

On time-consistent multi-horizon portfolio allocation

Simone Cerreia-Vioglio* Fulvio Ortu† Francesco Rotondi‡
Federico Severino§,¶

Working paper
October 30, 2020

Abstract

We analyse the problem of constructing multiple mean-variance portfolios over increasing investment horizons in stochastic interest rate markets. The traditional one-period mean-variance optimal portfolios of Hansen and Richard (1987) require the replication of two payoffs. When several maturities are considered, different payoffs have to be replicated each time, with an impact on transaction costs. Using martingale decomposition techniques and introducing a family of risk-adjusted measures linked to increasing maturities, we provide an intertemporal version of the traditional orthogonal decomposition of asset returns. This allows us to construct a multi-horizon mean-variance frontier that is time-consistent and requires the replication of solely two payoffs for all horizons under consideration. When transaction costs are taken into account, our time-consistent mean-variance frontier may outperform the traditional mean-variance optimal strategies in terms of Sharpe ratio. Some interesting examples of this fact come from long-term contracts as life annuities.

JEL Classification: G11, G12.

Keywords: return decomposition, multiple horizons, time consistency, mean-variance frontier, martingale pricing, stochastic interest rates.

*Università Bocconi, Milan, Department of Decision Sciences and IGIER: simone.cerreia@unibocconi.it.

†Università Bocconi, Milan, Department of Finance and IGIER: fulvio.ortu@unibocconi.it.

‡Università Bocconi, Milan, Department of Finance: rotondi.francesco@unibocconi.it

§Université Laval, Québec, Department of Finance, Insurance and Real Estate: federico.severino@fsa.ulaval.ca.

¶We thank Anna Battauz, Massimo Marinacci and participants at AMASES XLIII (University of Perugia, 2019) for useful comments. Federico Severino acknowledges the financial support of ERC (grant SDDM-TEA) and FRQSC (Fonds de recherche du Québec Société et culture).

1 Motivations and main results

The mean-variance approach for asset returns first rigorously formalized by the seminal work of Hansen and Richard (1987) is a cornerstone in the theory of portfolio allocation. Nevertheless, the orthogonal decomposition of returns proposed therein allows the characterization of the mean-variance frontier only at a single fixed time horizon T . Returns that lie on the mean-variance frontier at T generally do not exhibit this desirable property at any intermediate date t before T . Therefore, we propose a generalization of the traditional mean-variance approach that accounts for multi-horizon investment decisions. We also allow interest rates to be stochastic.

Specifically, we consider a continuous-time arbitrage-free market with finite horizon T , stochastic instantaneous rates and several risky securities that pay no dividends. We present all the results in a conditional setting, where we take into consideration two sources of randomness: prices of primary assets and instantaneous rates.

In order to decompose asset returns, we construct a space of forward prices, that we denote H_s^T , and a risk-adjusted measure Q^T , that we obtain from the forward measure. H_s^T is made by conditional forward martingales that can be univocally associated with returns. Moreover, this space is endowed with an inner product based on the conditional expectation of terminal values under Q^T . The overall structure is termed Hilbert module by Cerreia-Vioglio, Maccheroni, and Marinacci (2017). Interestingly, no-arbitrage prices feature an inner product representation in H_s^T , in agreement with the literature since Harrison and Kreps (1979). After decomposing the module H_s^T , in Corollary 4 we show that a return process $\{u_\tau(s)\}_{\tau \in [s, T]}$, where each $u_\tau(s)$ is the ratio of no-arbitrage prices π_τ/π_s , satisfies the orthogonal decomposition

$$u_\tau(s) = g_\tau(s) + \omega_s e_\tau(s) + n_\tau(s), \quad \tau \in [s, T]$$

in the spirit of Hansen and Richard (1987). Here $g(s)$ is the so-called *log optimal return*, $e(s)$ is the *mean excess return*, $n(s)$ is an additional zero-price return and ω_s is a random weight measurable at time s . All returns in the decomposition are (conditionally) orthogonal according to the measure Q^T . In addition, the associated mean-variance frontier in the period $[s, T]$ is made up of asset returns with null $n(s)$. A Two-fund Separation Theorem holds (Theorem 10) and so the frontier turns out to be spanned by $g(s)$ and the return $f(s)$ associated with a pure discount T -bond. In addition, it is possible to decompose returns also in any subperiod $[s, t]$ with $t < T$ by using a proper risk-adjusted measure Q^t associated with the horizon t .

The main advantage of our decompositions is time consistency. Since we decompose the whole forward martingale processes that define returns, the decompositions over different

temporal windows overlap. As a result, a time consistency property holds for our mean-variance returns: returns on the mean-variance frontier at time T are mean-variance returns at date t , too (Corollary 9). For example, a buy-and-hold one-year horizon mean-variance portfolio (according to Q^T) turns out to lie on the mean-variance frontier also at six-month horizon (according to Q^t). In fact, our mean-variance frontiers are spanned by the same two assets across a continuum of maturities. This feature is absent in the classical treatment of mean-variance portfolio selection, where second moments computed with respect to different information structures are usually incomparable.

The practical advantage implied by this feature is remarkable in terms of transaction costs, in particular replication costs. Consider, e.g., an investor that faces a multi-horizon portfolio allocation problem, namely, who has to meet N expected return targets at N subsequent maturities. She wants to achieve these expected returns by investing in N buy-and-hold portfolios while minimizing their variances. According to the standard mean-variance approach, she should solve N different optimization problems that would lead to N portfolios requiring the replication of 2 different assets each. Globally, she would need to replicate $2 \times N$ assets. On the contrary, according to our approach, she has to replicate only 2 securities as all the N portfolios she builds involve different units of the same assets (the ones with returns $g(s)$ and $f(s)$). The loss in terms of standard deviation of our investment strategies can be compensated by sizable savings on their implementation costs. Indeed, after incorporating transaction costs in the analysis, time-consistent mean-variance portfolios can display a higher Sharpe ratio than classical mean-variance portfolios. We illustrate several examples in Section 5, including life annuities.

Finally, similarly to Cochrane (2014), we provide a microeconomic foundation for our mean-variance frontier by showing that our mean-variance returns are optimal for a specific quadratic utility agent that solves a consumption-investment problem. In agreement with our theory, the arising optimal portfolio turns out to be time-consistent with respect to changes in the investment horizon.

The paper is organized as follows. Section 2 involves the theoretical instruments to study return dynamics. The section describes the no-arbitrage pricing framework, introduces the risk-adjusted measures Q^t and the module of forward prices H_g^t . We devote Section 3 to the orthogonal decomposition of asset returns at a fixed horizon while, in Section 4, we present our mean-variance frontier and its time consistency. We also formulate a Two-fund Separation Theorem and a β -pricing representation for excess returns. We devote Section 5 to several numerical simulations to illustrate time consistency and the reduction of transaction costs in different contexts. Section 6 describes the link between our mean-variance portfolios and the optimal allocation of a constrained investor with quadratic utility. Sec-

tion 7 concludes. Finally, the Appendix contains additional results and simulations, and all the proofs.

2 Framework and essentials

We describe the asset pricing framework and the essential tools for the intertemporal decomposition of returns: the risk-adjusted measures Q^t and the modules of forward prices H_s^t . We simultaneously introduce the notation of the paper.

2.1 Arbitrage-free market and numéraire changes

Fix $T > 0$ and consider the probability space (Ω, \mathcal{F}, P) , where P is the physical measure. We build a financial market on (Ω, \mathcal{F}, P) by considering two random processes $X = \{X_t\}_{t \in [0, T]}$ and $Y = \{Y_t\}_{t \in [0, T]}$. The process X is N -dimensional and consists of prices of N primary risky assets that generate the market, i.e. $X_t = [X_t^{(1)}, \dots, X_t^{(N)}]'$. On the contrary, Y is one-dimensional and represents the stochastic instantaneous interest rate. The money market account has value $e^{\int_0^t Y_\tau d\tau}$ at any time t . Moreover, pure discount bonds with any possible maturity and face value equal to 1 are traded.

At any instant t we define the vector of relative prices $Z_t = e^{-\int_0^t Y_\tau d\tau} X_t$ and consider the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$, where $\mathbb{F} = \{\mathcal{F}_t\}_{t \in [0, T]}$ is the filtration generated by the pair (Z, Y) . We assume that our price system satisfies the no-free lunch with vanishing risk (NFLVR) condition and that relative asset prices Z are semimartingale (Delbaen and Schachermayer, 1994). By the First Fundamental Theorem of Asset Pricing there exists a probability measure equivalent to P that makes Z a sigma-martingale (Delbaen and Schachermayer, 1998).¹ We assume that at least one of the sigma-martingale measures, denoted by Q , is an equivalent martingale measure, i.e. it makes Z a martingale process.

The Radon-Nikodym derivative of Q with respect to P is L_T and we define $L_t = \mathbb{E}_t[L_T]$ and $L_{t, T} = L_T/L_t$ at any time $t \in [0, T]$. As in Harrison and Kreps (1979), we assume that $e^{-\int_0^t Y_\tau d\tau} L_t$ belongs to $L^2(\mathcal{F}_t)$ for all t . We denote by $M = \{M_t\}_{t \in [0, T]}$ the strictly positive stochastic discount factor process associated with Q , i.e. $M_t = e^{-\int_0^t Y_\tau d\tau} L_t$. The related pricing kernel in the time interval $[t, T]$ is $M_{t, T} = M_T/M_t$.

We now consider a pure discount bond with maturity T and we denote its no-arbitrage price at time t by $\pi_t(1_T)$. The yield to maturity at time t is $r_t^T = -\log \pi_t(1_T)/(T - t)$ and r_T^T denotes the a.s. (finite) limit of r_t^T when t approaches T . By using as numéraire $\pi_t(1_T)$, we construct the forward measure with horizon T and we denote it by F (see Geman,

¹As explained by Émery (1980), Z turns out to be the martingale transform of some martingale, via an integrable predictable process.

El Karoui, and Rochet, 1995, for the theory of numéraire changes). This probability measure is equivalent to Q and we denote its Radon-Nikodym derivative with respect to P by G_T . Importantly, G_T belongs to $L^2(\mathcal{F}_T)$ because $e^{-\int_0^T Y_\tau d\tau} L_T$ is included in $L^2(\mathcal{F}_T)$. Moreover, we set $G_t = \mathbb{E}_t[G_T]$ for any t and we define $G_{t,T} = G_T/G_t$. See details in Appendix A.1.

Using the forward measure, the stochastic discount factor rewrites as $M_t = e^{r_t^T(T-t) - r_0^T T} G_t$ and the pricing kernel in any time interval $[s, t]$ with $s \leq t \leq T$ becomes

$$M_{s,t} = e^{r_t^T(T-t) - r_s^T(T-s)} G_{s,t}.$$

In addition, any attainable \mathcal{F}_T -measurable payoff h_T with finite $\mathbb{E}^F[|h_T|]$ has no-arbitrage price at time t given by

$$\pi_t(h_T) = e^{-r_t^T(T-t)} \mathbb{E}_t^F[h_T]. \quad (1)$$

We finally introduce some numéraire changes based on the investment horizon $t \in [0, T]$. We already defined the process $\{G_t\}_{t \in [0, T]}$ from the Radon-Nikodym derivative of the forward measure with respect to P . We take as numéraire the no-arbitrage price process of a (virtual) security generating the payoff G_t and we denote by Q^t the related *risk-adjusted measure*. Its Radon-Nikodym derivative with respect to the physical measure is

$$I_t^t = \frac{G_t^2}{\mathbb{E}[G_t^2]}.$$

For all $\tau \in [0, t]$ we define $I_\tau^t = \mathbb{E}_\tau[I_t^t]$ and $I_{\tau,t}^t = I_t^t/I_\tau^t$. Note that, by moving t , we obtain a continuum of risk-adjusted measures Q^t associated with increasing horizons. See the detailed derivation in Appendix A.1.

2.2 The Hilbert modules H_s^t and linear pricing functionals

We adopt the forward measure and we consider the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, F)$. We fix an instant $s \in [0, T]$ and develop some tools to deal with conditioning information in \mathcal{F}_s . We start with considering at any time $t \in [s, T]$ the conditional L^1 -space $L_s^1(\mathcal{F}_t) = \{f \in L^0(\mathcal{F}_t) : \mathbb{E}_s[|f|] \in L^0(\mathcal{F}_s)\}$. Cerreia-Vioglio, Kupper, Maccheroni, Marinacci, and Vogelpoth (2016) show that $L_s^1(\mathcal{F}_t)$ is an L^0 -module with the multiplicative decomposition $L_s^1(\mathcal{F}_t) = L^0(\mathcal{F}_s)L^1(\mathcal{F}_t)$.²

In our construction, we consider adapted processes that take values in $L_s^1(\mathcal{F}_t)$. An important role will be played by *conditional* (or *generalized*) *martingales*. We use this terminology for processes \hat{z} defined in the time interval $[s, t]$ with all the properties of

²Clearly, $L_s^1(\mathcal{F}_t)$ contains all functions f in $L^1(\mathcal{F}_t)$: in this case $\mathbb{E}_s[|f|] \in L^1(\mathcal{F}_s)$. In general, however, the conditional expectation is defined for random variables that are merely in $L^0(\mathcal{F}_t)$ as discussed, for instance, in Chapter II, §7 of Shiryaev (1996).

martingales except for integrability, which is replaced by the weaker condition $\mathbb{E}_s[|\hat{z}(\tau)|] \in L^0(\mathcal{F}_s)$ for all $\tau \in [s, t]$. See, for instance, Chapter VII, §1 of Shiryaev (1996). In finance a classical example of (conditional) martingale is provided by forward prices for the settlement date t (see Section 9.6 in Musiela and Rutkowski, 2005). With this property in mind, for any $t \in [s, T]$ we define the space

$$H_s^t = \left\{ \text{conditional } F\text{-martingale } \hat{z} : [s, t] \rightarrow L_s^1(\mathcal{F}_t), \mathbb{E}_s^{Q^t} [\hat{z}_t^2] \in L^0(\mathcal{F}_s) \right\},$$

H_s^t contains the forward prices of traded assets with square integrable terminal value under Q^t . Interestingly, H_s^t can be characterized in differential terms (see Proposition 2.4 in Marinacci and Severino, 2018, about weak time-derivatives). For our construction the relation between $H_s^{t_1}$ and $H_s^{t_2}$ with $t_1 \leq t_2$ is crucial: if \hat{z} belongs to $H_s^{t_2}$, then its restriction on $[s, t_1]$ belongs to $H_s^{t_1}$. Indeed, the conditional expectation of $\hat{z}_{t_1}^2$ is always defined as an extended real random variable and $\mathbb{E}_s[G_{t_1}^2 \hat{z}_{t_1}^2] \leq \mathbb{E}_s[G_{t_2}^2 \hat{z}_{t_2}^2]$.

Fixed $t \in [s, T]$, H_s^t is a pre-Hilbert module on the algebra $L^0(\mathcal{F}_s)$ when we define the outer product $\cdot : L^0(\mathcal{F}_s) \times H_s^t \rightarrow H_s^t$ and the L^0 -valued inner product $\langle \cdot, \cdot \rangle_s^t : H_s^t \times H_s^t \rightarrow L^0(\mathcal{F}_s)$ respectively by

$$a_s \cdot \hat{z} = a_s \hat{z}, \quad \langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s^{Q^t} [\hat{z}_t \hat{v}_t].$$

The homogeneity of the inner product with respect to \mathcal{F}_s -measurable variables, i.e.

$$\langle a_s \cdot \hat{z}, \hat{v} \rangle_s^t = a_s \langle \hat{z}, \hat{v} \rangle_s^t$$

for any \hat{z}, \hat{v} in H_s^t and a_s in $L^0(\mathcal{F}_s)$, is relevant for financial applications because it allows for contingent strategies in portfolio theory. Moreover, the inner product structure delivers a natural notion of orthogonality: two processes \hat{z}, \hat{v} in H_s^t are orthogonal when

$$\langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s^{Q^t} [\hat{z}_t \hat{v}_t] = 0.$$

Our inner product mimics the conditional structure of Hansen and Richard (1987), who built up a conditional asset pricing framework under the physical measure. Here we use a different probability measure and we work directly on the terminal values of forward martingale processes. These features drive our decomposition of returns and the related mean-variance frontier.

Importantly, H_s^t is a selfdual pre-Hilbert module or, more simply, a Hilbert module. Self-duality is the property that allows for an inner product representation of any L^0 -linear and bounded functional on H_s^t (see Definition 2 in Cerreia-Vioglio, Maccheroni, and Marinacci, 2017).

Proposition 1 H_s^t is a selfdual pre-Hilbert module on $L^0(\mathcal{F}_s)$.

Selfduality provides an inner product representation of linear pricing functionals, a fact which is consistent with the asset pricing literature: see Harrison and Kreps (1979), Ross (1978) and Hansen and Richard (1987) among the others.

To elucidate this point, fix $t = T$ and consider an \mathcal{F}_T -measurable payoff h_T with $\mathbb{E}_s^{Q^T}[h_T^2]$ in $L^0(\mathcal{F}_s)$. This payoff is the terminal value of the forward price process $\hat{h} = \{\hat{h}_\tau\}_{\tau \in [s, T]}$ in H_s^T defined by $\hat{h}_\tau = \mathbb{E}_\tau^F[h_T]$ at any τ . Hence, the no-arbitrage price of eq. (1) induces the L^0 -valued functional $\Pi_s : H_s^T \rightarrow L^0(\mathcal{F}_s)$ such that

$$\Pi_s : \hat{h} \mapsto e^{-r_s^T(T-s)} \hat{h}_s.$$

Π_s is a positive, L^0 -linear bounded functional and, in line with the selfduality of H_s^T , it is represented by the L^0 -valued inner product

$$\Pi_s(\hat{h}) = \left\langle \hat{z}, \hat{h} \right\rangle_s^T, \quad \hat{z}_\tau = \frac{e^{-r_s^T(T-s)} \mathbb{E}_s[G_T^2]}{G_s} \frac{1}{G_\tau}, \quad \tau \in [s, T]$$

for any $\hat{h} \in H_s^T$. Since $1/G$ belongs to H_s^T , also z does and the L^0 -valued inner product representation is well-defined. The process $1/G$ will play a fundamental role in our decomposition of excess returns. We will investigate the financial meaning of this object in Subsection 3.2.

3 Return decomposition

In this section we build the relation between returns and processes in H_s^t with $t \in [s, T]$. We, then, orthogonally decompose any H_s^t by exploiting the L^0 -valued inner product $\langle \cdot, \cdot \rangle_s^t$. As illustrated in Section 3.3 of Cerreia-Vioglio, Maccheroni, and Marinacci (2019), the decomposition of a Hilbert module needs topological conditions in order to be well-defined. Nevertheless, in case H is a selfdual L^0 -module and M is a finitely generated submodule, the decomposition $H = M \oplus M^\perp$ is well-posed (here M^\perp denotes the orthogonal complement of M in H). That is the case of our interest, because we deal with submodules generated by single returns, specifically $g(s)$ and $e(s)$ that we define in Subsections 3.2 and 3.3. Once the decomposition of modules is established in Theorem 3, we determine in Corollary 4 a risk-adjusted decomposition of asset returns. Our result parallels Hansen and Richard (1987) decomposition (which, on the contrary, exploits the physical measure) in a conditional setting with stochastic rates.

Remark. We assume that the forward measure is different from the physical one and that a security with terminal payoff $1/M_{s,T}$ is traded at any instant $t \in [s, T]$. Equivalently, there exists an admissible self-financing portfolio with value $1/M_{s,T}$ at time T .

3.1 Return definition

We consider the no-arbitrage price process induced by an attainable payoff h_T at time T such that $\mathbb{E}_s^{Q^T}[h_T^2] \in L^0(\mathcal{F}_s)$. We want to compute the return associated with h_T at any time t between s and T . This return is taken into account, for instance, by a trader that purchased a European derivative with payoff h_T at time s and is willing to sell it at time t . Formally, the return process $u(s) = \{u_\tau(s)\}_{\tau \in [s, T]}$ is the ratio of no-arbitrage prices, i.e. $u_\tau(s) = \pi_\tau(h_T)/\pi_s(h_T)$, it satisfies, for all $\tau \in [s, T]$, the relation $\mathbb{E}_s[M_{s, \tau} u_\tau(s)] = 1$ and $\mathbb{E}_s^{Q^T}[u_T^2] \in L^0(\mathcal{F}_s)$. Moreover, u is associated with the conditional martingale $\{M_{s, \tau} u_\tau(s)\}_{\tau \in [s, T]}$ whose value at time s is equal to 1. Thus, to develop our theory, we abstract from the previous example and provide the following definition of return.

Definition 2 We call return process any adapted process $u(s)$ such that the process $\hat{u}^T(s)$ defined, for all $\tau \in [s, T]$, by

$$\hat{u}_\tau^T(s) = \frac{M_{s, \tau}}{G_{s, \tau}} u_\tau(s) = \frac{\pi_s(1_T)}{\pi_\tau(1_T)} u_\tau(s) \quad (2)$$

belongs to H_s^T and $\hat{u}_s^T(s) = 1$.

Definition 2 relies on the bijection between returns and normalized conditional F -martingales (or, equivalently, forward prices) in H_s^T . Indeed, given any $\hat{u}^T(s) \in H_s^T$ with $\hat{u}_s^T(s) = 1$, by eq. (2) we can construct the process $\{u_\tau(s)\}_{\tau \in [s, T]}$ that satisfies $\mathbb{E}_s[M_{s, \tau} u_\tau(s)] = 1$ at any time τ as well as the square summability requirement that we mentioned before. Such a *return* materializes in the market when it can be replicated by an admissible self-financing portfolio of traded securities.

Example. Suppose that $\mathbb{E}_s[G_T^4]$ belongs to $L^0(\mathcal{F}_s)$ and consider a payoff at T that coincides with the pricing kernel $M_{s, T}$. This payoff is fundamental in the mean-variance decomposition provided by Hansen and Richard (1987). The related return process and the conditional martingale in H_s^T are

$$u_\tau(s) = \frac{\mathbb{E}_\tau[M_{\tau, T} M_{s, T}]}{\mathbb{E}_s[M_{s, T}^2]} = \frac{\mathbb{E}_\tau[G_T^2]}{M_{s, \tau} \mathbb{E}_s[G_T^2]}, \quad \hat{u}_\tau^T(s) = \frac{\mathbb{E}_\tau[G_T^2]}{G_{s, \tau} \mathbb{E}_s[G_T^2]}.$$

Return processes define conditional martingales also in any time subinterval $[s, t]$, with t preceding T . In fact, by defining $\hat{u}^t(s)$ as the restriction of $\hat{u}^T(s)$ on $[s, t]$, we have that $\hat{u}^t(s) \in H_s^t$ and $\hat{u}_s^t(s) = 1$.

3.2 The log optimal return $g(s)$

We define the submodule of H_s^T associated with zero-price payoffs (or excess returns)

$$\begin{aligned}\mathring{H}_s^T &= \{\hat{i}^T(s) \in H_s^T : \mathbb{E}_s[M_{s,T}\iota_T(s)] = 0\} \\ &= \{\hat{i}^T(s) \in H_s^T : \mathbb{E}_s[G_{s,T}\hat{i}_T^T(s)] = 0\} \\ &= \left\{ \hat{i}^T(s) \in H_s^T : \mathbb{E}_s^{Q^T} \left[\frac{1}{G_T} \hat{i}_T^T(s) \right] = 0 \right\} \\ &= \left\{ \hat{i}^T(s) \in H_s^T : \left\langle \frac{1}{G}, \hat{i}^T(s) \right\rangle_s^T = 0 \right\},\end{aligned}$$

where $\iota(s)$ and $\hat{i}^T(s)$ are related as above and $1/G$ belongs to H_s^T because G/G constitutes the martingale identically equal to 1. After a normalization, we define the process $\hat{g}^T(s)$ in H_s^T and the associated return process respectively by

$$\hat{g}_\tau^T(s) = \frac{G_s}{G_\tau}, \quad g_\tau(s) = \frac{1}{M_{s,\tau}}.$$

This return refers to an asset with terminal payoff $1/M_{s,T}$, which is assumed to be traded in the market. As expected, the process $1/G$ is the one that permits the inner product representation of pricing functionals described at the end of Subsection 2.2. Moreover, $1/M_{s,T}$ is the optimal terminal wealth of an investor that maximizes $\mathbb{E}[\log w_T]$ over all attainable wealth profiles (i.e. replicable payoffs) w_T , given an initial wealth equal to 1, that is $1 = \mathbb{E}[M_{s,T}w_T]$. See Chapter 20 of Björk (2004). Since the log optimal portfolio is the admissible self-financing strategy that generates the optimal wealth profile, we can refer to g as the *log optimal return*.

Observe also that any $\hat{i}^T(s) \in \mathring{H}_s^T$ satisfies $\hat{i}_s^T(s) = 0$ since $\hat{i}_s^T(s) = \mathbb{E}_s[G_{s,T}\hat{i}_T^T(s)] = 0$.

In addition, the module H_s^T orthogonally decomposes as

$$H_s^T = \text{span}_{L^0} \{\hat{g}^T(s)\} \oplus \mathring{H}_s^T.$$

When we restrict ourselves to a nearer horizon t , excess returns are associated with the submodule of H_s^t defined by

$$\mathring{H}_s^t = \{\hat{i}^t(s) \in H_s^t : \mathbb{E}_s[M_{s,t}\iota_t(s)] = 0\} = \left\{ \hat{i}^t(s) \in H_s^t : \left\langle \frac{1}{G}, \hat{i}^t(s) \right\rangle_s^t = 0 \right\}.$$

Consistently, $\hat{g}^t(s)$ is defined by the restriction of $\hat{g}^T(s)$ on $[s, t]$.

3.3 The mean excess return $e(s)$

Since the measures P and F do not coincide, G_T is different from 1. In this case, the process constantly equal to 1 in the time interval $[s, T]$ differs from $\hat{g}^T(s)$ and it belongs to

H_s^T because G is a P -martingale and $\mathbb{E}_s[G_T^2] \in L^0(\mathcal{F}_s)$. Such process is associated with the return process

$$f_\tau(s) = \frac{\pi_\tau(1_T)}{\pi_s(1_T)}$$

related to a zero-coupon bond with expiry T . This instrument can be considered as the riskless security in our stochastic rates setting. We formalize this intuition in Section 6 by solving the optimal allocation problem of a (possibly prudent) investor with quadratic utility.

Therefore, we are allowed to define $\hat{e}^T(s)$ as the orthogonal projection of 1 on the submodule \mathring{H}_s^T i.e.

$$\hat{e}^T(s) = \text{proj}_{\mathring{H}_s^T} 1,$$

meaning that $\hat{e}_s^T(s) = 0$ and $\mathbb{E}_s^{Q^T}[(1 - \hat{e}_T^T(s))\hat{i}_T^T(s)] = 0$ for all $\hat{i}^T(s)$ in \mathring{H}_s^T . Since the orthogonal projection of 1 on $\text{span}_{L^0}\{\hat{g}^T(s)\}$ is $\{G_s/G_\tau\}_{\tau \in [s, T]}$, we have $1 = \hat{e}^T(s) + G_s/G$ so that

$$\hat{e}^T(s) = 1 - \hat{g}^T(s).$$

Moreover, \mathring{H}_s^T decomposes as

$$\begin{aligned} \mathring{H}_s^T &= \text{span}_{L^0}\{\hat{e}^T(s)\} \oplus \left\{ \hat{n}^T(s) \in \mathring{H}_s^T : \mathbb{E}_s^{Q^T}[\hat{e}_T^T(s)\hat{n}_T^T(s)] = 0 \right\} \\ &= \text{span}_{L^0}\{\hat{e}^T(s)\} \oplus \left\{ \hat{n}^T(s) \in \mathring{H}_s^T : \mathbb{E}_s^{Q^T}[\hat{n}_T^T(s)] = 0 \right\} \end{aligned}$$

from the definition of $\hat{e}^T(s)$. Similarly to before, we define $e(s)$ by

$$e_\tau(s) = \frac{\pi_\tau(1_T)}{\pi_s(1_T)} \hat{e}_\tau^T(s) = f_\tau(s) - g_\tau(s),$$

which embodies the meaning of *mean excess return*.

At the shorter horizon t , we define the conditional martingale $\hat{e}^t(s)$ as

$$\hat{e}^t(s) = \text{proj}_{\mathring{H}_s^t} 1$$

with the orthogonality induced by Q^t , namely $\mathbb{E}_s^{Q^t}[(1 - \hat{e}_t^t(s))\hat{i}_t^t(s)] = 0$ for all $\hat{i}^t(s)$ in \mathring{H}_s^t . In addition, $\hat{e}^t(s)$ coincides with the restriction of $\hat{e}^T(s)$ on the time interval $[s, t]$, that is $\hat{e}^t(s) = 1 - \hat{g}^t(s)$. We show some useful properties of $\hat{g}^t(s)$ and $\hat{e}^t(s)$ in Lemma 17 in Appendix B.

3.4 Orthogonal decompositions of returns

The orthogonality in H_s^t allows us to determine a conditional decomposition of asset returns under the risk-adjusted measure Q^t . To achieve this goal, we start from the decomposition of conditional forward martingales.

Theorem 3 (Martingale decomposition) $\hat{u}^t(s)$ belongs to H_s^t and $\hat{u}_s^t(s) = 1$ if and only if there exist $\omega_s \in L^0(\mathcal{F}_s)$ and $\hat{n}^t(s) \in \hat{H}_s^t$ such that

$$\mathbb{E}_s^{Q^t} [\hat{n}_t^t(s)] = \mathbb{E}_s^{Q^t} [\hat{g}_t^t(s)\hat{n}_t^t(s)] = \mathbb{E}_s^{Q^t} [\hat{e}_t^t(s)\hat{n}_t^t(s)] = 0$$

and

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) + \hat{n}^t(s).$$

A straightforward application of Theorem 3 delivers an orthogonal decomposition of asset returns in the time window $[s, t]$, according to the risk-adjusted measure Q^t .

Corollary 4 (Return decomposition) $u(s)$ is a return in $[s, t]$ if and only if there exist $\omega_s \in L^0(\mathcal{F}_s)$ and $\hat{n}^t(s) \in \hat{H}_s^t$ such that

$$\mathbb{E}_s^{Q^t} [e^{r_t^T(T-t)}n_t(s)] = \mathbb{E}_s^{Q^t} [e^{2r_t^T(T-t)}g_t(s)n_t(s)] = \mathbb{E}_s^{Q^t} [e^{2r_t^T(T-t)}e_t(s)n_t(s)] = 0$$

with $n_\tau(s) = \hat{n}_\tau^t(s)G_{s,\tau}/M_{s,\tau}$ for all $\tau \in [s, t]$ and

$$u(s) = g(s) + \omega_s e(s) + n(s). \quad (3)$$

The proof of Theorem 3 exploits the definition of the projection coefficient ω_s in $L^0(\mathcal{F}_s)$, that turns out to be

$$\omega_s = \pi_s(1_T) \mathbb{E}_s [G_t^2] \frac{\mathbb{E}_s^{Q^t} [u_t(s)/\pi_t(1_T)] - \mathbb{E}_s^{Q^t} [g_t(s)/\pi_t(1_T)]}{\text{var}_s(G_t)}.$$

In particular, when asset returns are computed on the whole trading period $[s, T]$, we obtain the same decomposition of eq. (3) with

$$\mathbb{E}_s^{Q^T} [n_T(s)] = \mathbb{E}_s^{Q^T} [g_T(s)n_T(s)] = \mathbb{E}_s^{Q^T} [e_T(s)n_T(s)] = 0$$

and

$$\omega_s = \pi_s(1_T) \mathbb{E}_s [G_T^2] \frac{\mathbb{E}_s^{Q^T} [u_T(s)] - \mathbb{E}_s^{Q^T} [g_T(s)]}{\text{var}_s(G_T)},$$

which is reminiscent of the Sharpe ratio between $u_T(s)$ and $g_T(s)$ under the measure Q^T . In this case ω_s is uniquely determined by $\mathbb{E}_s^{Q^T} [u_T(s)]$.

4 Mean-variance returns

A simple outcome of Corollary 4 is that, among all return processes $u(s)$ in the period $[s, T]$, $g(s)$ is the one with minimum conditional second moment according to the measure Q^T . Indeed, any $u(s)$ satisfies

$$\mathbb{E}_s^{Q^T} [u_T^2(s)] = \mathbb{E}_s^{Q^T} [g_T^2(s)] + \omega_s^2 \mathbb{E}_s^{Q^T} [e_T^2(s)] + \mathbb{E}_s^{Q^T} [n_T^2(s)] \geq \mathbb{E}_s^{Q^T} [g_T^2(s)].$$

In general, our purpose is to provide a characterization of returns with minimum conditional variance, once the conditional expectation is fixed, under the measure Q^t related to the horizon into account. As it is well-known, mean-variance portfolio analysis (under the physical measure) has its roots in the seminal works by Markowitz (1952) and Tobin (1958) and had a huge development in the last decades. With respect to the existing literature, the convenience of our approach relies in the time consistency property that we describe in Subsection 4.1. Subsection 4.2 illustrates a Two-fund Separation Theorem and a β -pricing representation of excess returns.

Definition 5 *We say that a return process $u(s)$ is on the mean-variance frontier (or it is a mean-variance return) at time T when it minimizes $\text{var}_s^{Q^T}(u_T(s))$ for some given $\mathbb{E}_s^{Q^T}[u_T(s)]$ in $L^0(\mathcal{F}_s)$. In that case, we consistently say that the associated $\hat{u}^T(s)$ is a conditional mean-variance martingale in $[s, T]$. Similarly, mean-variance returns on the shortened time interval $[s, t]$ involve $\mathbb{E}_s^{Q^t}[u_t(s)/\pi_t(1_T)]$ and $\text{var}_s^{Q^t}(u_t(s)/\pi_t(1_T))$.*

We prove the results for a generic horizon t , beginning with conditional martingales.

Theorem 6 (Mean-variance martingales) *Consider $\hat{u}^t(s) \in H_s^t$ such that $\mathbb{E}_s^{Q^t}[\hat{u}_t^t(s)] = k_s$ for some $k_s \in L^0(\mathcal{F}_s)$. Among them, the conditional martingale that minimizes $\text{var}_s^{Q^t}(\hat{u}_t^t(s))$ is*

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) \quad \text{with} \quad \omega_s = \frac{k_s \mathbb{E}_s[G_t^2] - G_s^2}{\text{var}_s(G_t)}.$$

From Theorem 6 we easily deduce the characterization of mean-variance returns at t .

Corollary 7 (Mean-variance returns) *Consider returns $u(s)$ such that $\mathbb{E}_s^{Q^t}[u_t(s)/\pi_t(1_T)] = h_s$ for some $h_s \in L^0(\mathcal{F}_s)$. Among them, the return that minimizes $\text{var}_s^{Q^t}(u_t(s)/\pi_t(1_T))$ is*

$$u(s) = g(s) + \omega_s e(s) \quad \text{with} \quad \omega_s = \frac{h_s \pi_s(1_T) \mathbb{E}_s[G_t^2] - G_s^2}{\text{var}_s(G_t)}.$$

As expected, the mean-variance frontier associated with the maturity time T requires no discounting in first and second moments. In fact, among returns $u(s)$ such that $\mathbb{E}_s^{Q^T}[u_T(s)] = h_s$, the return

$$u(s) = g(s) + \omega_s e(s), \quad \omega_s = \frac{h_s \pi_s(1_T) \mathbb{E}_s[G_T^2] - G_s^2}{\text{var}_s(G_T)}.$$

minimizes $\text{var}_s^{Q^T}(u_T(s))$. As an example, consider the zero-coupon return process $f(s)$ on the period $[s, T]$, which satisfies

$$f(s) = g(s) + e(s).$$

By Corollary 7, $f(s)$ minimizes the conditional variance of any return $u(s)$ such that $\mathbb{E}_s^{Q^T}[u_T(s)] = 1/\pi_s(1_T)$.

Finally, note that at any horizon t mean-variance returns can be easily identified by their expectation under the physical measure. Indeed, if we fix $\mathbb{E}_s[u_t(s)] = \tilde{h}_s$, then the weight ω_s is univocally determined by

$$\omega_s = \frac{\tilde{h}_s - \mathbb{E}_s[g_t(s)]}{\mathbb{E}_s[e_t(s)]}. \quad (4)$$

4.1 Time consistency

A fundamental property of our approach to mean-variance portfolio analysis is time consistency. Indeed, if a return process belongs to our mean-variance frontier at date T , then it is on the mean-variance frontier at any other previous time t , too. This feature is ultimately due to the fact that the decomposition of Theorem 3 involves the whole return process in the time range $[s, T]$ and so there is a mechanical overlap with the decompositions built at shorter horizons.

The time consistency of portfolio or consumption choices is an old issue of economic theory. For example, a first distinction between precommitment and consistent planning can be retrieved in the seminal work by Strotz (1955). In addition, Mossin (1968) highlights the inconsistency of multiperiod mean-variance analysis because the quadratic utility does not satisfy the Bellman principle of optimality. These important issues are also discussed in Basak and Chabakauri (2010) and Czichowsky (2013).

In our framework the collection of risk-adjusted measures Q^t , when t varies from s to T , is the key tool for modeling the time consistency of returns. We first deal with conditional forward martingales.

Proposition 8 (Martingales time consistency) *Let $s \leq t \leq T$. If $\hat{u}^T(s)$ is a conditional mean-variance martingale in $[s, T]$, then $\hat{u}^t(s)$ is a conditional mean-variance martingale in $[s, t]$.*

We now establish time consistency of asset returns.

Corollary 9 (Returns time consistency) *Consider returns in the interval $[s, T]$. A mean-variance return at T is also a mean-variance return at any $t \in [s, T]$.*

From the standpoint of interpretation, we can set s as today and consider portfolios with maturity T of one year. Moreover, t may identify a six-month horizon from now. We build our six-month and one-year horizon mean-variance frontiers, based on the information

available today. However, Corollary 9 ensures that portfolios on the yearly frontier are also on the six-month frontier. This feature is absent in classical mean-variance analysis that exploits the decomposition of returns based on the physical measure. In fact, the standard construction does not provide any relation between the decompositions of returns at different horizons. On the contrary, the methodology that we propose relies on the decomposition of the underlying forward martingale processes and so return representations at different dates are interrelated. The practical benefit of our approach is that the arising mean-variance frontiers are generated by the same two processes across a multiplicity of horizons.

4.2 Additional properties

We state and prove a Two-fund Separation Theorem for our mean-variance frontier. Theorem 10 establishes in our setting the celebrated result by Merton (1972) and makes the implementation of the frontier easier by replacing the mean excess return $e(s)$ with the return $f(s)$ of a pure discount T -bond.

Theorem 10 (Two-fund Separation) *$u(s)$ is a mean-variance return if and only if*

$$u(s) = \alpha_s v(s) + (1 - \alpha_s) z(s)$$

for some mean-variance returns $v(s), z(s)$ and $\alpha_s \in L^0(\mathcal{F}_s)$. In particular,

$$u(s) = \alpha_s g(s) + (1 - \alpha_s) f(s)$$

where $\alpha_s = 1 - \omega_s$ and ω_s is obtained from Corollary 7.

In few words, $g(s)$ and $f(s)$ span the mean-variance frontier of returns at any horizon.

We finally provide a simple β -pricing representation for excess returns at time t . We employ $g(s)$ as benchmark return and we obtain an \mathcal{F}_s -measurable coefficient β_s .

Proposition 11 *Consider a return $u(s)$ and fix $t \in [s, T]$. Then,*

$$\mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} u_t(s) \right] - e^{r_s^T(T-s)} = \beta_s \left(\mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} g_t(s) \right] - e^{r_s^T(T-s)} \right) \quad (5)$$

where β_s in $L^0(\mathcal{F}_s)$ is

$$\beta_s = \text{cov}_s^{Q^t} \left(e^{r_t^T(T-t)} g_t(s), e^{r_t^T(T-t)} u_t(s) \right) / \text{var}_s^{Q^t} \left(e^{r_t^T(T-t)} g_t(s) \right).$$

When returns are computed on the whole time window $[s, T]$, the β -pricing representation reduces to

$$\mathbb{E}_s^{Q^T} [u_T(s) - f_T(s)] = \beta_s \mathbb{E}_s^{Q^T} [g_T(s) - f_T(s)]$$

where β_s in $L^0(\mathcal{F}_s)$ is

$$\beta_s = \frac{\text{cov}_s^{Q^T}(g_T(s), u_T(s))}{\text{var}_s^{Q^T}(g_T(s))}.$$

This coefficient is the one implied by the Two-Fund Separation Theorem about returns on the mean-variance frontier at the terminal date.

Appendix A.2 illustrates additional results about the mean-variance representation provided by Corollary 7. In particular, Proposition 15 describes the relation between the conditional and the unconditional version of the mean-variance frontier. Indeed, when unconditional second moments of returns are finite, unconditional mean-variance returns belong to the conditional frontier, too.

As for excess returns, one may also be interested in the shape of their mean-variance frontier per se. Proposition 16 shows that the mean-variance frontier of excess returns at time T is exactly spanned by $e(s)$. See again Appendix A.2.

5 Simulations: mean-variance optimization, time consistency and transaction costs

In this section we consider a multiperiod mean-variance portfolio problem in the time interval $[s, T]$. Interest rates are stochastic and only buy-and-hold investment strategies set at time s are allowed. Our investor may be thought as a manager or a company that aims at building portfolios with a target mean across a sequence of maturities t_1, t_2, \dots, t_N with $s < t_1 < t_2 < \dots < t_N = T$. Each of this portfolios must be optimal in terms of a mean-variance criterion in the time window under consideration. The need of designing such a term structure of portfolios may come from intertemporal hedging reasons due, e.g. to cashflow management or medium-term production plans. The asset allocation across multiple horizons is decided ex ante because of costly, or even forbidden, rebalancing. A detailed example in the context of life annuities is provided in Subsection 5.3.

Several dynamic extensions of mean-variance optimization have been proposed in the literature. Remarkable examples are given by Li and Ng (2000), Zhou and Li (2000) and Leippold, Trojani, and Vanini (2004) among the others. Several time periods are considered and the target return is usually achieved through the selection of suitable self-financing portfolios. However, differently from our approach, no investment target at intermediate horizons is considered.

In our problem, the investor builds a portfolio with return

$$\sum_{i=1}^N \lambda_s^{(i)} u_t^{(i)}(s)$$

at any $t \in [s, T]$, where each $\lambda_s^{(i)} \in L^0(\mathcal{F}_s)$ is the weight of the subportfolio i , i.e. the one with return $u^{(i)}(s)$, in the overall portfolio. Moreover, each return $u^{(i)}(s)$ solves

$$\min \text{var}_s \left(u_{t_i}^{(i)}(s) \right) \quad \text{sub} \quad \mathbb{E}_s \left[u_{t_i}^{(i)}(s) \right] = h_s^{(i)}$$

with $h_s^{(i)} \in L^0(\mathcal{F}_s)$ given, for $i = 1, \dots, N$. The weights $\lambda_s^{(i)}$ are positive, they sum up to 1 and, in case the overall portfolio is equally-weighted, $\lambda_s^{(i)} = 1/N$ for all i .

The unique solution to this optimization problem is achieved by subportfolios on the classical mean-variance frontier of Hansen and Richard (1987): at each date t_i

$$u_{t_i}^{(i)}(s) = \frac{M_{s,t_i}}{\mathbb{E}_s [M_{s,t_i}^2]} + \tilde{w}_s^{(i)} \left(1 - e^{-r_s^{t_i}(t_i-s)} \frac{M_{s,t_i}}{\mathbb{E}_s [M_{s,t_i}^2]} \right), \quad \tilde{w}_s^{(i)} \in L^0(\mathcal{F}_s).$$

By employing the return of zero-coupon bonds with expiry t_i , the Two-fund Separation Theorem permits to rewrite the classical mean-variance frontier as

$$u_{t_i}^{(i)}(s) = \tilde{\alpha}_s^{(i)} \frac{M_{s,t_i}}{\mathbb{E}_s [M_{s,t_i}^2]} + \left(1 - \tilde{\alpha}_s^{(i)} \right) f_{t_i}(s)$$

with $\tilde{\alpha}_s^{(i)} = 1 - \pi_s(1_{t_i})\tilde{w}_s^{(i)}$.³

At each maturity t_i , the initial implementation of a subportfolio that delivers the return $u^{(i)}(s)$ requires the replication, by self-financing portfolio strategies, of two payoffs at t_i : one coincides with the pricing kernel M_{s,t_i} , the other is the terminal value of a zero-coupon t_i -bond. Considering the whole sequence of maturities in the problem, $2 \times N$ payoffs need to be replicated in order to implement the optimal asset allocation (or, at least, N payoffs if all the needed pure discount bonds are traded in the market). As a result, depending on the severity of market incompleteness, the optimal solution may be unattainable. The requirement of replicating $2 \times N$ payoffs via existing securities may be too demanding in general. This fact worsen the practical implementation of mean-variance efficient portfolios, which already suffers from notorious pitfalls, as described by Michaud (1989), Best and Grauer (1991) and DeMiguel, Garlappi, and Uppal (2007), among the others. Indeed, the standard mean-variance portfolio weights are often unstable and sensitive to small changes in the estimate of returns moments. The practical implementation can, then, lead to suboptimal investments.

Hereby, we propose an alternative strategy by exploiting our time-consistent mean-variance frontier. Although theoretically suboptimal, our frontier requires solely the replication of two payoffs at T , whatever the number N of maturities involved (or only one

³Indeed, such pure-discount bonds belong to the frontier since

$$f_{t_i}(s) = \frac{M_{s,t_i}}{\mathbb{E}_s [M_{s,t_i}^2]} + \frac{1}{\pi_s(1_{t_i})} \left(1 - e^{-r_s^{t_i}(t_i-s)} \frac{M_{s,t_i}}{\mathbb{E}_s [M_{s,t_i}^2]} \right).$$

payoff if the zero-coupon T -bond is traded in the market). When asset replication is costly or difficult, this feature constitutes a sizable advantage, that may compensate the loss of mean-variance optimality with respect to the (possibly unattainable) classical solution. In particular, we consider subportfolios with returns

$$v_t^{(i)}(s) = g_t(s) + \omega_s^{(i)} e_t(s), \quad t \in [s, T],$$

where each $\omega_s^{(i)}$ is chosen so that the conditional expectation of $v_{t_i}^{(i)}(s)$ meets the target $h_s^{(i)}$ as in eq. (4). In line with Theorem 10, we build our subportfolios by exploiting the return $g(s)$ of the log optimal portfolio and the return $f(s)$ of a zero-coupon T -bond. These two financial instruments are employed at any intermediate maturity t_i , as a consequence of time consistency. We finally compare the performance of the two families of subportfolios $u^{(i)}(s)$ and $v^{(i)}(s)$ by considering the transaction costs and their impact on the Sharpe ratios.

Specifically, to quantify the advantage of using solely two securities in the time-consistent strategy, we assume that transaction costs are present. Such commissions are composed by trading and replication costs, similarly to Irle and Sass (2006).

Trading costs are constant for every asset unit and apply to both short and long positions. Their total amount is *proportional* to traded volumes. In our simulations, the implementation of each classical mean-variance subportfolio i generates the trading costs $c(|\tilde{\alpha}_s^{(i)}| + |1 - \tilde{\alpha}_s^{(i)}|)$ with $c > 0$. The analogous expression with $\alpha_s^{(i)}$ delivers the trading costs of the time-consistent mean-variance return $v^{(i)}(s)$.

As for the replication costs, we assume that the design of the replication strategies for $g_T(s)$ and $M_{s,t_i}/\mathbb{E}_s[M_{s,t_i}^2]$ at any maturity t_i entails a positive *fixed* cost C for any (possibly linearly independent) security. Therefore, the implementation of each classical mean-variance subportfolio i requires the additional expenditure of C . On the contrary, if we proportionally spread the replication cost of $g_T(s)$ across the horizons t_1, \dots, t_N , each time-consistent subportfolio i needs to bear the cost $\lambda_s^{(i)} C$. We attribute no cost to the possible replication of the pure discount bonds used in the two frontiers. As a result, each mean-variance optimal subportfolio i and each time-consistent subportfolio i have commissions, respectively,

$$C + c \left(\left| \tilde{\alpha}_s^{(i)} \right| + \left| 1 - \tilde{\alpha}_s^{(i)} \right| \right) \quad \text{and} \quad \lambda_s^{(i)} C + c \left(\left| \alpha_s^{(i)} \right| + \left| 1 - \alpha_s^{(i)} \right| \right). \quad (6)$$

Accordingly, the overall portfolios have commissions:

$$CN + c \sum_{i=1}^N \lambda_s^{(i)} \left(\left| \tilde{\alpha}_s^{(i)} \right| + \left| 1 - \tilde{\alpha}_s^{(i)} \right| \right) \quad \text{and} \quad C + c \sum_{i=1}^N \lambda_s^{(i)} \left(\left| \alpha_s^{(i)} \right| + \left| 1 - \alpha_s^{(i)} \right| \right).$$

In terms of risk/return trade-off, at any maturity t_i we consider a *modified Sharpe ratio* given by the difference of the Sharpe ratio and the ratio between transaction costs and

standard deviation. In this way, the expected return of each subportfolio i is reduced by the proper commissions of eq. (6):

$$\text{modified Sharpe ratio} = \text{Sharpe ratio} - \frac{\text{transaction costs}}{\text{standard deviation}}.$$

The modified Sharpe ratio can be negative even if the Sharpe ratio is positive. Interestingly, the modified Sharpe ratios can reverse the relations between the Sharpe ratios of the classical and the time-consistent mean-variance optimal strategies, making the time-consistent approach valuable. This happens in the simulations of Subsections 5.2 and 5.3. Subsection 5.1 precisely describes the market in which we set the simulations and Subsection 5.4 concludes with a comparison of measures Q , F and Q^T in the market under scrutiny.

5.1 Reference market

As in Appendix B of Brigo and Mercurio (2006), we assume that short-term rates move as in Vasicek (1977) model in the time interval $[s, T]$ with positive parameters k, θ, σ . Then, we consider a stock price X that follows a geometric Brownian motion with volatility $\eta > 0$, correlated with interest rates shocks. The instantaneous correlation between the two underlying Wiener processes is ϕ . We orthogonalize the two sources of randomness and consider, without loss of generality, the dynamics

$$\begin{cases} dX_t = X_t Y_t dt + \eta X_t \left[\phi dW_t^Q + \sqrt{1 - \phi^2} dZ_t^Q \right] \\ dY_t = k(\theta - Y_t) dt + \sigma dW_t^Q, \end{cases}$$

where W^Q and Z^Q are independent Wiener processes. A money market account with dynamics $dB_t = Y_t B_t dt$ is also present. A more general model with two risky stocks is illustrated in Appendix C.1.

Yields to maturity are affine, i.e. $r_t^T(T - t) = -A(t, T) + B(t, T)Y_t$, with

$$A(t, T) = \left(\theta - \frac{\sigma^2}{2k^2} \right) (B(t, T) - T + t) - \frac{\sigma^2}{4k} B^2(t, T)$$

and $B(t, T) = (1 - e^{-k(T-t)})/k$. The pure discount T -bond price at time t is function of t and Y_t , obtained from Itô's formula. Hence, beyond the money market account, the assets that generate the market are

$$\begin{cases} dX_t = X_t Y_t dt + \eta X_t \left[\phi dW_t^Q + \sqrt{1 - \phi^2} dZ_t^Q \right] \\ d\pi_t(1_T) = \pi_t(1_T) Y_t dt - \pi_t(1_T) B(t, T) \sigma dW_t^Q. \end{cases} \quad (7)$$

At the same time, under the physical measure,

$$\begin{cases} dX_t = X_t \mu_t^X dt + \eta X_t \left[\phi dW_t^P + \sqrt{1 - \phi^2} dZ_t^P \right] \\ d\pi_t(1_T) = \pi_t(1_T) \mu_t^P dt - \pi_t(1_T) B(t, T) \sigma dW_t^P, \end{cases}$$

where μ^X and μ^P are adapted processes. They are related to the drifts under Q via the bivariate process of *market price of risk* $[\nu^W, \nu^Z]'$ such that

$$\begin{bmatrix} dW_t^Q \\ dZ_t^Q \end{bmatrix} = \begin{bmatrix} \nu_t^W \\ \nu_t^Z \end{bmatrix} dt + \begin{bmatrix} dW_t^P \\ dZ_t^P \end{bmatrix}.$$

Specifically,

$$\begin{bmatrix} \eta\phi & \eta\sqrt{1-\phi^2} \\ -B(t, T)\sigma & 0 \end{bmatrix} \begin{bmatrix} \nu_t^W \\ \nu_t^Z \end{bmatrix} = \begin{bmatrix} \mu_t^X - Y_t \\ \mu_t^P - Y_t \end{bmatrix}$$

so that

$$\nu_t^W = -\frac{\mu_t^P - Y_t}{B(t, T)\sigma}, \quad \nu_t^Z = \frac{\mu_t^X - Y_t - \eta\phi\nu_t^W}{\eta\sqrt{1-\phi^2}}. \quad (8)$$

The Radon-Nikodym derivative of Q with respect to P is, at any $t \in [0, T]$,

$$L_t = e^{-\frac{1}{2} \int_0^t [(\nu_\tau^W)^2 + (\nu_\tau^Z)^2] d\tau - \int_0^t \nu_\tau^W dW_\tau^P - \int_0^t \nu_\tau^Z dZ_\tau^P}$$

and we assume that the Novikov condition is satisfied, that is $\mathbb{E}[e^{\frac{1}{2} \int_0^T [(\nu_t^W)^2 + (\nu_t^Z)^2] dt}]$ is finite. Moreover, we postulate that $\mu_t^P = (1 - \xi B(t, T)\sigma) Y_t$ for some $\xi > 0$ so that $\nu_t^W = \xi Y_t$, in line with the usual approach of Vasicek short-term rates. Finally, the dynamics of the pricing kernel are given by

$$dM_{s,t} = -Y_t M_{s,t} dt - \nu_t^W M_{s,t} dW_t^P - \nu_t^Z M_{s,t} dZ_t^P.$$

The parameters that we use in the simulations of the interest rate process are $k = 1$, $\theta = 0.05$ and $\sigma = 0.01$ with initial value $Y_0 = 0.02$, on a monthly time grid. Moreover, we set $\eta = 0.1$ and $\phi = 0.1$, and we assume that the drift of the stock price under the physical measure is $\mu_t^X = Y_t + 0.05$.

5.2 A six-horizon mean-variance optimization

In this set of simulations, we fix $s = 0$ and consider an equally-weighted portfolio over six maturities: $N = 6$ and $\lambda_s^{(i)} = 1/N$ for all $i = 1, \dots, 6$. We employ a monthly time grid and horizons t_1, \dots, t_6 associated with six subsequent semesters. We set the target means equal to $h_0^{(i)} = 1.06$ for $i = 1, \dots, 6$. In other words, we are assuming that the investor wants to obtain a 6% flat return at the end of each of six subsequent semesters. She plans to do so by investing in 6 equally weighted buy-and-hold subportfolios built at $s = 0$. We further assume that the cashflows obtained at the end of each semester from the liquidation of the related subportfolio are not re-invested. Implementation and transaction costs of the strategies are discussed later on.

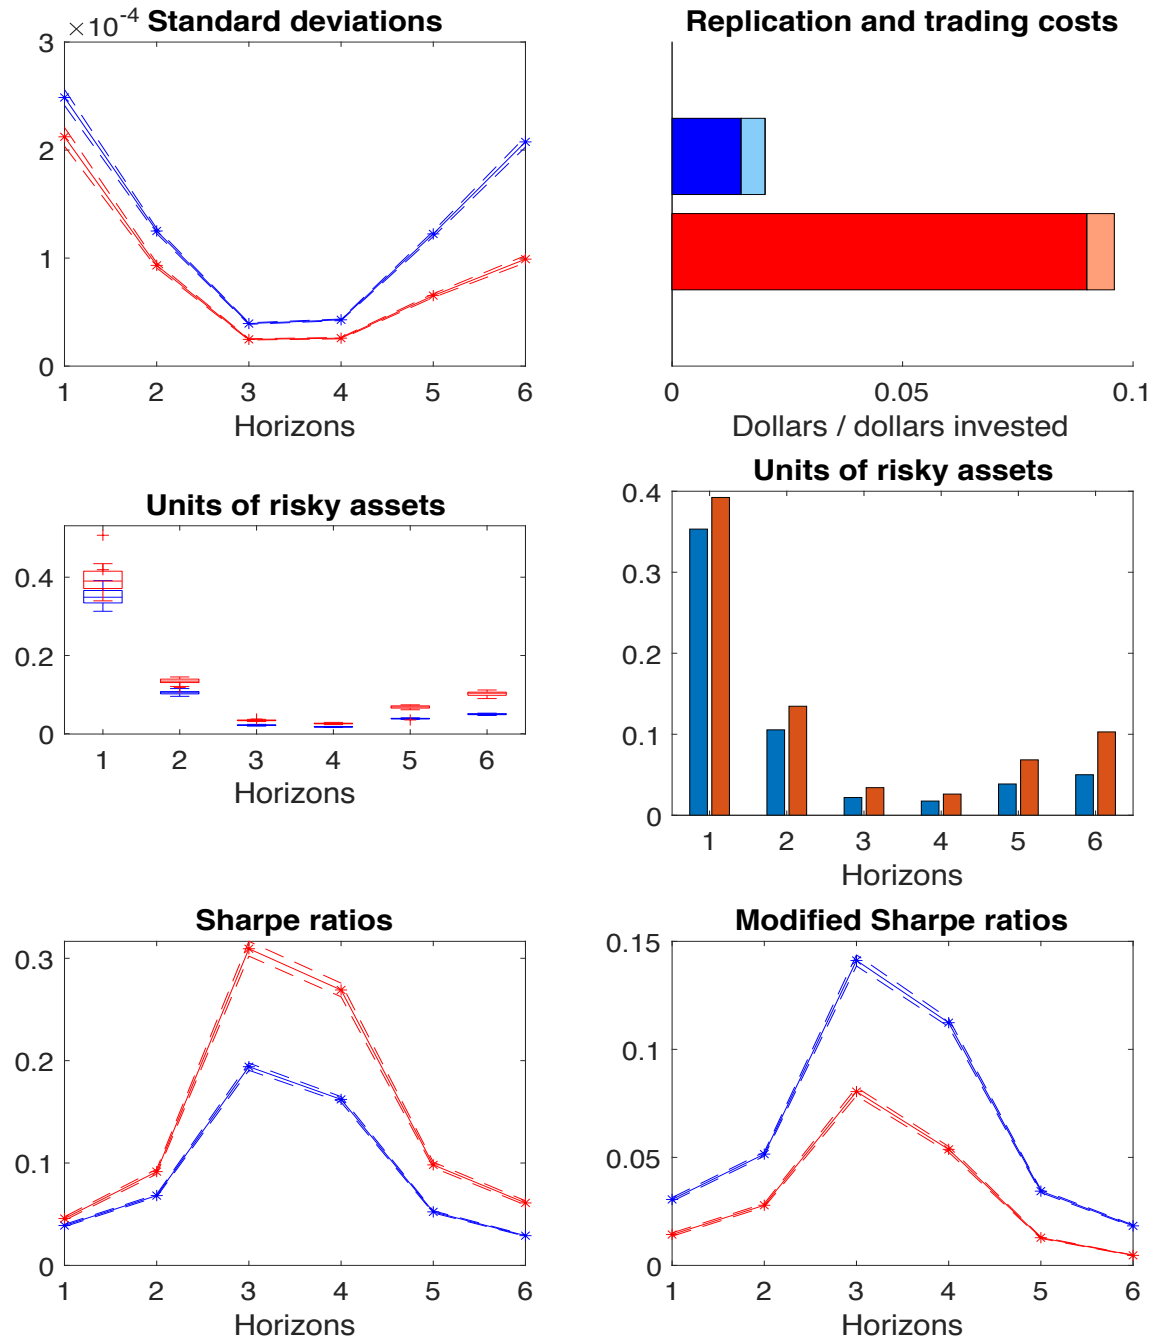


Figure 1: Red (resp. blue) lines, bars and boxes refer to the classical (resp. time-consistent) mean-variance solution for the problem of Subsection 5.2. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda_0^{(i)}$ for all $i = 1, \dots, 6$. 90% confidence intervals for these variables are represented. The top-right panel represents the transaction costs of the time-consistent portfolio (blue for replication costs, light blue for trading costs) and of the mean-variance portfolio (red for replication costs, light red for trading costs). Medium panels contain the box-and-whisker plot at 25th and 75th percentiles and the bar plot of loadings $|\alpha_s^{(i)}|$ and $|\tilde{\alpha}_s^{(i)}|$ at all horizons.

We simulate both the classical and the time-consistent multiperiod portfolios described above. We, then, repeat the exercise by employing, in total, 30 different seeds for the initial Gaussian random sampling to obtain a sample of averages and standard deviations of each subportfolio i with return $u^{(i)}(0)$ and horizon t_i , for $i = 1, \dots, 6$. Sharpe ratios are computed by using as reference risk-free securities pure discount bonds at increasing maturities. Results are summarized in Figure 5.2, where standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda_s^{(i)} = 1/N$. Every simulated subportfolio matches perfectly the target means $h_0^{(i)}$ at the proper horizon $i = 1, \dots, 6$. As predicted by the theory, classical mean-variance subportfolios display lower standard deviations with respect to our time-consistent asset allocations, whose advantage relies on a parsimonious implementation.

In our simulations the loadings of the time-consistent subportfolios are smaller than the ones of the classical mean-variance strategies, requiring to buy/sell fewer assets. We visualize this fact in the medium panels of Figure 5.2, where we plot the absolute values of $\alpha_s^{(i)}$ and $\tilde{\alpha}_s^{(i)}$ at each maturity t_i . The graphs depict the units of risky assets - i.e. the ones associated with $g_T(s)$ and $M_{s,t_i}/\mathbb{E}_s[M_{s,t_i}^2]$ respectively - contained in each subportfolio. The exposure to the risky securities is higher at maturities near in time. However, at any horizon, the loadings in the time-consistent subportfolios are lower than the ones in the classical subportfolios (with slightly lower dispersion). Consequently, the implementation of the portfolio with returns $v^{(i)}(s)$ involves narrower long (or short) positions, both in $g(s)$ and in $f(s)$, a valuable feature in case of short-selling constraints.

The medium panels of Figure 5.2 give also an idea of the magnitude of the transaction costs of both portfolios that we summarize in the top-right panel by setting $c = \$0.005$ and $C = \$0.015$. Under this assumption, by considering an initial investment of \$100, total transaction costs roughly amount to \$10 if the investor builds the portfolios according to the standard mean-variance frontier, and to \$2 if the investor exploits our time-consistent frontier.

The commission shrinkage of the time-consistent approach impacts the risk/return trade-off between the two strategies, as we can see in the bottom panels of Figure 5.2. Indeed, after including the transaction costs, the modified Sharpe ratio indicates that the time-consistent solution is the best performing. The excess standard deviation of the time-consistent portfolio is fully compensated by its reduced transaction costs (in particular, replication costs), as captured by the modified Sharpe ratio. This phenomenon depends on the model parameters (interest rate dynamics, stock price movements, number of maturities, target returns...) and, in general, it is relevant when a high number of maturities is taken into account.

5.3 A life annuity application

Still in the market of Subsection 5.1, we compare the time-consistent and the classical mean-variance approaches in the context of a life annuity.

Consider a life annuity payed with a lump sum at date 0 by a cohort of subscribers (see e.g. Chapter 5 in Bower, Gerber, Hickman, Jones, and Nesbitt, 1997). The annuity provides yearly payments to each subscriber until she dies. The insurance company invests the received capital in N subportfolios with increasing maturities that allow to meet the future payments. For example, we can assume that each subportfolio has target return $h_0^{(i)} = 1.05$ for $i = 1, \dots, N$ with $N = 20$ years.

The random variable *time-until-death* captures the difference between the insureds age at death and her age at subscription. It gives an idea of the potential length of the life annuity. We suppose that the cumulative distribution of time-until-death is $\mathcal{P}(t_i) = 1 - e^{-\gamma t_i^3}$ defined on the years $t_i = i$ for $i = 1, 2, \dots, 20$. This specification ensures a unimodal distribution with a peak at around ten years if we set $\gamma = 0.001$. Importantly, the weight of each subportfolio i depends on the proportion of survivors at the maturity-year t_i , i.e.

$$\lambda_0^{(i)} = \frac{1 - \mathcal{P}(t_i)}{\sum_{i=1}^{20} (1 - \mathcal{P}(t_i))}.$$

If the company aims at reducing the risk of each subportfolio, it can consider a (classical or time-consistent) mean-variance approach for each return $u_t^{(i)}(0)$ satisfying $\mathbb{E}[u_{t_i}^{(i)}(0)] = 1.05$ for $i = 1, \dots, 20$.

Similarly to Subsection 5.2, we scale standard deviations, Sharpe ratios and modified Sharpe ratios in the two approaches by the weights $\lambda_0^{(i)}$ for $i = 1, \dots, 20$. In so doing, we account for the amount of surviving subscribers at each horizon. As to transaction costs, we set $c = \$0.003$ and $C = \$0.006$. Results are summarized in Figure 5.3.

In the top-left panel of the figure, the excess standard deviation of time-consistent portfolios is more evident at intermediate horizons and vanishes when maturities approach 20 years, in agreement with the scaling induced by the time-until-death. The top-right panel highlights the difference in transaction costs between the two frontiers. The convenience of the time-consistent approach comes from the replication of one risky payoff instead of the $N = 20$ payoffs required by the classical mean-variance optimal strategies. The commission shrinkage affects the portfolio performance, as we can note from the Sharpe ratios and the modified Sharpe ratios in the bottom panels. Without considering the transaction costs, the standard mean-variance approach outperforms the optimal time-consistent strategy. Nevertheless, the introduction of the commissions reverses the conclusion: the optimal mean-variance portfolio turns out to have a lower (and sometimes negative) modified Sharpe

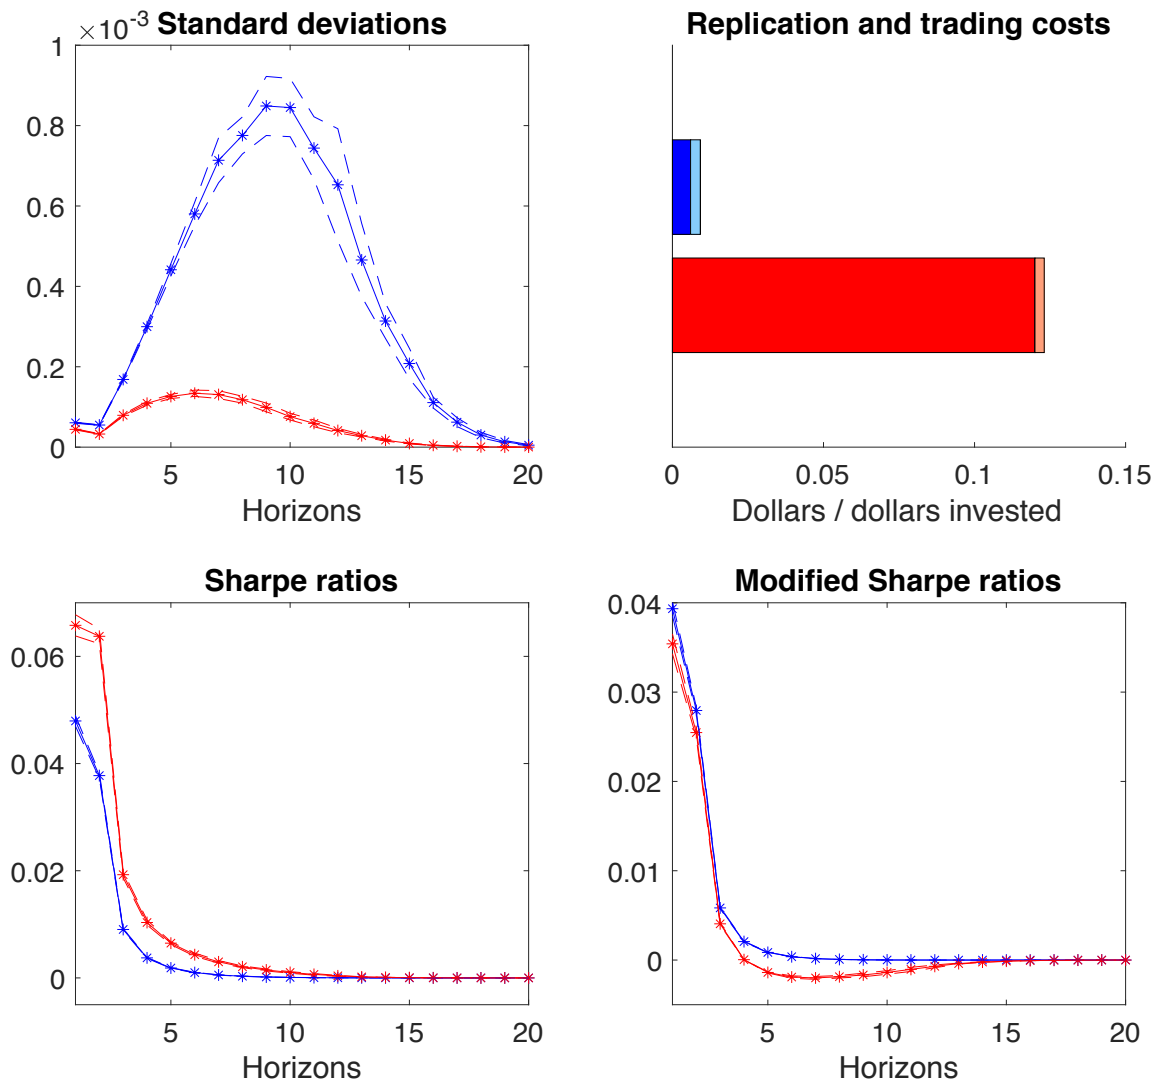


Figure 2: Red (resp. blue) lines, bars and boxes refer to the classical (resp. time-consistent) mean-variance solution for the life-annuity problem. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda_0^{(i)}$ for all $i = 1, \dots, 20$. 90% confidence intervals for these variables are represented. The top-right panel represents the transaction costs of the time-consistent portfolio (blue for replication costs, light blue for trading costs) and of the mean-variance portfolio (red for replication costs, light red for trading costs).

ratio. This effect is mostly due to the number of payments in the life annuity contract, which requires the replication of many risky securities.

5.4 Measure comparison

An important aspect of our theory is the comparison among the measures that we employ. We exploit the reference market of Subsection 5.1 to quantify the relations of Q , F and Q^T with respect to the physical measure in a practical context.

To measure the discrepancy between probabilities we use the Kullback-Leibler distance. This metric, also called *divergence* or *relative entropy*, is widely employed in the theory of finance. See, for instance, Frittelli (2000), Hansen and Sargent (2001), Alvarez and Jermann (2005) and Maccheroni, Marinacci, and Rustichini (2006) among the others. We assume that P is uniformly distributed on the interval $[0, 1]$ and we plot in Figure 3 the Kullback-Leibler distance between each of Q , F , Q^T and P over increasing horizons up to $T = 36$. All the divergences are increasing in the maturities under consideration, namely, the farther the horizon, the larger the distance between P and the three alternative measures considered. In particular, the (narrow) relative entropy between the risk-neutral and the physical measure is associated with the magnitude of the market reward-to-risk ratios. Furthermore, the forward measure turns out to be rather similar to P , whereas the risk-adjusted measure departs significantly from the other measures. Indeed, this is reflected also by the sizable differences between our portfolio weights and the ones derived through the traditional frontier.

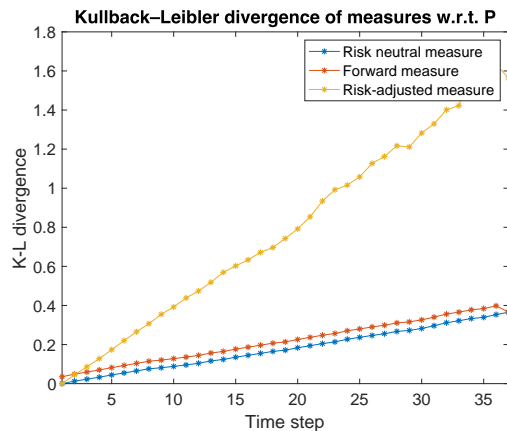


Figure 3: Kullback-Leibler distance of the measures Q , F and Q^T from the physical measure across the monthly time grid $t = 0, \dots, 36$.

To further assess the distances between these four probability measures, we look at the moments of the same random variable under the different measures. We first take into

account a random variable, say Θ , uniformly distributed on $[0,1]$ under P . We build a sample of Θ by applying the inverse transform method to the normally distributed realizations employed to generate the Gaussian increments of W^P . The uniform probability of each *state* is the inverse of the sample size. By using the Radon-Nikodym densities L_T , G_T and I_T^T , we compute the distribution of Θ under the risk-neutral measure, the forward measure and the risk-adjusted measure Q^T respectively. Then, we plot in the left panel of Figure 4 the estimated cumulative distribution functions of Θ with respect to the different probability measures. In addition, we estimate mean, variance, skewness and kurtosis under the different measures in Table 1.⁴ In particular, the four distributions differ slightly only with respect to higher moments and we notice that the risk-adjusted measure Q^T is the measure that mostly departs from P .

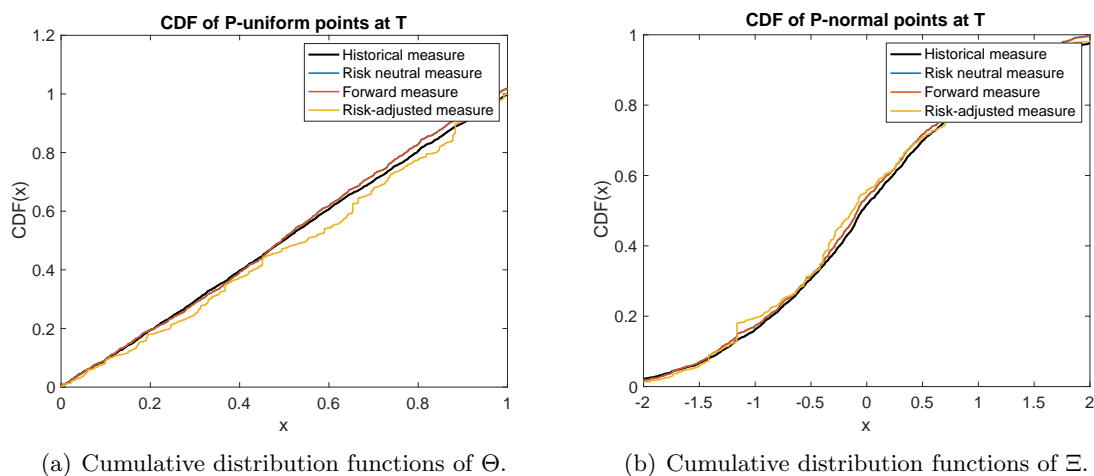


Figure 4: Estimated cumulative distribution functions of the random variables Θ and Ξ (on the left and on the right panel respectively) with respect to the measures P , Q , F and Q^T . Under P , Θ has uniform distribution on $[0, 1]$ while Ξ is a standard Normal.

⁴The benchmark values for the four moments are 0 for the mean and the skewness, $1/12 = 0.083$ for the variance and $9/5 = 1.8$ for the kurtosis.

| | Mean of Θ | Variance of Θ | Skewness of Θ | Kurtosis of Θ |
|-------|------------------|----------------------|----------------------|----------------------|
| P | 0.4979 | 0.0836 | 0.0186 | 1.8151 |
| Q | 0.5050 | 0.0814 | -0.0695 | 1.9993 |
| F | 0.5041 | 0.0817 | -0.0602 | 1.9830 |
| Q^T | 0.4828 | 0.0859 | 0.1196 | 1.7629 |

Table 1: First four centered moments estimated from the simulated sample of Θ with respect to the measures P , Q , F and Q^T .

| | Mean of Ξ | Variance of Ξ | Skewness of Ξ | Kurtosis of Ξ |
|-------|---------------|-------------------|-------------------|-------------------|
| P | -0.0069 | 1.0088 | -0.02087 | 2.9686 |
| Q | -0.0160 | 1.0364 | 0.01236 | 2.8009 |
| F | -0.0165 | 1.0352 | 0.01227 | 2.8053 |
| Q^T | -0.0487 | 0.9987 | 0.09260 | 2.7417 |

Table 2: First four centered moments estimated from the simulated sample of Ξ with respect to the measures P , Q , F and Q^T .

Finally, we repeat the analysis by considering a standard Gaussian variable Ξ under P . A sample for Ξ is already provided by the unitary increments of the Wiener process W^P . We depict the cumulative distributions in the right panel of Figure 4 and we record the estimated centered moments in Table 2. Results are qualitatively unchanged and we notice that, again, the moments under Q^T are the most different from the benchmark.

6 Mean-variance frontier and optimal investment

We provide a microeconomic foundation of the mean-variance frontier of returns described by Corollary 7. Similarly to Cochrane (2014) we show that optimal investments from date s to date T produce return processes that lie on our mean-variance frontier. In particular, such returns turn out to be a linear combination of returns $g(s)$ and $f(s)$ in agreement with Theorem 10. Moreover, an analogue of time consistency property of mean variance returns can be retrieved in optimal investment policies.

6.1 Optimal investment problem

We consider the optimization problem of an investor that decides her consumption policy $c = \{c_\tau\}_{\tau \in [s, T]}$. She is endowed with a positive initial wealth w_s in $L^0(\mathcal{F}_s)$ and receives an

exogenous income stream $i = \{i_\tau\}_{\tau \in [s, T]}$. The agent invests her initial wealth by selecting a payoff stream with value $w = \{w_\tau\}_{\tau \in [s, T]}$ and, at any instant τ , she consumes $c_\tau = i_\tau + w_\tau$. All processes are adapted. To make the investment affordable, w_s is required to satisfy the budget constraint

$$w_s = \mathbb{E}_s \left[\int_s^T M_{s,\tau} w_\tau d\tau \right].$$

The agent has an instantaneous quadratic utility

$$U(c_\tau) = -\frac{1}{2} (b_\tau - M_{s,\tau} c_\tau)^2,$$

where the process $b = \{b_\tau\}_{\tau \in [s, T]}$ defines a time-varying bliss point. Moreover, the investor deflates her consumption c_τ by exploiting the pricing kernel $M_{s,\tau}$. The intertemporal optimization problem to solve is

$$\max_c \mathbb{E}_s \left[\int_s^T U(c_\tau) d\tau \right] \quad \text{sub } w_s = \mathbb{E}_s \left[\int_s^T M_{s,\tau} w_\tau d\tau \right], \quad c_\tau = i_\tau + w_\tau.$$

The related reduced form is

$$\max_w \mathbb{E}_s \left[\int_s^T U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } w_s = \mathbb{E}_s \left[\int_s^T M_{s,\tau} w_\tau d\tau \right]. \quad (9)$$

Proposition 12 *If in Problem (9) the income stream is null and, for all $\tau \in [s, T]$, the bliss point is*

$$b_\tau = \frac{b_s \pi_\tau(1_T)}{T - s} M_{s,\tau}, \quad b_s \in L^0(\mathcal{F}_s),$$

then the optimal payoff stream defines the mean-variance return

$$\frac{(T - s)w_\tau^*}{w_s} = \frac{b_s \pi_s(1_T)}{w_s} f_\tau(s) + \left(1 - \frac{b_s \pi_s(1_T)}{w_s} \right) g_\tau(s).$$

Given some consumption c_τ , the parameter of relative risk aversion in the quadratic utility introduced above is

$$\gamma = -\frac{c_\tau U''(c_\tau)}{U'(c_\tau)} = \frac{c_\tau}{b_\tau - c_\tau}.$$

Under the assumptions of Proposition 12 about income and bliss points,

$$\frac{1}{\gamma} = \frac{b_s \pi_\tau(1_T) - c_\tau}{c_\tau}.$$

Therefore, in the optimal solution of the investment problem, by interpreting $w_s \pi_\tau(1_T) / \pi_s(1_T)$ as a special consumption, say \tilde{c}_τ , we have

$$1 - \frac{b_s \pi_s(1_T)}{w_s} = -\frac{b_s \pi_\tau(1_T) - \tilde{c}_\tau}{\tilde{c}_\tau}.$$

Moreover, by defining $\tilde{\gamma}$ through

$$\frac{1}{\tilde{\gamma}} = \frac{b_s \pi_\tau (1_T) - \tilde{c}_\tau}{\tilde{c}_\tau},$$

we can rewrite the return process associated with the optimal investment of Problem (9) as

$$\frac{(T-s)w_\tau^*}{w_s} = f_\tau(s) + \frac{1}{\tilde{\gamma}} (f_\tau(s) - g_\tau(s)).$$

Hence, the optimal portfolio is split within a unitary amount of zero-coupon bonds with maturity T and a risky allocation that depends on the agent's attitude toward risk. Optimal portfolio riskiness decreases as relative risk aversion rises. As expected, an infinitely risk-averse agent is willing to bear just the risks originated by the floating rates. Proposition 12 delves more deeply and specifies that the optimal choice of an infinitely risk-averse person is to buy only T -bonds. For instance, $f(s)$ is preferred to the return of any rolling strategy made by short-term bonds, that may constitute a safer investment in principle. This result justifies the use of zero-coupon T -bonds as the analogous of risk-free assets in a stochastic interest rates environment: they are optimal for infinitely risk-averse agents with quadratic utility over the investment period $[s, T]$.

The solidity for this approach relies on Wachter (2003) which shows that, in a multiperiod complete market, the utility of an infinitely risk-averse agent is maximized by a portfolio of bonds with the same maturity as the investment horizon, even if interest rates are random. In fact, Wachter (2003) provides a formalization of the *preferred habitat theory* that dates back to Modigliani and Sutch (1966). The conjecture of Modigliani and Sutch (1966) is, indeed, that a risk-averse person chooses assets with the same maturity as the end of the investment period, i.e. assets in the so-called maturity habitat.⁵

Therefore, under the assumptions of Proposition 12 the optimal consumption stream of an infinitely risk-averse investor coincide with the optimal wealth, that is

$$c_\tau = w_\tau = (T-s)w_s f_\tau(s), \quad \tau \in [s, T],$$

and it constitutes a very smooth consumption policy in a stochastic rates setting.

6.2 Time consistency of optimal cashflows

Inspired by the time consistency of our mean-variance frontier claimed in Corollary 9, we investigate whether a similar feature is kept in the optimal portfolio problem. Specifically, once Problem (9) is solved by a payoff stream $w^* = \{w_\tau^*\}_{\tau \in [s, T]}$ on the time interval $[s, T]$,

⁵Short- or long-living securities turn out to be inconvenient because they produce extra risks and investment costs that need to be compensated in some ways.

we assess whether the restriction of w^* is optimal on the subperiod $[s, t]$, too. In particular, we consider the problem

$$\max_w \mathbb{E}_s \left[\int_s^t U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } \tilde{w}_s = \mathbb{E}_s \left[\int_s^t M_{s,\tau} w_\tau d\tau \right], \quad (10)$$

where \tilde{w}_s is a given initial wealth in $L^0(\mathcal{F}_s)$.

Proposition 13 *Under the assumptions of Proposition 12, if w^* solves Problem (9) with initial wealth w_s , then it also solves Problem (10) with initial wealth*

$$\tilde{w}_s = \frac{t-s}{T-s} w_s.$$

The mean-variance return which is optimal on the investment period $[s, T]$ is still optimal on the subperiod $[s, t]$ for the same investor with a smaller initial endowment. The intuition behind the lower initial wealth is that the fraction $(t-s)/(T-s)$ of w_s is employed for buying the cashflow w^* on $[s, t]$. The remaining portion, namely $(T-t)/(T-s)$, is left for the last subinterval $[t, T]$. The nonlinear dependence of the optimal return from the initial endowment is actually a well-known issue for quadratic investment problems. See, for instance, Mossin (1968) for a deeper discussion.

An analogous reasoning of Proposition 13 shows that w^* is optimal also on the terminal subperiod $[t, T]$, according to

$$\max_w \mathbb{E}_s \left[\int_t^T U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } \hat{w}_s = \mathbb{E}_s \left[\int_t^T M_{s,\tau} w_\tau d\tau \right], \quad (11)$$

where \hat{w}_s belongs to $L^0(\mathcal{F}_s)$. Indeed, the following result holds.

Corollary 14 *Under the assumptions of Proposition 12, if w^* solves Problem (9) with initial wealth w_s , then it also solves Problem (11) with*

$$\hat{w}_s = \frac{T-t}{T-s} w_s.$$

Although Problem (11) involves the time window $[t, T]$, the conditional expectation in the objective function and in the budget constraint is taken at the previous date s . The pricing kernel is based on s as well. Accordingly, \hat{w}_s is \mathcal{F}_s -measurable and it represents the portion of initial wealth assigned to the final subperiod. The time consistency of w^* that we show requires, in fact, the same information set. This approach is known as *precommitment* in the language of Strotz (1955).

In general, if the decision were contingent at time t , a more profitable optimal investment would arise in the final time frame. Hence, our construction is consistent with a rational

inattention approach, as described in Sims (2003). Indeed, we implicitly assume that our investor makes a decision at time s for the whole period $[s, T]$ because she has a limited ability to process the incoming information at time t . In other words, observing the portfolio value at t may be costly and transaction costs may discourage changes in the investment policy. A more recent theory for optimal infrequent adjustments in portfolio selection is developed in Abel, Eberly, and Panageas (2013).

7 Conclusions

We obtain a conditional orthogonal decomposition of asset returns in the spirit of Hansen and Richard (1987) by employing a family of risk-adjusted measures derived from the forward measure. In addition, the associated mean-variance frontier features an important time consistency property, with practical advantages for multiperiod portfolio optimization in terms of replication costs. The whole construction lies within the linear pricing paradigm and it is consistent with the consumption-investment plan of an agent that maximizes a quadratic utility.

Introducing further specific dynamics of interest rates, beyond Vasicek model, may constitute an interesting avenue for future research. Such dynamics may convey special shapes of the mean-variance frontier that could improve the applicability of our construction in specific contexts.

References

- A.B. Abel, J.C. Eberly, and S. Panageas. Optimal inattention to the stock market with information costs and transactions costs. *Econometrica*, 81(4):1455–1481, 2013.
- F. Alvarez and U.J. Jermann. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016, 2005.
- S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23(8):2970–3016, 2010.
- M.J. Best and R.R. Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *Review of Financial Studies*, 4(2):315–342, 1991.
- T. Björk. *Arbitrage theory in continuous time*. Oxford University Press, 2004.
- N.L. Bower, H.U. Gerber, J.C. Hickman, D.A. Jones, and C.J. Nesbitt. *Actuarial mathematics*. The Society of Actuaries, Schaumburg, Illinois, 1997.

- D. Brigo and F. Mercurio. *Interest rate models-theory and practice: with smile, inflation and credit*. Springer-Verlag, Berlin Heidelberg, 2 edition, 2006.
- S. Cerreia-Vioglio, M. Kupper, F. Maccheroni, M. Marinacci, and N. Vogelpoth. Conditional L_p -spaces and the duality of modules over f -algebras. *Journal of Mathematical Analysis and Applications*, 444(2):1045–1070, 2016.
- S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci. Hilbert A-modules. *Journal of Mathematical Analysis and Applications*, 2017.
- S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci. Orthogonal decompositions in Hilbert A-modules. *Journal of Mathematical Analysis and Applications*, 470(2):846–875, 2019.
- J.H. Cochrane. *Asset pricing (Revised edition)*. Princeton University Press, Princeton, 2005.
- J.H. Cochrane. A mean-variance benchmark for intertemporal portfolio theory. *Journal of Finance*, 69(1):1–49, 2014.
- C. Czichowsky. Time-consistent mean-variance portfolio selection in discrete and continuous time. *Finance and Stochastics*, 17(2):227–271, 2013.
- F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1):463–520, 1994.
- F. Delbaen and W. Schachermayer. The fundamental theorem of asset pricing for unbounded stochastic processes. *Mathematische Annalen*, 312(2):215–250, 1998.
- V. DeMiguel, L. Garlappi, and R. Uppal. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953, 2007.
- M. Émery. Compensation de Processus VF non localement intégrables. In *Séminaire de Probabilités XIV 1978/79*, pages 152–160. Springer, Berlin, 1980.
- M. Frittelli. The minimal entropy martingale measure and the valuation problem in incomplete markets. *Mathematical Finance*, 10(1):39–52, 2000.
- H. Geman, N. El Karoui, and J-C Rochet. Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied probability*, 32(02):443–458, 1995.
- L.P. Hansen and S.F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, 1987.

- L.P. Hansen and T.J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.
- J.M. Harrison and D.M. Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic theory*, 20(3):381–408, 1979.
- A. Irle and J. Sass. Optimal portfolio policies under fixed and proportional transaction costs. *Advances in Applied Probability*, 38(4):916–942, 2006.
- M. Leippold, F. Trojani, and P. Vanini. A geometric approach to multiperiod mean variance optimization of assets and liabilities. *Journal of Economic Dynamics and Control*, 28(6):1079–1113, 2004.
- D. Li and W-L. Ng. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10(3):387–406, 2000.
- F. Maccheroni, M. Marinacci, and A. Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498, 2006.
- M. Marinacci and F. Severino. Weak time-derivatives and no-arbitrage pricing. *Finance and Stochastics*, 22(4):1007–1036, 2018.
- H. Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.
- R.C. Merton. An analytic derivation of the efficient portfolio frontier. *Journal of Financial and Quantitative Analysis*, 7(04):1851–1872, 1972.
- R.O. Michaud. The Markowitz optimization enigma: Is optimized optimal? *ICFA Continuing Education Series*, 1989(4):43–54, 1989.
- F. Modigliani and R. Sutch. Innovations in interest rate policy. *American Economic Review*, 56(1/2):178–197, 1966.
- J. Mossin. Optimal multiperiod portfolio policies. *Journal of Business*, 41(2):215–229, 1968.
- M. Musiela and M. Rutkowski. *Martingale methods in financial modelling*. Springer-Verlag, Berlin Heidelberg, 2 edition, 2005.
- S.A. Ross. A simple approach to the valuation of risky streams. *Journal of Business*, 51(3):453–475, 1978.
- A.N. Shiryaev. *Probability*. Springer-Verlag, New York, 2 edition, 1996.

- C.A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3): 665–690, 2003.
- R.H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1955.
- J. Tobin. Liquidity preference as behavior towards risk. *Review of Economic Studies*, 25(2):65–86, 1958.
- O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.
- J.A. Wachter. Risk aversion and allocation to long-term bonds. *Journal of Economic Theory*, 112(2):325–333, 2003.
- X.Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42(1):19–33, 2000.

A Complements of the theory

A.1 Forward measure and numéraire changes

The T -forward measure F is constructed by employing as numéraire the no-arbitrage price of a zero-coupon bond with maturity T . F is equivalent to the risk-neutral measure Q and its Radon-Nikodym derivative with respect to Q is

$$F_T = \frac{e^{-\int_0^T Y_\tau d\tau}}{\mathbb{E} \left[L_T e^{-\int_0^T Y_\tau d\tau} \right]} = e^{r_0^T T - \int_0^T Y_\tau d\tau}.$$

Moreover,

$$F_t = \mathbb{E}_t [L_{t,T} F_T] = e^{r_0^T T - r_t^T (T-t) - \int_0^t Y_\tau d\tau}$$

and we set $F_{t,T} = F_T/F_t$. The Radon-Nikodym derivative of F with respect to P is $G_T = F_T L_T$, which belongs to $L^2(\mathcal{F}_T)$. From $F_t = \mathbb{E}_t [L_{t,T} F_T]$, we have

$$G_t = \mathbb{E}_t [G_T] = \mathbb{E}_t [L_T F_T] = L_t F_t$$

and we define $G_{t,T} = G_T/G_t$.

If rates of interest are constant over time, $F_t = 1$ for any t so that $G_{t,T} = L_{t,T}$.

We now introduce some numéraire changes based on the investment horizon into consideration. Considering the maturity T , Q^T is the risk-adjusted measure associated with the no-arbitrage price process of a security generating the payoff G_T , employed as numéraire. Its Radon-Nikodym derivative with respect to F equals $G_T/\mathbb{E}[G_T^2]$ and, at previous times τ , it becomes

$$\frac{\mathbb{E}_\tau [G_T^2]}{G_\tau \mathbb{E} [G_T^2]}.$$

We denote by I_T^T the Radon-Nikodym density of Q^T with respect to P and by I_τ^T its conditional expectation at any previous time τ . They are obtained via multiplication by G_T and G_τ respectively:

$$I_T^T = \frac{G_T^2}{\mathbb{E} [G_T^2]}, \quad I_\tau^T = \frac{\mathbb{E}_\tau [G_T^2]}{\mathbb{E} [G_T^2]}.$$

Moreover, we define $I_{\tau,T}^T = I_T^T/I_\tau^T$ and so, for any random variable V , $\mathbb{E}_\tau^{Q^T} [V] = \mathbb{E}_\tau [I_{\tau,T}^T V]$.

In a similar fashion, on the restricted time frame $[s, t]$ with $t \leq T$, we define the measure Q^t . Its Radon-Nikodym derivative I_t^t with respect to P and its conditional values I_τ^t are

$$I_t^t = \frac{G_t^2}{\mathbb{E} [G_t^2]}, \quad I_\tau^t = \frac{\mathbb{E}_\tau [G_t^2]}{\mathbb{E} [G_t^2]}.$$

In addition, we define $I_{\tau,t}^t = I_t^t/I_\tau^t$ and so $\mathbb{E}_\tau^{Q^t} [V] = \mathbb{E}_\tau [I_{\tau,t}^t V]$ for any V as before.

A.2 Additional properties of the mean-variance frontier

The following result involves the relation between the conditional and the unconditional version of the mean-variance frontier. Indeed, unconditional mean-variance returns belong to the conditional frontier, too.

Proposition 15 *Consider returns $u(s)$ such that $\mathbb{E}^{Q^T}[u_T^2(s)] < +\infty$ and $\mathbb{E}^{Q^T}[u_T(s)] = h$ for some $h \in \mathbb{R}$. If, among them, $z(s)$ minimizes $\text{var}^{Q^T}(u_T(s))$, then $z(s)$ minimizes also $\text{var}_s^{Q^T}(u_T(s))$ among all $u(s)$ with $\mathbb{E}_s^{Q^T}[u_T(s)] = \mathbb{E}_s^{Q^T}[z_T(s)]$.*

Proof of Proposition 15. Suppose that there exists a return $u(s)$ such that $\mathbb{E}_s^{Q^T}[u_T(s)] = \mathbb{E}_s^{Q^T}[z_T(s)]$ and $\text{var}_s^{Q^T}(u_T(s)) < \text{var}_s^{Q^T}(z_T(s))$ a.s. Moreover, $\mathbb{E}^{Q^T}[u_T(s)] = \mathbb{E}^{Q^T}[z_T(s)] = h$ and

$$\mathbb{E}_s^{Q^T}[u_T^2(s)] - \left(\mathbb{E}_s^{Q^T}[u_T(s)]\right)^2 < \mathbb{E}_s^{Q^T}[z_T^2(s)] - \left(\mathbb{E}_s^{Q^T}[z_T(s)]\right)^2$$

so that $\mathbb{E}_s^{Q^T}[u_T^2(s)] < \mathbb{E}_s^{Q^T}[z_T^2(s)]$. This implies that $\mathbb{E}^{Q^T}[u_T^2(s)] < \mathbb{E}^{Q^T}[z_T^2(s)]$ and so

$$\begin{aligned} \text{var}^{Q^T}(u_T(s)) &= \mathbb{E}^{Q^T}[u_T^2(s)] - \left(\mathbb{E}^{Q^T}[u_T(s)]\right)^2 < \mathbb{E}^{Q^T}[z_T^2(s)] - \left(\mathbb{E}^{Q^T}[z_T(s)]\right)^2 \\ &= \text{var}^{Q^T}(z_T(s)), \end{aligned}$$

which is a contradiction. ■

The converse implication of Proposition 15 does not hold, as highlighted by Cochrane (2005). However, an equivalent claim to conditional mean-variance returns can be derived by employing unconditional *scaled returns*, or payoffs of managed portfolios. Rephrasing Chapter 8 of Cochrane (2005) in our framework, a scaled return is obtained by multiplying a return process $u(s)$ by some (fixed) weight α_s in $L^0(\mathcal{F}_s)$. The orthogonal decomposition of $\alpha_s u(s)$ is immediate once the decomposition of $u(s)$ is achieved in H_s^T . In addition, scaled returns allow to write an unconditional formulation of the conditional mean-variance frontier. By considering returns $u(s)$ such that $\mathbb{E}_s^{Q^T}[u_T(s)] = h_s$ as in Corollary 7 and denoting by $z(s)$ the mean-variance return in $[s, T]$, we have

$$\text{var}_s^{Q^T}(z_T(s)) \leq \text{var}_s^{Q^T}(u_T(s)), \quad \mathbb{E}_s^{Q^T}[u_T(s)] = h_s$$

if and only if

$$\text{var}^{Q^T}(\alpha_s z_T(s)) \leq \text{var}^{Q^T}(\alpha_s u_T(s)), \quad \mathbb{E}_s^{Q^T}[u_T(s)] = h_s \quad \forall \alpha_s \in L^0(\mathcal{F}_s).$$

The former equivalence is due to the definition of conditional expectation. Conditioning may be dropped when managed portfolios are considered. However, the last formulation does not refer to the unconditional mean-variance frontier because the knowledge of $\mathbb{E}^{Q^T}[u_T(s)]$ is not sufficient for the equivalence.

We now focus on the mean-variance frontier of excess returns. Specifically, the mean-variance frontier of excess returns at time T is spanned by $e(s)$.

Proposition 16 Consider excess returns $\iota(s)$ such that $\mathbb{E}_s^{Q^T}[\iota_T(s)] = h_s$ for some $h_s \in L^0(\mathcal{F}_s)$. Among them, the excess return that minimizes $\text{var}_s^{Q^T}(\iota_T(s))$ is

$$\iota(s) = w_s e(s), \quad w_s = \frac{h_s \pi_s(1_T) \mathbb{E}_s[G_T^2]}{\text{var}_s(G_T)}.$$

Proof of Proposition 16. Any excess return $\iota(s)$ is associated with a process $\hat{\iota}(s) \in \mathring{H}_s^T$ defined by $\hat{\iota}_\tau^T(s) = \iota_\tau(s) M_{s,\tau} / G_{s,\tau}$. Since

$$\mathring{H}_s^T = \text{span}_{L^0} \{ \hat{e}^T(s) \} \oplus \left\{ \hat{n}^T(s) \in \mathring{H}_s^T : \mathbb{E}_s^{Q^T}[\hat{\iota}_T^T(s) \hat{n}_T^T(s)] = 0 \right\},$$

$\hat{\iota}^T(s)$ decomposes as

$$\hat{\iota}^T(s) = w_s \hat{e}^T(s) + \hat{n}^T(s), \quad w_s = \frac{e^{-r_s^T(T-s)} h_s \mathbb{E}_s[G_T^2]}{\text{var}_s(G_T)}.$$

The proper expression of w_s is derived as in the proof of Theorem 3. As a result, the excess return $\iota(s)$ decomposes as $\iota(s) = w_s e(s) + n(s)$ with $\mathbb{E}_s^{Q^T}[n_T(s)] = \mathbb{E}_s^{Q^T}[e_T(s)n_T(s)] = 0$.

In addition, $\text{var}_s^{Q^T}(\iota_T(s)) = \mathbb{E}_s^{Q^T}[\iota_T^2(s)] - h_s^2$ for all excess returns $\iota(s)$ into consideration and

$$\begin{aligned} \mathbb{E}_s^{Q^T}[\iota_T^2(s)] &= w_s^2 \mathbb{E}_s^{Q^T}[e_T^2(s)] + \mathbb{E}_s^{Q^T}[n_T^2(s)] + 2w_s \mathbb{E}_s^{Q^T}[e_T(s)n_T(s)] \\ &= w_s^2 \mathbb{E}_s^{Q^T}[e_T^2(s)] + \mathbb{E}_s^{Q^T}[n_T^2(s)] \geq w_s^2 \mathbb{E}_s^{Q^T}[e_T^2(s)]. \end{aligned}$$

As a consequence, $\text{var}_s^{Q^T}(\iota_T(s)) \geq \text{var}_s^{Q^T}(w_s e_T(s))$, i.e. $\text{var}_s^{Q^T}(\iota_T(s))$ is minimized by the excess return characterized by $n(s) = 0$. ■

B Proofs

Lemma 17 (i) $\mathbb{E}_s[G_t \hat{g}_t^t(s)] = G_s$.

(ii) $\langle \hat{g}^t(s), \hat{u}^t(s) \rangle_s^t = \langle \hat{g}^t(s), \hat{g}^t(s) \rangle_s^t = G_s^2 / \mathbb{E}_s[G_t^2]$ for any $\hat{u}(s) \in H_s^t$ such that $\hat{u}_s^t(s) = 1$.

(iii) $\langle \hat{g}^t(s), \hat{g}^t(s) \rangle_s^t / \mathbb{E}_s[G_t \hat{g}_t^t(s)] = G_s / \mathbb{E}_s[G_t^2]$.

(iv) $\mathbb{E}_s^{Q^t}[(\hat{e}_t^t(s))^2] = \mathbb{E}_s^{Q^t}[\hat{e}_t^t(s)] = \text{var}_s(G_t) / \mathbb{E}_s[G_t^2]$.

Proof of Lemma 17.

(i)

$$\mathbb{E}_s[G_t \hat{g}_t^t(s)] = \mathbb{E}_s[G_t \hat{g}_t^T(s)] = \mathbb{E}_s \left[G_t \frac{G_s}{G_t} \right] = G_s.$$

(ii) All $\hat{u}^t(s) \in H_s^t$ such that $\hat{u}_s(s) = 1$ satisfy

$$\mathbb{E}_s[G_t^2] \langle \hat{g}^t(s), \hat{u}^t(s) \rangle_s^t = \mathbb{E}_s[G_t^2] \mathbb{E}_s^{Q^t}[\hat{g}_t^t(s) \hat{u}_t^t(s)] = G_s \mathbb{E}_s[G_t \hat{u}_t^t(s)] = G_s^2.$$

The same result holds when $\hat{u}^t(s)$ is chosen to be $\hat{g}^t(s)$.

(iii) It follows from (i) and (ii).

(iv) Since $\hat{e}^t(s) = \text{proj}_{H_s^t} 1$, for any $\hat{t}^t(s) \in H_s^t$, we have $\mathbb{E}^{Q^t} [(1 - \hat{e}_t^t(s)) \hat{t}^t(s)] = 0$. Then, the first equality follows when $\hat{t}^t(s) = \hat{e}^t(s)$. As for the second one,

$$\begin{aligned} \mathbb{E}_s [G_t^2] \mathbb{E}_s^{Q^t} [\hat{e}_t^t(s)] &= \mathbb{E}_s [G_t^2 \hat{e}_t^t(s)] = \mathbb{E}_s [G_t^2 \hat{e}_t^T(s)] \\ &= \mathbb{E}_s [G_t^2 (1 - \hat{g}_t^T(s))] = \mathbb{E}_s [G_t^2] - \mathbb{E}_s \left[G_t^2 \frac{1}{G_t} \right] G_s \\ &= \mathbb{E}_s [G_t^2] - G_s^2 = \mathbb{E}_s [G_t^2] - (\mathbb{E}_s [G_t])^2. \end{aligned}$$

■

Proof of Proposition 1

The algebra $L^0(\mathcal{F}_s)$ is endowed with the pointwise sum and product between random variables. The outer product $\cdot : L^0(\mathcal{F}_s) \times H_s^t \rightarrow H_s^t$ is well-defined because, for any $a_s \in L^0(\mathcal{F}_s)$ and $\hat{z} \in H_s^t$, $a_s \hat{z}$ belongs to H_s^t too.

Moreover, for each $a_s, b_s \in L^0(\mathcal{F}_s)$ and $\hat{z}, \hat{v} \in H_s^t$ the following properties hold.

- (1) $a_s \cdot (\hat{z} + \hat{v}) = a_s \cdot \hat{z} + a_s \cdot \hat{v}$.
- (2) $(a_s + b_s) \cdot \hat{z} = a_s \cdot \hat{z} + b_s \cdot \hat{z}$.
- (3) $a_s \cdot (b_s \cdot \hat{z}) = (a_s b_s) \cdot \hat{z}$.
- (4) If e_s denotes the \mathcal{F}_s -measurable random variable equal to one, $e_s \cdot \hat{z} = \hat{z}$.

These features make H_s^t a module over $L^0(\mathcal{F}_s)$.

Now consider the inner product $\langle \cdot, \cdot \rangle_s^t : H_s^t \times H_s^t \rightarrow L^0(\mathcal{F}_s)$. For all $\hat{z} \in H_s^t$, $\mathbb{E}_s^{Q^t} [\hat{z}_t^2] \in L^0(\mathcal{F}_s)$. Therefore, by Footnote 3 in Hansen and Richard (1987), $\langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s^{Q^t} [\hat{z}_t \hat{v}_t]$ belongs to $L^0(\mathcal{F}_s)$.

In addition, for each $a_s \in L^0(\mathcal{F}_s)$ and $\hat{z}, \hat{v}, \hat{w} \in H_s^t$ the following properties are satisfied.

- (5) $\langle \hat{z}, \hat{z} \rangle_s^t = \mathbb{E}_s^{Q^t} [\hat{z}_t^2] = \mathbb{E}_s [G_t^2 \hat{z}_t^2] / \mathbb{E}_s [G_t^2] \geq 0$ with equality if and only if $G_t \hat{z}_t = 0$. This implies that, for any $\tau \in [s, t]$, $\mathbb{E}_\tau [G_t \hat{z}_t] = G_\tau \hat{z}_\tau = 0$. As a result, $\hat{z} = 0$.
- (6) $\langle \hat{z}, \hat{v} \rangle_s^t = \langle \hat{v}, \hat{z} \rangle_s^t$.
- (7) $\langle \hat{z} + \hat{v}, \hat{w} \rangle_s^t = \langle \hat{z}, \hat{w} \rangle_s^t + \langle \hat{v}, \hat{w} \rangle_s^t$.
- (8) $\langle a_s \cdot \hat{z}, \hat{v} \rangle_s^t = a_s \mathbb{E}_s^{Q^t} [\hat{z}_t \hat{v}_t] = a_s \langle \hat{z}, \hat{v} \rangle_s^t$.

As a result, H_s^t is a pre-Hilbert module.

We now prove that H_s^t is selfdual. First, note that $L^0(\mathcal{F}_s)$ is endowed with the Lévy metric $d(f, g) = \mathbb{E}^F[\min\{|f - g|, 1\}]$ for all $f, g \in L^0(\mathcal{F}_s)$. As described in Cerreia-Vioglio, Maccheroni, and Marinacci (2017), in a pre-Hilbert L^0 -module a metric, denoted by d_H ,

is given by the composition of d with the L^0 -valued norm induced by the L^0 -valued inner product. Hence, the d_H distance between two processes u, v in H_s is

$$d_H(\hat{z}, \hat{v}) = d\left(\sqrt{\langle \hat{z} - \hat{v}, \hat{z} - \hat{v} \rangle_s^t}, 0\right) = \mathbb{E}^F \left[\min \left\{ \sqrt{\mathbb{E}_s^{Q^t} \left[(\hat{z}_t - \hat{v}_t)^2 \right]}, 1 \right\} \right].$$

Since the selfduality of a pre-Hilbert L^0 -module is equivalent to the d_H -completeness (see Theorem 5 in Cerreia-Vioglio, Maccheroni, and Marinacci, 2017), we establish this property in H_s^t . In addition, we observe that the metric d_H actually involves just terminal values \hat{z}_t and \hat{v}_t and so $d_H(\hat{z}, \hat{v})$ actually coincides with the distance between random variables \hat{z}_t, \hat{v}_t belonging to the L^0 -module $L_s^2(\mathcal{F}_t, Q^t) = \{f \in L^0(\mathcal{F}_t) : \mathbb{E}_s^{Q^t}[f^2] \in L^0(\mathcal{F}_s)\}$, which is complete: see Theorem 7 in Cerreia-Vioglio, Kupper, Maccheroni, Marinacci, and Vogelpoth (2016). This fact makes d_H -completeness of H_s^t straightforward.

Therefore, consider a Cauchy sequence $\{\hat{z}^{(n)}\}_{n \in \mathbb{N}} \subset H_s^t$: for all $\varepsilon > 0$ there is $N_\varepsilon \in \mathbb{N}$ such that, for all $n, m > N_\varepsilon$,

$$d_H\left(\hat{z}^{(n)}, \hat{z}^{(m)}\right) = \mathbb{E}^F \left[\min \left\{ \sqrt{\mathbb{E}_s^{Q^t} \left[\left(\hat{z}_t^{(n)} - \hat{z}_t^{(m)}\right)^2 \right]}, 1 \right\} \right] < \varepsilon.$$

Thus, we obtain a Cauchy sequence $\{\hat{z}_t^{(n)}\}_{n \in \mathbb{N}} \subset L_s^2(\mathcal{F}_t, Q^t)$, which is complete. As a result, this sequence has limit $\hat{z}_t \in L_s^2(\mathcal{F}_t, Q^t)$. From \hat{z}_t we define the process $\hat{z} = \{\hat{z}_\tau\}_{\tau \in [s, t]}$ by setting $\hat{z}_\tau = \mathbb{E}^F[\hat{z}_t]$. This process is a conditional F -martingale and belongs to H_s^t . To assess this fact, we check that $\mathbb{E}_s^F[|\hat{z}_\tau|] \in L^0(\mathcal{F}_s)$ for all τ .

Since any $|\hat{z}_\tau|$ is non-negative, its conditional expectation is always defined as an extended real random variable. Moreover, the conditional Cauchy-Schwartz' inequality guarantees that

$$\mathbb{E}_s^F[|\hat{z}_\tau|] \leq \mathbb{E}_s^F[|\hat{z}_t|] = \frac{\mathbb{E}_s[G_t |\hat{z}_t|]}{G_s} \leq \frac{\sqrt{\mathbb{E}_s[G_t^2 \hat{z}_t^2]}}{G_s} = \frac{\sqrt{\mathbb{E}_s^{Q^t}[\hat{z}_t^2] \mathbb{E}_s[G_t^2]}}{G_s},$$

where the last quantity belongs to $L^0(\mathcal{F}_s)$. Consequently, $\mathbb{E}_s^F[|\hat{z}_\tau|] \in L^0(\mathcal{F}_s)$ for all $\tau \in [s, t]$. We, then, determined a process $\hat{z} \in H_s^t$ such that

$$d_H\left(\hat{z}^{(n)}, \hat{z}\right) = \mathbb{E}^F \left[\min \left\{ \sqrt{\mathbb{E}_s^{Q^t} \left[\left(\hat{z}_t^{(n)} - \hat{z}_t\right)^2 \right]}, 1 \right\} \right]$$

is arbitrarily small. Since $\hat{z}^{(n)}$ goes to \hat{z} in d_H , H_s^t is d_H -complete and so selfdual.

Proof of Theorem 3

Let $\hat{u}^t(s)$ be defined by the relation $\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) + \hat{n}^t(s)$ with $\omega_s \in L^0(\mathcal{F}_s)$ and $\hat{n}^t(s) \in \hat{H}_s^t$. The process $\hat{u}^t(s) \in H_s^t$ because it is a linear combination of three processes in H_s^t . Moreover,

$$\hat{u}_s^t(s) = \hat{g}_s^t(s) + \omega_s \hat{e}_s^t(s) + \hat{n}_s^t(s) = 1 + 0 + 0 = 1$$

since $\hat{e}^t(s)$ and $\hat{n}^t(s)$ are included in \mathring{H}_s^t .

Conversely, consider any process $\hat{u}^t(s)$ in H_s^t with $\hat{u}_s^t(s) = 1$. Note that $\hat{u}^t(s) - \hat{g}^t(s)$ belongs to H_s^t and in particular to \mathring{H}_s^t because

$$\mathbb{E}_s [G_t (\hat{u}_t^t(s) - \hat{g}_t^t(s))] = \mathbb{E}_s [G_t \hat{u}_t^t(s)] - \mathbb{E}_s [G_t \hat{g}_t^t(s)] = G_s - G_s = 0.$$

Define the projection coefficient $\omega_s \in L^0(\mathcal{F}_s)$ by

$$\begin{aligned} \omega_s &= \frac{\mathbb{E}_s^{Q^t} [(\hat{u}_t^t(s) - \hat{g}_t^t(s)) \hat{e}_t^t(s)]}{\mathbb{E}_s^{Q^t} [(\hat{e}_t^t(s))^2]} = \frac{\mathbb{E}_s^{Q^t} [\hat{u}_t^t(s)] - \mathbb{E}_s^{Q^t} [\hat{g}_t^t(s)]}{\mathbb{E}_s^{Q^t} [\hat{e}_t^t(s)]} \\ &= \frac{\mathbb{E}_s^{Q^t} [\hat{u}_t^t(s)] \mathbb{E}_s [G_t^2] - G_s^2}{\text{var}_s(G_t)}, \end{aligned}$$

where last equalities are due to the definition of $\hat{e}^t(s)$ and Lemma 17 (iv). Define also the process $\hat{n}^t(s) = \hat{u}^t(s) - \hat{g}^t(s) - \omega_s \hat{e}^t(s)$, which belongs to \mathring{H}_s^T because both $\hat{u}^t(s) - \hat{g}^t(s)$ and $\hat{e}^t(s)$ are in \mathring{H}_s^T . In addition,

$$\mathbb{E}_s^{Q^t} [\hat{g}_t^t(s) \hat{n}_t^t(s)] = \mathbb{E}_s^{Q^t} [\hat{g}_t^t(s) \hat{u}_t^t(s)] - \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s))^2] - \omega_s \mathbb{E}_s^{Q^t} [\hat{g}_t^t(s) \hat{e}_t^t(s)] = 0$$

because $\mathbb{E}_s^{Q^t} [\hat{g}_t^t(s) \hat{u}_t^t(s)] = \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s))^2]$ by Lemma 17 and $\hat{g}^t(s)$ and $\hat{e}^t(s)$ belong to orthogonal submodules. Furthermore,

$$\mathbb{E}_s^{Q^t} [\hat{e}_t^t(s) \hat{n}_t^t(s)] = \mathbb{E}_s^{Q^t} [\hat{e}_t^t(s) (\hat{u}_t^t(s) - \hat{g}_t^t(s))] - \omega_s \mathbb{E}_s^{Q^t} [(\hat{e}_t^t(s))^2] = 0$$

by the expression of ω_s . By the definition of \hat{e}^t , $\mathbb{E}_s^{Q^t} [\hat{e}_t^t(s) \hat{n}_t^t(s)] = \mathbb{E}_s^{Q^t} [\hat{n}_t^t(s)] = 0$.

Proof of Corollary 4

The result follows from the relation between returns and martingales of eq. (2).

Proof of Theorem 6

Each conditional martingale $\hat{u}^t(s)$ satisfies the decomposition provided by Theorem 3

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) + \hat{n}^t(s), \quad \omega_s = \frac{k_s \mathbb{E}_s [G_t^2] - G_s^2}{\text{var}_s(G_t)}.$$

Moreover,

$$\text{var}_s^{Q^t} (\hat{u}_t^t(s)) = \mathbb{E}_s^{Q^t} [(\hat{u}_t^t(s))^2] - \left(\mathbb{E}_s^{Q^t} [\hat{u}_t^t(s)] \right)^2 = \mathbb{E}_s^{Q^t} [(\hat{u}_t^t(s))^2] - k_s^2.$$

We note that

$$\begin{aligned} \mathbb{E}_s^{Q^t} [(\hat{u}_t^t(s))^2] &= \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s) + \hat{n}_t^t(s))^2] \\ &= \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s))^2] + \mathbb{E}_s^{Q^t} [(\hat{n}_t^t(s))^2] \\ &\quad + 2\mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s)) \hat{n}_t^t(s)]. \end{aligned}$$

By Theorem 3, $\mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s)) \hat{n}_t^t(s)] = 0$ and so

$$\begin{aligned} \mathbb{E}_s^{Q^t} [(\hat{u}_t^t(s))^2] &= \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s))^2] + \mathbb{E}_s^{Q^t} [(\hat{n}_t^t(s))^2] \\ &\geq \mathbb{E}_s^{Q^t} [(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s))^2]. \end{aligned}$$

Therefore, $\text{var}_s^{Q^t}(\hat{u}_t^t(s))$ is minimized by the conditional martingale with $\hat{n}_t^t(s) = 0$.

Proof of Corollary 7

The relation $u_t(s) = e^{r_s^T(T-s) - r_t^T(T-t)} \hat{u}_t^t(s)$ guarantees that

$$\mathbb{E}_s^{Q^t} [\hat{u}_t^t(s)] = e^{-r_s^T(T-s)} \mathbb{E}_s^{Q^t} [e^{r_t^T(T-t)} u_t(s)] = e^{-r_s^T(T-s)} h_s$$

and

$$\text{var}_s^{Q^t}(\hat{u}_t^t(s)) = e^{-2r_s^T(T-s)} \text{var}_s^{Q^t}(e^{r_t^T(T-t)} u_t(s)).$$

Hence, the claim is an immediate consequence of Theorem 6 with $k_s = e^{-r_s^T(T-s)} h_s$.

Proof of Proposition 8

We show that, if $\hat{z}^T(s)$ is a conditional mean-variance martingale in $[s, T]$, then $\hat{z}^t(s)$ is a conditional mean-variance martingale in $[s, t]$.

Suppose that $\hat{z}^T(s)$ minimizes $\text{var}_s^{Q^T}(\hat{u}_T^T(s))$ among all conditional martingales in H_s^T with $\mathbb{E}_s^{Q^T}[\hat{u}_T^T(T)] = k_s$ for a given $k_s \in L^0(\mathcal{F}_s)$. By Theorems 3 and 6,

$$\hat{z}^T(s) = \hat{g}^T(s) + \omega_s \hat{e}^T(s), \quad \omega_s = \frac{k_s \mathbb{E}_s[G_T^2] - G_s^2}{\text{var}_s(G_T)}$$

and $\mathbb{E}_s^{Q^T}[\hat{g}_T^T(s) \hat{e}_T^T(s)] = 0$. The decomposition on $[s, T]$ induces a decomposition on $[s, t]$ for the conditional martingale $\hat{z}^t(s)$ obtained by restricting $\hat{z}^T(s)$ on $[s, t]$:

$$\hat{z}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s).$$

Moreover, $\mathbb{E}_s^{Q^t}[\hat{g}_t^t(s) \hat{e}_t^t(s)] = 0$ and so we retrieve the orthogonal decomposition of $\hat{z}^t(s)$ provided by Theorem 3 in the time window $[s, t]$. Indeed, since $\hat{z}_t^t(s) = \hat{z}_t^T(s)$, we have

$$\begin{aligned} \mathbb{E}_s[G_t^2] \mathbb{E}_s^{Q^t}[\hat{g}_t^t(s) \hat{e}_t^t(s)] &= \mathbb{E}_s[G_t^2 \hat{g}_t^t(s) \hat{e}_t^t(s)] = \mathbb{E}_s[G_t^2 \hat{g}_t^T(s) \hat{e}_t^T(s)] \\ &= \mathbb{E}_s\left[G_t^2 \frac{G_s}{G_t} \frac{G_t - G_s}{G_t}\right] = 0. \end{aligned}$$

In addition,

$$\omega_s = \frac{k_s \mathbb{E}_s[G_T^2] - G_s^2}{\text{var}_s(G_T)} = \frac{h_s \mathbb{E}_s[G_t^2] - G_s^2}{\text{var}_s(G_t)},$$

where $h_s \in L^0(\mathcal{F}_s)$ is identified by k_s through

$$h_s = \frac{1}{\mathbb{E}_s[G_t^2]} \left\{ G_s^2 + \frac{\text{var}_s(G_t)}{\text{var}_s(G_T)} (k_s \mathbb{E}_s[G_T^2] - G_s^2) \right\}.$$

Then, Theorem 6 ensures that $\hat{z}^t(s)$ minimizes $\text{var}_s^{Q^t}(\hat{u}_t^t(s))$ among all conditional martingales in H_s^t such that $\mathbb{E}_s^{Q^t}[\hat{u}_t^t(s)] = h_s$. In other words, $\hat{z}^t(s)$ is a conditional mean-variance martingale at time t .

Proof of Corollary 9

Suppose that the return $z(s)$ minimizes $\text{var}_s^{Q^T}(u_T(s))$ among all returns with $\mathbb{E}_s^{Q^T}[u_T(s)] = k_s$ for some given $k_s \in L^0(\mathcal{F}_s)$. By Corollaries 4 and 7,

$$z(s) = g(s) + \omega_s e(s), \quad \omega_s = \frac{k_s e^{-r_s^T(T-s)} \mathbb{E}_s[G_T^2] - G_s^2}{\text{var}_s(G_T^T)},$$

with $\mathbb{E}_s^{Q^T}[g_T(s)e_T(s)] = 0$. The former decomposition holds algebraically at time t in $[s, T]$ too, where, in addition $\mathbb{E}_s^{Q^t}[e^{2r_t^T(T-t)}g_t(s)e_t(s)] = 0$. Hence, by uniqueness of the decomposition, we obtain the same result that we get by decomposing $z(s)$ in the time range $[s, t]$ as prescribed by Corollary 4. Furthermore,

$$\omega_s = \frac{k_s e^{-r_s^T(T-s)} \mathbb{E}_s[G_T^2] - G_s^2}{\text{var}_s(G_T)} = \frac{h_s e^{-r_s^T(T-s)} \mathbb{E}_s[G_t^2] - G_s^2}{\text{var}_s(G_t)},$$

where $h_s \in L^0(\mathcal{F}_s)$ is determined by k_s through the relation

$$h_s = \frac{1}{\mathbb{E}_s[G_t^2]} \left\{ e^{r_s^T(T-s)} G_s^2 + \frac{\text{var}_s(G_t)}{\text{var}_s(G_T)} (k_s \mathbb{E}_s[G_T^2] - e^{r_s^T(T-s)} G_s^2) \right\}.$$

By Corollary 7 in $[s, t]$, this means that $z(s)$ minimizes $\text{var}_s^{Q^t}(e^{r_t^T(T-t)}u_t(s))$ among all returns with $\mathbb{E}_s^{Q^t}[e^{r_t^T(T-t)}u_t(s)] = h_s$. Hence $z(s)$ is a mean-variance return at time t , too.

Proof of Theorem 10

Suppose that $u(s)$ is a mean-variance return. Then, Corollary 7 guarantees that $u(s) = g(s) + \omega_s e(s)$ for some $\omega_s \in L^0(\mathcal{F}_s)$. Consider another mean-variance return $v(s)$ that decomposes as $v(s) = g(s) + \tilde{\omega}_s e(s)$ with $\tilde{\omega}_s$ different from zero in $L^0(\mathcal{F}_s)$. Hence, $e(s) = (v(s) - g(s))/\tilde{\omega}_s$ and so

$$u(s) = g(s) + \frac{\omega_s}{\tilde{\omega}_s} v(s) - \frac{\omega_s}{\tilde{\omega}_s} g(s) = \alpha_s v(s) + (1 - \alpha_s) z(s),$$

where $\alpha_s = \omega_s/\tilde{\omega}_s$ and $z(s) = g(s)$, which is on the mean-variance frontier, too.

Conversely, assume that a return process $u(s)$ satisfies the decomposition $u(s) = \alpha_s v(s) + (1 - \alpha_s) z(s)$ with $v(s), z(s)$ mean-variance returns, i.e. $v(s) = g(s) + \tilde{\omega}_s e(s)$ and $z(s) = g(s) + \bar{\omega}_s e(s)$ for some $\tilde{\omega}_s, \bar{\omega}_s \in L^0(\mathcal{F}_s)$. It follows that

$$u(s) = \alpha_s (g(s) + \tilde{\omega}_s e(s)) + (1 - \alpha_s) (g(s) + \bar{\omega}_s e(s)) = g(s) + (\alpha_s \tilde{\omega}_s + (1 - \alpha_s) \bar{\omega}_s) e(s)$$

and so, by Corollary 7, $u(s)$ is on the mean-variance frontier.

By considering the return process $f(s)$, we get $u(s) = \alpha_s g(s) + (1 - \alpha_s) f(s)$. Since $f(s) = g(s) + e(s)$, we immediately obtain that $\alpha_s = 1 - \omega_s$.

Proof of Proposition 11

Since $u(s)$ is a return process,

$$\begin{aligned} 1 &= \mathbb{E}_s [M_{s,t} u_t(s)] = \mathbb{E}_s [G_t^2] \frac{e^{-r_s^T(T-s)}}{G_s} \mathbb{E}_s^{Q^t} \left[\frac{e^{r_t^T(T-t)}}{G_t} u_t(s) \right] \\ &= \mathbb{E}_s [G_t^2] \frac{e^{-r_s^T(T-s)}}{G_s} \\ &\quad \cdot \left\{ \text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} u_t(s) \right) + \mathbb{E}_s^{Q^t} \left[\frac{1}{G_t} \right] \mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} u_t(s) \right] \right\}. \end{aligned}$$

As $\mathbb{E}_s^{Q^t} [1/G_t] = G_s / \mathbb{E}_s [G_t^2]$, we obtain

$$\frac{G_s}{\mathbb{E}_s [G_t^2]} \left\{ \mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} u_t(s) \right] - e^{r_s^T(T-s)} \right\} = -\text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} u_t(s) \right).$$

By choosing $u(s) = g(s)$, we deduce

$$\frac{G_s}{\mathbb{E}_s [G_t^2]} \left\{ \mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} g_t(s) \right] - e^{r_s^T(T-s)} \right\} = -\text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} g_t(s) \right).$$

As a result, in the previous representation

$$\begin{aligned} &\frac{G_s}{\mathbb{E}_s [G_t^2]} \left\{ \mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} u_t(s) \right] - e^{r_s^T(T-s)} \right\} \\ &= -\text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} g_t(s) \right) \frac{\text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} u_t(s) \right)}{\text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} g_t(s) \right)} \\ &= \beta_s \frac{G_s}{\mathbb{E}_s [G_t^2]} \left\{ \mathbb{E}_s^{Q^t} \left[e^{r_t^T(T-t)} g_t(s) \right] - e^{r_s^T(T-s)} \right\}, \end{aligned}$$

where

$$\beta_s = \text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} u_t(s) \right) \Big/ \text{cov}_s^{Q^t} \left(\frac{1}{G_t}, e^{r_t^T(T-t)} g_t(s) \right)$$

belongs to $L^0(\mathcal{F}_s)$. Therefore, we get eq. (5). In addition, the coefficient β_s can be rewritten as in the claim.

Proof of Proposition 12

The Lagrangian function is

$$\mathcal{L} = \mathbb{E}_s \left[\int_s^T U(i_\tau + w_\tau) - \lambda_s M_{s,\tau} w_\tau d\tau \right] + \lambda_s w_s$$

with $w_s \in L^0(\mathcal{F}_s)$. Note that \mathcal{L} is a function of λ_s and $w_\tau(\omega)$ for all instants $\tau \in [s, T]$ and states $\omega \in \Omega$. The first-order condition implies that (at any time and in any state) $U'(i_\tau + w_\tau) - \lambda_s M_{s,\tau} = 0$. Therefore,

$$w_\tau = (U')^{-1}(\lambda_s M_{s,\tau}) - i_\tau = \frac{b_\tau}{M_{s,\tau}} - \frac{\lambda_s}{M_{s,\tau}} - i_\tau = \frac{b_\tau}{M_{s,\tau}} - \lambda_s g_\tau(s) - i_\tau,$$

thanks to the quadratic utility. The constraint over w_s delivers

$$\begin{aligned} w_s &= \mathbb{E}_s \left[\int_s^T M_{s,\tau} \left(\frac{b_\tau}{M_{s,\tau}} - i_\tau \right) d\tau \right] - \lambda_s \mathbb{E}_s \left[\int_s^T M_{s,\tau} g_\tau(s) d\tau \right] \\ &= \mathbb{E}_s \left[\int_s^T M_{s,\tau} \left(\frac{b_\tau}{M_{s,\tau}} - i_\tau \right) d\tau \right] - \lambda_s (T - s) \end{aligned}$$

and so

$$\lambda_s = \frac{1}{T - s} \mathbb{E}_s \left[\int_s^T M_{s,\tau} \left(\frac{b_\tau}{M_{s,\tau}} - i_\tau \right) d\tau \right] - \frac{w_s}{T - s}.$$

As a result,

$$w_\tau = \frac{b_\tau}{M_{s,\tau}} - i_\tau - \left(\frac{1}{T - s} \mathbb{E}_s \left[\int_s^T M_{s,\tau} \left(\frac{b_\tau}{M_{s,\tau}} - i_\tau \right) dt \right] - \frac{w_s}{T - s} \right) g_\tau(s)$$

and we denote it by w_τ^* . Under the assumptions about income and bliss points,

$$\begin{aligned} w_\tau^* &= \frac{b_s \pi_\tau(1_T)}{T - s} - \left(\frac{1}{(T - s)^2} \mathbb{E}_s \left[\int_s^T e^{-r_\tau^T (T - \tau)} M_{s,\tau} b_s d\tau \right] - \frac{w_s}{T - s} \right) g_\tau(s) \\ &= \frac{b_s \pi_s(1_T)}{T - s} \frac{\pi_\tau(1_T)}{\pi_s(1_T)} - \left(\frac{b_s}{(T - s)^2} \pi_s(1_T) \mathbb{E}_s \left[\int_s^T G_{s,\tau} d\tau \right] - \frac{w_s}{T - s} \right) g_\tau(s) \\ &= \frac{b_s \pi_s(1_T)}{T - s} f_\tau(s) - \left(\frac{b_s \pi_s(1_T)}{T - s} - \frac{w_s}{T - s} \right) g_\tau(s). \end{aligned}$$

Consequently, the optimal payoff stream is associated with the return

$$\frac{(T - s)w_\tau^*}{w_s} = \frac{b_s \pi_s(1_T)}{w_s} f_\tau(s) - \left(\frac{b_s \pi_s(1_T)}{w_s} - 1 \right) g_\tau(s),$$

which is on the mean-variance frontier by Theorem 10.

Proof of Proposition 13

Following the same steps as in the proof of Proposition 12, the Lagrange multiplier is

$$\lambda_s = \frac{1}{t-s} \mathbb{E}_s \left[\int_s^t M_{s,\tau} \left(\frac{b_\tau}{M_{s,\tau}} - i_\tau \right) d\tau \right] - \frac{\tilde{w}_s}{t-s}.$$

Therefore, the optimal payoff stream is

$$\begin{aligned} w_\tau &= \frac{b_s \pi_\tau(1_T)}{T-s} - \left(\frac{1}{(T-s)(t-s)} \mathbb{E}_s \left[\int_s^T e^{-r_\tau(T-\tau)} M_{s,\tau} b_s d\tau \right] - \frac{\tilde{w}_s}{t-s} \right) g_\tau(s) \\ &= \frac{b_s \pi_s(1_T)}{T-s} \frac{\pi_\tau(1_T)}{\pi_s(1_T)} - \left(\frac{b_s}{(T-s)(t-s)} \pi_s(1_T) \mathbb{E}_s \left[\int_s^T G_{s,\tau} d\tau \right] - \frac{\tilde{w}_s}{t-s} \right) g_\tau(s) \\ &= \frac{b_s \pi_s(1_T)}{T-s} f_\tau(s) - \left(\frac{b_s \pi_s(1_T)}{T-s} - \frac{w_s}{T-s} \right) g_\tau(s). \end{aligned}$$

and it coincides with the one prescribed by Proposition 12.

C Additional simulations

C.1 Reference market with two stocks

We provide a generalization of the reference market of Subsection 5.1 by allowing for two risky stocks. We, then, repeat the simulations of Subsection 5.2 with 6 maturities. Generalizations with a higher number of assets can be developed in a similar way.

In the system of equations (7) under the measure Q , we consider an additional Wiener process V^Q , independent of W^Q and Z^Q and a novel stock price S_t with volatility $\kappa > 0$. The parameter ψ provides the instantaneous correlation between the new stock and the zero-coupon bond, while χ gives the instantaneous correlation with the old stock:

$$\begin{cases} dS_t = S_t Y_t dt + \kappa S_t \left[\psi dW_t^Q + \frac{\chi - \phi\psi}{\sqrt{1-\phi^2}} dZ_t^Q + \sqrt{1-\psi^2 - \frac{(\chi-\phi\psi)^2}{1-\psi^2}} dV_t^Q \right] \\ dX_t = X_t Y_t dt + \eta X_t \left[\phi dW_t^Q + \sqrt{1-\phi^2} dZ_t^Q \right] \\ d\pi_t(1_T) = \pi_t(1_T) Y_t dt - \pi_t(1_T) B(t, T) \sigma dW_t^Q. \end{cases}$$

The orthogonal shocks dW_t^Q , dZ_t^Q and dV_t^Q come from the Cholesky factorization of the 3×3 correlation matrix of the original Brownian motions.

The market price of risk is the multivariate process $[\nu^W, \nu^Z, \nu^V]'$ with the first two entries as in eq. (8) and

$$\nu_t^V = \frac{\mu_t^S - Y_t - \kappa \psi \nu_t^W - \frac{\chi - \phi\psi}{\sqrt{1-\phi^2}} \kappa \nu_t^Z}{\kappa \sqrt{\frac{\phi^2 - 2\phi\psi\chi + \psi^2 + \chi^2 - 1}{\phi^2 - 1}}},$$

where μ^S is the adapted drift process of dS_t/S_t under the physical measure. The Radon-Nikodym derivative of Q with respect to P , the Novikov condition and the pricing kernel

dynamics are modified to accommodate the extra component in the market price of risk. The other assumptions and the parameter choices of Subsection 5.1 are kept. In addition, we set $\kappa = 0.15$, $\psi = 0.1$, $\chi = -0.3$ and $\mu_t^S = Y_t + 0.08$.

We, then, repeat the six-semester mean-variance optimization of Subsection 5.2 with the constants $c = \$0.002$ for trading costs and $C = \$0.02$ for replication costs. Results are displayed in Figure C.1, where we represent (scaled) standard deviations, (scaled) Sharpe ratios and (scaled) modified Sharpe ratios across maturities, transaction costs and units of risky assets in each subportfolio, where risky assets coincide with the log optimal portfolio (in the time-consistent approach) and the portfolio replicating the pricing kernel (in the classical frontier). As the modified Sharpe ratio shows, in this simulation the time-consistent approach outperforms the standard mean-variance optimization when replication and trading costs are taken into account.

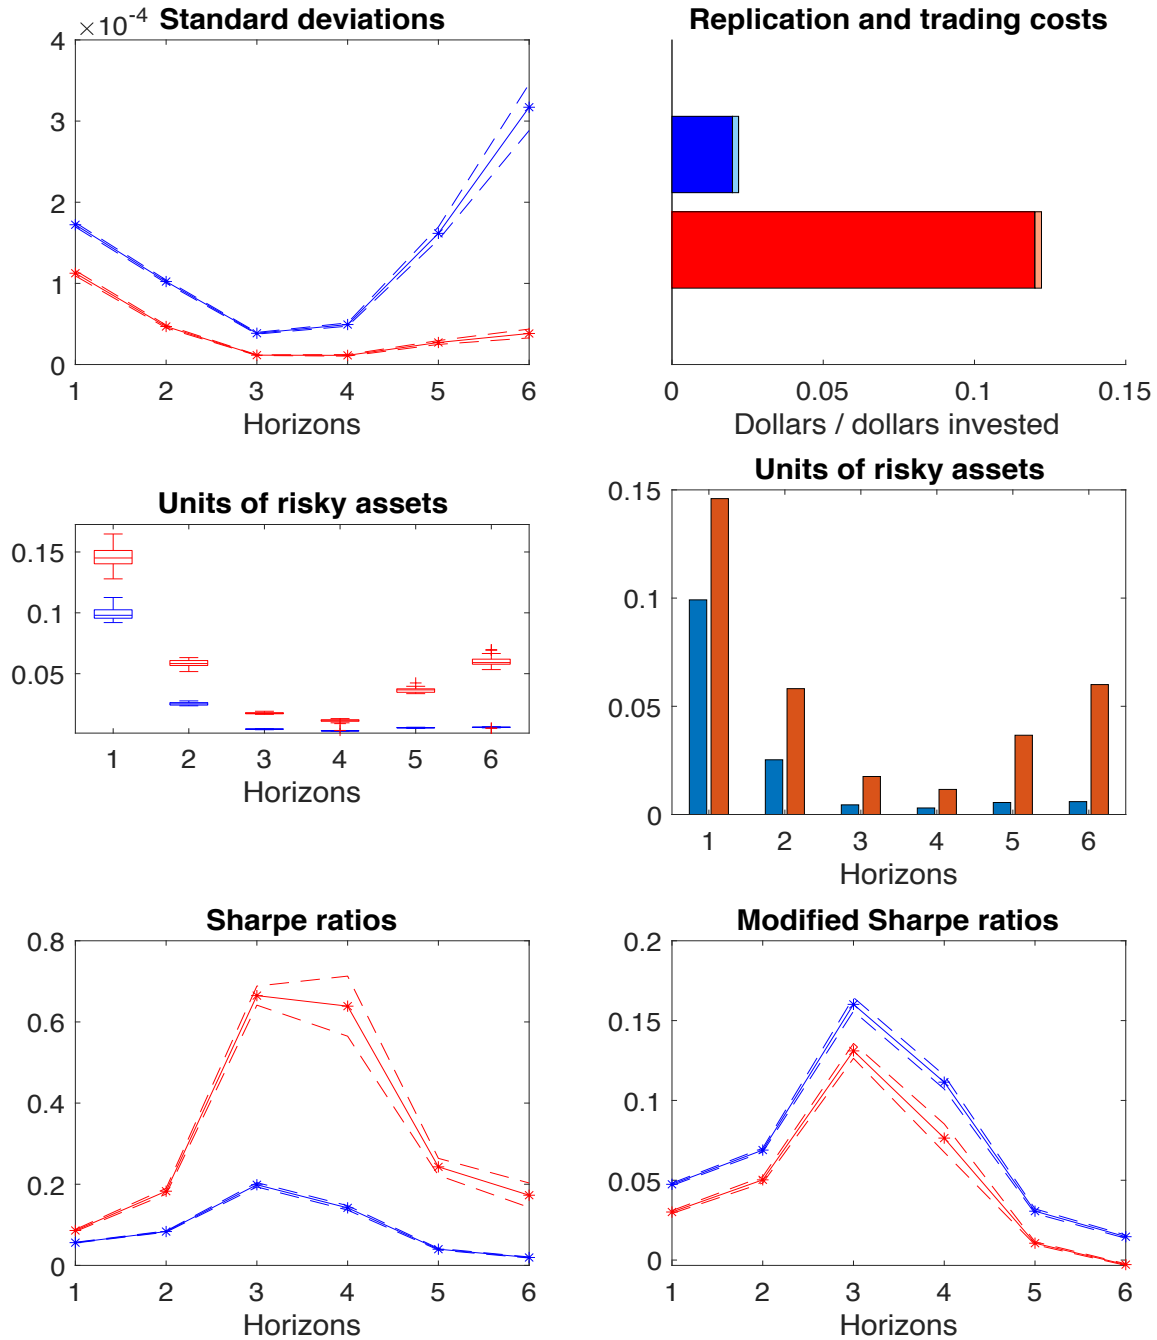


Figure 5: Red (resp. blue) lines, bars and boxes refer to the classical (resp. time-consistent) mean-variance solution for the problem of Subsection C.1. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda_0^{(i)}$ for all $i = 1, \dots, 6$. 90% confidence intervals for these variables are represented. The top-right panel represents the transaction costs of the time-consistent portfolio (blue for replication costs, light blue for trading costs) and of the mean-variance portfolio (red for replication costs, light red for trading costs). Medium panels contain the box-and-whisker plot at 25th and 75th percentiles and the bar plot of loadings $|\alpha_s^{(i)}|$ and $|\tilde{\alpha}_s^{(i)}|$ at all horizons.