at raising expected future wealth. (Basak & Chabakauri, 2010) It is, however, not fully clear how the investor balances target means and variances at such intermediate time points according to the interpretations

wealth, as derived in appendix A7, there is no reason to assume that this result, which is generally stated for the single-period case would not carry over the the dynamic programming case. This will further be discussed in section 4.

[36]Possibly, this framing of the problem is the reason for a lack of reflection on issues of Stochastic Dominance in connection with the pre-commitment solution in the literature. Johnstone & Lindley (2013) discuss the representability of the optimal Mean-Variance choice in the (μ, σ) plane, which holds under certain conditions. As will be discussed in section 4, however, these conditions do not appear to be satisfied when using the Martingale Method, as evidenced by the resulting stochastic strategy and non-normally distributed wealth process (see section 4). This does not imply, of course, that the optimal strategies, as derived in the literature, would be incorrect. However, it implies that the juxtaposition of pre-commitment solutions and dynamically consistent solutions is of particular interest not only because of issues of time-inconsistency, but also because of the former's problematic relationship to Stochastic Dominance.

in the liteature. In any case, according to the mechanics of the Martingale Method, concern for target expected return and variance would have to be intermediated by the current value process of the optimal terminal wealth profile, as derived in (82) in appendix A5. The depri-oritization of variance reduction in favor of raising expected returns would be consistent with failure to meet criteria of Second-Order Stochastic Dominance (Hadar & Russel, 1969; Levy, 1990; Machina & Rothschild, 1990; Johnstone & Lindley, 2013, Hanoch & Levy, 1969), as will be discussed further in section 4.

Another way to interpret Mean-Variance pre-commitment solutions is via the delta-hedge, which is part of the Martingale Method. This may also provide a more unified framework of interpretation for optimal CARA-, CRRA - and Mean-Variance pre-commitment strategies.
The optimal terminal wealth levels chosen by the CRRA (47) -, CARA (65) and Mean-Variance (80) investors are all inversely related to movements in the stochastic discount factor, which is reasonable, as a low realization of the stochastic discount factor implies a good state of the world, whereas a high stochastic discount factor reflects a bad state of the world (see e.g. Cochrane, 2005). However, while for the CRRA - and CARA investors optimal terminal wealth levels are strictly convex functions (for $\gamma > 1$) of the stochastic discount factor, the Mean-Variance investor's optimal wealth levels react linearly to changes in the state price deflator.
Based on solutions in (47) and (65) together with the stochastic discount factor in (23), the CRRA investor's optimal terminal wealth, thus, reacts exponentially to the underlying accumulated Brownian shock, whereas for the CARA investor, this relationship is linear. This is intuitive in two ways. First, the deterministic fractions for the CRRA investor and amounts invested for the CARA investor, respectively, imply the given patterns in terminal wealth. If investment fractions are constant and asset returns accrue exponentially, then terminal wealth must vary exponentially with the accumulated Brownian shocks. In turn, if present discounted amounts invested in the stock are constant [37], then at intermediate time points, excess returns in the stock are deposited in the risk-free asset, whereas after negative shocks to the asset price, the individual sells the bond and buys the stock to compensate for the fall in risky investment. As a result, optimal terminal wealth varies linearly with the accumulated Brownian shocks, as returns are not reinvested in the stock, but deposited in the risk-free asset, which grows at the same rate as present discounted future wealth. Vice versa, and more closely related to the mechanics of the delta-hedge underlying the Martingale Method, the given optimal terminal wealth levels imply the given hedging strategies. The CRRA investor's optimal terminal wealth in (47) is of the form ce^{bW_T}, with constants $c, b > 0$. With the investor choosing investment fractions and returns accruing exponentially, it is intuitive that the optimal hedge of this wealth profile is deterministic, and it follows from the given parameter setting that the fractions of the hedging portfolio are moreover constant. In turn, the CARA investor's optimal wealth profile, which varies linearly with accumulated Brownian shocks, is hedged by a strategy that posts excess returns on the risky asset in the risk-free, which accrues the same return as present discounted future wealth, so that the amount invested in the stock grows deterministically at rate r.
The Mean-Variance investor's optimal terminal wealth in (80) is given by a function of accumulated Brownian shocks of the form $a - ce^{-bW_T}$, with constants $a, b, c > 0$. Thus, optimal terminal wealth varies more with W_T at low levels of the sum of Brownian shocks. If we imagine the Mean-Variance investor at a short time-increment before the terminal time point, then with constant investment opportunities, the delta of the hedging strategy is, thus, larger, if the sum of Brownian shocks to date is low. At the same time, a low sum of Brownian shocks to date

[37] Using the risk-free rate, see discussion in section 3.2.2.

also implies that the current absolute stock exposure may be low, so that the exposure to the risky asset must be increases to hedge a given spread of wealth levels across states of the world in the future. Thus, from the mechanics of the underlying delta-hedge, potentially large swings in investment fractions may be expected, especially at wealth levels close to zero.

In turn, static investment strategies effectively reduce the investment horizon to one single period. As within the given market setting, the single-period Mean-Variance optimization problem is equivalent to the CARA investor's single period portfolio choice problem (see e.g. Munk, 2017), the static, buy-and hold strategies of the two investors coincide. Thus, as in (34), derived in appendix A8 for the CARA investor, the Mean-Variance investor's static strategy is given by

$$\theta_0^{MV,ST} = \frac{2(\mu - r)}{\delta\sigma^2(e^{rT} + 1)}. \tag{39}$$

Finally, the investor's myopic strategy is derived in appendix A12, following the procedure proposed by Chen & Zhou (2022) and specializing to the present setting. The strategy is given by

$$\theta_t^{MV,MY} = \frac{1}{\delta}\frac{(\mu - r)}{\sigma^2} e^{-r(T-t)+\lambda^2(T-t)}. \tag{40}$$

The resemblance between the myopic and and pre-commitment solutions is due to the latter's derivation on the basis of the former, as seen in appendix A12.

4 Evaluation of strategies based on certainty equivalents

To study Mean-Variance approximations, we now evaluate the Mean-Variance investor's strategies under CARA - and CRRA utilities. Furthermore, to study the commitment problems related to time-inconsistency, the Mean-Variance investor's strategies are also evaluated under her own objective function. Direct comparisons on the basis of utility[38] are not possible, as utility functions are only defined up to positive monotone transformations (see e.g. Mas-Colell et al., 1995). Thus, the use of Certainty Equivalents has previously been suggested instead to compare the value of various portfolio strategies to different investors (Kallberg & Ziemba, 1979; Pulley, 1985) In particular, we follow Balter et al. (2021) in converting the achieved certainty equivalent into certainty equivalent returns, that is the certain (guaranteed) returns that would make investors indifferent between receiving this fixed return and holding the risky prospect, respectively (ibid). On the one hand, we will use the certainty equivalent growth rate, as in Balter et al. (2021), that is

$$CE^G = \frac{1}{\hat{T}} \log\left(\frac{CE}{X_0}\right), \tag{41}$$

whereas $\hat{T}$ will be a variable time of evaluation, not necessarily equal to maturity T, and CE be understood as the achieved certainty equivalent from time zero to time $\hat{T}$. In the case of $CE < X_0$ it will often be more convenient, and if $CE < 0$ it will be necessary, to work with certainty equivalent total gross returns instead, that is

$$CE^R = \frac{CE}{X_0}. \tag{42}$$

This section starts by comparing investment fractions, amounts invested, as well as the implied terminal wealth profile and wealth dynamics in order to gain further understanding of the portfolio strategies through which the investors aim to maximize their respective value criterion. Next, certainty equivalents and certainty equivalent returns will be computed. Following Balter et al. (2021), where possible, the sources of sub-optimality for various strategies will be more closely assessed, with the aim of identifying the effect of over - and underinvestment and investment scheduling, respectively. The results will also speak to the value - as well as the difficulty of commitment for the Mean-Variance investor. Finally, it will be shown that time-inconsistency of the Mean-Variance investor's objective function is value enhancing when it comes to approximating the overall optimal CARA and CRRA investment strategy, respectively.

As the investment fractions of various strategies discussed are stochastic, the analysis will be based on simulations. Investors will be assumed to have initial wealth $X_0 = 100$. Drift and volatility terms of the stock will be set to $\mu = 0.08$ and $\sigma = 0.2$, respectively, whereas the risk-free rate will be set to $r = 0.02$. Finally, we will assume the initial values $S_0 = 10$ and $B_0 = 1$ for the stock and bond, respectively. Given that the stock follows a Geometric Brownian Motion, it may be simulated without bias by simulating the analytical solution of the respective price process, freeing up computation time to improve precision. We, hence, set the number of scenarios to 1 000 000. While the number of trajectories may seem excessive at first, it will

[38] Here, utility refers to the utility derived by the CARA - and CRRA investors. The 'value' derived by the Mean-Variance investor as through her objective function in fact delivers a comparable 'measure of value', as the objective function of a risky bet is equal to the certainty equivalent derived by the Mean-Variance investor.

be discussed below and shown in appendix A13 that a large amount of precision is required for Monte-Carlo means to be sufficiently close to the true underlying means. The investment horizon is initially set to $T = 10$ years and will subsequently be increased to $T = 30$ to evaluate certainty equivalent returns in the longer run. Moreover, we discretize using 120 and 360 time-steps, respectively, which yields monthly rebalancing opportunities. Finally, given the initial wealth, the risk aversion parameters will be set to $\gamma = 3$ and $\alpha = \delta = \frac{\gamma}{X_0} = 0.03$. Empirical estimates of the respective risk-aversion parameters vary widely (Elminejad et al., 2022) and without further adaptations in the given utility functions, the assumed asset-price parameters could not be justified on basis of the assumed risk-aversion parameter if we also assume that these utility function are to some degree representative and asset prices are consumption-based and reflective of these utility functions (Mehra & Prescott, 1985, Cochrane, 2001). Hence, given the inherent inconsistency therein, we set risk-aversion parameters such that they lead to reasonable attitudes towards risk within the context of the present discussion. To illustrate, given $X_0 = 100$ and $\gamma = 3$, the CRRA investor would pay a price of 20 for a gamble that would increase her wealth by one half to 150 or yield zero gain, both with probability 1/2. Moreover, she would require compensation of 33 to hold the risk of having her wealth halved to 50 or gaining 50 to reach 150, each with equal probability. Similarly, the CARA investor would pay 16 for the first gamble and would require compensation of 29 to hold the risk from the second gamble. These results appear broadly reasonable.

4.1 Comparison of risk-profiles

4.2 Average investment fractions and invested amounts

Figures 1 and 2 show the average fractions of financial wealth invested in the stock across trajectories at each time $t \in [0, 10]$, whereas Figure 4 shows the resulting amounts invested in the risky asset, for each of the strategies considered. Figure 3 shows the median investment amounts for comparison. The behavior of the myopic, consistently planning and pre-committed CARA - and CRRA investors follows immediately from the discussions in section 3. For all except for the static buy-and-hold strategy, the CRRA investor chooses constant investment fractions, leading to increasing average amounts invested in the stock. For the same strategies, the CARA investor chooses deterministically increasing absolute amounts invested in the stock. Given that amounts invested in the stock investment increase at rate r and that the stock earns positive expected excess returns, the chosen amounts invested lead to decreasing average portfolio fractions invested in the stock over time. For both investors, in turn, the static investment strategy leads to increasing average fractions - and consequently amounts - invested in the stock over time, which is again due to the stock's positive expected excess returns. The latter invested amounts also increase more aggressively over time than for the consistent planning, pre-commitment and myopic strategies, again due to expected excess returns on the stock together with a lack of rebalancing.
The Mean-Variance investor's static and dynamically consistent strategies coincide with those of the CARA investor, respectively, as evidenced by both amounts invested and investment fractions. The Mean-Variance investor's pre-commitment and myopic strategies, in turn, are characterized by considerable stochasticity and sudden large changes in portfolio positioning in

terms of investment fractions, as can be seen in Figures 1 and 2. Overall, the investment fractions and amounts invested in the risky asset according to both strategies follow a downward trend [39], starting at a slightly leveraged position a little higher than full investment in the risky asset and falling to roughly the same level as that preferred by the consistently planning Mean - Variance - and CARA investors at the end of the investment horizon. The latter result is intuitive, as the Mean-Variance investor is a CARA investor at heart and wishes to have at least a similar [40] absolute exposure to the risky asset at the terminal time point as the CARA investor. Median risky investment fractions in Figure 3 confirm a clear overall downward trend and show that average fractions are significantly influenced by rare and extreme portfolio re-positioning[41]. In spite of finishing at similar investment levels as the pre-committed CARA investor, it follows from the discussion in section 3 that the pre-committed Mean-Variance investor chooses a fundamentally different terminal wealth profile as compared to the pre-committed CARA investor, which results in significantly different investment strategies prior to the terminal time point.
Direct economic interpretation of (38) would suggest that the downward sloping schedule for investment in the risky asset is due to the asset's expected excess return, which incentivizes the investor to invest large fractions of total wealth in the stock at the beginning of the investment horizon.[42] This effect is larger over longer (remaining) investment horizons. As time progresses, the Mean-Variance investor reduces exposure to the risky asset over time, hence, arguably first prioritizing expected excess returns and then increasingly shifting focus to variance reduction to achieve roughly the same terminal wealth level as the pre-committed CARA investor. This is similar in spirit to a common theme in the life-cycle investment literature, which suggests that exposure to the risky asset should be larger the longer the investment horizon in order to benefit from expected excess returns while being able to cushion negative shocks by positive returns over time [43] (see e.g. Bovenberg et al., 2007). The downward sloping schedule for risky investment would also be in line with direct economic interpretation of the Mean-Variance optimization criterion. Given positive expected excess returns, initial wealth is (significantly) below target expected wealth, with this gap being - on average - reduced over time. Hence, it seems intuitive that the Mean-Variance investor would aim at closing this gap by larger risky investment at the beginning of the investment horizon (Basak & Chabakauri, 2010; Bajeux-Besnainou & Portait, 1998; Richardson, 1989), with this goal being gradually replaced by greater focus on variance reduction as average financial wealth increases over time. We would call this a 'forward-looking' direct economic interpretation, as it reflects a picture of an investor looking ahead at a given time and comparing her present wealth to target expected wealth.
As seen particularly clearly in Figure 4, the pre-committed Mean-Variance investor's risk appetite is considerably higher than her risk-appetite exhibited through her consistent-planning strategy or the risk-appetite shown in the CARA investor's pre-commitment or consistent planning strategies, respectively. This would also be in line with a violation of Second-order stochastic dominance by the pre-committed Mean-Variance investor, in the sense that she would weigh

[39] Richardson (1989) also finds the respective downward trend, whereas based on his chosen parameter setting, the Mean-Variance investor even initially uses leverage to further increase stock holdings beyond 100 per cent of financial wealth. However, the author does not further analyse the mechanisms behind the Mean-Variance investor's optimal choice driving this downward trend.

[40] The ultimately desired exposure might not be equal to that preferred by the CARA investor, as the two investors' actual risk aversion may differ due to the discussed imperfections in the parameterization discussed in section 2.

[41] See also discussion below.

[42] See the multiplicative term $e^{\frac{1}{2}\lambda^2(T-t)}$, which results from subtraction of the exponents in $e^{\lambda^2(T-t)}$ and M_t in (38).

[43] As discussed above, lifecycle investment tends to be analyzed on the basis of utility functions rather than Mean-Variance criteria. It is, hence, interesting that a similar theme emerges in both investment problems.

losses and gains equally, much in contrast to expected utility investors with strictly concave utility functions, such as CARA investors, whose dis-utility from a given loss outweighs the utility from an equally-sized gain (see Hadar & Russel, 1969; Hannoch & Levy, 1969; Machina & Rothschild, 1990). However, in a setting of normally distributed returns on the assets, as given by the Geometric Brownian Motion for the stock in (1), and normally distributed terminal wealth prior to application of a particular investment strategy, as given in (92), it is not entirely clear whether a 'forward looking' Mean-Variance investor would indeed violate Second-Order Stochastic Dominance. Let us be more precise and define a 'forward looking' Mean-Variance investor as an investor, who maximizes her current value from the mean and variance of future wealth via linear - and non-stochastic - portfolio weights. Given normally distributed asset returns - or wealth processes - such a 'forward looking' Mean-Variance investor satisfies Second-Order stochastic dominance (see discussion in Johnstone & Lindley, 2013). The static and dynamically consistent Mean-Variance optimization problems are 'forward looking' in this sense, as evidenced by the respective solutions[44] and the resulting terminal wealth processes.[45] (see discussion in Johnstone & Lindley, 2013; Hanoch & Levy, 1969) The pre-committed Mean-Variance investor, in turn, is not. In essence, the Martingale Method is inherently 'backward looking', determining an optimal terminal wealth profile first and then hedging it on the basis of financial assets in the market, which leads to the particular stochastic portfolio strategy in (38) and the resultant non-normally distributed terminal financial wealth, which, in turn, prevent her from qualifying as 'forward looking' in the above sense[46]. Thus, the pre-committed Mean-Variance investor may not adhere to Second-Order Stochastic Dominance (see discussion in Johnstone & Lindley, 2013), in contrast to the consistently planning Mean-Variance investor, for example.

This suggests that the significant differences between pre-committed CARA - and Mean-Variance investment, as well as between pre-committed and dynamically consistent Mean-Variance investment, in Figures 1 to 4, should be explained based on a 'backward-looking' interpretation of the Mean-Variance pre-commitment strategy[47], taking into account the particular mechanics of the Martingale Method, rather than a direct economic interpretation based on the strategy or the Mean-Variance objective alone, which appears to be forward-looking in character.

Given that the dynamically consistent strategies of the CARA - and Mean-Variance investors coincide, it might be argued that mechanics of the Martingale-Method and the resulting failure of the Mean-Variance pre-commitment strategy to satisfy Second-Order Stochastic Dominance are indeed pivotal for the dramatic deviation of the latter strategy from the alternatives available to the Mean-Variance investor. This does not prevent an interpretation of the respective value process X_t^* to be hedged, as given in (82), in terms of the current value of the expected future target value, which, in turn, results from the optimally chosen terminal wealth profile. The latter interpretation would be in line with Basak & Chabakauri (2010), Bajeux-Besnainou

[44]See appendices A7 and A8, respectively. Furthermore, the static Mean-Variance strategy, for the very reason of bering 'forward looking', can be computed as the static CARA strategy (see Munk, 2017; Johnstone & Lindley, 2013).

[45]See also Johnstone & Lindley's (2013) discussion that the forward-looking Mean-Variance investor's choices may be represented in the (μ, σ) plane, implying the existence of indifference curves, which is not necessarily the case without further conditions in the case of non-normally distributed portfolios (or wealth processes, that is).

[46]Neither can stochastic investment decisions reasonably be seen as 'forward looking' in the general understanding of the word, as the strategies are not adapted to the information set at the decision time.

[47]While the CARA investor's pre-commitment strategy is, likewise, computed via the Martingale Method, she applies her utility function to the optimization over terminal wealth levels in the first step, which in turn implies that her choice over terminal wealth levels should not lead to the same violations of Second-Order Stochastic Dominance as for the Mean-Variance investor.

& Portait (1998) and Richardson (1989). But arguably, the underlying mechanisms leading to the unique character of the Mean-Variance pre-commitment policy may only be seen with reference to the 2-step procedure according to the Martingale method, focus on the choice of terminal wealth profiles per se in step 1, the non-normality of terminal wealth, which results from the greater range of investment strategies available as per this method, and the resulting potential violation of Second-Order Stochastic Dominance by the Mean-Variance investor's pre-commitment strategy.

Turning to an explanation based on the hedging strategy, as per the Martingale Method, optimal future wealth levels tend to be (significantly) above the initial wealth level due to expected excess returns. As discussed in section 3, the resulting 'delta' of the hedging strategy is relatively large, which, in turn, leads to larger stock investment in the hedging portfolio. At the same time, the variance of the stock over longer time-periods is, likewise, higher, so that large changes in optimal terminal wealth levels may be hedged by large, but relatively stable stock investment. As wealth increases on average over time, the 'delta' of optimal terminal wealth is reduced, so that optimal risky asset exposure decreases, while becoming more volatile. In a sense, the Mean-Variance investor first hedges via levels of investment and later hedges via increasingly large changes in the composition of her hedging portfolio.

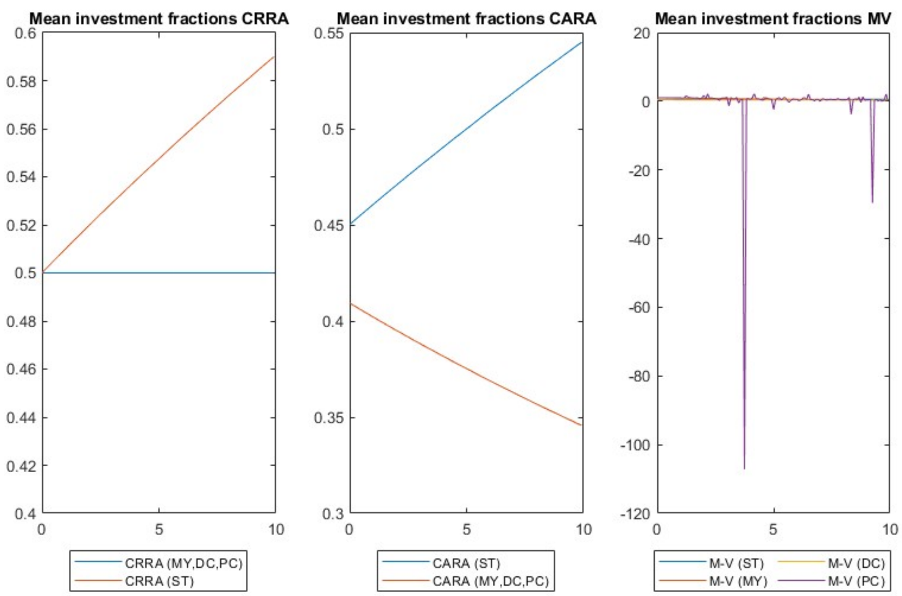

Figure 1: Average investment fractions, stock investment, across trajectories. X-axis: time, Y-axis: investment fraction. Parameters: T=10, m=1 000 000, n=120.

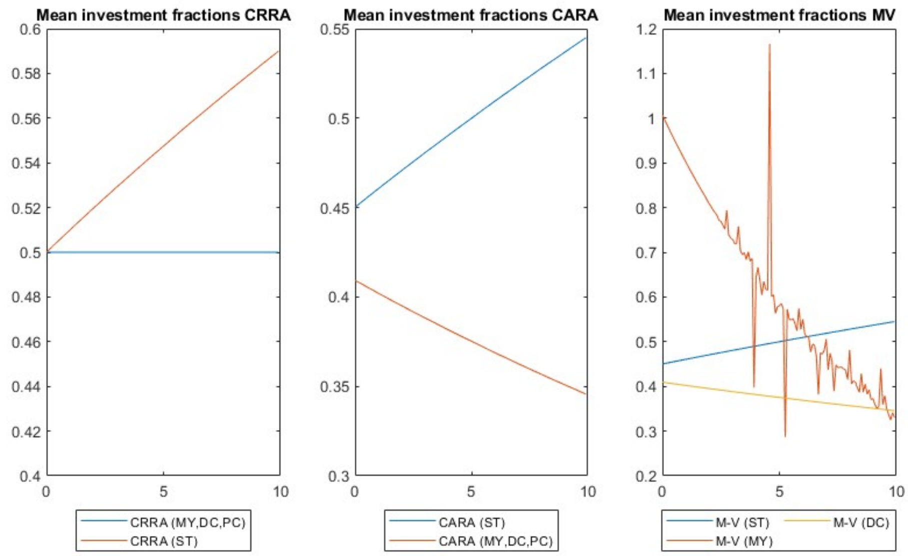

Figure 2: Average investment fractions, stock investment, across trajectories, Mean-Variance pre-commitment strategy excluded. X-axis: time, Y-axis: investment fraction. Parameters: T=10, m=1 000 000, n=120.

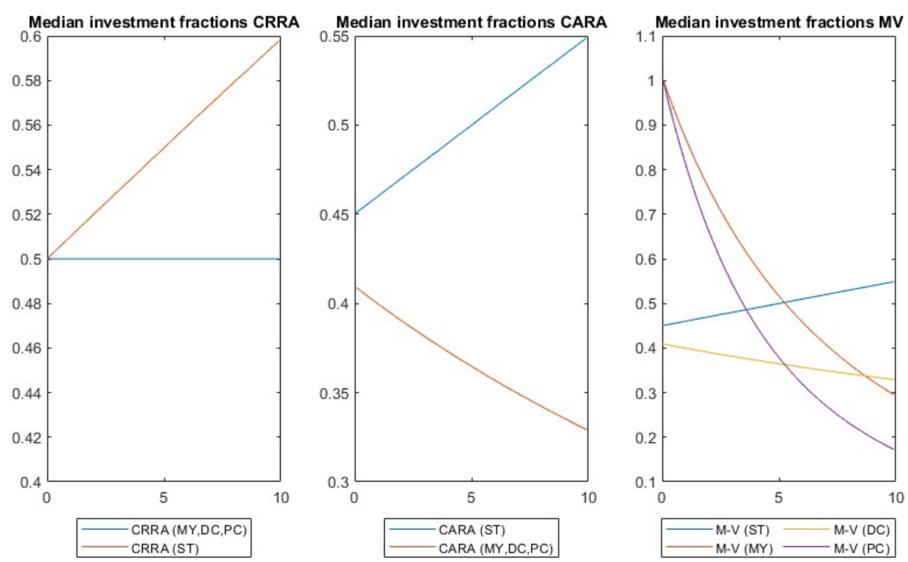

Figure 3: Median investment fractions, stock investment. X-axis: time, Y-axis: investment fraction. Parameters: T=10, m=1 000 000, n=120.

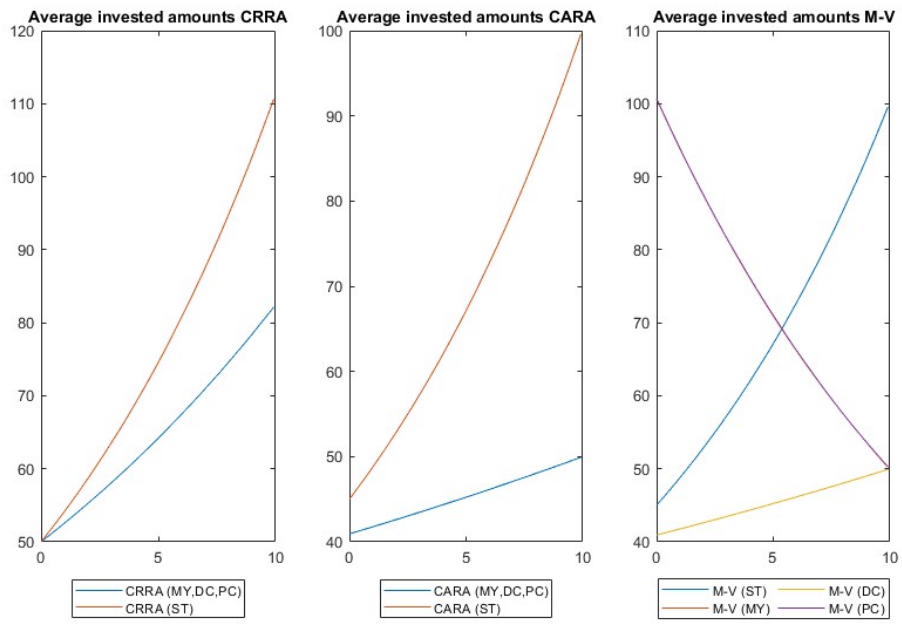

Figure 4: Average amounts invested in the stock across trajectories. X-axis: time, Y-axis: invested amount. Parameters: T=10, m=1 000 000, n=120.

To further analyze the large and sudden adjustments in investment fractions, as per the Mean-Variance investor's pre-commitment and myopic strategies, we perform another set of simulations of financial markets with the same parameter setting but using a coarser grid of 10 time steps, in order to allow for changes in key variables to become more readily visible. Figures 5 and 6 together provide a closer look at the six trajectories, which lead to the most extreme choices of portfolio weights in pursuit of the pre-commitment strategy. As for the myopic strategy, we confine ourselves to assessing more closely the two trajectories with the largest and smallest portfolio weight, respectively, in Figure 7, given that extremes are less likely under that strategy.

Figures 5 to 7 show that the most important determinant for large investment weights (in absolute terms) is wealth becoming close to zero, which is reasonable, because both strategies are primarily determined as invested amounts, whereas invested fractions are then obtained after division by wealth. In all cases, wealth becomes close to zero after a series of negative Brownian shocks. Furthermore, in all cases, both positive and negative weights lead to large long positions in the stock, so that negative weights do not signify genuine short positions, but merely positive invested amounts, which lead to negative fractions after division by negative wealth. In all cases, there is a correlative short position in the bond, so that the investor borrows at the risk-free rate to invest in the stock. This is in line with the Mean-Variance investor's character as a CARA investor at heart, who wishes to have a constant absolute exposure to the risky asset, irrespective of current wealth level. (compare Merton, 1969; Pratt, 1964; Eeckhoudt et al., 2005) As a result, the apparently extreme portfolio weights do not translate into extreme amounts invested in the stock per se, even though relative to her current wealth, these overall moderate levels of stock investment would still seem extreme. As a result, if portfolio constraints are not a concern and if we are willing to assume that markets would readily lend the investor 20 000 times her current wealth to invest an amount in the stock, which would overall seem reasonable for a general investor with average financial wealth, then the investor's behavior should rather be analyzed according to amounts invested, rather than investment fractions. We maintain, however, that borrowing without equity is in general difficult, especially if the intended use is for risky investment. We will, thus, maintain a focus on investment weights, but keep in mind that these extreme weights translate into overall reasonable amounts invested. Finally, we note that the Mean-Variance investor's gamble, that is either to increase expected returns given low current wealth (Richardson, 1989; Basak & Chabakauri, 2010) or to hedge a large delta of optimal terminal wealth relative to the hedging protfolio with an outsized risky position relative to current wealth, pays off in the majority of cases, at least in terms of terminal wealth on the selected trajectories.

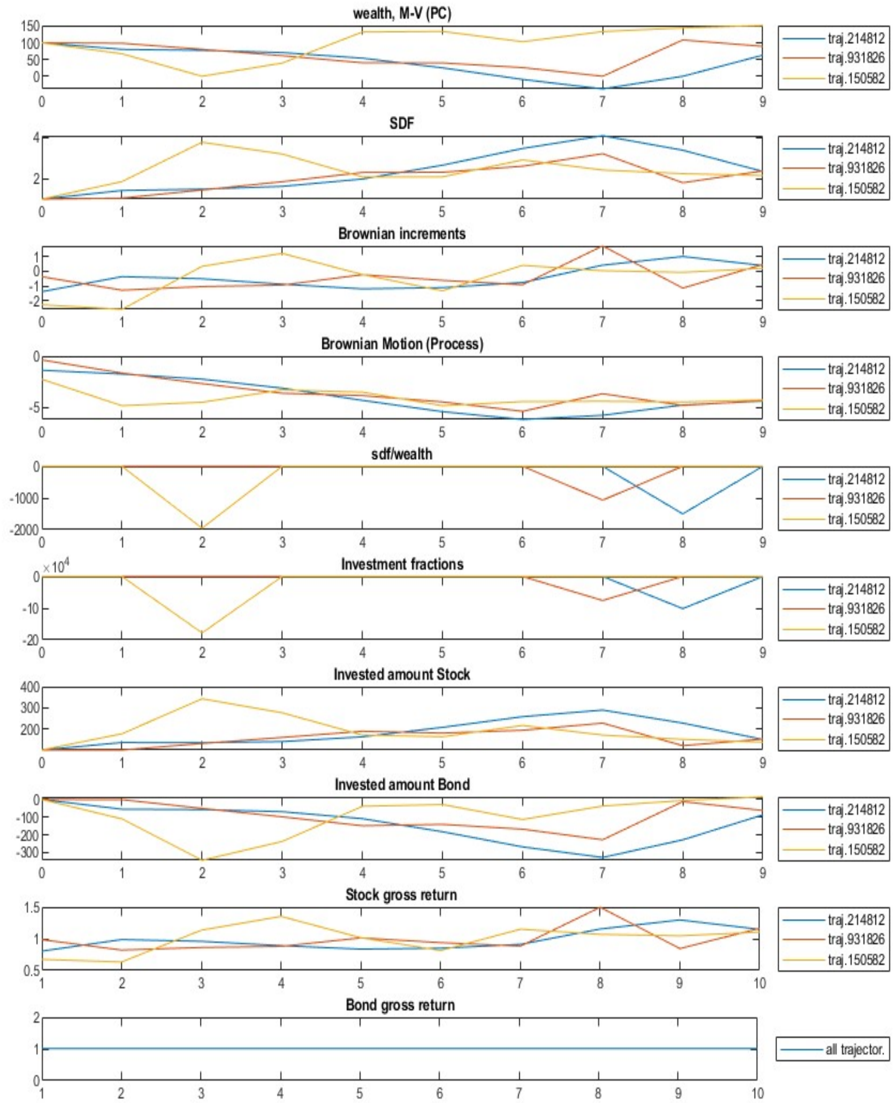

Figure 5: Large negative weights, M-V (PC) investor. X-axis: time, Y-axis: see plot heading. Parameters: T=10, m=1 000 000, n=10.

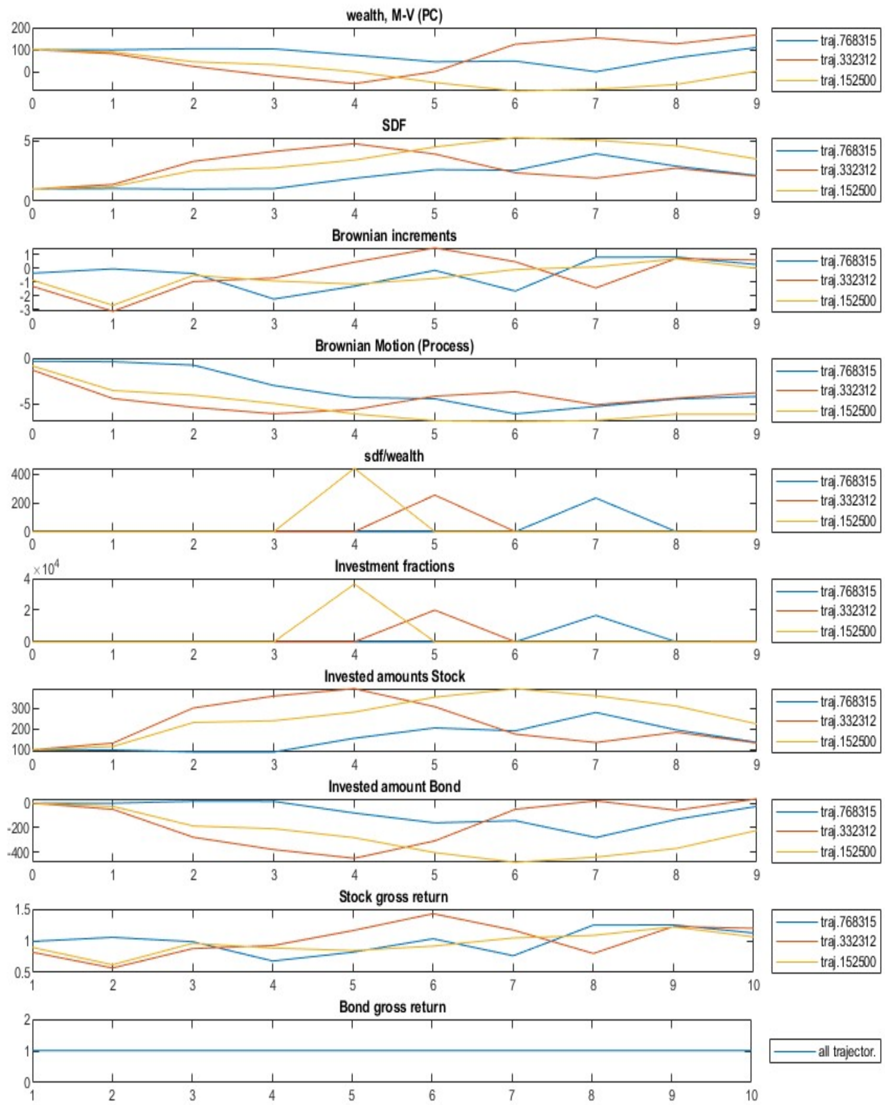

Figure 6: Large positive weights, M-V (PC) investor. X-axis: time, Y-axis: see plot heading. Parameters: T=10, m=1 000 000, n=10.

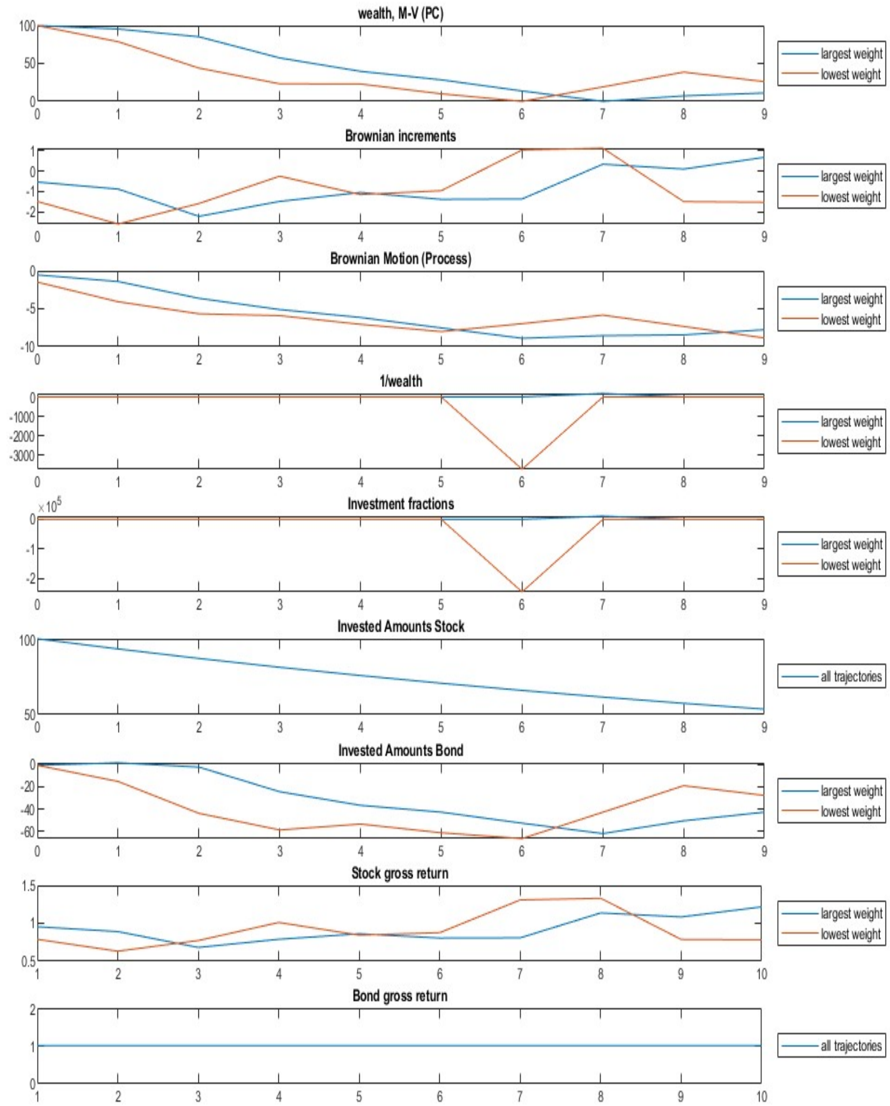

Figure 7: Large negative and positive weights, M-V (MY) investor. X-axis: time, Y-axis: see plot heading. Parameters: T=10, m=1 000 000, n=10.

Finally, for robustness, we re-run the original set of simulations, but use the average wealth level achieved by the CRRA investor as a reference level for parameterization. The CRRA investor's wealth level is chosen, as this is exogenous to the choices of the CARA - and Mean-Variance investors, which are, in turn, influenced by their own level of risk-aversion. The CRRA investor's average wealth level is given by $\bar{X} = 130.3651$. As seen in Figure 8, there are no significant qualitative changes in the results. The investment schedules for the CARA - and Mean-Variance investors are slightly shifted upwards, whereas the schedules for their static strategies, as well as the investment pattern of the myopic and pre-committed Mean-Variance investor, become steeper. Both reflect an overall lower risk-aversion on part of the CARA - and Mean-Variance investor, as compared to the base-case. As investment patterns remain broadly unchanged, however, the remainder of the analysis will be performed on the basis of the initial paramterization.

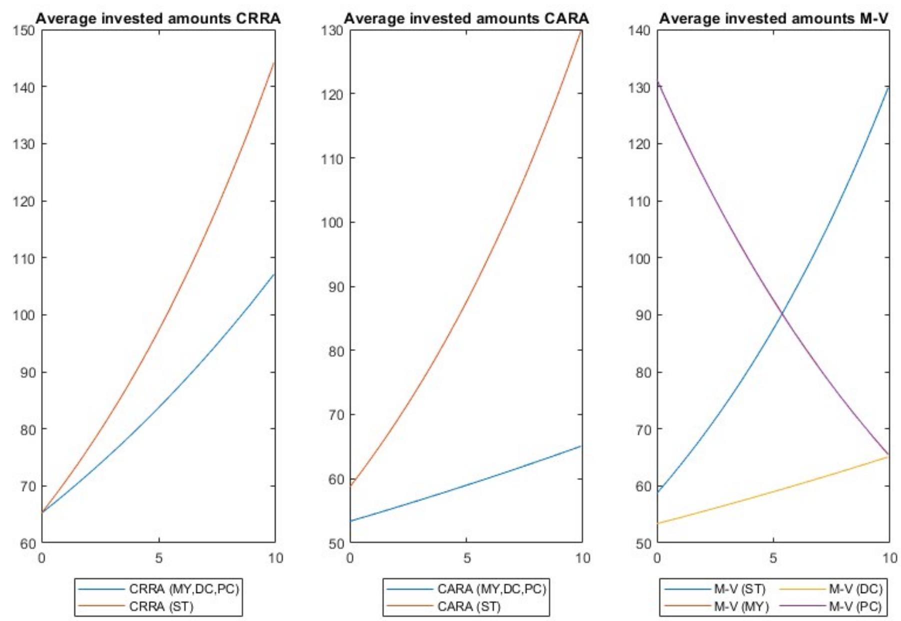

Figure 8: Average invested amounts across trajectories, reference wealth level X=130.3651. X-axis: time, Y-axis: amounts invested. Parameter setting: T=10, m=1 000 000, n-120.

4.3 Certainty equivalents, investment fractions and implied wealth processes

We now assess whether the differences in investment patterns - and amounts lead to differences in value from the respective risky prospects for the individual investors. Thus, we first discuss differences in the induced wealth processes, that is in particular means and standard deviations of terminal wealth and the variability of wealth along the trajectories. We then assess certainty equivalents from terminal wealth. Having witnessed the large positive and negative investment fractions chosen by the myopic and pre-committed Mean-Variance investors, we also impose ex-post portfolio constraints and assess their impact on certainty equivalents from the respective strategies.

We see from Figures 9 to 12 that the wealth resulting from the CRRA investor's myopic, precommitment and consistent planning strategies is slightly more variable across trajectories than for the respective strategies of the CARA investor. This coincides with a greater variability in amounts invested in the stock (with the amount invested being deterministic in the case of the CARA investor).

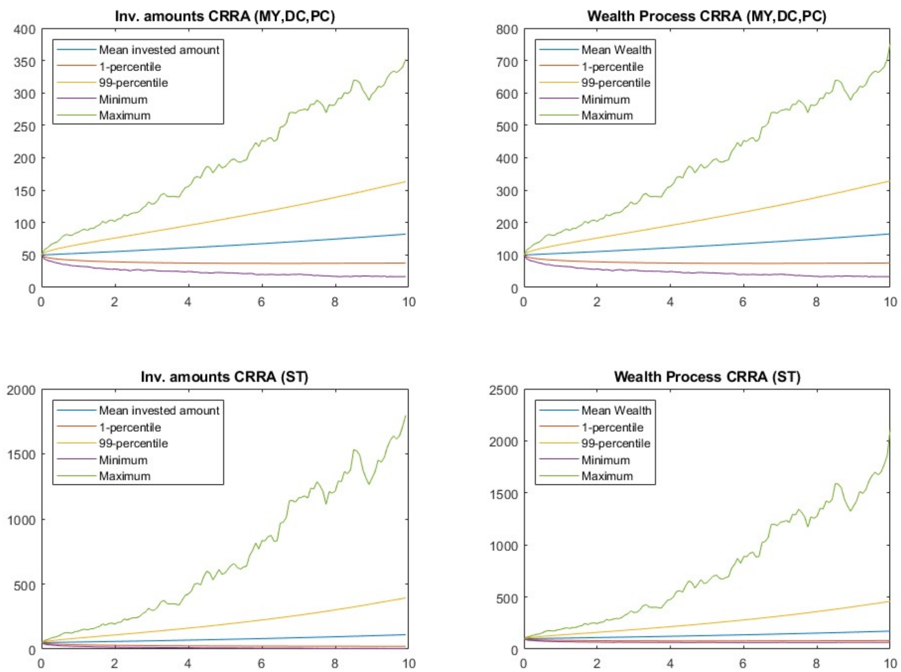

Figure 9: Invested amounts and wealth processes, CRRA investor. X-axis: time, Y-axis: invested amounts (left) & wealth (right). Parameters: m=1 000 000, n=120, T=10.

At the same time, wealth stays firmly positive for the CRRA investor at any time, whereas it eventually turns negative in rare instances for the CARA investor[48]. Hence, constant investment fractions provide greater protection against downside risks, whereas constant invested amounts lead to smaller overall variability of wealth. This is intuitive, because the CRRA investor scales down risky investment as wealth decreases, whereas the CARA investor does not increase risky investment past a certain amount - indeed, it does not vary at all - as wealth increases. For all three investors, static strategies lead to considerably greater variability of wealth than any other strategies, which is intuitive, because the static investor's strategy is to follow the whims of financial markets without changing course along the way.
Finally, as the two strategies coincide, the Mean-Variance investor's consistent-planning strat-

[48]The probability of negative realizations of wealth is of course dependent on the initial wealth assumed, so that the latter probability would be significantly larger if initial wealth was assumed to be considerably lower.

egy leads to the same amounts invested and wealth profile as the respective CARA strategy. The Mean-Variance investor's pre-commitment strategy, in turn, leads to significantly higher downside risk with respect to wealth than the respective CARA strategy, while invested amounts also vary considerably. Apart from a few extreme negative realizations, wealth under the pre-commitment strategy is considerably less variable for the Mean-Variance investor than for the CRRA investor, in turn. As will also be seen from the tables below, her terminal wealth is, however, more variable than for CARA investors under the respective strategy. Thus, as compared to the CRRA investor, the pre-committed Mean-Variance investor appears to adopt the overall greater focus on variance-reduction of the CARA investor. At the same time, in contrast to the CARA investor, she is willing to accept considerably greater downside risk in rare events for the benefit of a greater upside potential under the pre-commitment strategy. This seems broadly in line with the violation of Second-Order stochastic dominance discussed above. Similar reasoning applies to analysis of the Mean-Variance investor's myopic strategy.

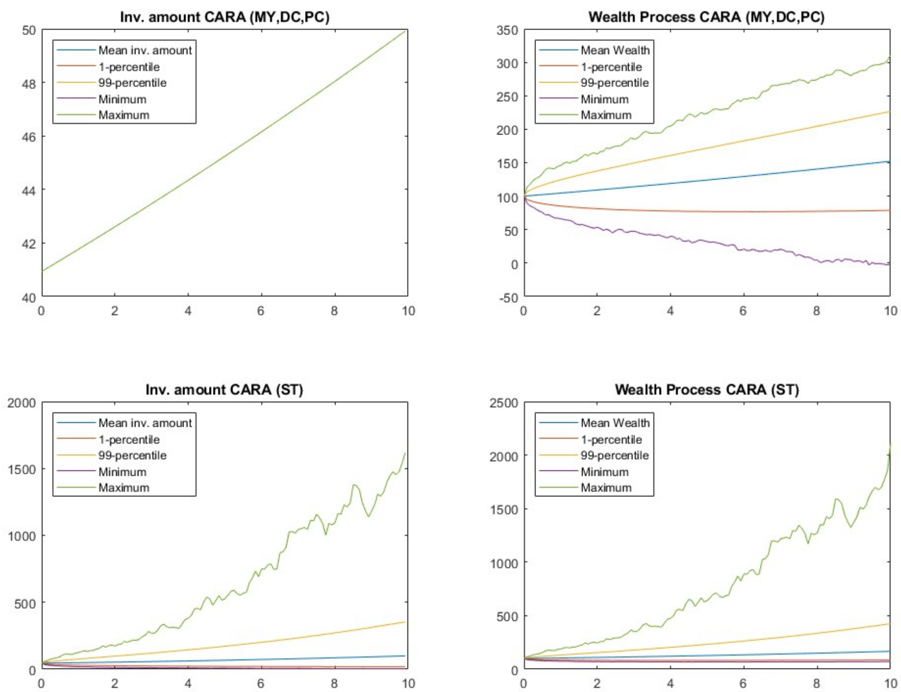

Figure 10: Invested amounts and wealth processes, CARA investor. X-axis: time, Y-axis: invested amounts (left) & wealth (right). Parameters: m=1 000 000, n=120, T=10.

Tables 1 to 3 below confirm that the pre-commitment strategy is, indeed the overall optimal strategy for all three investors. Even though comparison is across different risky prospects resulting from different strategies in this case, the Mean-Variance investor derives a greater certainty equivalent from her pre-commitment strategy than the CARA - and CRRA investors. The Mean-Variance investor also takes considerably larger risk under her pre-commitment strategy than under her consistent-planning strategy. Moreover, under the pre-commitment strategy, she achieves her greater certainty equivalent with significantly larger risk and larger average terminal wealth than the CARA investor. At the same time, under their respective pre-commitment strategies, the Mean-Variance investor appears to derive greater value than the CRRA investor, due to greater focus on variance reduction[49], as the difference in means of terminal wealth is relatively small, whereas the difference in standard deviations is larger than between the CARA - and Mean-Variance investors. Finally, the extreme portfolio re-positioning in terms of investment fractions, as seen above, is sufficiently rare so that the standard deviation of terminal wealth under the Mean-Variance pre-commitment strategy still stays below the standard deviation of terminal wealth under the respective strategy of the CRRA investor.

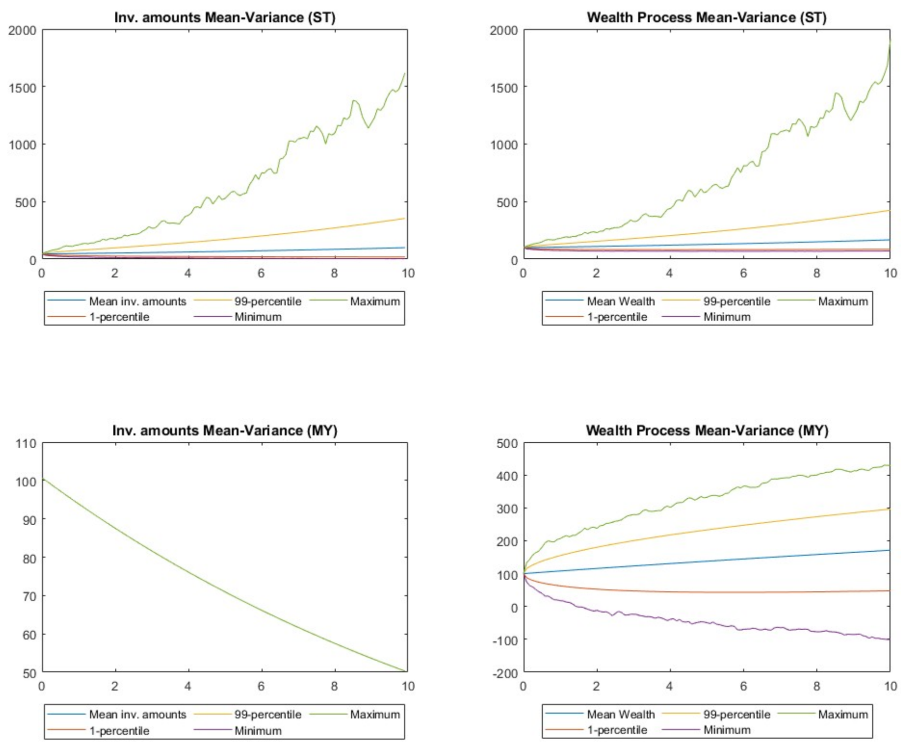

Figure 11: Invested amounts and wealth processes, M-V investor, static and myopic strategies. X-axis: time, Y-axis: invested amounts (left) & wealth (right). Parameters: m=1 000 000, n=120, T=10.

[49] Of course, as an expected utility investor, the CRRA investor does not target variance-reduction per se, but rather changes in expected utility. Variance-reduction is, thus, not itself a target but an effect. The same holds for the CARA expected utility investor above.

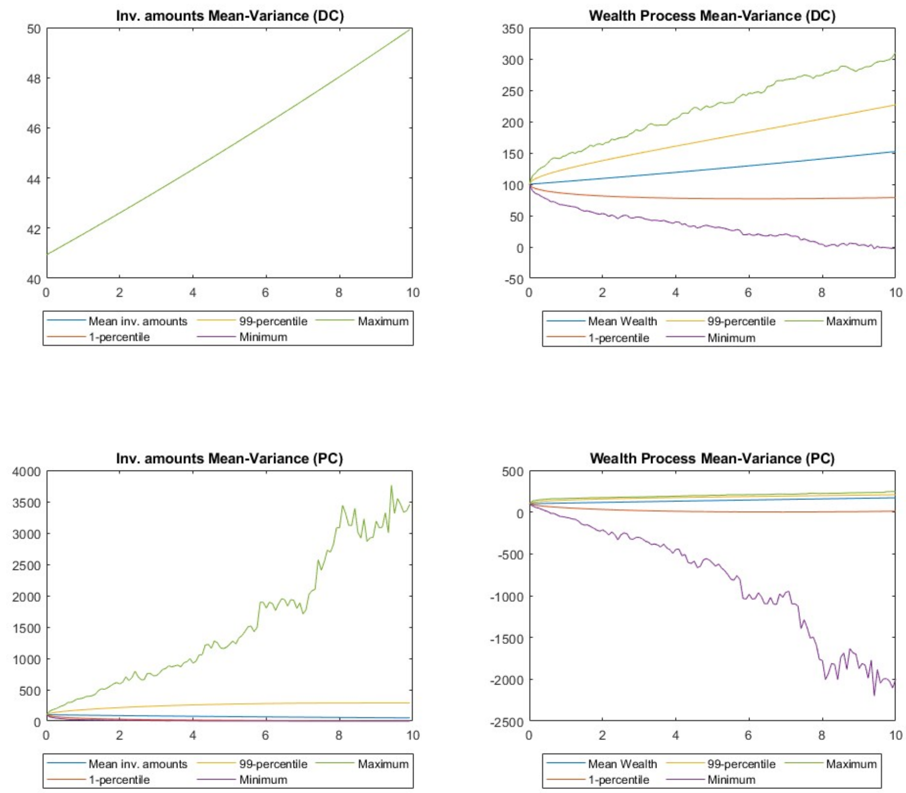

Figure 12: Invested amounts and wealth processes, M-V investor, dynamically consistent and pre-commitment strategies. X-axis: time, Y-axis: invested amounts (left) & wealth (right). Parameters: m=1 000 000, n=120, T=10.

μ_{MC} $(\sigma/\sqrt{n})$	ST	MY	DC	PC
M-V investor	167.38 (0.07)	171.12 (0.05)	152.22 (0.03)	171.14 (0.04)
CARA investor	167.38 (0.07)	152.9 (0.03)	152.9 (0.03)	152.9 (0.03)
CRRA investor	172.39 (0.08)	164.95 (0.05)	164.95 (0.05)	164.95 (0.05)

Table 1: Mean of terminal wealth achieved, unconstrained investor. Parameters: m=1 000 000, n=120, T=10.

$\sigma(X_T)$	ST	MY	DC	PC
M-V investor	70.51	53.51	31.81	40.84
CARA investor	70.51	31.81	31.81	31.81
CRRA investor	78.32	53.74	53.74	53.74

Table 2: Standard deviation of terminal wealth achieved, unconstrained investor. Parameters: m=1 000 000, n=120, T=10.

$CE(X_T)$	ST	MY	DC	PC
M-V investor	92.81	128.17	137.04	146.12
CARA investor	134.65	137.15	137.15	137.15
CRRA investor	141.06	141.92	141.92	141.92

Table 3: Certainty equivalents, unconstrained investor. Parameters: m=1 000 000, n=120, T=10.

4.4 The effect of portfolio constraints on certainty equivalents

As can be seen from Tables 4 to 9, portfolio constraints have a significant effect on the terminal wealth profile achieved by Mean-Variance under her pre-commitment strategy and a minor effect under her myopic strategy. As expected, results for the remaining strategies of all three investors remain unchanged[50]. We first restrict portfolio weights to $\pi \in [-1, 2]$, which implies that the investor may at most borrow 100% of her current wealth and invest at most 200% of her current wealth in the risky asset. Next, we constrain investment weights to $\pi \in [0, 1]$, which implies that the investor cannot borrow at all and may, thus, at most invest 100% of her current wealth in the risky asset. As regards the Mean-Variance investor's pre-commitment strategy, we see that the first set of portfolio constraints leads to a drop in the certainty equivalent achievable under Mean-Variance preferences. This drop is due to both a smaller mean terminal wealth achieved and a larger standard derivation, respectively. Interestingly, subsequent further

[50]As seen from Figure 25 in Appendix A15, together with Figure 4, toward the end of the investment horizon, as a small number of wealth trajectories for the CARA investor turn negative, the CARA investor likewise exhibits large and sudden changes in portfolio weights, akin to the Mean-Variance investor. However, they are less extreme than for the Mean-Variance investor and only occur on very rare occasions. Thus, even though portfolio constraints have been imposed on the CARA investor alike in the analysis of Tables 4 to 9, no significant changes as regards terminal wealth are found for the CARA investor.

μ_{MC} $(\sigma/\sqrt{n})$	ST	MY	DC	PC
M-V investor	167.38 (0.07)	171.04 (0.05)	152.22 (0.03)	166.32 (0.04)
CARA investor	167.38 (0.07)	152.22 (0.03)	152.22 (0.03)	152.22 (0.03)
CRRA investor	172.39 (0.08)	164.95 (0.05)	164.95 (0.05)	164.95 (0.05)

Table 4: Mean of terminal wealth achieved, $\pi \in [-1, 2]$. Parameters: m=1 000 000, n=120, T=10.

$\sigma(X_T)$	ST	MY	DC	PC
M-V investor	47.4	53.61	31.81	42.01
CARA investor	70.9	31.81	31.81	31.81
CRRA investor	78.32	53.74	53.74	53.74

Table 5: Standard deviation of terminal wealth achieved, $\pi \in [-1, 2]$. Parameters: m=1 000 000, n=120, T=10.

$CE(X_T)$	ST	MY	DC	PC
M-V investor	93.1	127.93	137.04	139.85
CARA investor	134.65	137.15	137.15	137.15
CRRA investor	141.06	141.92	141.92	141.92

Table 6: Certainty equivalents, $\pi \in [-1, 2]$. Parameters: m=1 000 000, n=120, T=10.

portfolio constraints do not lead to a further significant drop in the certainty equivalent for the Mean-Variance investor under this strategy, which may be due to the dampening effect of the portfolio constraints on the standard deviation of terminal wealth (see comparison of result in Tables 5 and 8). The effects on the myopic Mean-Variance investor are comparatively minor both in terms of means and standard deviations of terminal wealth and certainty equivalents, respectively. This suggests that ex-post portfolio constraints reduce optimal value achievable for the Mean-Variance investor in so far as it constrains her ability to take large risky investments when wealth is close to zero, in order to increase expected wealth to target expected wealth, or improve the delta hedge with respect to the optimal terminal wealth profile, respectively (see discussion in section 3). While her certainty equivalent from terminal wealth drops below that of the CRRA investor, the effect of portfolio constraints is not dramatic either, especially taking into account that the second set of portfolio constraints renders any borrowing impossible. Finally, we note that in both the constrained and the unconstrained cases, the static CARA - and CRRA strategy performs remarkably well, when comparing the respective certainty equivalents with the certainty equivalents under the investors' pre-commitment strategies. This provides further support for the argument in section 3 that within the Black-Scholes setting, there is little in terms of inter-temporal linkages to exploit, neither in terms of investment opportunities, nor in terms of preferences over time, that would yield considerable advantages to consistent-planning or pre-commitment strategies for those two investors. This is in contrast to the Mean-Variance investor, who derives considerable gains in terms of certainty equivalents from the effort of devising dynamically-consistent or pre-commitment strategies. The reasons behind this result will further be assessed below when analyzing the effects of over - and underinvestment, and investment scheduling on certainty equivalent returns.

μ_{MC} $(\sigma/\sqrt{n})$	ST	MY	DC	PC
M-V investor	167.38 (0.07)	170.38 (0.05)	152.22 (0.03)	162.13 (0.04)
CARA investor	167.38 (0.07)	152.22 (0.03)	152.22 (0.03)	152.22 (0.03)
CRRA investor	172.39 (0.08)	164.95 (0.05)	164.95 (0.05)	164.95 (0.05)

Table 7: Mean of terminal wealth achieved, $\pi \in [0,1]$. Parameters: m=1 000 000, n=120, T=10.

$\sigma(X_T)$	ST	MY	DC	PC
M-V investor	47.4	53.51	31.81	38.62
CARA investor	70.9	31.81	31.81	31.81
CRRA investor	78.32	53.74	53.74	53.74

Table 8: Standard deviation of terminal wealth achieved, $\pi \in [0,1]$. Parameters: m=1 000 000, n=120, T=10.

$CE(X_T)$	ST	MY	DC	PC
M-V investor	93.1	127.44	137.04	139.76
CARA investor	134.65	137.14	137.14	137.14
CRRA investor	141.06	141.92	141.92	141.92

Table 9: Certainty equivalents, $\pi \in [0,1]$. Parameters: m=1 000 000, n=120, T=10.

4.5 Certainty equivalent Returns from Mean-Variance Approximations

Having assessed the certainty equivalents that each investor derives from their own investment strategies, as well as the wealth dynamics resulting from the underlying risky prospects, we now assess the quality of Mean-Variance approximations more closely. Thus, we assess the certainty equivalent returns that the CARA - and CRRA investors may derive from (optimal) Mean-Variance strategies relative to the certainty equivalent returns that the investors could derive from their own pre-commitment strategies, respectively. As a further benchmark, we will also use static CARA - and CRRA strategies, even though, as noted above, the relative underperformance of static strategies for the CARA - and CRRA investors is not particularly large, due to the lack of inter-temporal linkages in the given setting. Nevertheless, if the goal is to approximate optimal expected utility strategies, doubt may be cast on the usefulness of the optimal pre-commitment Mean-Variance strategy, which may be quite costly in terms of trading costs, the human resources to devise the strategy, as well as the safeguards necessary to ensure commitment, if the strategy cannot beat a static expected utility strategy in terms of certainty-equivalent performance.

We assess certainty equivalent returns over a longer time-horizon of $T = 30$ years and using

1 000 000 trajectories, as well as 360 time steps, which yields a similar discretization as before. However, we constrain analysis to yearly performance measurement in terms of certainty equivalent returns for computational reasons. Appendix A13 underlines the importance of using both a large number of trajectories, as well as a 'larger' number of time steps for discretization, especially when assessing certainty equivalent returns over a longer time-horizon. Figure 22 shows that with 30 time steps, the theoretically optimal Mean-Variance pre-commitment strategy, would appear to yield lower certainty equivalent total gross returns than the dynamically consistent strategy. At the same time, the returns appear quite variable at the long end. Moreover, the Mean-Variance pre-commitment strategy would appear to yield a higher certainty equivalent return for the CRRA investor that her own, theoretically optimal pre-commitment strategy at the long end. Increasing the number of trajectories (Figure 23), certainty equivalent returns from the Mean-Variance pre-commitment strategy for both investors become much smoother and less volatile at the long end, the CRRA investor obtains the highest certainty equivalent return from her own pre-commitment strategy and the Mean-Variance investor's pre-commitment strategy appears close to optimal again under her own preferences. Finally, increasing also the number of time steps in the simulations to 360, the Mean-Variance investor's pre-commitment strategy again appears optimal, as seen in Figure 24. For the sake of completeness, Tables 10 and 11 show the improvements as regards precision. These results suggest that a relatively fine grid may be required for Mean-Variance pre-commitment strategies, which obtain the theoretical optimum for the Mean-Variance investor in continuous time, to appear optimal in a discretized setting. While the literature on discrete time multi-period (dynamic) mean-variance optimization is beyond the scope of this thesis, within the scope of the present analysis this suggests that considerable computation power may be required to use continuous-time Mean-Variance strategies in practice.

Figure 13 shows the underperformance of the pre-commitment Mean-Variance strategy under CRRA preferences, relative to the CRRA investor's own pre-commitment and static strategies. Given the vastly different optimal investment schedules and the relative disregard to downside risk by the Mean-Variance investor, as discussed in sections (4.2) and (4.3), it is not surprising that the CRRA investor dislikes the risky prospect proposed by the Mean-Variance investor. With an investment horizon of T=30 years, the gap in initial optimal investment between the two investors at the beginning of the investment horizon can be expected to be particularly large, leading to risks that the CRRA investor is unwilling to bear. The relatively strong improvement of the Mean-Variance investor's strategy as portfolio constraints are introduced (see lines MVPC(-1,2), MVPC(0,1), MVDC(-1,2) and MVDC(0,1), with numbers in brackets denoting the portfolio constraints discussed in section 4.4) in the eyes of the CRRA investor bears witness to this. This is again in line with failure on part of the Mean-Variance investor's pre-commitment strategy to meet criteria of Second-Order stochastic dominance. Moreover, the consistently planning Mean-Variance strategy underperforms both the static and the pre-committed CRRA strategies as well. As discussed in section 3, the dynamically consistent Mean-Variance policy is itself a 'constrained optimum' and, hence, yields sub-optimal outcomes, even for the Mean-Variance investor herself. While the strategy underperforms even relative to the static CRRA strategy, the Mean-Variance investor's consistent-planning strategy performs better in the eyes of the CRRA investor than the pre-committed Mean-Variance strategy. With reference to section 4.2, the investment schedule of the dynamically consistent Mean-Variance strategy resembles that of the optimal CRRA pre-commitment strategy more closely. Finally, as discussed above, the static CRRA strategy does not vastly underperform the investor's pre-commitment strategy in terms of certainty-equivalent gross returns.

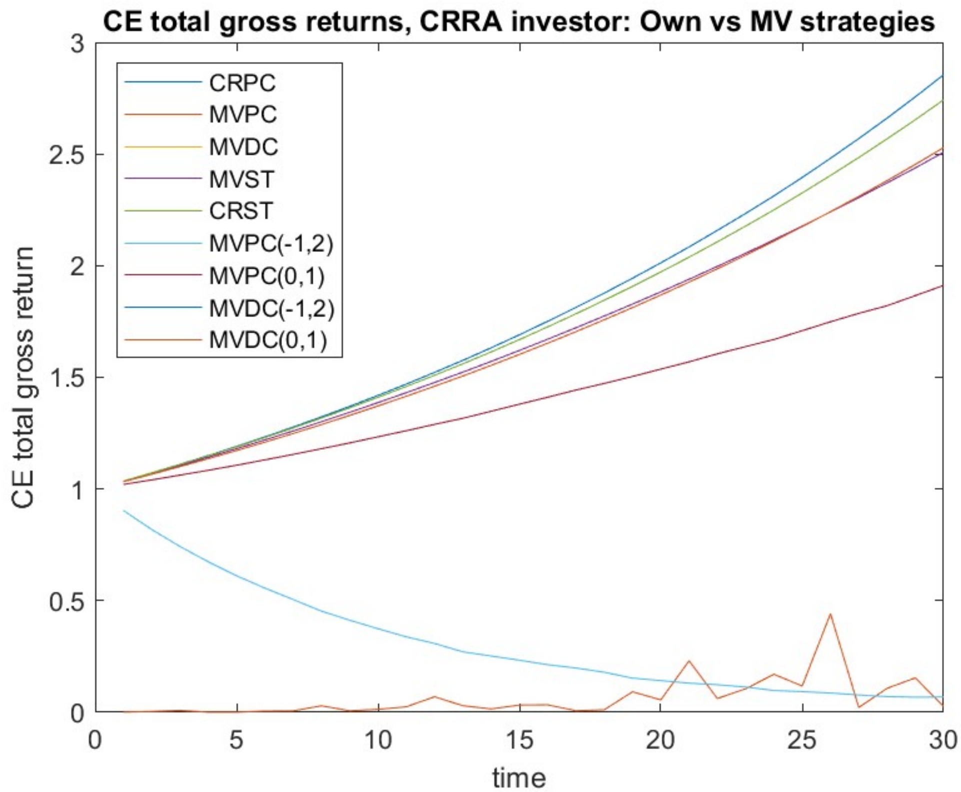

Figure 13: CE total gross returns for the CRRA investor from using static, dynamic and pre-commitment Mean-Variance strategies in terms of CE growth rate. The benchmark is given by the CRRA pre-commitment solution. Additional benchmark: CRRA investor's static strategy. Parameters: T=30, m=1000000,n=360, yearly performance measurement.

Given that all certainty equivalents (and certainty equivalent gross returns) on part of the CRRA investor remain positive across trajectories, certainty equivalent growth rates may also be derived. Figure 14 shows the losses in terms of certainty equivalent growth rates for the CRRA investor from using static CRRA - and constrained pre-commitment and consistent-planning Mean-Variance strategies, respectively, relative to the optimal pre-commitment CRRA strategy. Again, we see that the certainty equivalent growth rate lost due to use of the static CRRA strategy is almost equal to zero, whereas it is close to zero for the dynamically consistent Mean-Variance strategies under both sets of portfolio constraints[51]. In turn, losses in terms of annual certainty equivalent growth rates from using pre-committed Mean-Variance strategies become significantly larger as portfolio constraint are eased.

[51] Eventually, portfolio constraints also become binding for the dynamically consistent Mean-Variance investor, as seen in Figure 25 in appendix A14. However, as seen from the results in section 4.4, neither wealth processes nor certainty equivalents obtained change significantly due to the introduction of portfolio constraints for the dynamically consistent Mean-Variance investor (or, by extension, for the dynamically consistent CARA investor). The lack of significant consequences for the consistently - planning Mean-Variance investor is also suggested by the overlap of the two respective lines in Figure 14.

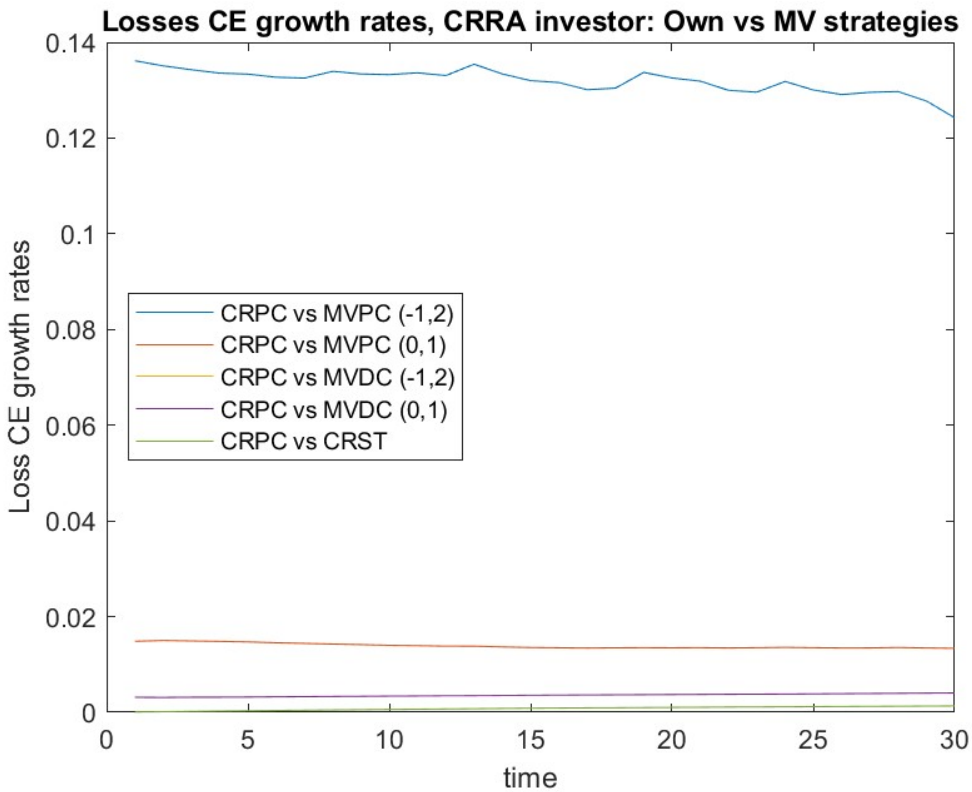

Figure 14: Losses in terms of CE growth rates for the CRRA investor from using (constrained) pre-commitment (PC) and consistent-planning (DC) strategies, measured against the CRRA pre-commitment strategy as a benchmark. The static CRRA strategy is also included for further comparison. Parameters: T=30, m=1000000,n=360, yearly performance measurement.

The extent of initial over-investment on part of the pre-committed Mean-Variance investor relative to the CRRA investor's preferences is expected to be larger the longer the investment horizon. Thus, we assess whether the Mean-Variance pre-commitment strategy would fare better in the eyes of the CRRA investor if a shorter investment horizon is assumed. However, Figure 15 suggests that while certainty total equivalent gross returns do not drop to zero right away as before at intermediate time-points of performance measurement, the patterns of underperformance remain generally unchanged and the Mean-Variance pre-commitment strategy fares by far the worst in the eye of the CRRA investor.

In summary, the CRRA investor sees optimal pre-committed Mean-Variance investment as a (burdensome) risk, rather than a (beneficial) investment. As measured in terms of certainty equivalents, the investor would essentially pay almost her entire initial wealth to avoid having to carry the risks that come with the Mean-Variance pre-commitment strategy.

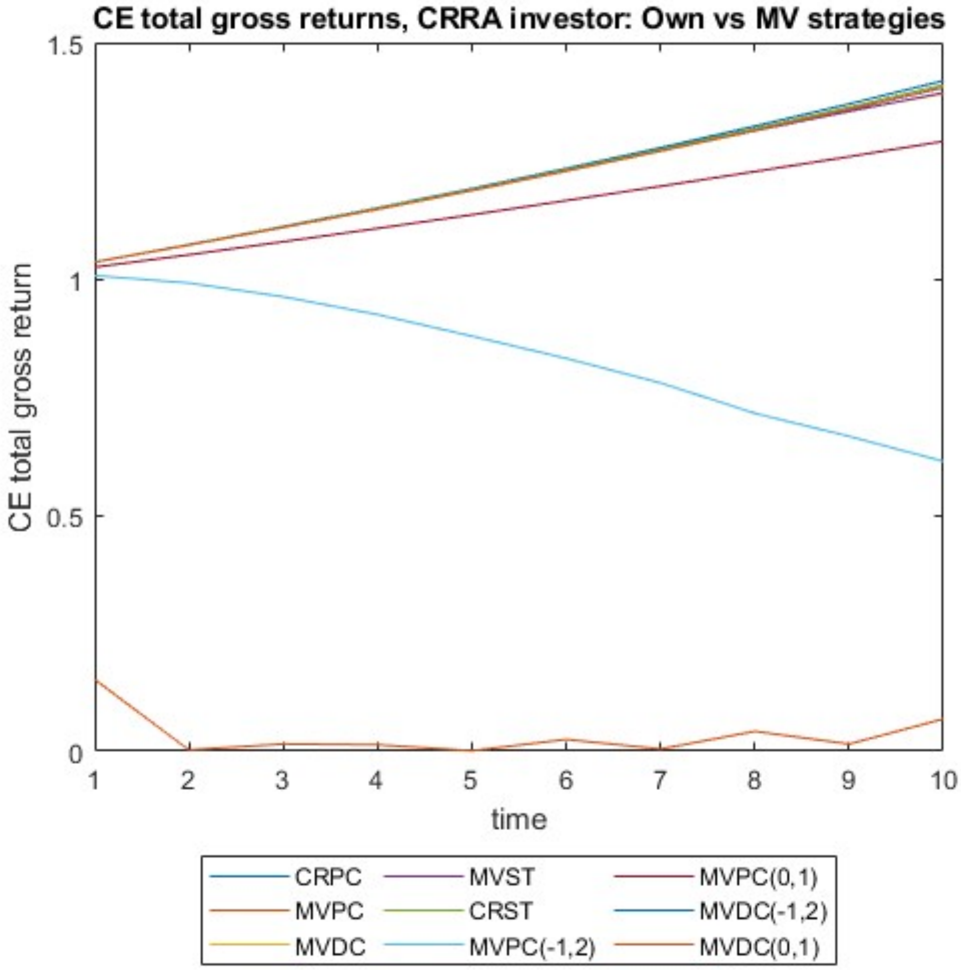

Figure 15: CE total gross returns for the CRRA investor from CRRA pre-commitment strategies and in comparison to various (constrained) Mean-Variance strategies and the CRRA investor's own static strategy. The shortening of the investment horizon does not increase the relative appeal of the Mean-Variance pre-commitment strategy to the CRRA investor. Parameters: T=30, m=1000000,n=360, yearly performance measurement.

As for the CARA investor, Figure 16 shows that her Certainty-Equivalent gross returns from the Mean-Variance pre-commitment strategy, as based on Monte-Carlo simulations of future wealth levels, wanders off far into negative territory. As shown in appendix A14, however, the expected utility for the CARA investor is, in this case, not defined. While this prevents further analysis in this case, it is in itself an interesting result.

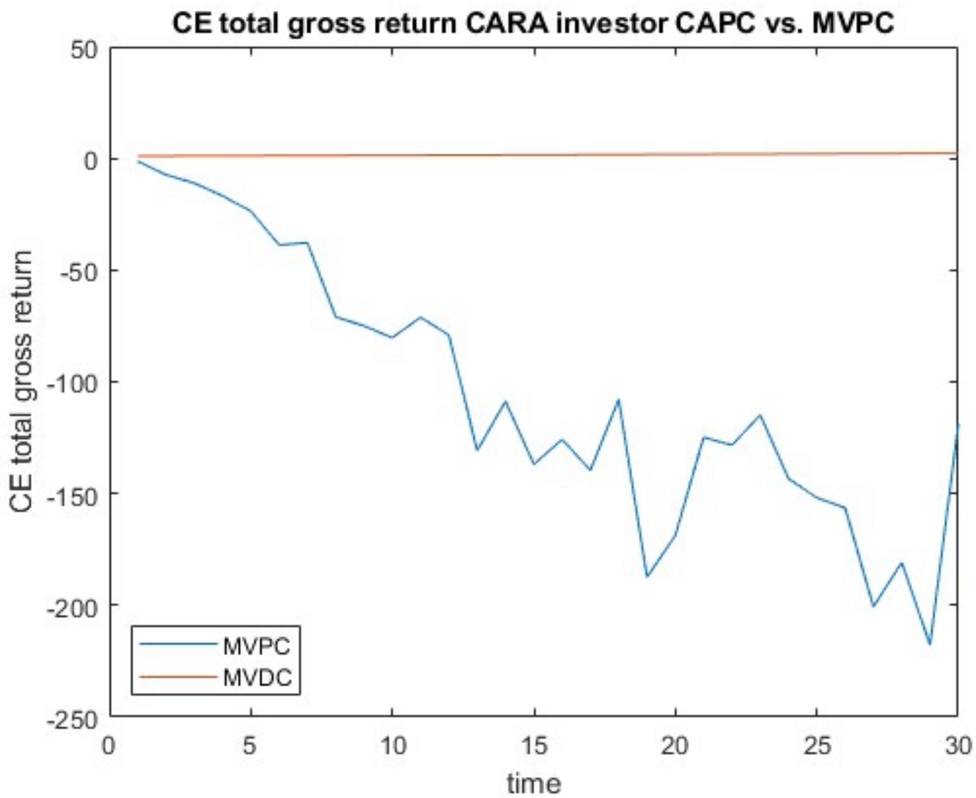

Figure 16: CE total gross returns for CARA investor from Mean-Variance pre-commitment and dynamically consistent strategies. Parameters: T=30, m=1000000,n=360, yearly performance measurement.

Next, similar to Balter & Schweizer (2021), we aim to investigate the relative effect of over-investment and different investment scheduling on the (dis-)utility from the Mean-Variance pre-commitment strategy relative to CRRA preferences. This is complicated by the fact that investment weights of the pre-committed Mean-Variance investor are stochastic and by the degree of volatility in these weights. To nevertheless obtain a rough indication of the relative effects, we choose the median Mean-Variance pre-commitment investment fractions as a better indication of the 'average' investment fraction chosen. The CRRA investor invests slightly larger average fractions in the risky asset for investment horizons ranging from one to ten years than the Mean-Variance investor under the respective pre-commitment strategies, with the average difference reaching 5.2% at a horizon of T=10 years.[52] Thus, we construct a 'raised Mean-Variance pre-commitment strategy' (MVPC raised), which shows the same investment schedule as the original Mean-Variance pre-commitment strategy, but has the same average investment fraction as the CRRA investor's strategy. Likewise, we construct a 'reduced CRRA pre-commitment strategy', which has the same investment schedule as the CRRA investor's optimal pre-commitment strategy, but the same average investment fraction as the Mean-Variance pre-commitment strategy. Figure 17 shows the performance in terms of certainty equivalent gross returns of the above strategies for maturities ranging from T=1 to T=10. The

[52]Figure 3 provides further illustration that the difference in average fractions is indeed quite small.

graph suggests that the effect of over-investment per se is not pivotal in explaining the CRRA investor's disutility from the Mean-Variance pre-commitment strategy. The reduced CRRA pre-commitment strategy performs almost as well as the investors original pre-commitment strategy. On the other hand, the raised Mean-Variance pre-commitment strategy performs remarkably well under CRRA preferences as well, so that the schedule under the latter strategy does not by itself cause the large loss in value witnessed before. Thus, it is the combination of over-investment and differences in investment-scheduling related to the Mean-Variance pre-commitment strategy, which causes the loss in value to the CRRA investor. From the discussion above, it would follow that the disutility for the CRRA investor is due to the large downside risks, which the Mean-Variance pre-commitment strategy entails. These risks are in particular due to the considerable overweight in the risky asset at the beginning of the investment horizon pursuant to this strategy, which, in turn, causes both the positive difference in 'average' investment weights and investment scheduling.

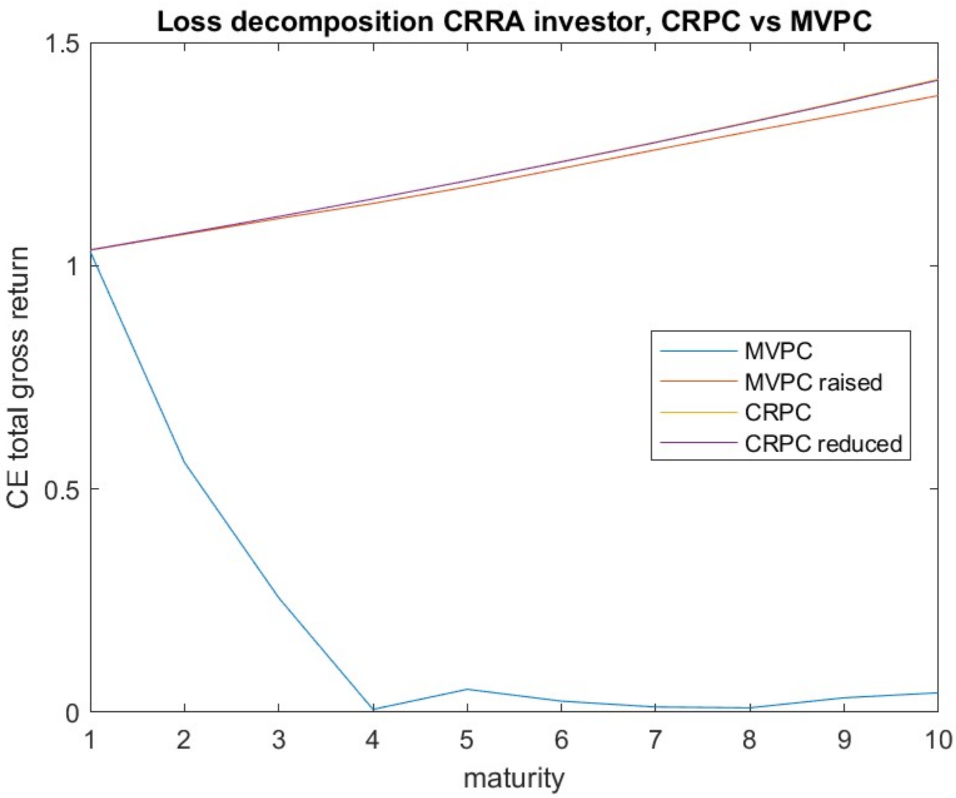

Figure 17: Comparison of CE growth rates for CRRA pre-commitment (PC) and Mean-Variance PC. For heuristic loss decomposition, CRPC reduced denotes a strategy that has the same investment pattern as the CRRA investor's pre-commitment strategy, but whose mean investment fraction is equal to the median investment fraction of the Mean-Variance PC strategy. Vice versa, the raised Mean-Variance investor's pre-commitment strategy (MVPC raised) has the same pattern as the Mean-Variance PC strategy, but the same average investment level as the CRRA investor's PC strategy. At the end of the investment horizon, the average CRPC strategy and median Mean-Variance PC investment level are almost the same. Parameters: T=1:10, m=1000000,n=120, yearly performance measurement.

4.6 Certainty equivalent Returns for the Mean-Variance investor

Setting again the investment horizon at $T = 30$ years and measuring performance in terms of certainty equivalent returns at intermediate time-points, we see in Figure 18 that it is only after more than 20 years that the Mean-Variance investor's pre-commitment strategy outperforms her dynamically consistent (and other) strategies. This points at large problems with respect to commitment, as up to this point, yearly performance measurement would suggest to the Mean-Variance investor to abandon her pre-commitment strategy in favor of her consistent-planning - or even her static investment strategy. Nevertheless, at the terminal time point, the pre-commitment strategy leads to almost twice the certainty equivalent total gross return as compared to the consistent-planning strategy. That is to say, while commitment may be problematic along the way, there are vast benefits to commitment for the Mean-Variance investor. The discussed portfolio constraints negate these benefits, however, with the dynamically consistent strategy and static strategies outperforming the constrained pre-commitment strategies at each point in time over the investment horizon of 30 years.

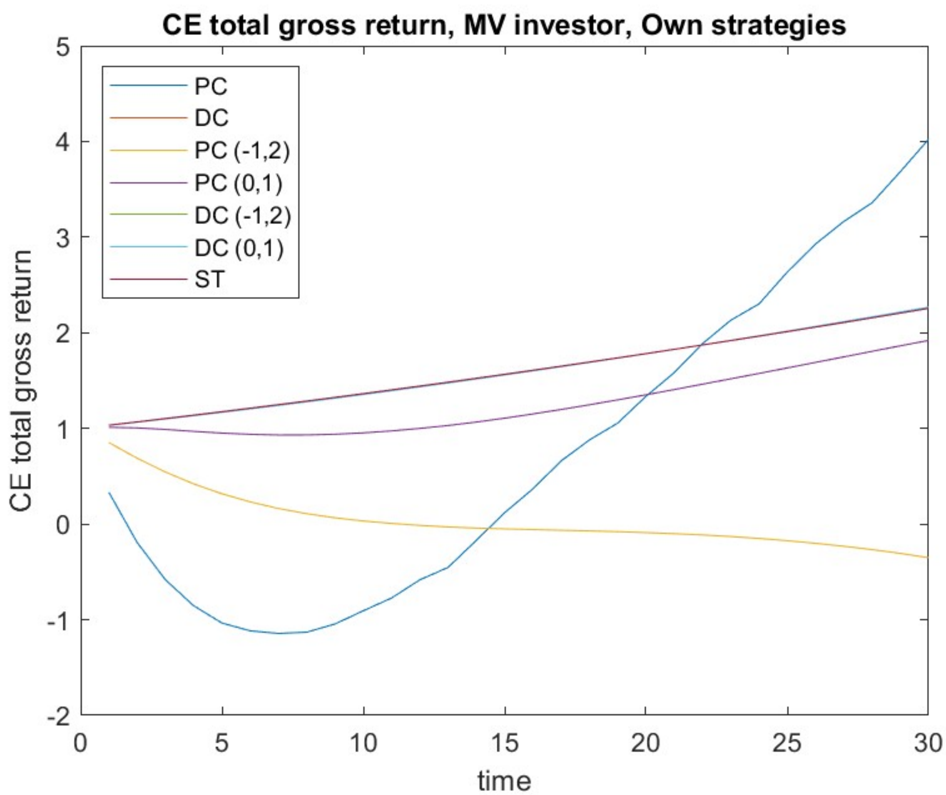

Figure 18: Certainty equivalent gross returns for Mean-Variance objective function and from Mean-Variance pre-commitment (PC, portfolio constrains in brackets), consistent-planning (DC, portfolio constraints in brackets) and static strategies. Parameters: T=30, m=1000000,n=360, yearly performance measurement.

To better link these results to the simulation results in section 4.3, we re-compute the certainty equivalent total gross returns from Figure 18 for an investment horizon of $T = 10$ years. As shown in Figure 19, with a shorter overall investment horizon, the constrained pre-commitment strategies still outperform the dynamically consistent strategy at the end of the investment horizon, confirming simulation results in section 4.3 . Thus, in the eyes of the Mean-Variance investor, portfolio constraints are particularly costly in the longer term. Figure 19 also shows, however, that commitment remains both problematic at intermediate time points and beneficial at the end of the investment horizon, even though the incentive to deviate in terms of the difference of certainty equivalents total gross returns at intermediate time points is smaller if the investment horizon is shorter[53].

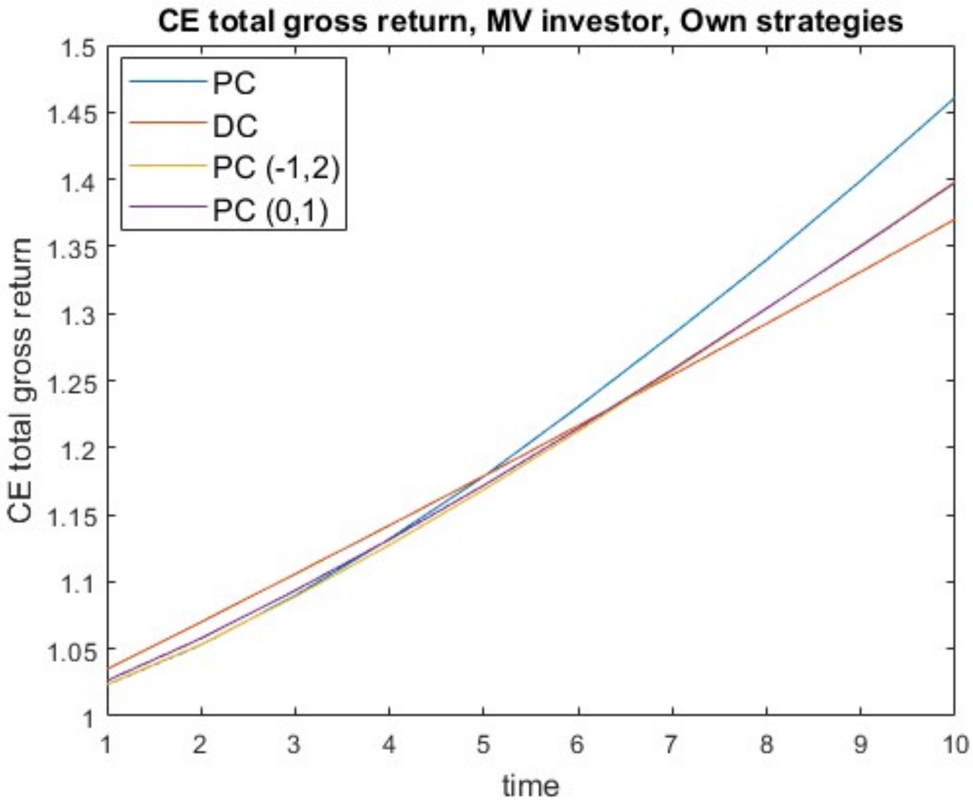

Figure 19: Re-assessment of the value of commitment in Figure 18 for the Mean-Variance investor for shorter time horizon. Parameters: T=10, m=1000000,n=120, yearly performance measurement.

With the exception of the myopic strategy (see Figure 24) in appendix A13, all of the Mean-Variance investor's strategies lead to positive certainty equivalents. Thus, Figure 20 shows the (annual) certainty equivalent growth rates that are derived from the certainty equivalent gross returns in Figure 19. Thus, we see more clearly that it takes 5 years for the Mean-Variance pre-commitment strategy to outperform the static or dynamically consistent strategies, respectively, whereas outperformance at the end of the investment horizon is roughly half a percentage point of certainty equivalent growth rate. It is, thus, questionable whether the Mean-Variance investor could be tempted to accept the behavioral constraints at intermediate time-points by the promised 'excess return' at the end of the investment horizon.
It follows that significant outperformance of the Mean-Variance pre-commitment strategy requires a longer-investment horizon, whereas the incentive to deviate likewise increases as the investment horizon is extended.

[53]Figure 26 in Appendix A15 includes also the constrained dynamically consistent Mean-Variance strategies, which have been omitted in Figure 19 to avoid cluttering.

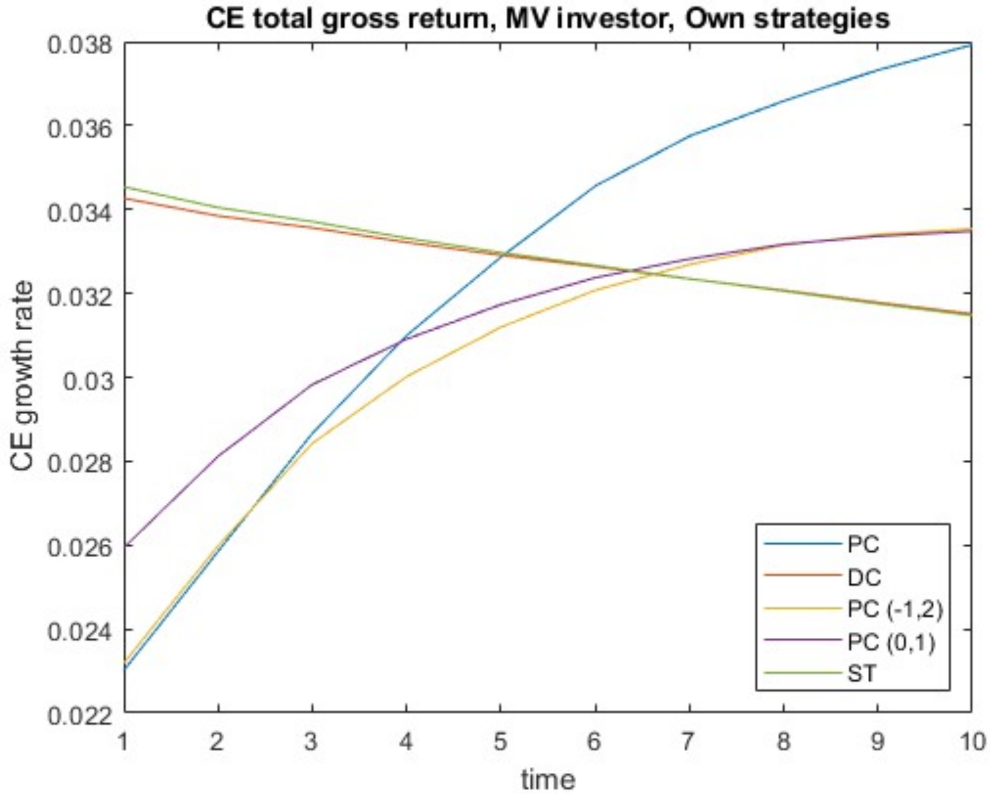

Figure 20: Comparison of CE growth rates for the Mean-Variance investor from (constrained) Mean-Variance pre-commitment (PC), consistent-planning (DC) and static Mean-Variance strategies.

Finally, we aim to decompose the relative losses in terms of certainty equivalent growth rates for the Mean-Variance investor from her dynamically consistent strategy relative to her pre-commitment strategy, similar to the discussion concerning the CRRA investor in section 4.5 . The pre-committed Mean-Variance investor incurs considerable losses in value if she is forced to reduce the average investment amount, while retaining the investment schedule of her pre-commitment strategy. At the same time, her losses from changing the investment pattern to the schedule under the dynamically consistent strategy, while retaining the average invested amount at the current level under the pre-commitment strategy, her losses incurred are significantly smaller. This suggests that the invested amount is crucial to the pre-committed Mean-Variance investor. The average difference in invested fractions is also slightly larger as compared to the case of the CRRA investor above, reaching 7.2% at a ten-year investment horizon. At the same time, the investor also gains significantly in terms of certainty equivalent growth rates if she changes her investment schedule from the reduced pre-commitment strategy to the DC strategy, while keeping the average risky investment at the average level under the former strategy, which, in turn, suggests that re-scheduling alone has a large effect on the certainty equivalent growth rate as well. Thus, starting off at the pre-commitment solution, the effect of a re-scheduling is smaller than the effect of reduction of the average investment fraction, which stands in contrast to results in Figure 17 concerning the CRRA investor. However, in turn, optimal invested amounts remain strongly linked to optimal investment schedules under the respective optimization problem.

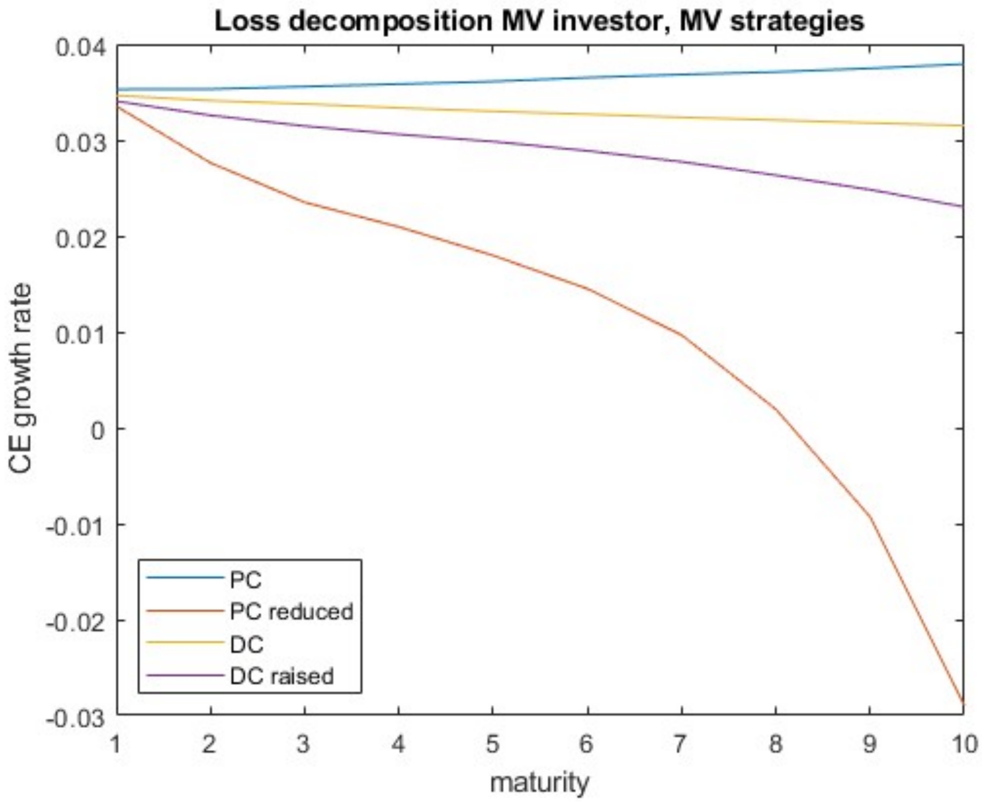

Figure 21: Comparison of CE growth rates for Mean-Variance pre-commitment PC and consistent planning DC strategies. For heuristic loss decomposition PC-DC, DC raised denotes a strategy that has the same investment pattern as DC, but whose mean investment fraction is equal to the median investment fraction of the PC strategy. Parameters: T=1:10, m=1000000,n=120, yearly performance measurement.

5 Conclusion

The quality of Mean-Variance approximations in continuous time and the single period or (myopic) multi-period discrete time case differs considerably. While the literature has found that optimal Mean-Variance investment in the latter case tends to approximate optimal policies under expected utility criteria quite closely, the optimal continuous time Mean-Variance policy, as given by the pre-commitment solution, performs very poorly in approximating the optimal pre-commitment policies under a CRRA expected utility criterion, even in a setting of Gaussian returns. The results presented suggest that CRRA investors could do better by following a static investment strategy under their own criterion rather than having their financial wealth invested according to the Mean-Variance pre-commitment policy. The discussion traces these results back to the mechanics of the Martingale approach as the underlying solution method for the computation of continuous time pre-commitment policies, and the Mean-Variance investor's failure to satisfy criteria of Second-Order stochastic dominance under this approach. With respect to the latter, it is argued that dynamically consistent investor, much like static investors, maximize expected utility or value in a forward-looking manner, whereas optimal investment under the Martingale approach is essentially backward looking from an optimal terminal wealth profile, which has been computed state-by-state. As a result, the safeguards for Second-Order stochastic dominance that are in place in the single-period case when maximizing a mean-variance criterion under Gaussian returns do not apply to the the computation of optimal pre-commitment solutions under the Martingale approach. The Mean-Variance pre-commitment solution differs significantly from all strategies of the CARA and CRRA investors, as well as her own static and dynamically consistent strategies, respectively, in terms of amounts invested and investment schedule. Moreover, it yields a stochastic strategy, whereas the former are deterministic and at most time-varying. According to the results presented, the CRRA investor would essentially be willing to spend almost her entire wealth to insure against having to hold the risk represented by Mean-Variance pre-commitment strategies. Furthermore, the analysis suggests that this is due to the considerable downside risk that the pre-committed Mean-Variance investor is willing to take, which causes large losses in terms of certainty equivalents for the CRRA investor. This is consistent with the pre-committed Mean-Variance investor's failure with respect to Second-Order stochastic dominance. Expected utility and respective certainty equivalents of the Mean-Variance pre-commitment strategy under CARA utility could not be computed, as the underlying mean was found not to exist. Furthermore, the costs of time-inconsistency of the Mean-Variance criterion to the Mean-Variance investor in continuous time was assessed in terms of certainty equivalent returns. While costly for the Mean-Variance investor, the criterion's time-inconsistency turns out to be value enhancing with respect to Mean-Variance approximations of optimal policies under the CRRA criterion in indirect ways. This is due to the greater similarity of dynamically consistent Mean-Variance strategies, to which the Mean-Variance investor may have to revert if commitment becomes impossible, bear much greater similarity to the strategies chosen by CRRA expected utility maximizers. In particular, the dynamically consistent Mean-Variance strategy leads to much smaller downside risk as compared to the Mean-Variance pre-commitment solution. Given the at first unexpectedly bad performance of Mean-Variance pre-commitment strategies in approximating CRRA expected utility strategies in settings of Gaussian returns, interesting further research would include the study of continuous-time Mean-Variance approximations in alternative financial market settings. As the present research suggests, however, considerable computing power may be required especially if longer investment horizons are to be investigated. Moreover, additional sources of time-inconsistency[54] could be studied as to their effect on the relative performance of dynamically consistent and optimal pre-commitment Mean-Variance strategies in approximating optimal policies under expected utility criteria. Finally, the potential failure of the Mean-Variance investor to satisfy criteria of Second-Order stochastic dominance under a Martingale approach should be more thoroughly studied, especially as to the juxtaposition with her apparent adherence to these economic efficiency criteria under the dynamic programming approach.

[54] See e.g. Björk et al., 2014 and Balter et al., 2021.

A Appendix

A.1 Optimal pre-commitment strategy, CRRA investor, amounts invested

Following[55] derivations in Munk (2017) and Björk (2020) and specializing to the case at hand, the Lagrangian for the constrained static optimization problem of maximizing (25) subject to (26) is given by

$$\mathcal{L} = \mathbb{E}\left[\frac{X_T^{1-\gamma}}{1-\gamma}\right] + \psi\left(X_0 - \mathbb{E}\left[M_T X_T\right]\right), \tag{43}$$

with $\gamma > 1$. Optimization is performed by state of the world, that is determining for each state of the world $\omega \in \Omega$ at time T the optimal value $X_T(\omega)$. The first-order conditions are, respectively,

$$X_T^{-\gamma} - \psi M_T = 0 \iff X_T = \psi^{-\frac{1}{\gamma}} M_T^{-\frac{1}{\gamma}}, \tag{44}$$

and

$$X_0 - \mathbb{E}[M_T X_T] = 0 \iff X_0 = \mathbb{E}[M_T X_T]. \tag{45}$$

Plugging (44) into (45) and solving for the Lagrangian multiplier ψ yields

$$X_0 = \mathbb{E}\left[M_T \psi^{-\frac{1}{\gamma}} M_T^{-\frac{1}{\gamma}}\right] \iff \psi = X_0^{-\gamma}\left(\mathbb{E}\left[M_T^{1-\frac{1}{\gamma}}\right]\right)^{\gamma}. \tag{46}$$

Plugging (46) back into (44) yields

$$X_T^* = \left(X_0^{-\gamma}\left(\mathbb{E}\left[M_T^{1-\frac{1}{\gamma}}\right]\right)^{\gamma}\right)^{-\frac{1}{\gamma}} M_T^{-\frac{1}{\gamma}} = \frac{X_0 M_T^{-\frac{1}{\gamma}}}{\mathbb{E}\left[M_T^{1-\frac{1}{\gamma}}\right]}. \tag{47}$$

The market consistent value of future time T - optimal wealth at any time $t \in [0, T]$ is given by the process X_t^* as follows, whereas M_t is the state-price deflator process as given in (23):

$$X_t^* = \frac{1}{M_t}\mathbb{E}_t\left[M_T X_T^*\right] = \mathbb{E}_t\left[\frac{M_T}{M_t}\frac{X_0 M_T^{-\frac{1}{\gamma}}}{\mathbb{E}\left[M_T^{1-\frac{1}{\gamma}}\right]}\right] = \frac{M_t^{-\frac{1}{\gamma}} X_0}{\mathbb{E}\left[M_T^{1-\frac{1}{\gamma}}\right]}\mathbb{E}_t\left[\left(\frac{M_T}{M_t}\right)^{1-\frac{1}{\gamma}}\right]. \tag{48}$$

[55] The derivations included in this appendix and many of the appendices below build on the existing knowledge as to the solution of the above problems. This includes in particular textbook treatments of these models as included in Björk (2020), Cvitanic & Zapatero (2004) and Munk (2017) and Shreve (2010).

Thus, the process X_t^* is given by a time-dependent function of the stochastic discount factor M_t, $X_t^* = f_t(M_t)$. Hence, by Ito's lemma the differential form of the process may be written as

$$dX_t^* = ...dt + \frac{\partial}{\partial M_t} X_t^* dM_t = ...dt + \frac{1}{\gamma} X_t^* \lambda dW_t. \tag{49}$$

The pay-off $X_T^*(\omega)$ is adapted to the filtration $\mathcal{F}(T)$ generated by the Brownian Motion W_t at time T and is, thus, also spanned by the basic financial assets in our market. Moreover, the Black-Scholes financial market admits the risk-neutral probability measure $\mathbb{Q}$ defined by the given probability measure $\mathbb{P}$ together with the Radon-Nikodym derivative

$$\xi_T = \exp\left(-\frac{1}{2}\lambda^2 T - \lambda W_T\right), \tag{50}$$

with the respective Radon-Nikodym process, as given by (24) (see e.g. Shreve, 2010).

Measure $\mathbb{Q}$ is equivalent to measure $\mathbb{P}$ and under this measure, the discounted stock price is a martingale such that $S(t) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r(T-t)}S(T)\right]$. Moreover, given the particular market setting, the price of risk λ is determined uniquely by $\lambda = \frac{(\mu - r)}{\sigma}$ and, hence, so are the Radon-Nikodym derivative and the Radon-Nikodym process, respectively. As a result, we may use the First - and Second Fundamental Theorem of Asset Pricing to determine the unique price of any pay-off to be received at time T and given by a random variable $\tilde{X}_T$ adapted to $\mathcal{F}_T$ in either of two equivalent ways:

$$\tilde{X}_t = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r(T-t)}\tilde{X}_T\right] = \mathbb{E}_t\left[\frac{d\mathbb{Q}}{d\mathbb{P}} e^{-r(T-t)}\tilde{X}_T\right] = \mathbb{E}_t\left[e^{-r(T-t)}\frac{\xi_T}{\xi_t}\tilde{X}_T\right] \tag{51}$$

and

$$\tilde{X}_t = \mathbb{E}_t\left[\frac{M_T}{M_t}\tilde{X}_T\right] = \mathbb{E}_t\left[e^{-r(T-t)-\frac{1}{2}\lambda^2(T-t)-\lambda(W_T - W_t)}\tilde{X}_T\right] = \mathbb{E}_t\left[e^{-r(T-t)}\frac{\xi_T}{\xi_t}\tilde{X}_T\right], \tag{52}$$

respectively, whereas M_t is again the state-price deflator process as given in (23). (See e.g. Shreve, 2010) Hence, with the value of the optimally chosen pay-off X_T^* at time t, X_t^*, being determined by (48), it follows from (51) and (52) that under measure $\mathbb{Q}$ the process X_t^* is expected to grow exponentially at constant rate r. Finally, on the basis of Girsanov's theorem, we note that under the particular measure $\mathbb{Q}$ as defined above, the process

$$\tilde{W}_t = W_t + \lambda t \tag{53}$$

is a Brownian Motion. (see e.g. Shreve, 2010) Thus, from (49) together with (51) - (53) it follows that

$$dX_t^* = ...dt + \frac{1}{\gamma} X_t^* \lambda dW_t = X_t^* r dt + \frac{1}{\gamma} X_t^* \lambda d\tilde{W}_t. \tag{54}$$

Moreover, the differential form of the wealth process dX_t resulting from the replicating portfolio, as chosen by the investor, given in (5) can be re-written as

$$
\begin{aligned}
dX_t &= X_t r dt + X_t \pi_t (\mu - r) dt + X_t \pi_t \sigma dW_t \\
&= X_t r dt + X_t \pi_t \sigma \left(\frac{(\mu - r)}{\sigma} dt + dW_t \right) \\
&= X_t r dt + X_t \pi_t \sigma d\tilde{W}_t .
\end{aligned}
\tag{55}
$$

We could go one step further from here, discount both processes in (54) and (55) with the risk-free rate and apply the Martingale Representation Theorem to determine the investment fraction π_t such that the wealth process (55) matches process (54) in all states of the world, subsequently using the equivalence of measures $\mathbb{P}$ and $\mathbb{Q}$ such that if the processes match in all states of the world under measure $\mathbb{Q}$ then they also match in all states of the world under measure $\mathbb{P}$. (Shreve, 2010) However, we may also simply note that two processes, which are driven by the same Brownian Motion, which agree at an initial point in time and which have the same drift and volatility term, respectively, are equal to each other at each time after the initial point in time. (Baxter & Rennie, 2012) Due to the budget constraint in optimization problem (43), the processes described by (54) and (55) are indeed required to have the same value, $X_0 = \bar{X}_0$, at the initial point in time. Moreover, they share the same drift, so that we may determine the perfect hedge π_t by matching volatility terms in (54) and (55). Thus,

$$
\pi_t^* = \pi^* = \frac{\psi}{\gamma\sigma} = \frac{(\mu - r)}{\gamma\sigma^2}. \tag{56}
$$

Again, by equivalence of measures $\mathbb{P}$ and $\mathbb{Q}$, the hedge works perfectly under the physical probability measure as well, which concludes the argument. In equating volatility terms in (54) and (55) to obtain (56), we replaced X_t^* by X_t, as for each incremental step forward at any time $t \in [0, T]$, $X_t^* = X_t$ is ensured by the fact that $X_0 = X_0^* = \bar{X}_0$ and all increments leading up to time t are matched. We may also note with reference to (54) and (55) that, as is commonly known, a change in measure is a change in the drift, leaving volatility terms unaffected. (see e.g. Shreve, 2010) Thus, in future instances, in which the Martingale Method is applied, we will skip the underlying argument as regards changes in measure and rest assured that within the given market setting we can transform the given price - and value processes such that they have the same deterministic drift, before matching volatility terms.

A.2 Optimal pre-commitment strategy, CRRA investor, relative weights

The derivation of the value X_t^* of the optimal pay-off X_T^* in equations (43) to (49) is agnostic as to whether investment strategies are specified in terms of relative fractions or amounts invested. If we now propose to use absolute amounts invested to specify the exposure to the risky asset, the value process of the replicating portfolio is given by $dX_t = X_t r dt + \theta_t(\mu - r)dt + \theta\sigma dW_t$. Again denoting $\tilde{W}_t$ a Brownian Motion under measure $\mathbb{Q}$ it follows from the same reasoning as in Appendix A1 that

$$
dX_t = X_t r dt + \theta_t \sigma d\tilde{W}_t = X_t^* r dt + \frac{1}{\gamma} X_t^* \lambda d\tilde{W}_t = dX_t^* \tag{57}
$$

and

$$X_0 = X_0^*, \tag{58}$$

so that

$$\theta_t^* = X_t \frac{\lambda}{\gamma\sigma} = X_t \frac{(\mu - r)}{\gamma\sigma^2}. \tag{59}$$

A.3 Optimal pre-commitment strategy, CARA investor, relative weights

The Lagrangian for the static constrained optimization problem faced by the CARA investor is given by

$$\mathcal{L} = \mathbb{E}\left[-\frac{1}{\alpha}e^{-\alpha X_T}\right] + \psi\left(X_0 - \mathbb{E}[M_T X_T]\right), \tag{60}$$

whereas maximizing again by state of the world at time T the first-order conditions are, respectively,

$$e^{-\alpha X_T} - \psi M_T = 0 \iff e^{-\alpha X_T} = \psi M_T, \tag{61}$$

and

$$X_0 - \mathbb{E}[M_T X_T] = 0 \iff X_o = \mathbb{E}[M_T X_T]. \tag{62}$$

From (61), we have that

$$X_T = \frac{-\log(\psi M_T)}{\alpha}, \tag{63}$$

whereas plugging (63) into (62) and solving for ψ yields

$$\begin{aligned} \alpha X_0 =& \ \mathbb{E}\left[M_T(-\log(\psi M_T))\right] \\ =& \ -\mathbb{E}\left[M_T \log(M_T)\right] - \mathbb{E}\left[M_T \log(\psi)\right] \\ \iff \quad -\mathbb{E}[M_T]\log(\psi) =& \ \alpha X_0 + \mathbb{E}[M_T \log(M_T)] \\ \iff \quad \psi =& \ \exp\left[\frac{\alpha X_0 + \mathbb{E}[M_T \log(M_T)]}{\mathbb{E}[-M_T]}\right], \end{aligned} \tag{64}$$

whereas the third equality follows from the fact that the marginal value of a relaxation of the budget constraint at time $t = 0$ is known at that time and, hence, its logarithm can be taken out of the expectation. Given initial wealth X_0, we now set $\psi = \bar{\psi}$. Therefore, the optimal terminal wealth is given by

$$X_T^* = \frac{-\log(\bar{\psi} M_T)}{\alpha}. \tag{65}$$

From this we derive the value process X_t^* for all $t \in [0, T]$, and its respective differential form as follows:

$$\begin{aligned} X_t^* &= -\frac{1}{M_t} \frac{\mathbb{E}_t[M_T \log(\bar{\psi} M_T)]}{\alpha} \\ &= -\frac{1}{\alpha M_t} \mathbb{E}_t[M_T \log(\bar{\psi})] - \frac{1}{\alpha M_t} \mathbb{E}_t[M_T \log(M_T)] \\ &= -\frac{\log(\bar{\psi})}{\alpha M_t} M_t e^{-r(T-t)} - \frac{1}{\alpha M_t} \mathbb{E}_t[M_T \log(M_T)], \end{aligned} \tag{66}$$

whereas the first term in the final expression is time-dependent, but non-stochastic after cancellation of M_t in the numerator and denominator. The expectation in the second term is computed using the Law of the Unconscious Statistician and by integration with respect to the standard normal density. Given the stochastic discount factor in (23), we may split the expectation in the second term into two parts:

$$\begin{aligned} \mathbb{E}_t[M_T \log(M_T)] &= M_t \mathbb{E}_t \left[e^{(-r(T-t) - \frac{1}{2}\lambda^2 (T-t) - \lambda(W_T - W_t))} \right] \left(\log(M_t) - r(T-t) - \frac{1}{2}\lambda^2 (T-t) \right) \\ &\quad - M_t \mathbb{E}_t \left[e^{(-r(T-t) - \frac{1}{2}\lambda^2 (T-t) - \lambda(W_T - W_t))} \lambda (W_T - W_t) \right] \\ &= M_t e^{-r(T-t)} \left(\log(M_t) - r(T-t) - \frac{1}{2}\lambda^2 (T-t) \right) \\ &\quad - M_t e^{-r(T-t) - \frac{1}{2}\lambda^2 (T-t)} \int_{-\infty}^{\infty} e^{-\lambda \sqrt{T-t} z} \lambda \sqrt{T-t} z \frac{1}{\sqrt{2\pi}} e^{-\frac{-z^2}{2}} dz, \end{aligned} \tag{67}$$

with $z \sim \mathcal{N}(0, 1)$. Completing the square to compute the integral in the second term, we obtain

$$\begin{aligned} \int_{-\infty}^{\infty} e^{-\lambda \sqrt{T-t} z} \lambda \sqrt{T-t} z \frac{1}{\sqrt{2\pi}} e^{-\frac{-z^2}{2}} dz &= e^{\frac{1}{2}\lambda^2 (T-t)} \lambda \sqrt{T-t} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{(z + \lambda\sqrt{T-t})^2}{2}} z \, dz \\ &= e^{\frac{1}{2}\lambda^2 (T-t)} \lambda \sqrt{T-t} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{\tilde{z}^2}{2}} (\tilde{z} - \lambda \sqrt{T-t}) \, d\tilde{z} \\ &= e^{\frac{1}{2}\lambda^2 (T-t)} \lambda \sqrt{T-t} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{\tilde{z}^2}{2}} \tilde{z} d\tilde{z} \\ &\quad - e^{\frac{1}{2}\lambda^2 (T-t)} \lambda^2 (T-t) \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{\tilde{z}^2}{2}} \, d\tilde{z} \\ &= -e^{\frac{1}{2}\lambda^2 (T-t)} \lambda^2 (T-t), \end{aligned} \tag{68}$$

whereas the second equality follows from the change of variable $\tilde{z} = z + \lambda\sqrt{T-t} \iff z = \tilde{z} - \lambda\sqrt{T-t}$ and the final equality from taking the two integrals, which yields the expectation of a standard normal variable, that is equal to zero, and the integral over the range of the standard normal density, which is equal to one, respectively. Plugging the result in (68) back into (67) yields

$$\begin{aligned}\mathbb{E}_t[M_T \log(M_T)] &= M_t e^{-r(T-t)}\left(\log(M_t) - r(T-t) - \frac{1}{2}\lambda^2(T-t)\right) + M_t e^{-r(T-t)}\lambda^2(T-t) \\ &= M_t e^{-r(T-t)}\left(\log(M_t) - r(T-t) + \frac{1}{2}\lambda^2(T-t)\right).\end{aligned} \tag{69}$$

Plugging (69) back into (66), we obtain

$$\begin{aligned}X_t^* &= -\frac{\log(\bar{\psi})}{\alpha}e^{-r(T-t)} - \frac{1}{M_t}\frac{\mathbb{E}_t[M_T \log(\bar{\psi}M_T)]}{\alpha} = \\ &= -\frac{\log(\bar{\psi})}{\alpha}e^{-r(T-t)} - \frac{1}{\alpha}e^{-r(T-t)}\left(\log(M_t) - r(T-t) + \frac{1}{2}\lambda^2(T-t)\right).\end{aligned} \tag{70}$$

Thus, the process X_t^* is a time-dependent function of the stochastic discount factor M_t, $X_t = f_t(M_t)$, so that by Ito's lemma, its differential form can be written as

$$\begin{aligned}dX_t^* &= ...dt - \frac{1}{\alpha}e^{-r(T-t)}\frac{1}{M_t}dM_t \\ &= ...dt + \frac{1}{\alpha}e^{-r(T-t)}\left(rdt + \lambda dW_t\right) = ...dt + \frac{1}{\alpha}e^{-r(T-t)}\lambda dW_t\end{aligned} \tag{71}$$

Again, as in Appendix A1, we find the replicating portfolio by re-writing the value process (71) and the wealth process in (5) using process $\tilde{W}_t$ as defined in (53), which is a Brownian Motion under the risk-neutral measure $\mathbb{Q}$ defined in Appendix A1, and equating volatility terms, that is

$$dX_t^* = X_t^* rdt + \frac{1}{\alpha}e^{-r(T-t)}\lambda d\tilde{W}_t = X_t rdt + X_t \pi_t \sigma_t d\tilde{W}_t = dX_t,$$

with

$$X_0^* = X_0,$$

so that

$$\pi_t^* = \frac{\lambda}{\alpha\sigma_t X_t}e^{-r(T-t)} = \frac{(\mu - r)}{\alpha\sigma_t^2 X_t}e^{-r(T-t)}.$$

A.4 Optimal pre-commitment strategy, CARA investor, amounts invested

As in the discussion in appendices A1 and A2, and given the results and discussion in appendix A3, we now replicate wealth process in (71) using the wealth process in (6) and the respective initial condition:

$$dX_t^* = X_t^* r dt + \frac{1}{\gamma} e^{-r(T-t)} \lambda d\tilde{W}_t = X_t r dt + \theta_t \sigma_t d\tilde{W}_t = dX_t,$$

with

$$X_0^* = X_0,$$

so that

$$\theta_t^* = \frac{\lambda}{\alpha \sigma_t} e^{-r(T-t)} = \frac{(\mu - r)}{\alpha \sigma_t^2} e^{-r(T-t)}.$$

A.5 Optimal pre-commitment strategy, Mean-Variance investor, amounts invested

As suggested by Basak & Chabakauri (2012), the optimal pre-commitment portfolio strategy for the Mean-Variance investor is computed by the following constrained optimization problem

$$\mathcal{L} = \mathbb{E}[X_T] - \frac{\delta}{2} \mathrm{Var}[X_T] + \psi \left(X_0 - \mathbb{E}[M_T X_T] \right). \tag{72}$$

The respective first-order conditions, taken again by state of the world at time T, are given by

$$1 - \delta X_T + \delta \mathbb{E}[X_T] - \psi M_T = 0 \tag{73}$$

and

$$X_0 - \mathbb{E}[M_T X_T] = 0, \tag{74}$$

respectively. To compute (73), we note that

$$\mathrm{Var}[X_T] = \mathbb{E}[(X_T - \mathbb{E}[X_T])^2] = \int_{\omega \epsilon \Omega} (X_T(\omega) - \mathbb{E}[X_T])^2 d\mathbb{P}(\omega), \tag{75}$$

so that taking the first-order derivative of the variance with respect to X_T for each state of the world separately effectively means that we condition on the state of the world $\omega \epsilon \Omega$ in (75) before taking the derivative with respect to X_T. Thus, the first order derivative for all states of the world $\omega = \bar{\omega} \epsilon \Omega$ is given by

$$\frac{d}{dX_T}\text{Var}[X_T]|_{\omega=\bar{\omega}\epsilon\Omega} = 2X_T - 2\mathbb{E}[X_T].$$

From (73), it follows that

$$X_T = \frac{1}{\delta}\left(1 + \delta\mathbb{E}[X_T] - \psi M_T\right). \tag{76}$$

Noting that both the Lagrangian multiplier and expectation of future optimal wealth are unknowns, we first take the expectation on both sides of (76) to solve for ψ to obtain

$$\delta\mathbb{E}[X_T] = 1 + \delta\mathbb{E}[X_T] - \psi e^{-rT} \iff \psi = e^{rT}. \tag{77}$$

Plugging back into (76), we obtain

$$X_T = \frac{1}{\delta}\left(1 + \delta\mathbb{E}[X_T] - e^{rT}M_T\right). \tag{78}$$

X_T must be chosen such that the budget constraint (74) is satisfied, which in this case implies a condition on the $\mathbb{E}[X_T]$. Plugging (78) into the budget constraint (74) and solving for $E[X_T]$ yields

$$\begin{aligned}
X_0 &= \mathbb{E}\left[M_T\frac{1}{\delta}\left(1 + \delta\mathbb{E}[X_T] - e^{rT}M_T\right)\right] \\
&= \mathbb{E}[M_T]\frac{1}{\delta} + \mathbb{E}[M_T]\mathbb{E}[X_T] - \frac{1}{\delta}\mathbb{E}\left[M_T^2\right]e^{rT} \\
&= e^{-rT}\frac{1}{\delta} + e^{-rT}\mathbb{E}[X_T] - \frac{1}{\delta}e^{-2rT+\lambda^2T+rT} \\
\iff \mathbb{E}[X_T] &= X_0e^{rT} - \frac{1}{\delta} + \frac{1}{\delta}e^{\lambda^2T}.
\end{aligned} \tag{79}$$

Plugging (79) back into (78) yields

$$\begin{aligned}
X_T^* &= \frac{1}{\delta}\left(1 + \delta\left(X_0e^{rT} - \frac{1}{\delta} + \frac{1}{\delta}e^{\lambda^2T}\right) - e^{rT}M_T\right) \\
\iff X_T^* &= X_0e^{rT} + \frac{1}{\delta}e^{\lambda^2T} - \frac{1}{\delta}M_Te^{rT}.
\end{aligned} \tag{80}$$

As before, we next derive the market-consistent value of pay-off X_T^* at any time $t \in [0, T]$:

$$\begin{aligned}
X_t^* &= \frac{1}{M_t}\mathbb{E}_t[M_TX_T] = \frac{1}{M_t}\mathbb{E}_t\left[M_TX_0e^{rT} + \frac{1}{\delta}e^{\lambda^2T}M_T - M_T^2\frac{1}{\delta}e^{rT}\right] \\
&= \frac{1}{M_t}\left(M_te^{-r(T-t)}X_0e^{rT} + \frac{1}{\delta}e^{\lambda^2T}M_te^{-r(T-t)} - \frac{1}{\delta}M_t^2e^{-2r(T-t)+\lambda^2(T-t)+rT}\right) \\
&= e^{-r(T-t)}X_0e^{rT} + \frac{1}{\delta}e^{\lambda^2T}e^{-r(T-t)} - \frac{1}{\delta}M_te^{-2r(T-t)+\lambda^2(T-t)+rT},
\end{aligned} \tag{81}$$

which is a time-dependent function of M_t. As above, given the differential form for the stochastic discount factor, $dM_t = M_t(-rdt - \lambda dW_t)$, the differential of the process X_t^* is again computed via Ito's lemma and given by

$$\begin{aligned} dX_t^* = df_t(M_t) &= .....dt + \frac{\partial f_t(M_t)}{\partial M_t} M_t \lambda dW_t \\ &= ....dt + \frac{1}{\delta} e^{-2r(T-t)+\lambda^2(T-t)+rT} M_t \lambda dW_t. \end{aligned} \tag{82}$$

Denoting again $\tilde{W}_t$ as defined in (53) a Brownian Motion under measure $\mathbb{Q}$ as defined above, we may rewrite processes (82) and (6), respectively, and match them, as before, which yields

$$dX_t^* = X_t^* rdt + \frac{1}{\delta} e^{-2r(T-t)+\lambda^2(T-t)+rT} M_t \lambda d\tilde{W}_t = X_t rdt + \theta_t \sigma d\tilde{W}_t = dX_t, \tag{83}$$

with

$$X_0^* = X_0.$$

Equating volatility terms, we obtain the perfect hedge

$$\begin{aligned} \theta_t^* \sigma &= \frac{1}{\delta} 2 e^{-2r(T-t)+\lambda^2(T-t)+rT} M_t \lambda \\ \iff \theta_t^* &= \frac{1}{\delta} \frac{(\mu - r)}{\sigma^2} e^{-r(T-t)+\lambda^2(T-t)+rt} M_t, \end{aligned} \tag{84}$$

which coincides with the result for the pre-commitment solution as stated by Basak & Chabakauri (2012).

A.6 Optimal pre-commitment strategy, Mean-Variance investor, relative weights

Following the discussion in appendixes A1 and A2, as well as the results in appendix A5, the relative investment weights are found by replicating

$$dX_t^* = X_t^* rdt + \frac{1}{\delta} e^{-2r(T-t)+\lambda^2(T-t)+rT} M_t \lambda d\tilde{W}_t \tag{85}$$

by wealth process

$$dX_t = X_t rdt + X_t \pi_t \sigma d\tilde{W}_t, \tag{86}$$

with

$$X_0^* = X_0.$$

Equating volatility terms in (85) and (86), we obtain

$$\theta_t^* = \frac{1}{\delta}\frac{(\mu - r)}{\sigma^2}\frac{M_t}{X_t}e^{-r(T-t)+\lambda^2(T-t)+rt}.$$

A.7 Consistent-planning strategy, Mean-Variance investor

Following Björk's (2020, 2017) notation, specializing Basak & Chabakauri's (2012) results to our current setting and filling in the gaps in the derivation, let $\theta_s^{t,Y_t} = g(s, Y_s) : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}, s \in [t,T]$ be a control law, mapping for all time periods $s \in [t,T]$ the realization Y_s of the state variable into an investment strategy in the real numbers. The state variable remains an unspecified and generic Markovian state variable. The natural state variable within the Black-Scholes market setting would be given by wealth itself, however, the results in this section show that both the value function of the dynamic optimization problem and the optimal policy θ_t is independent of current wealth. Based on the model setting, however, it is clear that the stock price S_t may serve as a Markovian state variable, as the filtration generated by S_t, $\mathcal{F}_t^S$ is exactly the same as the filtration generated by wealth X_t, $\mathcal{F}_t^X$, so that the two are informationally equivalent.
The superscripts in θ_s^{t,Y_t} denote that the policy described by the control law is devised based on observation Y_t at time t, whereas through the codomain of the function g any portfolio constraints are assumed away. Moreover, where sub - and superscripts are suppressed, θ describes a variable entering into a function as an argument, without any assumption as to the optimality of the given investment choice. If only subscripts are suppressed, the $\theta^{Y_t,t}$ denotes the control law as devised at time $t \in [0,T]$ upon observation of state variable Y_t, providing the investment strategy for the entire remaining time horizon. As in previous sections, policy θ is assumed to denote the investment in the risky asset, the stock, in terms of of amounts invested. Trading is, moreover, assumed to be continuous in time and frictionless. Furthermore, let $J(t, S_t, X_t, \theta)$ be the agent's mean-variance value function [56] at time t from terminal wealth X_T^θ, given investment policy θ and conditional on the realization of the stock value S_t and wealth X_t at time t. Hence,

$$J(t, S_t, X_t, \theta) = \mathbb{E}_t\left[X_T^\theta\right] - \frac{\delta}{2}\mathrm{Var}_t\left[X_T^\theta\right]. \tag{87}$$

The respective (optimal) value function, in turn is given by

$$V(t, X_t) = \sup_{\theta^{t,Y_t}} J(t, S_t, X_t, \theta). \tag{88}$$

To derive the Hamilton-Jacobi-Bellman (HJB) equation, we note that

[56] That is, we make a distinction between value function and optimal value function here (compare Björk 2020).

$$
\begin{aligned}
V(t,S_t,X_t) &= \mathbb{E}_t\left[X_T^{\theta^{t,Y_t}}\right] - \frac{\delta}{2}\mathrm{Var}_t\left[X_T^{\theta^{t,Y_t}}\right] \\
&= \mathbb{E}_t\left[E_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right] - \frac{\delta}{2}\left(\mathbb{E}_t\left[\mathrm{Var}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right] + \mathrm{Var}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right]\right) \\
&= \mathbb{E}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right] - \frac{\delta}{2}\mathrm{Var}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right] - \frac{\delta}{2}\mathrm{Var}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right] \\
&= \mathbb{E}_t\left[V\left(t+\tau, S_{t+\tau}, X_{t+\tau}^{\theta^{t,Y_t}}\right)\right] - \frac{\delta}{2}\mathrm{Var}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right].
\end{aligned} \tag{89}
$$

Thus, due to the conditional variance decomposition, there is a second term on the right hand side of (89), which already suggests that the optimal value function is not the optimum over expected future value functions. Hence, we spell this out and first show that Bellman's principle of optimality does not hold in the present case. To see this, we follow the standard procedure (Björk, 2020, Basak & Chabakauri, 2012) in assuming a small time increment τ and two optimal strategies given different starting times at which these strategies are determined, that is θ^{t,Y_t} and $\theta^{t+\tau,Y_{t+\tau}}$. We will intuitively show why θ_s^{t,Y_t}, for $s \in [t+\tau, T]$ is not equal to $\theta_s^{t+\tau,Y_{t+\tau}}$, for $s \in [t+\tau, T]$. This result follows immediately from (89) and considering that

$$
\begin{aligned}
V(t,X_t,S_t) &= \mathbb{E}_t\left[J\left(t+\tau, S_{t+\tau}, X_{t+\tau}^{\theta^{t,Y_t}}, \theta^{t,Y_t}\right)\right] - \frac{\delta}{2}\mathrm{Var}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t,Y_t}}\right]\right] \\
&\neq \sup_{\{\theta_s\}_{s=t}^{t+\tau}}\left\{\mathbb{E}_t\left[J\left(t+\tau, S_{t+\tau}, X_{t+\tau}^{\theta_s}, \theta^{t+\tau,Y_{t+\tau}}\right)\right] - \frac{\delta}{2}\mathrm{Var}_t\left[\mathbb{E}_{t+\tau}\left[X_T^{\theta^{t+\tau,Y_{t+\tau}}}\right]\right]\right\}.
\end{aligned} \tag{90}
$$

The inequality is due to the fact that at time $t+\tau$, when strategy $\theta^{t+\tau,Y_{t+\tau}}$ is determined, the second term after the inequality sign is equal to zero, as the conditional expectation at time $t+\tau$ is in $\mathcal{F}_{t+\tau}$ and, hence, known at that time, so that its conditional variance is equal to zero, whereas the term is not equal to zero and, thus, taken into account by the investor at time t, when strategy θ^{t,Y_t} is determined. As a result, $\theta_s^{t+\tau,Y_{t+\tau}} \neq \theta_s^{t+\tau,Y_{t+\tau}}$, for $s \in [t+\tau, T]$. Hence, the conditional variance decomposition in (89) leads to an adjustment term, which in turn leads to time inconsistency in (90). (Basak & Chabakauri, 2012, Björk & Murgoci, 2010, Strotz, 1956)

Based on the second term on the right hand side of (89) and (90), Basak & Chabakauri (2012) suggest an adjustment term to take into account the incentive to deviate with respect to the investment strategy over time. As a result of the investor taking into account her own incentive to deviate over time, a time-consistent strategy in terms of a constrained dynamic optimum may be derived by use of dynamic programming. (compare Chen & Zhou, 2022)

Given that $\mathbb{E}_t\left[X_T^{\theta^{t,Y_t}}\right] \in \mathcal{F}_t$, $\mathrm{Var}_t\left[\mathbb{E}_t\left[X_T^{\theta^{t,Y_t}}\right]\right] = 0$ and $\mathrm{Cov}_t\left[\mathbb{E}_t\left[X_T^{\theta^{t,Y_t}}\right], \mathbb{E}_{t+\tau}[X_T^{\theta^{t,Y_t}}\right] = 0$, so that we may subtract the conditional expectation within the conditional variance - term in (90). It is, thus, also innocent to impose the optimal strategy θ^{t,S_t} in this conditional expectation at time t in this instance and no circularity in the economic argument is introduced when we ask the economic agent to find the optimal $\theta_s, s \in [t, t+\tau]$, while at the same time imposing in the adjustment term of equation (90) that an optimal strategy for this time period is already known. Hence, accepting sub-optimality as discussed and setting the optimal (achievable) value function at time t equal to its (downward) adjusted expected future value in (90), we obtain

$$V(t, S_t, X_t) = \sup_{\{\theta_s\}_{s=t}^{t+\tau}} \left\{ \mathbb{E}_t \left[V(t+\tau, S_{t+\tau}, X_{t+\tau}) \right] - \frac{\delta}{2} \mathrm{Var}_t \left[\mathbb{E}_{t+\tau} \left[X_T^{\theta^{t+\tau, Y_{t+\tau}}} \right] - \mathbb{E}_t \left[X_T^{\theta^{t, Y_t}} \right] \right] \right\}. \tag{91}$$

Next, we may solve for X_T from the stochastic differential equation in (6). Given

$$dX_t = X_t r dt + \theta_t (\mu - r) dt + \theta_t \sigma dW_t,$$

we assume the function $Y_t = f(t, X_t) = e^{-rt} X_t$ and obtain, by Ito's lemma,

$$\begin{aligned} df(t, X_t) &= -re^{-rt} X_t dt + e^{-rt} dX_t \\ &= -re^{-rt} X_t dt + e^{-rt} \left(X_t r dt + \theta_t (\mu - r) dt + \theta_t \sigma dW_t \right) \\ &= e^{-rt} \theta_t (\mu - r) dt + e^{-rt} \theta_t \sigma dW_t \\ \iff e^{-rT} X_T &= e^{-rt} X_t + \int_t^T e^{-rs} \theta_s (\mu - r) ds + \int_t^T e^{-rs} \theta_s \sigma dW_s \\ \iff X_T &= e^{r(T-t)} X_t + \int_t^T e^{r(T-s)} \theta_s (\mu - r) ds + \int_t^T e^{r(T-s)} \theta_s \sigma dW_s, \end{aligned} \tag{92}$$

so that

$$\mathbb{E}_t[X_T^\theta] = e^{r(T-t)} X_t + \mathbb{E}_t \left[\int_t^T e^{r(T-s)} \theta_s (\mu - r) ds \right]. \tag{93}$$

This holds for all $t \in [0, T]$, so that $\mathbb{E}_{t+\tau}[X_T^\theta] = e^{r(T-(t+\tau))} X_{t+\tau} + \mathbb{E}_{t+\tau} \left[\int_{t+\tau}^T e^{r(T-s)} \theta_s (\mu - r) ds \right]$. Hence, plugging the latter two results into (91) and noting that

$$V(t+\tau, S_{t+\tau}, X_{t+\tau}) = \mathbb{E}_{t+\tau} \left[X_T^{\theta^{t+\tau, Y_{t+\tau}}} \right] - \frac{\delta}{2} \mathrm{Var}_{t+\tau} \left[X_T^{\theta^{t+\tau, Y_{t+\tau}}} \right], \tag{94}$$

we obtain

$$\begin{aligned} V(t, S_t, X_t) = \sup_{\{\theta_s\}_{s=t}^{t+\tau}} \mathbb{E}_t \left[\mathbb{E}_{t+\tau} \left[X_T^{\theta^{t+\tau, Y_{t+\tau}}} \right] - \frac{\delta}{2} \mathrm{Var}_{t+\tau} \left[X_T^{\theta^{t+\tau, Y_{t+\tau}}} \right] \right] \\ -\frac{\delta}{2} \mathrm{Var}_t \left(X_{t+\tau} e^{(T-(t+\tau))} + \int_{t+\tau}^T e^{r(T-s)} \theta_s^{t+\tau, Y_{t+\tau}} (\mu - r) ds - X_t e^{(T-t)} - \int_t^T e^{r(T-s)} \theta_s^{t, Y_t} (\mu - r) ds \right). \end{aligned} \tag{95}$$

As is also noted by Basak & Chabakauri (2012), the optimal policy θ^{t, Y_t} is independent of current wealth due to the separability of the objective function (87) in current wealth. This can be seen by plugging (92) and (93) into (87), as

$$J(t,S_t,X_t,\theta) = e^{r(T-t)}X_t + \mathbb{E}_t\left[\int_t^T e^{r(T-s)}\theta_s(\mu-r)ds\right] - \frac{\delta}{2}\mathrm{Var}_t\left[e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right], \tag{96}$$

whereas given that $X_t \in \mathcal{F}_t$, its conditional variance at time t is zero. Hence, we note that θ^{t,Y_y}, which maximizes $J(t,S_t,X_t,\theta)$, is independent of X_t. Moreover it then follows that

$$V(t,S_t,X_t) = J(t,S_t,X_t,\theta^{t,Y_t}) = e^{r(T-t)}X_t + \tilde{V}(t,S_t), \tag{97}$$

for all $t \in [0,T]$. Using this result for $V(t+\tau, S_{t+\tau}, X_{t+\tau})$, we see that the latter is only affected by θ_t^{t,Y_t} via $e^{r(T-(t+\tau))}X_{t+\tau}$. Moreover, the conditional variance-term in (95) is only affected by θ_t^{t,Y_t} via the differential $X_{t+\tau}e^{(T-(t+\tau))} - X_t e^{(T-t)}$, as $\theta^{t+\tau,Y_{t+\tau}}$ already reflects the future optimal strategy, which is in turn independent of future wealth levels, and the term including θ^{t,Y_t} is only an artificial term added for computational purposes, which does not affect the result, as discussed above. Hence, the conditional variance term in (95) is affected by θ_t^{t,Y_t} only via the wealth-differential. As is commonly done, we now subtract $V(t,S_t,W_t)$ from both sides in (95) and let τ go to zero to obtain,

$$0 = \sup_{\theta_t}\mathbb{E}_t\left[d\tilde{V}(t,S_t) + d\left(X_t e^{r(T-t)}\right)\right] - \frac{\delta}{2}\mathrm{Var}_t\left[d\left(\int_t^T e^{r(T-s)}\theta_s^{t,Y_t}(\mu-r)ds\right) + d\left(X_t e^{r(T-t)}\right)\right]. \tag{98}$$

By Ito's lemma, we obtain

$$\begin{aligned} d(X_t e^{r(T-t)}) &= d\tilde{f}(t,X_t) = -re^{r(T-t)}X_t dt + e^{r(T-t)}(X_t r dt + \theta_t(\mu-r)dt + \theta_t\sigma dW_t) \\ &= e^{r(T-t)}\theta_t(\mu-r)dt + e^{r(T-t)}\theta_t\sigma dW_t \\ \implies E_t[d(X_t e^{r(T-t)})] &= e^{r(T-t)}\theta_t(\mu-r)dt \end{aligned} \tag{99}$$

so that after dropping all terms that do not depend on θ_t, plugging in results for $dX_t e^{r(T-t)}$ and its expectation from (99) into (98) and taking the conditional variance, we obtain

$$0 = \sup_{\theta_t}\left\{e^{r(T-t)}\theta_t(\mu-r)dt - \frac{\delta}{2}e^{2r(T-t)}\theta_t^2\sigma^2 dt\right\}. \tag{100}$$

From the respective first-order condition it follows that

$$e^{r(T-t)}(\mu-r)dt - \delta e^{2r(T-t)}\theta_t\sigma^2 dt = 0 \iff \theta_t^* = \frac{e^{-r(T-t)}(\mu-r)}{\delta\sigma^2}. \tag{101}$$

A.8 Static investment strategy, CARA investor

Let the optimization problem for the static exponential utility (CARA) investor be given by

$$\max_{\bar{\theta}} \mathbb{E}\left[-e^{-\alpha X_T}\right]$$

subject to

$$dX_t = (rX_t + \bar{\theta}(\mu - r))dt + \bar{\theta}_t \sigma dW_t,$$

with $X_0 = \bar{X}_o$. Choice variable $\bar{\theta}$ is a static, in the sense of buy-and-hold, exposure in terms of amounts invested in the stock. We first solve for wealth at the terminal time point, X_T,

$$\begin{aligned} X_T &= e^{rT}\bar{X}_0 + \int_0^T e^{r(T-s)}\bar{\theta}(\mu - r)ds + \int_0^T e^{r(T-s)}\bar{\theta}\sigma dW_s \\ &= e^{rT}\bar{X}_0 + \frac{(e^{rT} - 1)\bar{\theta}(\mu - r)}{r} + \int_0^T e^{r(T-s)}\bar{\theta}\sigma dW_s, \end{aligned} \tag{102}$$

so that taking the exponential and subsequently the expectation over the resulting log-normal random variable, we obtain

$$\begin{aligned} \mathbb{E}[u(X(T)] &= -exp\left[-\alpha\left(e^{rT}\bar{X}_0 + \frac{(e^{rT} - 1)\bar{\theta}(\mu - r)}{r}\right) + \frac{\alpha^2}{2}\int_0^T e^{2r(T-s)}\bar{\theta}^2\sigma^2 ds\right] \\ &= -exp\left[-\alpha\left(e^{rT}\bar{X}_0 + \frac{(e^{rT} - 1)\bar{\theta}(\mu - r)}{r}\right) - \frac{\alpha^2}{4r}\left(\bar{\theta}^2\sigma^2\left(1 - e^{2rT}\right)\right)\right]. \end{aligned} \tag{103}$$

As above, due to the monotone transformation given by the exponential, we maximize the criterion by minimizing the exponent, which yields the following first-order condition:

$$FOC: -\frac{\alpha}{r}(\mu - r)\left(e^{rT} - 1\right) - \frac{\alpha^2}{2r}\bar{\theta}\sigma^2\left(1 - e^{2rT}\right) = 0$$

Solving for $\bar{\theta}$ yields

$$\begin{aligned} \bar{\theta}^* &= \frac{2(\mu - r)(e^{rT} - 1)}{\alpha\sigma^2(e^{2rT} - 1)} \\ &= \frac{2(\mu - r)}{\alpha\sigma^2(e^{rT} + 1)}. \end{aligned} \tag{104}$$

A.9 Static investment strategy, CRRA investor

The CRRA investor's optimization problem over all static strategies $\bar{\pi}$ in terms of relative fractions over the time horizon $[0, T]$ is given by

$$\max_{\bar{\pi}} \mathbb{E}_t \left[\frac{X(T)^{1-\gamma}}{1-\gamma} \right] \tag{105}$$

subject to

$$dX_t = X_t(r + \bar{\pi}_t(\mu - r))dt + X_t \bar{\pi} \sigma dW_t,$$

with $X_0 = \bar{X}_0$. Solving for $X(T)$ using a log-transformation and Ito's lemma, we obtain

$$X(T) = X(0) \exp \left[rT + \bar{\pi}(\mu - r)T - \frac{1}{2}\bar{\pi}^2\sigma^2 T + \bar{\pi}\sigma W(T) \right].$$

Plugging the solution back into (105) yields

$$\max_{\bar{\pi}} \mathbb{E}_t \left[\frac{X(0)^{1-\gamma}}{1-\gamma} \exp \left((1-\gamma) \left(\left(r + \bar{\pi}(\mu - r) - \frac{1}{2}\bar{\pi}^2\sigma^2 \right) T + +\bar{\pi}\sigma W(T) \right) \right) \right],$$

so that after taking the expectation, we obtain

$$= \max_{\bar{\pi}} \frac{X(0)^{1-\gamma}}{1-\gamma} \exp \left((1-\gamma) \left(\left(r + \bar{\pi}(\mu - r) - \frac{1}{2}\bar{\pi}^2\sigma^2 \right) T + (1-\gamma)^2 \frac{1}{2}\bar{\pi}^2\sigma^2 T \right) \right).$$

Again, we maximize the entire criterion by maximizing the exponent, from which, after cancelling terms, we obtain the following first-order condition and solution, respectively:

$$FOC : (\mu - r) - \bar{\pi}\sigma^2 + (1-\gamma)\bar{\pi}\sigma^2 = 0,$$

whereas solving for $\bar{\pi}$ yields

$$\bar{\pi} = \frac{(\mu - r)}{\gamma\sigma^2}.$$

A.10 Derivation Arrow-Pratt measure of absolute risk aversion

For the sake of completeness, this section derives the Arrow-Pratt measure of absolute risk aversion (Pratt, 1964) for the CRRA investor (for the CARA investor, the derivation is analogous) in (15), in order to underline the role of and dependence on the reference wealth level X_{t^R}. As in Pratt (1964), let us assume a zero-mean risk $\tilde{Z}$ with known variance σ^2 and let us denote the risk premium π. The Arrow-Pratt measure of absolute risk aversion is based on the following first - and second-order Taylor expansions (and expectations thereof) of utility of future wealth around X_{t^R} (Pratt, 1964):

$$
\begin{aligned}
\mathbb{E}_t\left[u^{CR}(X_{t^R}+\tilde{Z})\right] &\approx \mathbb{E}_t\left[u^{CR}(X_{t^R})+u'^{CR}(X_{t^R})\tilde{Z}+\frac{1}{2}u''^{CR}(X_{t^R})\tilde{Z}^2\right] \\
&= u^{CR}(X_{t^R})+\frac{1}{2}u''^{CR}(X_{t^R})\sigma^2
\end{aligned}
\tag{106}
$$

and

$$
u^{CR}(X_{t^R}-\pi) \approx u^{CR}(X_{t^R}) - u'^{CR}(X_{t^R})\pi. \tag{107}
$$

Equating (106) and (107), as per the definition of the risk-premium (Pratt, 1964), we obtain

$$
\pi \approx -\frac{1}{2}\sigma^2\frac{u''^{CR}(X_{t^R})}{u'^{CR}(X_{t^R})}, \tag{108}
$$

which underlines that the Arrow-Pratt measure of absolute risk aversion uses a reference level at the point in time when a decision is taken by the investor in order to approximate the risk premium by a function of changes of utility as wealth fluctuates around the given reference level. This also illustrates why there may be sizeable errors in using initial wealth at time $t^R = 0$ as a reference level in a multi-period (possibly dynamic) investment problem.

A.11 Myopic investment strategies, CARA investor

Schweizer et al.'s (2021) strategy is to derive the optimal precommitment strategy at each time point $t \in [0, T]$ via the investment strategy that maximizes the investor's certainty equivalent at that point in time. The CARA investor's certainty equivalent is given by

$$
CE_t^{CA} = U^{CA^{-1}}\mathbb{E}_t\left[U^{CA}(X_T)\right]. \tag{109}
$$

The utility function for the CARA investor is given by $U^{CA} = -e^{-\alpha X_T}$, so that

$$
U^{CA} = -e^{-\alpha X_T} = c \iff U^{CA^{-1}}(c) = -\frac{1}{\alpha}ln(-c) = X_T. \tag{110}
$$

Moreover, assuming that the CARA investor optimizes over absolute invested amounts and that the respective wealth dynamics are given by (6), wealth at time T conditional on wealth at time t being fixed at X_t is given by

$$
X_T = e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu-r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s. \tag{111}
$$

Hence, it follows from (109) and (110) that

$$
\begin{aligned}
CE_t^{CA} &= -\frac{1}{\alpha} ln\left(\mathbb{E}\left[\exp\left(-\alpha\left(e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu - r)ds + \int_t^T e^{r(T-s)}\theta_s\sigma dW_s\right)\right)\right]\right) \\
&= -\frac{1}{\alpha} ln\left(\exp\left(-\alpha\left(e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu - r)ds\right) + \frac{1}{2}\alpha^2\int_t^T e^{2r(T-s)}\theta_s^2\sigma^2 ds\right)\right) \\
&= e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu - r)ds - \frac{1}{2}\alpha\int_t^T e^{2r(T-s)}\theta_s^2\sigma^2 ds.
\end{aligned}
\tag{112}
$$

As in Schweizer et al.'s argument concerning CRRA investors, we may now use the fact that CARA investors' optimal precommitment strategy is time-dependent but deterministic to simplify the maximization of the certainty equivalent C_t^{CA} in the following way. Defining

$$
m_s = e^{r(T-s)}\theta_s \tag{113}
$$

and plugging back into (112), we obtain

$$
CE_t^{CA} = e^{r(T-t)}X_t + \int_t^T m_s(\mu - r)ds - \frac{1}{2}\alpha\int_t^T m_s^2\sigma^2 ds. \tag{114}
$$

Maximizing (114) over $\{m_s\}_{s=t}^T$, we note that the certaint equivalent is maximized if $\int_t^T m_s^2 ds$ is minimized for any κ such that $\int_t^T m_s ds = \kappa$. The latter constrained optimization problem,

$$
\min_{\{m_s\}_{s=t}^T} \int_t^T m_s^2 ds \tag{115}
$$

$$
s.t. \int_t^T m_s ds = \kappa, \tag{116}
$$

yields $m_s = \frac{\kappa}{(T-t)}$, which is, hence, constant, as in Schweizer et al. (2021). As a results we may set $m_s = \bar{m}$ when maximizing (114) over $\{m_s\}_{s=t}^T$, so that

$$
\max_{\bar{m}} CE_t^{CA} = \max_{\bar{m}} \left\{e^{r(T-t)}X_t + \bar{m}(T-t)(\mu - r) - \frac{1}{2}\alpha\bar{m}^2\sigma^2(T-t)\right\}, \tag{117}
$$

from which it follows after taking the first order condition and solving for $\bar{m}$ that

$$
\bar{m} = \frac{(\mu - r)}{\alpha\sigma^2}. \tag{118}
$$

From (113) it then follows that $\theta_s = e^{-r(T-s)}\frac{(\mu-r)}{\alpha\sigma^2}$, for all $s \in [t, T]$.

A.12 Myopic investment strategy, Mean-Variance investor

Similar to the constrained optimization problem for the Mean-Variance investor under the pre-commitment strategy, the myopic Mean-Variance investor solves the following constrained optimization problem at each point in time $t \in [0, T]$:

$$\max_{X_T} \left\{ \mathbb{E}_t[X_T] - \frac{\delta}{2} \mathrm{Var}_t[X_T] \right\} \tag{119}$$

subject to the budget constraint

$$\frac{1}{M_t} \mathbb{E}_t[M_T X_T] = X_t. \tag{120}$$

The Lagrangian for the maximization problem is given by

$$\mathcal{L} = \mathbb{E}_t[X_T] - \frac{\delta}{2} \mathrm{Var}_t[X_T] + \psi \left(X_t - \frac{1}{M_t} \mathbb{E}_t[M_T X_T] \right), \tag{121}$$

whereas the respective first-order conditions, taken by state of the world, are

$$1 - \delta X_T + \delta \mathbb{E}_t[X_T] - \psi \frac{M_T}{M_t} = 0 \tag{122}$$

and

$$X_t - \frac{1}{M_t} \mathbb{E}_t[M_T X_T] = 0, \tag{123}$$

respectively. To compute (122), we note that

$$\mathrm{Var}_t[X_T] = \mathbb{E}_t \left[(X_T - \mathbb{E}_t[X_T])^2 \right] = \int\limits_{\omega \in \Omega : X_t = \tilde{X}_t} (X_T(\omega) - \mathbb{E}_t[X_T])^2 \, d\mathbb{P}(\omega), \tag{124}$$

so that taking the first-order derivative of the variance with respect to X_T for each state of the world separately effectively means that we condition on the state of the world $\omega \epsilon \Omega$ in (124) before taking the derivative. Thus, the first order derivative conditional on state of the world $\bar{\omega} \epsilon \Omega$ is given by

$$\frac{d}{dX_T} \mathrm{Var}_t[X_T] |_{\omega = \bar{\omega} \epsilon \Omega} = 2X_T - 2\mathbb{E}_t[X_T].$$

From (122), it follows that

$$X_T = \frac{1}{\delta} \left(1 + \delta \mathbb{E}_t[X_T] - \psi \frac{M_T}{M_t} \right). \tag{125}$$

The optimal value for X_T by state of the world is yet to be determined, so that its expectation is still unknown, as is the value of the Lagrangian multiplier ψ. Thus, we first take the conditional expectation on both sides of (125) to solve for ψ to obtain

$$\delta\mathbb{E}_t[X_T] = 1 + \delta\mathbb{E}_t[X_T] - \psi e^{-r(T-t)} \iff \psi = e^{r(T-t)}. \tag{126}$$

Plugging back into (125), we obtain

$$X_T = \frac{1}{\delta}\left(1 + \delta\mathbb{E}_t[X_T] - e^{r(T-t)}\frac{M_T}{M_t}\right). \tag{127}$$

X_T must be chosen such that the budget constraint (123) is satisfied. This also implies a condition on $\mathbb{E}_t[X_T]$, which depends on the optimal choice of X_T by state of the world and the known distribution $\mathbb{P}$, so that we may determine $E[X_T]$ as implied by the optimal choice for terminal wealth by plugging (127) into the budget constraint (123) and solving for $E[X_T]$:

$$\begin{aligned} X_t &= \frac{1}{M_t}\mathbb{E}_t\left[M_T\frac{1}{\delta}\left(1 + \delta\mathbb{E}_t[X_T] - e^{r(T-t)}\frac{M_T}{M_t}\right)\right] \\ &= \frac{1}{\delta}\mathbb{E}_t\left[\frac{M_T}{M_t}\right] + \mathbb{E}_t\left[\frac{M_T}{M_t}\right]\mathbb{E}_t[X_T] - \frac{1}{\delta}\mathbb{E}_t\left[\frac{M_T^2}{M_t^2}\right]e^{r(T-t)} \\ &= e^{-r(T-t)}\frac{1}{\delta} + e^{-r(T-t)}\mathbb{E}_t[X_T] - \frac{1}{\delta}e^{-2r(T-t)+\lambda^2(T-t)+r(T-t)} \\ \iff \mathbb{E}_t[X_T] &= X_t e^{r(T-t)} - \frac{1}{\delta} + \frac{1}{\delta}e^{\lambda^2(T-t)}, \end{aligned} \tag{128}$$

whereas the third equality follows from

$$\mathbb{E}_t\left[\frac{M_T^2}{M_t^2}\right] = \mathbb{E}_t\left[e^{\left(-2r(T-t)-\lambda^2(T-t)-2\lambda(W_T-W_t)\right)}\right] = e^{-2r(T-t)+\lambda^2(T-t)}$$

and

$$\mathbb{E}_t\left[\frac{M_T}{M_t}\right] = \mathbb{E}_t\left[e^{-r(T-t)-\frac{1}{2}\lambda^2(T-t)-\lambda(W_T-W_t)}\right] = e^{-r(T-t)}.$$

Plugging (128) back into (127) yields

$$\begin{aligned} X_T &= \frac{1}{\delta}\left(1 + \delta\left(X_0 e^{r(T-t)} - \frac{1}{\delta} + \frac{1}{\delta}e^{\lambda^2(T-t)}\right) - e^{r(T-t)}\frac{M_T}{M_t}\right) \\ \iff X_T^* &= X_t e^{r(T-t)} + \frac{1}{\delta}e^{\lambda^2(T-t)} - \frac{1}{\delta}e^{r(T-t)}\frac{M_T}{M_t}. \end{aligned} \tag{129}$$

To compute the prefect hedge, which will ultimately yield the optimal investment strategy for the investor standing at decision time t, we first derive the market consistent value of the optimal terminal wealth X_T at any time $t^R \in [t, T]$:

$$
\begin{aligned}
X^*_{t^R} &= \frac{1}{M_{t^R}}\mathbb{E}_{t^R}[M_T X^*_T] \\
&= \frac{1}{M_{t^R}}\mathbb{E}_{t^R}\left[M_T X_t e^{r(T-t)} + \frac{1}{\delta}e^{\lambda^2(T-t)}M_T - \frac{M_T}{M_t}^2\frac{1}{\delta}e^{r(T-t)}\right] \\
&= \frac{1}{M_{t^R}}\left(M_{t^R}e^{-r(T-t^R)}X_t e^{r(T-t)} + \frac{1}{\delta}e^{\lambda^2(T-t)}M_{t^R}e^{-r(T-t^R)} - \frac{1}{\delta M_t}M^2_{t^R}e^{-2r(T-t^R)+\lambda^2(T-t^R)+r(T-t)}\right) \\
&= e^{-r(t^R-t)}X_t + \frac{1}{\delta}e^{\lambda^2(T-t)}e^{-r(T-t^R)} - \frac{1}{\delta}\frac{M_{t^R}}{M_t}e^{-r(T-t^R)+\lambda^2(T-t^R)+r(t^R-t)} \\
&= f_t(M_{t^R}).
\end{aligned}
\tag{130}
$$

Given the differential form for the stochastic discount factor, $dM_t = M_t(-rdt - \lambda dW_t)$ for any time t, the differential of the process X_{t^R} at time $t^R \in [t, T]$ following the strategy as determined at time t is given by

$$
\begin{aligned}
dX^*_{t^R} = df_t(M_{t^R}) &= .....dt + \frac{\partial f_t(M_{t^R})}{\partial M_t}\frac{M_{t^R}}{M_t}\lambda dW_{t^R} \\
&= ....dt + \frac{1}{\delta}e^{-r(T-t^R)+\lambda^2(T-t^R)+r(t^R-t)}\frac{M_{t^R}}{M_t}\lambda dW_{t^R}
\end{aligned}
\tag{131}
$$

Rewriting, as above, processes (131) and (6) using $\tilde{W}_t$, a Brownian Motion under measure $\mathbb{Q}$, we note that, as above, both processes show the same drift term, so that the perfect hedge may be computed by equating the volatility terms of the differential form of the processes in (131) and (6), respectively (See discussion in Appendix A1). Hence, we obtain

$$
\begin{aligned}
\theta^*_{t^R}\sigma &= \frac{1}{\delta}e^{-r(T-t^R)+\lambda^2(T-t^R)+r(t^R-t)}\frac{M_{t^R}}{M_t}\lambda\theta_{t^R}\sigma \\
\iff \theta^*_{t^R} &= \frac{1}{\delta}\frac{(\mu-r)}{\sigma^2}e^{-r(T-t^R)+\lambda^2(T-t^R)+r(t^R-t)}\frac{M_{t^R}}{M_t}.
\end{aligned}
\tag{132}
$$

For the sake of completeness, the derivation of relative investment weights follows the same steps as in appendices A2, A4 and A6, hence, we merely state the result here, referring for derivations to the appendices above:

$$
\pi^*_{t^R} = \frac{1}{\delta}\frac{(\mu-r)}{\sigma^2}\frac{M_{t^R}}{M_t X_{t^R}}e^{-r(T-t^R)+\lambda^2(T-t^R)+r(t^R-t)}. \tag{133}
$$

Let us assume that strategy in (132) is applied over a short time-interval, that is for all $t^R \in [t, t+\tau]$. If the length of the interval, τ, goes to zero, then $t^R \to t$, so that the solution for θ_t as in (132), however for all $t \in [0, T]$, becomes

$$
\theta^*_t = \frac{1}{\delta}\frac{(\mu-r)}{\sigma^2}e^{-r(T-t)+\lambda^2(T-t)}. \tag{134}
$$

A.13 Certainty equivalent total gross returns and growth rates

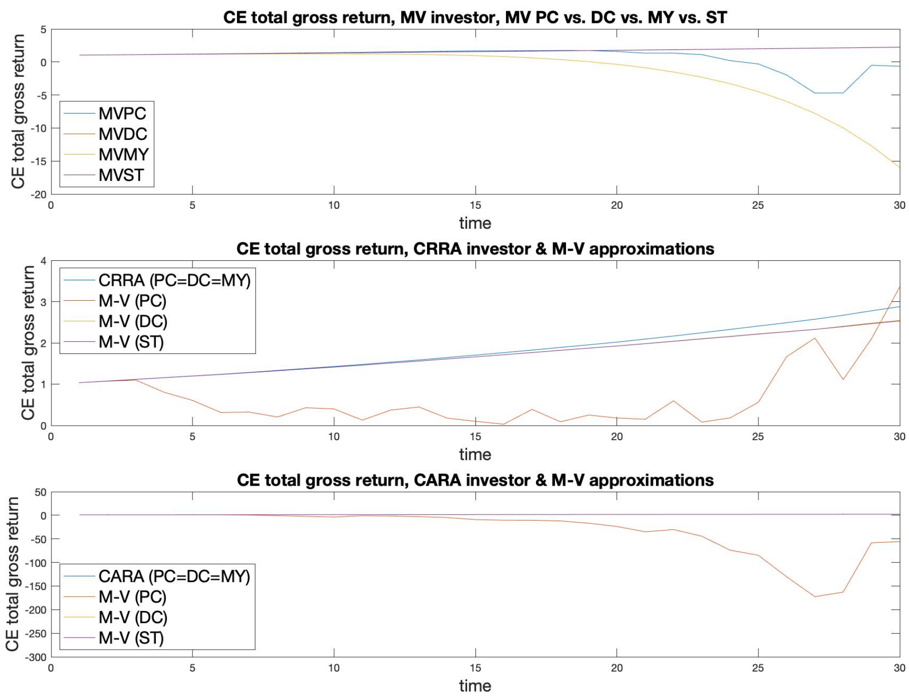

Figure 22: CE total gross returns, T=30, n=30, m=10000

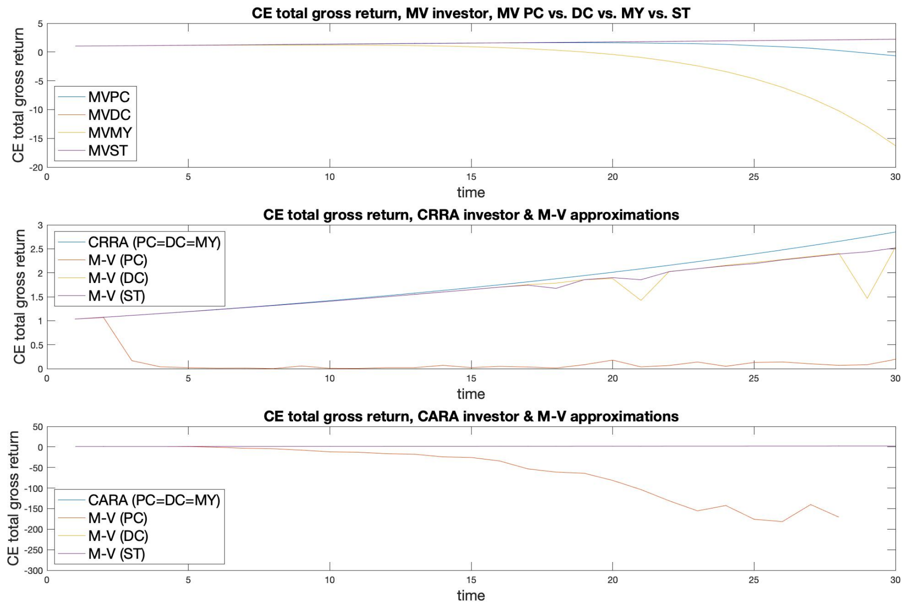

Figure 23: CE total gross returns, T=30, n=30, m= 1000000. Computed Certainty Equivalent of pre-commitment Mean-Variance strategy drops to minus infinity, whereas the underlying mean is not defined (see discussion in main body of the text).

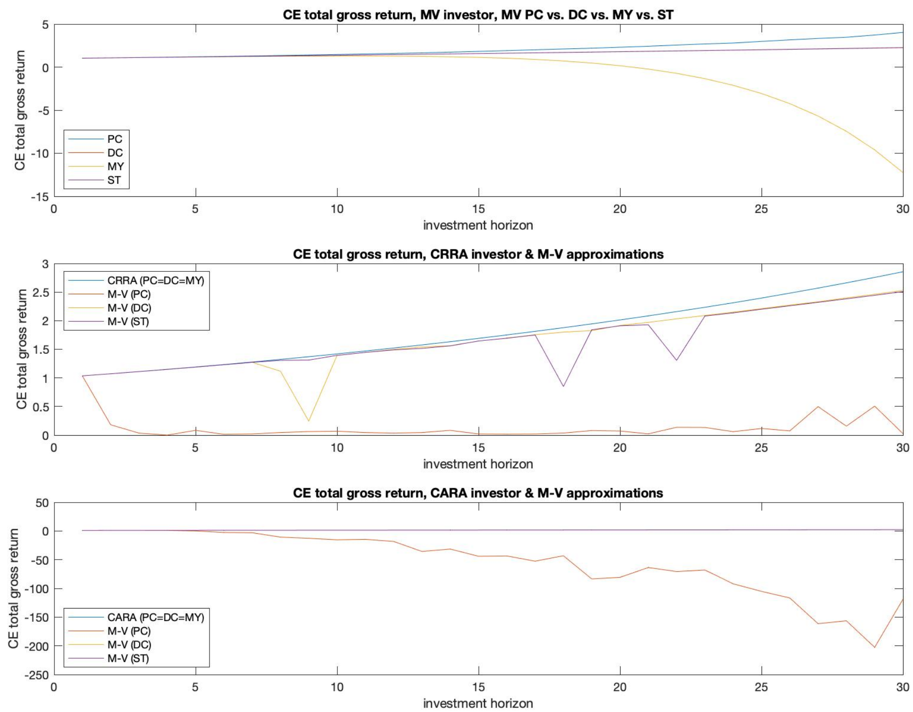

Figure 24: CE total gross returns. Parameters: T=30, n=360, m= 1000000, yearly performance measurement

μ_{MC} $(\sigma/\sqrt{n})$	ST	MY	DC	PC
M-V investor	507.27 (0.6)	647.4 (0.35)	272.33 (0.05)	647.8 (0.13)
CARA investor	505.27 (0.6)	272.33 (0.05)	272.33 (0.05)	272.33 (0.05)
CRRA investor	668.9 (0.84)	448.17 (0.27)	448.17 (0.27)	448.17 (0.27)

Table 10: Monte Carlo Mean of wealth at T=30, with m=1000000, n=360

μ_{MC} $(\sigma/\sqrt{n})$	ST	MY	DC	PC
M-V investor	511.16 (5.78)	655.23 (3.53)	273.1 (0.55)	649.69 (1.65)
CARA investor	511.16 (6.78)	273.1 (0.55)	273.1 (0.55)	273.1 (0.55)
CRRA investor	646.38 (8.16)	451.23 (2.66)	451.23 (2.66)	451.23 (2.66)

Table 11: Monte Carlo Mean of wealth at T=30, with m=10000, n=360

A.14 Proof Divergence Expected CARA Utility from MV-PC strategy

As seen from time t, terminal wealth X_T is given by

$$X_T = e^{r(T-t)}X_t + \int_t^T e^{r(T-s)}\theta_s(\mu - r)ds + \int_t^T e^{r(T-s)}\theta_s \sigma dW_s. \tag{135}$$

with

$$\theta_s^* = \frac{1}{\delta}\frac{(\mu - r)}{\sigma^2}M_s e^{-r(T-s)+\lambda^2(T-s)+rs} \tag{136}$$

and

$$M_s = \exp\left(-rs - \frac{1}{2}\lambda^2 s - \lambda W_s\right). \tag{137}$$

The question is whether $\mathbb{E}_t[\exp(-\alpha X_T)]$ exists. From (135) and (136), at any time t wealth after one time increment is normally distributed, as θ_s^* is an adapted process and W_s and dW_s are independent. However, the distribution of wealth after more than one time-increment is not normal (and neither log-normal). Hence, we cannot use any results directly based on knowledge of the normal or log-normal distributions to check whether $\mathbb{E}_t[\exp(-\alpha X_T)]$ exists. Thus, we need to evaluate the expectation based on the joint distribution of the Brownian increments that lead up to wealth X_T. Hence, if we denote $\Pi = \{t_0, t_t, ..., t_n\}$ a partition of T, $(W_{t_0}, W_{t_1}, ..., W_{t_n})$ the vector of the values of the standard Brownian motion at each time point t_i, $\Delta W_{t_j} = W_{t_{j+1}} - W_{t_j}$ for $j = 0, 1, ..., n-1$ the Brownian increments with $\Delta W_{t_j} = z_j\sqrt{\Delta t_j}$, with $z_j \sim \mathcal{N}(0,1)$ for all $j = 0, ..., n-1$ and $z = (z_0\sqrt{\Delta t_0}, ..., z_{n-1}\sqrt{\Delta t_{n-1}})$, and $\Sigma = \text{diag}(\Delta_{t_0}, ..., \Delta_{t_{n-1}})$ the Variance-Covarianve matrix of the Brownian increments, then the integral underlying the expectation above exists if

$$\int \cdots \int_{-\infty}^{\infty} \exp\left(-\alpha \lim_{\|\Pi\|\to 0} \sum_{j=0}^{n-1} \exp\left(-\lambda \sum_{i=0}^{j-1} z_i\sqrt{\Delta t_i}\right) - \alpha \lim_{\|\Pi\|\to 0} \sum_{j=0}^{n-1} \exp\left(-\lambda \sum_{i=0}^{j-1} z_i\sqrt{\Delta t_i}\right) z_j\sqrt{\Delta t_j}\right) \times \frac{\|\Sigma\|^{-\frac{1}{2}}}{\sqrt{2\pi}} \exp\left(z^T\Sigma^{-1}z\right) dz_1 dz_2 .... dz_{n-1} \tag{138}$$

exists (that is, abstracting from constants or deterministic functions of time). However, sums of normal random variables are normally distributed and products of log-normal random variables are log-normally distributed. Moreover, the potential problem locations are at $z_i \to \pm\infty$, so we also abstract from the Variance of the Brownian increments Δ_t. Thus, the above integral exists if the following integral exists

$$\int_{-\infty}^{\infty} \exp\left(-\alpha\exp(-\lambda x) - \alpha\exp(-\lambda x)x\right)\exp\left(-\frac{x^2}{2}\right) dx \tag{139}$$

The problem location is at $\lim_{x\to-\infty}$, hence, we split up the integral and focus on the integral

$$\int_{-\infty}^{0} \exp\left(-\alpha\exp(-\lambda x) - \alpha\exp(-\lambda x)x\right)\exp\left(-\frac{x^2}{2}\right) dx \tag{140}$$

, which after a simple change of variable (z=-x) is equal to

$$\int_0^{\infty} \exp\left(-\alpha \exp(\lambda z) + \alpha \exp(\lambda z) z\right) \exp\left(-\frac{z^2}{2}\right) dx. \tag{141}$$

We now show that the integrand is infinite at the upper limit and bounded from below by test function f(z)=z, the integral of which is infinite. First, we rewrite the integrand as

$$\begin{aligned}
\exp\left(-\alpha \exp(\lambda z) + \alpha \exp(\lambda z) z\right) exp\left(-\frac{z^2}{2}\right) &= \frac{\exp\left(\alpha \exp\left(\lambda z\right) z\right)}{\exp\left(\alpha \exp\left(\lambda z\right)\right) \exp\left(\frac{z^2}{2}\right)} \\
&= \frac{\exp\left(\alpha \exp\left(\lambda z\right)\right)^{z}}{\exp\left(\alpha \exp\left(\lambda z\right)\right) \exp\left(\frac{z^2}{2}\right)} \\
&= \frac{\exp\left(\alpha \exp\left(\lambda z\right)\right)^{z-1}}{\exp\left(\frac{z^2}{2}\right)} \\
&= \frac{\exp\left(\alpha \exp\left(\lambda z\right)\right)^{z-2}}{\frac{\exp\left(\frac{z^2}{2}\right)}{\exp(\alpha \exp(\lambda z))}}.
\end{aligned} \tag{142}$$

The denominator goes to zero as $z \to \infty$, so that there is a number $k \in \mathbb{R}$ s.t. $\frac{\exp\left(\frac{z^2}{2}\right)}{\exp(\alpha \exp(\lambda z))} \leq 1$, for $z \geq k$. Moreover, for $z \geq 3$, the numerator goes to infinity and is larger than z (for $\alpha, \lambda > 0$). Thus, for $l = max\{3, k\}$, the integral in (141) is bounded from below by

$$\int_l^{\infty} z dz, \tag{143}$$

which diverges to ∞.

A.15 Further graphs for comparison

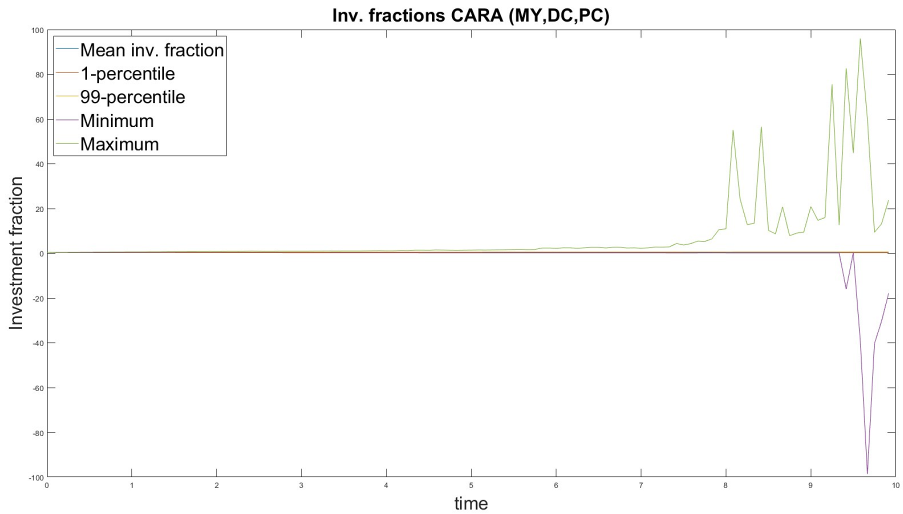

Figure 25: Investment fractions (stock) investment, CARA investor. Parameters: T=10, m=1 000 000, n=120.

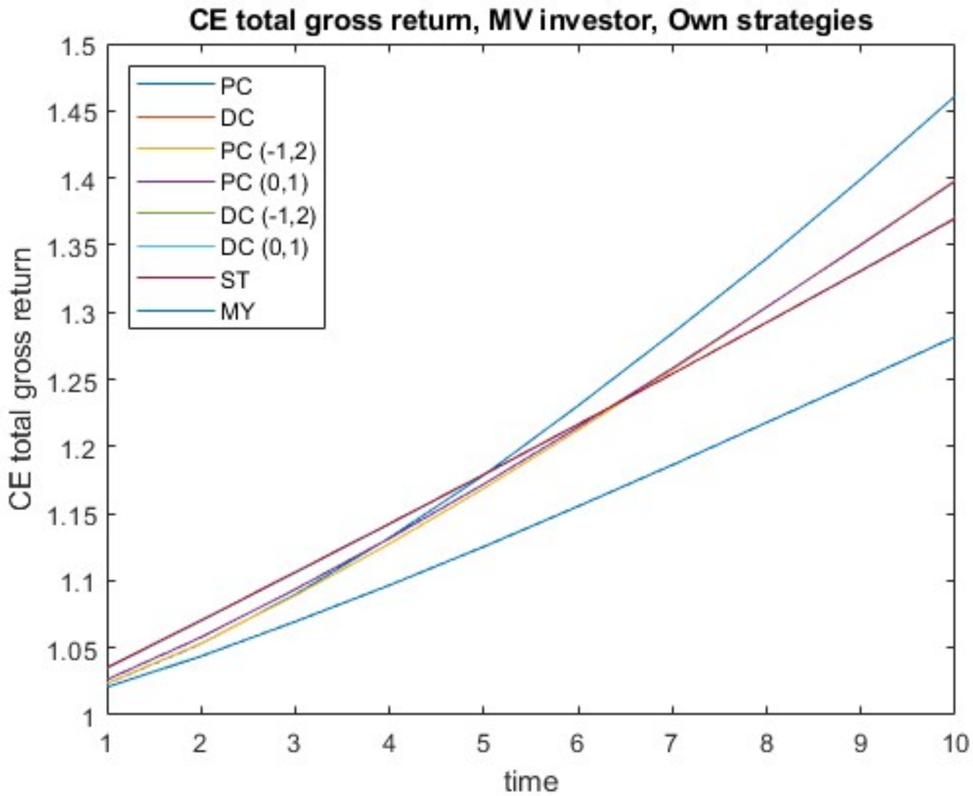

Figure 26: CE total gross return for Mean-Variance investors and from various (constrained) Mean-Variance strategies, including myopic strategy. Parameters: T=10, m=1000000,n=360, yearly performance measurement.

B References

Allais, M., Allais Paradox, in Eatwell, J., Milgate, M., Newman, P. (Eds.), Utility and Probability, The Macmillan Press Limited, 1990, pp. 251-253.

Arrow, K.J., The Theory of Risk-Bearing: Small and Great Risks, Journal of Risk and Uncertainty, Vol.12, No.2/3, 1996, pp.103-111.

Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., Coherent Measures of Risk, Mathematical Finance, Vol.9, No.3, 1999, pp.203-228.

Bajeux-Besnainou, I., Portait, R., Dynamic Asset Allocation in a Mean-Variance Framework, Management Science, Vol. 44, No.11, 1999, pp.579-595.

Baron, D.P., In the Utility-Theoretic Foundations of Mean-Variance Analysis, The Journal of Finance, Vol. 32, No.5, 1977, pp.1683-1697.

Basak, S., Chabakauri, G., Dynamic Mean-Variance Asset Allocation, The Review of Financial Studies, Vol.23, No.8, 2010, pp.2970-3016.

Baxter, M., Rennie, A., Financial Calculus, Cambridge University Press, 2012.

Balter, A.G., Mahayni, A., Schweizer, N., Time-consistency of optimal investment under smooth ambiguity, European Journal of Operational Research, Vol. 293, 2021, pp.643-657.

Bernoulli, D., Exposition of a new theory on the measurement of risk, Econometrica, Vol. 22, No.1, 1954, pp. 23-36.

Bielecki, T.R., Jin, Hanqing, Pliska, S.R., Zhou, X.Y., Continuous-time Mean-Variance Portfolio Selection with Bankrupticy Prohibition, Mathematical Finance, Vol. 15, No.2, 2005, pp.213-244.

Black, F., Scholes, M., The Pricing of Options and Corporate Liabilities, The Journal of Political Economy, Vol. 81, No.3, 1973, pp.637-654.

Björk, T., Arbitrage Theory in Continuous Time, Oxford University Press, 2020.

Björk, T., Murgoci, A., A General Theory of Markovian Time Inconsistent Stochastic Control Problems, SSRN Working Paper, 2010.

Björk, T., Murgoci, A., Zhou, X.Y., Mean-Variance Portfolio Optimization with State-Dependent Risk Aversion, Mathematical Finance, Vol. 24, No. 1, 2014, pp.1-24.

Björk, T., Khapko, M., Murgoci, A., On time-inconsistent stochastic control in continuous time, Finance and Stochastics, Vol.21, 2017, pp.331-360.

Borch, K., A Note on uncertainty and Indifference Curves, The Review of Economic Studies, Vol.36, No.1, 1969, pp.1-4.

Bovenberg, A.L., Koijen, R.S.J., Nijman, T.E., Teulings, C.N., Saving and investing over the life cycle and the role of collective pension funds, De Economist, Vol. 155, No. 5, 2007, pp.347-415.

Briys, E., Eeckhoudt, L., Loubergé, H., Endogenous Risks and the Risk Premium, Theory and Decision, Vol. 26, 1989, pp.37-46.

Chen, L., Zhou, X.Y., Naïve Markovitz Policies, Working Paper 2022, available under https://arxiv.org/abs/2212.07516.

Cochrane, J., Asset Pricing, Princeton University Press, 2001.

Cox, J.C., Huang, C.-F., Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion process, Journal of Economic Theory, Vol. 49, 1989, pp.33-83.

Cvitanic, J., Zapatero, F., Introduction to the Economics and Mathematics of Financial Markets, The MIT Press, 2004.

Cvitanic, J., Lazrak, A., Wang, T., implications of the Sharpe ratio as a performance measure in multi-period settings, Journal of Economic Dynamics and Control, Vol. 32, 2008, pp.1622-1649.

Eeckhoudt, L., Gollier, C., Schlesinger, H., Economic and Financial Decisions under Risk, Princeton University Press, 2005.

Elminejad, A., Havranek, T., Zuzuna, I., Relative Risk Aversion: A Meta-Analysis, ZBW Working Paper, 2022, available under https://www.econstor.eu/bitstream/10419/260586/1/RRA.pdf

Figlewski, S., Freund, S., The pricing of convexity risk and time decay in options markets, Journal of Banking and Finance Vol. 18, 1994, pp.73-91.

Hadar, J., Russell, W.R., Rules for Ordering Uncertain Prospects, The American Economic Review, Vol.59, No.1, 1969, pp.25-34.

Harkansson, N.H., Capital Growth and the Mean-Variance Approach to Portfolio Selection, The Journal of Financial and Quantitative Analysis, Vol.6, No.1, 1971, pp.517-557.

Hanoch, G., Levy, H., The Efficiency Analysis of Choices Involving Risk, The Review of Economic Studies, Vol 36, No.3, 1969, pp.335-346.

Hodder, J.E., Jackwerth, J.C., Kolokolova, O., Improved Portfolio Choice Using Second-Order Stochastic Dominance, Review of Finance, Vol.19, 2015, pp.1623-1647.

Johnstone, D., Lindley, D., Mean-Variance and Expected Uitlity: The Borch Paradox, Statictial Science, Vol. 28, No.2, 2013, pp.223-237.

Kallberg, J.G., Ziemba, W.T., On the robustness of the Arrow-Pratt risk aversion measure, Economics Letters, Vol. 2, 1979, pp.21-26.

Karni, E., State-dependent preferences, in Utility and Probability, Eatwell, J., Milgate, M., Newman, P. (Eds.), The Macmillan Press Limited, 1990, pp. 79-96.

Kroll, Y., Levy, H., Markowitz, H.M., Mean-variance Versus Direct Utility Maximization, The Journal of Finance, Vol. 39, No.1, 1984, pp.47-61.

Laffont, J.J., Martimort, D., The Theory of Incentives: The Principal-Agent Model, Princeton University Press, 2002.

Lassance, N., Reconciling mean-variance portfolio theory with non-Gaussian returns, European Journal of Operational Research, Vol.297, No.2, 2022, pp.729-740.

Levy, H., Stochastic dominance, Stochastic Dominance, Investment Decision Making under Uncertainty, Springer, 2016.

Levy, H., Stochastic Dominance, in Eatwell, J., Milgate, M., Newman, P. (Eds.), Utility and Probability, The Macmillan Press Limited, 1990, pp. 251-253.

Levy, H., Markowitz, H.M., Approximating Expected Utility by a Function of Mean and Variance, The American Economic Review, Vol.69, No.3, 1979, pp.308-317.

Li, X., Zhou, X.Y., Lim, A.E.B., Dynamic Mean-Variance Portfolio Selection with No-Shorting Constraints, Siam Journal on Control and Optimization, Vol. 40, No.5, 2002, pp.1540-1555.

Machina, M.J., Expected Utility Hypothesis, in Utility and Probability, Eatwell, J., Milgate, M., Newman, P. (Eds.), The Macmillan Press Limited, 1990, pp. 79-96.

Machina, M.J., Rothschild, M., Risk, in Eatwell, J., Milgate, M., Newman, P. (Eds.), Utility and Probability, The Macmillan Press Limited, 1990, pp. 227-138.

Malavesi, M., Ortobelli Lozza, S., Trück, S., Second order of stochastic dominance efficiency vs mean variance efficiency, European Journal of Operational Research, Vol.290, 2021, pp.1192-1206.

Markowitz, H., Portfolio Selection, The Journal of Finance, Vol. 7, No.1, 1952, pp.77-91.

Markowitz, H., Mean-variance approximations to expected utility, European Journal of Operational Research, Vol.234, 2014, pp.346-355.

Mas-Colell, A., Whinston, M.D., Green, J.R., Microeconomic Theory, Oxford University Press, 1995.

Mehra, R., Prescott, E.C., The Equity Premium, A Puzzle, Journal of Monetary Economics, Vol. 15, 1985, pp.145-161.

Merton, R.C., Lifetime Portfolio Selection under Uncertainty: The continuous - time case, The Review of Economics and Statistics, Vol. 51, No.3, 1969, pp.247 - 257.

Merton, R.C., Optimum Consumption and Portfolio Rules in a Continuous-Time Model, Journal of Economic Theory, Vol.3, 1971, pp.373-413.

Metselaar, L., Zwaneveld, P., van Ewijk, C., Reforming Occupational Pensions in the Netherlands: Contract and Intergenerational Aspects, De Economist, Vol. 170, 2022, pp.7-36.

Munk, C., Dynamic Asset Allocation, Lecture Notes, 2017, available under https://sites.google.com/view/clausmunk/teaching

Pedersen, L.H., Efficiently Inefficient, Princeton University Press, 2015.

Pliska, S.R. A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios, Mathematics of Operations Research, Matjematics of Operations Research, Vol.11, No.2, 1986, pp.371-382.

Polemarchakis, H.M., State preference approach, in Utility and Probability, Eatwell, J., Milgate, M., Newman, P. (Eds.), The Macmillan Press Limited, 1990, pp. 248-251.

Pratt, J.W., Risk Aversion in the Small and in the Large, Econometrica, Vol. 32, No.1/2, 1964, pp. 122-136.

Pulley, L.B., A General Mean-Variance Approximation to Expected Utility for Short Holding Periods, The Journal of Financial and Quantitative Analysis, Vol.16, No.3, 1981, pp.361-373.

Pulley, L.B., Mean-Variance Approximations to Expected Logarithmic Utility, Operations Research, Vol. 31, No.4, 1983, pp.685-696.

Pulley, L.B., Mean-Variance versus Direct Utility Maximization: A Comment, The Journal of Finance, Vol.40, No.2, 1985, pp.601-602.

Reid, D.W., Tew, B.V., Mean-Variance versus Direct Utility Maximization: A Comment, The Journal of Finance, Vol. 31, No.5, 1986, pp.1177-1179.

Richardson, H.R., A Minimum Variance Result in Continuous Trading Portfolio Optimization, Management Science, Vol. 35, No.9, 1989, pp.1045-1055.

Rothschild, M., Stiglitz, J.E., Increasing Risk: I. A Definition, Journal of Economic Theory, No.2, 1970, pp.225-243.

Rothschild, M., Stiglitz, J.E., Increasing Risk II: Its Economic Consequences, Journal of Economic Theory, No.3, 1971, pp.66-84.

Rubinstein, M., The Strong Case for the Generalized Logarithmic Utility Model as the Premier Model of Financial Markets, The Journal of Finance, Vol.31, No.2, 1976, pp.551-571.

Sharpe, W., Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, The Journal of Finance, Vol. 19, No.3, 1964, pp.425-442.

Shreve, S.E., Stochastic Calculus for Finance II, Springer, 2010.

Simaan, Y., The opportunity cost of mean-variance choice under estimation risk, European Journal of Operational Research, Vol.234, No.2, 2014, pp.382-391.

Strotz, R.H., Myopia and Inconsistency in Dynamic Utility Maximization, The Review of Economic Studies, Vol. 23, No.3, 1955, pp.165-180.

Wang, Zengwu, Xia, Jianming, Zhang, Lihong, Optimal investment for an insurer: The martingale approach, Insurance: Mathematics and Economics, Vol.40, 2007, pp.322-334.

Yaari, E.M., Some Remarks on Measures of Risk Aversion and on Their Use, Journal of Economic Theory, Vol. 1, 1969, pp.315-329.

Zhao, Y., Ziemba, W.T., Mean-Variance versus Expected Utility in Dynamic Investment Analysis. Working Paper, 2002.

Zhou, X.Y., Li, D., Continuous-Time Mean-Variance Portfolio Selection: A Stochastic LQ Framwork, Applied Mathematical Optimization, Vol. 42, 2000, pp.19-33.