On time-consistent multi-horizon portfolio allocation

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

September 12, 2021

Abstract

We analyze the problem of constructing multiple mean-variance portfolios over increasing investment horizons in continuous-time arbitrage-free stochastic interest rate markets. The traditional one-period mean-variance optimization of Hansen and Richard (1987) requires the replication of a risky payoff for each investment horizon. When many maturities are considered, a large number of payoffs must be replicated, with an impact on transaction costs. In this paper, we orthogonally decompose the whole processes defined by asset returns to obtain a mean-variance frontier generated by the same two securities across a multiplicity of horizons. Our risk-adjusted mean-variance frontier rests on the martingale property of the returns discounted by the log-optimal portfolio and features a time-consistency property. The outcome is that the replication of a single risky payoff is required to implement such frontier at any investment horizon. As a result, when transaction costs are taken into account, our risk-adjusted mean-variance frontier may outperform the traditional mean-variance optimal strategies in terms of Sharpe ratio. Realistic numerical examples show the improvements of our approach in medium- or long-term cashflow management, when a sequence of target returns at increasing investment horizons is considered.

JEL Classification: G11, G12.

[*]Università Bocconi, Milan (Italy), Department of Decision Sciences and IGIER: simone.cerreia@unibocconi.it.

[†]Università Bocconi, Milan (Italy), Department of Finance and IGIER: fulvio.ortu@unibocconi.it.

[‡]Università degli Studi di Padova, Padova (Italy), Department of Mathematics and Università Bocconi, Milan (Italy), Department of Finance: francesco.rotondi@math.unipd.it.

[§]✉ Université Laval, Québec (Canada), Department of Finance, Insurance and Real Estate: federico.severino@fsa.ulaval.ca.

[¶]We thank Anna Battauz, Marie-Claude Beaulieu, Claudio Fontana, Fabio Girardi, Massimo Marinacci and participants at AMASES XLIII Annual Meeting (University of Perugia, 2019), Goethe Universität Frankfurt, 10^{th} General AMaMeF Conference (2021), 37^{th} AFFI International Conference (2021), CORS Annual Conference (2021), 29^{th} EFMA Annual Conference (2021) and Econometric Society European Meeting (2021) for useful comments. S. Cerreia-Vioglio acknowledges the financial support of ERC (grant SDDM-TEA). F. Severino acknowledges the financial support of FRQSC (Fonds de recherche du Québec – Société et culture).

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

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. Indeed, frontiers at different horizons are generated by different securities: a risky security associated with the pricing kernel on the time window under consideration and a riskless asset. Moreover, relevant computational issues are present: optimal 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. See, e.g. Michaud (1989), Best and Grauer (1991) and DeMiguel, Garlappi, and Uppal (2007).

These issues worsen in a *multi-horizon* allocation problem where a sequence of N maturities, each associated with a target expected return, are considered. This kind of problem is of utmost importance for insurance companies, pension funds and any financial intermediary managing long-term multiple cashflows, such as annuities. We can consider, e.g., an investor that wants to meet N expected return targets at N subsequent horizons by investing in N buy-and-hold portfolios which attain the targets with the minimum possible variance each (the proper formalization is in Section 5). According to the standard mean-variance approach, the investor should solve N different optimization problems that would lead to N optimal portfolios requiring the replication of one different risky security each: globally, the investor would need to replicate N risky assets. When transaction costs (in particular, replication costs) are taken into account, the resulting portfolio allocation may be inconvenient.

To tackle these issues, we propose a generalization of the traditional mean-variance optimization by decomposing the whole return process of each traded security on the larger time interval under consideration. We obtain a *risk-adjusted* mean-variance frontier which is spanned by the same two securities (a risky one and a riskless one) at any time horizon. This feature implies a drastic reduction of replication costs that can benefit multi-horizon investment decisions. Considering again the multi-horizon problem above, according to our risk-adjusted approach, the investor has to replicate only one risky security as all the N built portfolios involve different units of the same assets. Comparing with the classical mean-variance approach, the loss in terms of standard deviation of risk-adjusted investment strategies can be compensated by sizable savings on their implementation costs. Indeed, after incorporating transaction costs in the analysis, risk-adjusted mean-variance portfolios can display a higher Sharpe ratio than classical mean-variance portfolios. We illustrate several examples in Section 5, in the contexts of fixed-income markets and life annuities.

To give a snapshot of our construction, we consider a continuous-time arbitrage-free market with finite horizon T, stochastic interest rates and, possibly, a bunch of risky securities. Pure discount bonds with any expiry are traded, too, as well as the log-optimal portfolio (details in Subsection 2.1). All the results are presented 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, in Subsection 2.2 we construct the space H_s^T of conditional martingales obtained by discounting asset returns by the value of the log-optimal portfolio. Specifically, H_s^T is endowed with an inner product based on the conditional expectation of martingale terminal values. 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 3 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

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

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, with an orthogonality condition implied by the structure of H_s^T. In addition, the associated *risk-adjusted* mean-variance frontier in the period $[s,T]$ is made up of asset returns with null $n(s)$ (see Corollary 6). A Two-fund Separation Theorem holds (Theorem 9) and so the frontier turns out to be spanned by $g(s)$ and the return $f(s)$ associated with a pure discount T-bond. Importantly, it is possible to decompose returns also in any subperiod $[s,t]$ with $t \leqslant T$ in an analogous way. Note that the use of the log-optimal portfolio and the risk-neutral variance of asset returns shares some similarities with Martin and Wagner (2019).

The main advantage of our decompositions is time consistency. Since we decompose the process that defines returns over the longest horizon, restrictions on closer horizons naturally obtain. Moreover, the orthogonality relations at different horizons ensure that a time consistency property holds for our mean-variance returns: returns on the risk-adjusted mean-variance frontier at time T are risk-adjusted mean-variance returns at date t, too (Corollary 8). For example, a buy-and-hold one-year horizon risk-adjusted mean-variance portfolio turns out to lie on the risk-adjusted mean-variance frontier also at the six-month horizon. In fact, our risk-adjusted mean-variance frontiers are spanned by the same two assets across a continuum of horizons, a crucial property for the practitioners. 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.

Similarly to Cochrane (2014), we provide in Section 6 a microeconomic foundation of our risk-adjusted 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 different investment horizon. Finally, the Appendix contains some complements of the theory and additional simulations.

1.1 Additional related literature

As it is well-known, one-period mean-variance portfolio analysis has its roots in the seminal works by Markowitz (1952), Tobin (1958) and Sharpe (1964) and the abstract formalization is provided by Hansen and Richard (1987). The development in the last decades has been huge and its summary goes beyond our scope. Interestingly, multi-period 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. However, differently from our multi-horizon approach, intermediate dates are only useful for rebalancing purposes, and no intermediate target is considered.

Our risk-adjusted mean-variance frontier features a time consistency property that allows to generate optimal returns via the same two securities across a sequence of investment horizons. Time consistency of portfolio or consumption choices is an old issue of economic theory. 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). Van Staden, Dang, and Forsyth (2018) provide a detailed summary of the two literature streams – one related to precommittment, the other to time consistency. The mean-variance theory proposed here is peculiar in this respect. Indeed, it shares some aspects of both streams: our risk-adjusted mean-variance frontier features time consistency and our application to multi-horizon portfolio allocation lies within the precommitment paradigm (the problem is solved ex ante and the investor never changes the plan). The addressed problem is, in fact, peculiar and different from the ones treated in the literature because we are considering multiple investment targets at increasing horizons.

As we already mentioned, our mean-variance theory is designed in a conditional framework. For a comparison between conditional and unconditional mean-variance optimization, one can refer to Ferson and Siegel (2001), where mean-variance optimization problems in the presence of conditioning information are discussed. It is also worth mentioning the parallel literature stream about the use of conditional information for the mean-variance frontier of stochastic discount factors. Starting from the celebrated dual result of Hansen and Jagannathan (1991), conditional variance bounds on pricing kernels have been illustrated by Bekaert and Liu (2004), Ferson and Siegel (2003) and Gallant, Hansen, and Tauchen (1990), among the others. See, e.g. the review in Favero, Ortu, Tamoni, and Yang (2020).

2 Framework and essentials

We describe the asset pricing framework and the essential tools for the intertemporal decomposition of returns. We simultaneously introduce the notation of the paper.

2.1 Arbitrage-free market and numéraire changes

Fix $T > 0$ and consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$, where P is the *physical measure* and the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t\in[0,T]}$ satisfies the usual conditions. The adapted process $Y = \{Y_t\}_{t\in[0,T]}$ represents the stochastic instantaneous rate. The money market account has value $e^{\int_0^t Y_\tau d\tau}$ at any time t. Pure discount bonds with any possible maturity and face value equal to 1 are traded. Additional risky securities, with adapted price processes, can be present in the market, too. Moreover, we assume that the *log-optimal portfolio* (or *growth optimal portfolio*) is traded in the market. This is the self-financing portfolio that maximizes the expected utility on the terminal wealth (at time T) of a log-utility investor with initial wealth equal to 1 (see Chapter 20 of Björk, 2009).

If X_t is the price process of a traded security at time t, we obtain its relative price Z_t by discounting X_t with the money market account: $Z_t = e^{-\int_0^t Y_\tau d\tau} X_t$. We assume that all relative asset prices are semimartingales and there exists an equivalent martingale measure Q (the *risk-neutral measure*). The market is, then, arbitrage-free. The Radon-Nikodym derivative of Q with respect to P on $\mathcal{F}_T$ is $L_T = dQ/dP$ and we define $L_t = \mathbb{E}_t[L_T]$ and $L_{t,T} = L_T/L_t$ at any time $t \in [0, T]$. Here, $\mathbb{E}_t$ denotes the conditional expectation with respect to $\mathcal{F}_t$ under the measure P. 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 $M_{t,T} = M_T/M_t$. Importantly, the value of the log-optimal portfolio at any time $t \in [0, T]$ is $M_{0,t}^{-1}$. Moreover, the log-optimal portfolio works as *numéraire portfolio*, meaning that prices discounted by the log-optimal portfolio are martingales under P. See Long (1990) and Section 26.9 in Björk (2009).

We now consider a pure discount T-bond 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 (or T-forward measure) and we denote it by F (see Geman, 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 on $\mathcal{F}_T$ by $G_T = dF/dP$. 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.

Using the T-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 $s \leqslant t \leqslant T$ 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

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

2.2 The Hilbert modules H_s^t and linear pricing functionals

We consider the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$. 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 $L_s^1(\mathcal{F}_t) = L^0(\mathcal{F}_s)L^1(\mathcal{F}_t)$.[1]

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 martingales except for integrability, which is replaced by the weaker condition $\mathbb{E}_s[|\hat{z}(\tau)|] \in$

[1]Clearly, $L_s^1(\mathcal{F}_t)$ 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).

$L^0(\mathcal{F}_s)$ for all $\tau \in [s,t]$. See, for instance, Chapter VII, §1 of Shiryaev (1996).[2] With this property in mind, for any $t \in [s,T]$ we define the space

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

H_s^t contains the price processes discounted by the log-optimal portfolio with the appropriate square-integrability condition. Interestingly, H_s^t can be characterized in differential terms: see Proposition 2.4 in Marinacci and Severino (2018) about weak time-derivatives and Subsection 2.4 in Severino (2021). For our construction the relation between $H_s^{t_1}$ and $H_s^{t_2}$ with $t_1 \leqslant t_2$ 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 $\mathbb{E}_s[\hat{z}_{t_1}^2] \leqslant \mathbb{E}_s[\hat{z}_{t_2}^2]$.

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 $\cdot : L^0(\mathcal{F}_s) \times H_s^t \to H_s^t$ and the L^0-valued inner product $\langle\ ,\ \rangle_s^t : H_s^t \times H_s^t \to L^0(\mathcal{F}_s)$ respectively by

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

The inner product homogeneity 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 $\langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s\left[\hat{z}_t \hat{v}_t\right] = 0$.

Our inner product mimics the conditional structure of Hansen and Richard (1987) and Gallant, Hansen, and Tauchen (1990), who employ a conditional asset pricing framework under the physical measure. Here we will apply such an approach on the martingale processes induced by discounted prices under risk neutrality.

Importantly, H_s^t is a selfdual pre-Hilbert module or, more simply, a Hilbert module. Selfduality 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 $L^0(\mathcal{F}_s)$.*

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 $\cdot : L^0(\mathcal{F}_s) \times H_s^t \to H_s^t$ 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.

[2] For example, when the forward measure is employed, (conditional) martingales are provided by forward prices for the settlement date t (see Section 9.6 in Musiela and Rutkowski, 2005).

(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 $\langle\ ,\ \rangle_s^t : H_s^t \times H_s^t \to L^0(\mathcal{F}_s)$. For all $\hat{z} \in H_s^t$, $\mathbb{E}_s[\hat{z}_t^2] \in L_s^0(\mathcal{F}_s)$. Therefore, by Footnote 3 in Hansen and Richard (1987), $\langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s[\hat{z}_t \hat{v}_t]$ 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) $\langle \hat{z}, \hat{z} \rangle_s^t = \mathbb{E}_s[\hat{z}_t^2] \geqslant 0$ with equality if and only if $\hat{z}_t = 0$. This implies that, for any $\tau \in [s, t]$, $\mathbb{E}_\tau[\hat{z}_t] = \hat{z}_\tau = 0$. 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) $\langle a_s \cdot \hat{z}, \hat{v} \rangle_s^t = a_s \mathbb{E}_s[\hat{z}_t \hat{v}_t] = a_s \langle \hat{z}, \hat{v} \rangle_s^t$.

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 $d(f, g) = \mathbb{E}[\min\{|f - g|, 1\}]$ 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

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

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 $L_s^2(\mathcal{F}_t) = \{f \in L^0(\mathcal{F}_t) :\ \mathbb{E}_s[f^2] \in L^0(\mathcal{F}_s)\}$, 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 $n, m > N_\varepsilon$,

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

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

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

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

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, consider an $\mathcal{F}_t$-measurable payoff h_t with $\mathbb{E}_s[M_{s,t}^2 h_t^2]$ in $L^0(\mathcal{F}_s)$. Consider, then, the process of prices discounted by the log-optimal portfolio, i.e. $\hat{h} = \{\hat{h}_\tau\}_{\tau\in[s,t]}$ defined by $\hat{h}_\tau = M_{s,\tau}\pi_\tau(h_t)$. Such process belongs to H_s^t and, in particular, $\hat{h}_s = \pi_s(h_t) = \mathbb{E}_s[M_{s,t}h_t]$. Hence, the no-arbitrage price of eq. (1) induces the L^0-valued functional $\Pi_s :\ H_s^t \to L^0(\mathcal{F}_s)$ such that

$$\Pi_s : \quad \hat{h} \mapsto \hat{h}_s.$$

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

$$\Pi_s\left(\hat{h}\right) = \left\langle \hat{g}^t(s), \hat{h}\right\rangle_s^t, \qquad \hat{g}_\tau^t(s) = 1 \quad \forall \tau \in [s,t] \tag{2}$$

for any $\hat{h} \in H_s^t$. The constant conditional martingale $\hat{g}^t(s)$ clearly belongs to H_s^t. This process will play a fundamental role in our decomposition of excess returns and its financial meaning is related to the log-optimal portfolio, as we discuss in Subsection 3.2.

3 Return decomposition

In this section we build the relation between asset returns and conditional martingales in H_s^t with $t \in [s,T]$. We orthogonally decompose any H_s^t by exploiting the L^0-valued inner product $\langle\ ,\ \rangle_s^t$ and, as a consequence, we retrieve a decomposition of returns. 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). This is the case of our interest, because we deal with submodules generated by single return processes, 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 2, we determine in Corollary 3 a decomposition of asset returns. Our result parallels Hansen and Richard (1987) decomposition but it exploits a different orthogonality condition inspired by the martingale processes induced by asset returns.

3.1 Return definition

Consider the time τ between s and T. In our theory, a *return* of a traded asset at time τ is an $\mathcal{F}_\tau$-measurable random variable $u_\tau(s)$ satisfying

$$\mathbb{E}_s\left[M_{s,\tau}u_\tau(s)\right] = 1 \quad \forall \tau \in [s,T] \qquad \text{and} \qquad \mathbb{E}_s[M_{s,T}^2 u_T^2] \in L^0(\mathcal{F}_s). \tag{3}$$

The related *return process* is the adapted process $u(s) = \{u_\tau(s)\}_{\tau \in [s,T]}$. To be precise, when dealing with $t \in [s,T]$, we will call *return process in* $[s,t]$ the restriction of $u(s)$ on the time interval $[s,t]$.

As an example, we can consider an attainable payoff h_T at time T such that $\mathbb{E}_s[M_{s,T}^2 h_T^2] \in L^0(\mathcal{F}_s)$. At each $\tau \in [s,T]$, the return is the ratio of no-arbitrage prices $u_\tau(s) = \pi_\tau(h_T)/\pi_s(h_T)$ and the relations in (3) are fulfilled.

Importantly, by discounting returns by the values of the log-optimal portfolio, we obtain a conditional martingale, that we denote by $\hat{u}^T(s)$, which belongs to H_s^T. In particular, $\hat{u}^T(s)$ satisfies

$$\hat{u}_\tau^T(s) = M_{s,\tau}u_\tau(s) \quad \forall \tau \in [s,T] \qquad \text{and} \qquad \hat{u}_s^T(s) = 1. \tag{4}$$

Hence, asset returns are mapped into conditional martingales in H_s^T. Moreover, return processes define conditional martingales also in any time subinterval $[s,t]$ with $t \leqslant T$. Indeed, we define $\hat{u}^t(s)$ in H_s^t as the restriction of $\hat{u}^T(s)$ on $[s,t]$:

$$\hat{u}_\tau^t(s) = M_{s,\tau}u_\tau(s) \quad \forall \tau \in [s,t] \qquad \text{and} \qquad \hat{u}_s^t(s) = 1. \tag{5}$$

Example 1. Consider a zero-coupon bond with expiry T. In this case, by the relation in (4), the return process and the associated conditional martingale in H_s^T are

$$f_\tau(s) = \frac{\pi_\tau(1_T)}{\pi_s(1_T)}, \qquad \hat{f}_\tau^T(s) = G_{s,\tau} \qquad \forall \tau \in [s,T]. \tag{6}$$

Example 2. 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 $M_{s,T}$ (see Subsection 2.1 for the definition of G_T). This payoff is fundamental in the mean-variance decomposition of Hansen and Richard (1987). By the previous relations, the related return process and the conditional martingale in H_s^T are given by

$$u_\tau(s) = \frac{\mathbb{E}_\tau\left[M_{\tau,T} M_{s,T}\right]}{\mathbb{E}_s\left[M_{s,T}^2\right]} = \frac{\mathbb{E}_\tau\left[G_T^2\right]}{M_{s,\tau}\mathbb{E}_s\left[G_T^2\right]}, \qquad \hat{u}_\tau^T(s) = \frac{\mathbb{E}_\tau\left[G_T^2\right]}{\mathbb{E}_s\left[G_T^2\right]} \qquad \forall \tau \in [s,T].$$

3.2 The log-optimal return *g(s)*

Fix $t \in [s,T]$. We define the submodule of H_s^t associated with zero-price payoffs (or excess returns)

$$\begin{aligned} \mathring{H}_s^t &= \left\{\hat{\iota}^t(s) \in H_s^t : \quad \mathbb{E}_s\left[M_{s,t}\iota_t(s)\right] = 0\right\} \\ &= \left\{\hat{\iota}^t(s) \in H_s^t : \quad \mathbb{E}_s\left[\hat{\iota}_t^t(s)\right] = \hat{\iota}_s^t(s) = 0\right\} \\ &= \left\{\hat{\iota}^t(s) \in H_s^t : \quad \left\langle \hat{g}^t(s), \hat{\iota}^t(s)\right\rangle_s^t = 0\right\}, \end{aligned}$$

where $\iota(s)$ and $\hat{\iota}^t(s)$ are linked by the relation in (5) and $\hat{g}^t(s)$ is defined in eq. (2). Precisely, the process $\hat{g}^t(s)$ in H_s^t and the associated return process $g(s)$ are respectively defined by

$$\hat{g}_\tau^t(s) = 1, \qquad g_\tau(s) = \frac{1}{M_{s,\tau}} \qquad \forall \tau \in [s,t].$$

As expected, the process $\hat{g}^t(s)$ is the one that permits the inner product representation of pricing functionals described at the end of Subsection 2.2. Moreover, $g(s)$ is the return process of the log-optimal portfolio. Hence, we refer to $g(s)$ as the *log-optimal return.*

In addition, the module H_s^t orthogonally decomposes as

$$H_s^t = \operatorname{span}_{L^0}\left\{\hat{g}^t(s)\right\} \oplus \mathring{H}_s^t.$$

3.3 The mean excess return *e(s)*

Fix again $t \in [s,T]$. From the definition of $\hat{f}^T(s)$ in eq. (6), we consider the conditional martingale $\hat{f}^t(s)$ associated with the pure discount T-bond and we define $\hat{e}^t(s)$ as the orthogonal projection of $\hat{f}^t(s)$ on the submodule $\mathring{H}_s^t$, namely

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

meaning that $\hat{e}_s^t(s) = 0$ and $\mathbb{E}_s[(\hat{f}_t^t(s) - \hat{e}_t^t(s))\hat{\imath}_t^t(s)] = 0$ for all $\hat{\imath}^t(s)$ in $\mathring{H}_s^t$. Since the orthogonal projection of $\hat{f}^t(s)$ on $\text{span}_{L^0}\{\hat{g}^t(s)\}$ is $\hat{g}^t(s)$, we have $\hat{f}^t(s) = \hat{e}^t(s) + \hat{g}^t(s)$ so that $\hat{e}_\tau^t(s) = G_{s,\tau} - 1$ for all $\tau \in [s,t]$. Moreover, $\mathring{H}_s^t$ decomposes as

$$\begin{aligned}
\mathring{H}_s^t &= \text{span}_{L^0}\left\{\hat{e}^t(s)\right\} \oplus \left\{\hat{n}^t(s) \in \mathring{H}_s^t : \quad \mathbb{E}_s\left[\hat{e}_t^t(s)\hat{n}_t^t(s)\right] = 0\right\} \\
&= \text{span}_{L^0}\left\{\hat{e}^t(s)\right\} \oplus \left\{\hat{n}^t(s) \in \mathring{H}_s^t : \quad \mathbb{E}_s\left[G_{s,t}\hat{n}_t^t(s)\right] = \mathbb{E}_s^F\left[\hat{n}_t^t(s)\right] = 0\right\}
\end{aligned}$$

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

$$e_\tau(s) = f_\tau(s) - g_\tau(s) \qquad \forall \tau \in [s,t], \tag{7}$$

which embodies the meaning of *mean excess return*. Such return is attainable because the log-optimal portfolio and the zero-coupon T-bond are traded.

3.4 Orthogonal decompositions of returns

The orthogonality in H_s^t implies an orthogonal decomposition of conditional martingales and, in turns, of asset returns. To achieve this goal, we start from the decomposition of conditional martingales.

Theorem 2 (Martingale decomposition) *Given $t \in [s,T]$, $\hat{u}^t(s)$ 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 $\hat{n}^t(s) \in \mathring{H}_s^t$ such that*

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

and

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

Proof of Theorem 2. We first show that

$$\mathbb{E}_s\left[\left(\hat{e}_t^t(s)\right)^2\right] = \mathbb{E}_s\left[G_{s,t}\hat{e}_t^t(s)\right] = var_s\left(G_{s,t}\right).$$

Indeed, since $\hat{e}^t(s) = \text{proj}_{\mathring{H}_s^t} \hat{f}^t(s)$, for any $\hat{\iota}^t(s) \in \mathring{H}_s^t$, we have $\mathbb{E}[(\hat{f}_t^t(s) - \hat{e}_t^t(s))\hat{\iota}^t(s)] = 0$. Then, the first equality follows when $\hat{\iota}^t(s) = \hat{e}^t(s)$. As for the second one,

$$\mathbb{E}_s\left[G_{s,t}\hat{e}_t^t(s)\right] = \mathbb{E}_s\left[G_{s,t}^2 - G_{s,t}\right] = \mathbb{E}_s\left[G_{s,t}^2\right] - 1 = \frac{\mathbb{E}_s\left[G_t^2\right]}{G_s^2} - \left(\frac{\mathbb{E}_s\left[G_t\right]}{G_s}\right)^2 = \frac{var_s\left(G_t\right)}{G_s^2}.$$

Now, 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 $\hat{n}^t(s) \in \mathring{H}_s^t$. 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 $\hat{n}^t(s)$ belong to $\mathring{H}_s^t$.

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 $\mathbb{E}_s[\hat{u}_t^t(s) - \hat{g}_t^t(s)] = 1 - 1 = 0$. Define the projection coefficient $\omega_s \in L^0(\mathcal{F}_s)$ by

$$\omega_s = \frac{\mathbb{E}_s\left[\left(\hat{u}_t^t(s) - \hat{g}_t^t(s)\right)\hat{e}_t^t(s)\right]}{\mathbb{E}_s\left[\left(\hat{e}_t^t(s)\right)^2\right]} = \frac{\mathbb{E}_s\left[G_{s,t}\hat{u}_t^t(s)\right] - 1}{\mathbb{E}_s\left[G_{s,t}\hat{e}_t^t(s)\right]} = \frac{\mathbb{E}_s\left[G_{s,t}\hat{u}_t^t(s)\right] - 1}{var_s\left(G_{s,t}\right)},$$

where last equalities are due to the definition of $\hat{e}^t(s)$ and its properties. 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,

$$\mathbb{E}_s\left[\hat{g}_t^t(s)\hat{n}_t^t(s)\right] = \mathbb{E}_s\left[\hat{g}_t^t(s)\hat{u}_t^t(s)\right] - \mathbb{E}_s\left[\left(\hat{g}_t^t(s)\right)^2\right] - \omega_s\mathbb{E}_s\left[\hat{g}_t^t(s)\hat{e}_t^t(s)\right] = 1 - 1 - 0 = 0$$

because $\hat{g}^t(s)$ and $\hat{e}^t(s)$ belong to orthogonal submodules. Furthermore,

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

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

A straightforward application of Theorem 2 delivers an orthogonal decomposition of asset returns in the time window $[s, t]$.

Corollary 3 (Return decomposition) *Let $t \in [s, T]$. If $u(s)$ is a return process in $[s, t]$, there exist $\omega_s \in L^0(\mathcal{F}_s)$ and $\hat{n}^t(s) \in \mathring{H}_s^t$ such that*

$$\mathbb{E}_s\left[M_{s,t}^2 g_t(s)n_t(s)\right] = \mathbb{E}_s\left[M_{s,t}^2 e_t(s)n_t(s)\right] = \mathbb{E}_s^F\left[M_{s,t}n_t(s)\right] = \mathbb{E}_s\left[M_{s,t}n_t(s)\right] = 0$$

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

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

Proof of Corollary 3. The return process $u(s)$ in $[s,t]$ can be associated with the conditional martingale $\hat{u}^t(s) \in H_s^t$ by the relation in (5). Then, the result follows directly from Theorem 2. In addition, $\mathbb{E}_s[M_{s,t}n_t(s)] = 0$ because $\hat{n}^t(s)$ belongs to $\mathring{H}_s^t$. ■

The proof of Theorem 2 exploits the definition of the projection coefficient ω_s in $L^0(\mathcal{F}_s)$, that turns out to be

$$\omega_s = \frac{\mathbb{E}_s^F[M_{s,t}u_t(s)] - 1}{var_s(G_{s,t})}. \tag{8}$$

Hence, ω_s depends on the expected return discounted by the log-optimal portfolio under the T-forward measure.

4 Risk-adjusted mean-variance returns

Following the standard mean-variance approach logic, we define risk-adjusted mean-variance returns in our setting. Here we decompose the whole processes of returns discounted by the log-optimal portfolio to obtain the time consistency property that we describe in Subsection 4.1. We, then, illustrate a useful Two-fund Separation Theorem.

Definition 4 *Fixed $t \in [s,T]$, we say that a return process $u(s)$ is on the* risk-adjusted mean-variance frontier *(or it is a* risk-adjusted mean-variance return*) in $[s,t]$ when it minimizes $var_s(M_{s,t}u_t(s))$ for some given $\mathbb{E}_s^F[M_{s,t}u_t(s)]$ in $L^0(\mathcal{F}_s)$. In that case, we say that the conditional martingale in H_s^t associated to $u(s)$ via the relation in* (5) *is a* conditional mean-variance martingale in $[s,t]$. *Such conditional martingale minimizes $var_s(\hat{u}_t^t(s))$ for the given $\mathbb{E}_s^F[\hat{u}_t^t(s)]$.*

As already mentioned, the use of expected returns under the measure F comes from the expression of portfolio weights in eq. (8). We begin with the characterization of conditional mean-variance martingales.

Theorem 5 (Mean-variance martingales) *Let $t \in [s,T]$. Consider $\hat{u}^t(s) \in H_s^t$ with $\hat{u}_s^t(s) = 1$ such that $\mathbb{E}_s^F[\hat{u}_t^t(s)] = k_s$ for some $k_s \in L^0(\mathcal{F}_s)$. Among them, the conditional martingale that minimizes $var_s(\hat{u}_t^t(s))$ is*

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) \qquad \textit{with} \qquad \omega_s = \frac{k_s - 1}{var_s(G_{s,t})}.$$

Proof of Theorem 5. Each conditional martingale $\hat{u}^t(s) \in H_s^t$ with $\hat{u}_s^t(s) = 1$ and $\mathbb{E}_s[\hat{u}_t^t(s)] = k_s$ satisfies the decomposition provided by Theorem 2:

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) + \hat{n}^t(s), \qquad \omega_s = \frac{k_s - 1}{var_s(G_{s,t})}.$$

Moreover, $var_s(\hat{u}^t_t(s)) = \mathbb{E}_s[(\hat{u}^t_t(s))^2] - (\mathbb{E}_s[\hat{u}^t_t(s)])^2 = \mathbb{E}_s[(\hat{u}^t_t(s))^2] - 1$. We note that

$$\begin{aligned}\mathbb{E}_s\left[\left(\hat{u}^t_t(s)\right)^2\right] &= \mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s) + \hat{n}^t_t(s)\right)^2\right] \\ &= \mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s)\right)^2\right] + \mathbb{E}_s\left[\left(\hat{n}^t_t(s)\right)^2\right] + 2\mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s)\right)\hat{n}^t_t(s)\right].\end{aligned}$$

By Theorem 2, $\mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s)\right)\hat{n}^t_t(s)\right] = 0$ and so

$$\mathbb{E}_s\left[\left(\hat{u}^t_t(s)\right)^2\right] = \mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s)\right)^2\right] + \mathbb{E}_s\left[\left(\hat{n}^t_t(s)\right)^2\right] \geqslant \mathbb{E}_s\left[\left(\hat{g}^t_t(s) + \omega_s\hat{e}^t_t(s)\right)^2\right].$$

Therefore, $var_s(\hat{u}^t_t(s))$ is minimized by the conditional martingale with $\hat{n}^t(s) = 0$. ■

From Theorem 5 we easily deduce the characterization of risk-adjusted mean-variance returns in $[s, t]$.

Corollary 6 (Risk-adjusted mean-variance returns) *Let $t \in [s, T]$. Consider return processes $u(s)$ in $[s, t]$ such that $\mathbb{E}^F_s[M_{s,t}u_t(s)] = h_s$ for some $h_s \in L^0(\mathcal{F}_s)$. Among them, the return process that minimizes $var_s(M_{s,t}u_t(s))$ is*

$$u(s) = g(s) + \omega_s e(s) \qquad \textit{with} \qquad \omega_s = \frac{h_s - 1}{var_s\left(G_{s,t}\right)}.$$

Proof of Corollary 6. Consider the conditional martingales in H^t_s associated with returns via the relation in (5). We have that $\mathbb{E}^F_s\left[\hat{u}^t_t(s)\right] = h_s$ and $var_s(\hat{u}^t_t(s)) = var_s(M_{s,t}u_t(s))$. Hence, the claim is an immediate consequence of Theorem 5 with $k_s = h_s$. ■

As an example, consider the zero-coupon T-bond return process $f(s)$ in $[s, T]$. By eq. (7), such return process satisfies $f(s) = g(s) + e(s)$ and so, by Corollary 6, $f(s)$ minimizes the conditional variance of any $M_{s,T}u_T(s)$ with $\mathbb{E}^F_s[M_{s,T}u_T(s)] = \pi_s(1_T)\mathbb{E}[G^2_{s,T}]$.

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

$$\omega_s = \frac{\tilde{h}_s - \mathbb{E}_s\left[g_t(s)\right]}{\mathbb{E}_s\left[e_t(s)\right]} = \frac{\tilde{h}_s - \mathbb{E}_s\left[g_t(s)\right]}{\mathbb{E}_s\left[f_t(s)\right] - \mathbb{E}_s\left[g_t(s)\right]}. \tag{9}$$

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 the risk-adjusted mean-variance frontier in $[s, T]$, then it is also on the risk-adjusted mean-variance frontier in $[s, t]$ for any $t \leqslant T$. This feature is ultimately due to the fact that the decomposition of Theorem 2 involves the whole process of the return discounted by the log-optimal portfolio in the time range $[s, T]$ and so there is a mechanical overlap with the decompositions built at shorter horizons. There is, then, an interplay between such overlaps and the definitions of orthogonality at different horizons.

We first deal with conditional mean-variance martingales. For simplicity, we express the result by using the time indices t and T but it clearly holds for any $t_1, t_2 \in [s, T]$ with $t_1 \leqslant t_2$.

Proposition 7 (Mean-variance martingales time consistency) *Let $t \in [s, T]$. 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 $[s, t]$.*

Proof of Proposition 7. 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 $var_s(\hat{u}^T_T(s))$ among all conditional martingales in H^T_s with $\mathbb{E}^F_s[\hat{u}^T_T(s)] = k_s$ for a given $k_s \in L^0(\mathcal{F}_s)$. By Theorems 2 and 5,

$$\hat{z}^T(s) = \hat{g}^T(s) + \omega_s \hat{e}^T(s), \qquad \omega_s = \frac{k_s - 1}{var_s\,(G_{s,T})}$$

and $\mathbb{E}_s[\hat{g}^T_T(s)\hat{e}^T_T(s)] = 0$. 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, $\mathbb{E}_s[\hat{g}^t_t(s)\hat{e}^t_t(s)] = 0$ because $\mathring{H}^t_s$ is orthogonal to $\mathrm{span}_{L^0}\{\hat{g}^t(s)\}$ and so we retrieve the orthogonal decomposition of $\hat{z}^t(s)$ provided by Theorem 2 in the time window $[s, t]$. In addition, $\omega_s = (k_s - 1)/var_s(G_{s,T}) = (h_s - 1)/var_s(G_{s,t})$, where $h_s \in L^0(\mathcal{F}_s)$ is univocally determined by k_s. Then, Theorem 5 ensures that $\hat{z}^t(s)$ minimizes $var_s(\hat{u}^t_t(s))$ among all conditional martingales in H^t_s such that $\mathbb{E}^F_s[\hat{u}^t_t(s)] = h_s$. In other words, $\hat{z}^t(s)$ is a conditional mean-variance martingale in $[s, t]$. ∎

We now establish the time consistency property of risk-adjusted mean-variance returns.

Corollary 8 (Risk-adjusted mean-variance returns time consistency) *Let $t \in [s, T]$. A risk-adjusted mean-variance return in $[s, T]$ is also a risk-adjusted mean-variance return in $[s, t]$.*

Proof of Corollary 8. Suppose that the return process $z(s)$ in $[s, T]$ minimizes $var_s(M_{s,T} u_T(s))$ among all return processes with $\mathbb{E}^F_s[M_{s,T} u_T(s)] = k_s$ for some given $k_s \in L^0(\mathcal{F}_s)$. By Corollaries 3 and 6,

$$z(s) = g(s) + \omega_s e(s), \qquad \omega_s = \frac{k_s - 1}{var_s\,(G_{s,T})},$$

with $\mathbb{E}_s[M^2_{s,T} g_T(s) e_T(s)] = 0$. The former decomposition holds algebraically at any time in $[s, t]$ and $\mathbb{E}_s[M^2_{s,t} g_t(s) e_t(s)] = 0$. 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 3. Furthermore, $\omega_s = (k_s - 1)/var_s(G_{s,T}) = (h_s - 1)/var_s(G_{s,t})$, where $h_s \in L^0(\mathcal{F}_s)$ is univocally determined by k_s. By Corollary 6, this means that $z(s)$ minimizes $var_s(M_{s,t}u_t(s))$ among all return processes with $\mathbb{E}_s^F[M_{s,t}u_t(s)] = h_s$. Hence $z(s)$ is a risk-adjusted mean-variance return in $[s,t]$, too. ■

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 risk-adjusted mean-variance frontiers, based on the information available today. Corollary 8 ensures that portfolios on the yearly frontier are also on the six-month frontier. This feature is absent in classical mean-variance analysis. 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 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.

We finally state and prove a Two-fund Separation Theorem for our risk-adjusted mean-variance frontiers. Theorem 9 establishes in our setting the celebrated result by Merton (1972) and makes the implementation of the frontiers easier by replacing the mean excess return $e(s)$ with the return $f(s)$ of the pure discount T-bond.

Theorem 9 (Two-fund Separation) *Given $t \in [s,T]$, $u(s)$ is a risk-adjusted mean-variance return in $[s,t]$ if and only if*

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

for some risk-adjusted mean-variance returns $v(s), z(s)$ in $[s,t]$ 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 6.

Proof of Theorem 9. Suppose that $u(s)$ is a risk-adjusted mean-variance return in $[s,t]$. Then, Corollary 6 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 $v(s)$ in $[s,t]$ 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 risk-adjusted mean-variance frontier, too.

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

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

and so, by Corollary 6, $u(s)$ is on the risk-adjusted mean-variance frontier in $[s,t]$.

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 $\alpha_s = 1 - \omega_s$.

■

In few words, $g(s)$ and $f(s)$ span the risk-adjusted mean-variance frontiers of asset returns at any horizon under consideration.

Note that our mean-variance optimization does not directly employ expectation and variance of asset returns $u(s)$: we rather use $\mathbb{E}_s^F[M_{s,t} u_t(s)]$ and $var_s(M_{s,t} u_t(s))$ to define risk-adjusted mean-variance returns in $[s,t]$. This approach actually conceals a measure change under which returns are evaluated. The whole theory can be rewritten by introducing a collection of risk-adjusted measures $\{Q^t\}_{t\in[s,T]}$ equivalent to P, each associated with a precise horizon t. Specifically, the Radon-Nikodym derivative of Q^t with respect to P on $\mathcal{F}_t$ is given by $dQ^t/dP = G_t^2/\mathbb{E}[G_t^2]$. The same risk-adjusted mean-variance frontier of Theorem 9 arises but the derivation is more complex than the one presented here: the measures Q^t affect the inner product, the orthogonality relations and so on.

5 Simulations: multi-horizon mean-variance optimization

To ease the notation and for the sake of interpretability, in this section we fix $s = 0$, we omit the s subscript whenever possible and we denote return processes by u instead of $u(s)$.

As sketched in Section 1, we consider a multi-horizon mean-variance portfolio problem in the time interval $[0,T]$, where only buy-and-hold investment strategies set at time 0 are allowed. Our investor may be thought as a manager or a company that aims at building portfolios with target expected returns across a sequence of maturities $t_1, t_2 \ldots, t_N$ with $0 < t_1 < t_2 < \cdots < t_N = T$. Each of these portfolios must be optimal in terms of the mean-variance criterion in its specific time horizon. The need to design such a term structure of portfolios may come from multi-horizon 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.

Specifically, the investor builds a portfolio with return process

$$\sum_{i=1}^{N} \lambda^{(i)} u^{(i)},$$

where each $\lambda^{(i)} \in \mathbb{R}$ is the weight of the sub-portfolio i, i.e. the one with return process $u^{(i)}$, in the overall portfolio. Each $u^{(i)}$ is properly a return process in $[0, t_i]$ and the position of the sub-portfolio i is liquidated at time t_i. Moreover, each $u^{(i)}$ solves

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

with $h^{(i)} \in \mathbb{R}$ given, for $i = 1, \ldots, N$. The weights $\lambda^{(i)}$ are positive, they sum up to 1 and, in case the overall portfolio is equally-weighted, $\lambda^{(i)} = 1/N$ for all i.

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

$$u_{t_i}^{(i)} = \frac{M_{0,t_i}}{\mathbb{E}\left[M_{0,t_i}^2\right]} + \widetilde{w}^{(i)} \left(1 - e^{-r_0^{t_i}(t_i - 0)} \frac{M_{0,t_i}}{\mathbb{E}\left[M_{0,t_i}^2\right]}\right), \qquad \widetilde{w}^{(i)} \in \mathbb{R}.$$

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 in $[0, t_i]$ as

$$u_{t_i}^{(i)} = \widetilde{\alpha}^{(i)} \frac{M_{0,t_i}}{\mathbb{E}\left[M_{0,t_i}^2\right]} + \left(1 - \widetilde{\alpha}^{(i)}\right) f_{t_i}$$

with $\widetilde{\alpha}^{(i)} = 1 - \pi_0(1_{t_i})\widetilde{w}^{(i)}$.[3]

For each horizon t_i, the initial implementation of the sub-portfolio delivering the return process $u^{(i)}$ in $[0, t_i]$ requires the replication, by self-financing portfolio strategies, of the payoff at t_i that coincides with the pricing kernel M_{0,t_i}. Considering the whole sequence of maturities in the problem, N payoffs need to be replicated in order to implement the mean-variance optimal asset allocation. Depending on the severity of market incompleteness, the optimal solution may require costly approximations or may even be unattainable.

Hereby, we propose an alternative strategy by exploiting our risk-adjusted mean-variance frontier. Although theoretically suboptimal, our frontier requires the replication of a single payoff at T (the one of log-optimal portfolio), whatever the number N of horizons involved.

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

$$f_{t_i} = \frac{M_{0,t_i}}{\mathbb{E}\left[M_{0,t_i}^2\right]} + \frac{1}{\pi_s\left(1_{t_i}\right)} \left(1 - e^{-r_0^{t_i}(t_i - 0)} \frac{M_{0,t_i}}{\mathbb{E}\left[M_{0,t_i}^2\right]}\right).$$

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 classical solution. From Corollary 6, for any $i = 1, \ldots, N$, we consider a sub-portfolio with return process

$$v^{(i)} = g + \omega^{(i)} e,$$

where $\omega^{(i)}$ is chosen so that the conditional expectation of $v^{(i)}_{t_i}$ meets the target $h^{(i)}$ as in eq. (9). By Theorem 9, we build our sub-portfolios by exploiting the return process g of the log-optimal portfolio and the return process f of the zero-coupon T-bond. These two financial instruments are employed for any intermediate horizon t_i, as a consequence of time consistency. We finally compare the performance of the two families of sub-portfolios with returns $u^{(i)}$ and $v^{(i)}$, respectively, by considering the transaction costs and their impact on the Sharpe ratios.

Specifically, we assume that transaction costs are present in the market and, similarly to Irle and Sass (2006), they are composed by trading and replication costs.

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 sub-portfolio i generates the trading costs $c(|\widetilde{\alpha}^{(i)}| + |1 - \widetilde{\alpha}^{(i)}|)$ with $c > 0$. The analogous expression with $\alpha^{(i)} = 1 - \omega^{(i)}$ delivers the trading costs of the risk-adjusted mean-variance return $v^{(i)}$.

As for the replication costs, we assume that the design of the replication strategies for g_T and $M_{0,t_i}/\mathbb{E}[M^2_{0,t_i}]$ at any horizon t_i entails a positive *fixed* cost C for any (possibly linearly independent) security. Therefore, the implementation of each classical mean-variance sub-portfolio i requires the additional expenditure of C. On the contrary, if we proportionally spread the replication cost of g_T across the horizons $t_1, \ldots, t_N$, each time-consistent sub-portfolio i needs to bear the cost $\lambda^{(i)} C$. As a result, each mean-variance optimal sub-portfolio i and each risk-adjusted mean-variance sub-portfolio i have commissions, respectively,

$$C + c\left(\left|\widetilde{\alpha}^{(i)}\right| + \left|1 - \widetilde{\alpha}^{(i)}\right|\right) \qquad \text{and} \qquad \lambda^{(i)} C + c\left(\left|\alpha^{(i)}\right| + \left|1 - \alpha^{(i)}\right|\right). \tag{10}$$

Accordingly, the overall portfolios have commissions:

$$CN + c\sum_{i=1}^{N} \lambda^{(i)}\left(\left|\widetilde{\alpha}^{(i)}\right| + \left|1 - \widetilde{\alpha}^{(i)}\right|\right) \qquad \text{and} \qquad C + c\sum_{i=1}^{N} \lambda^{(i)}\left(\left|\alpha^{(i)}\right| + \left|1 - \alpha^{(i)}\right|\right).$$

In terms of risk/return trade-off, at any horizon t_i we consider a *modified Sharpe ratio* given by the difference of the Sharpe ratio and the ratio between transaction costs (as percentage of the initial capital) and standard deviation. In this way, the expected return of each sub-portfolio i is reduced by the proper commissions of eq. (10):

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

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 risk-adjusted mean-variance optimal strategies, making the risk-adjusted approach valuable. This happens in the simulations of Subsections 5.2 and 5.3. Subsection 5.1 describes the market in which we set such simulations.

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 $[0, T]$ 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

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

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 B.

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

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

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

$$\begin{cases} dX_t = X_t Y_t \; dt + \eta X_t \left[\phi dW_t^Q + \sqrt{1-\phi^2} dZ_t^Q\right] \\ d\pi_t\left(1_T\right) = \pi_t\left(1_T\right) Y_t dt - \pi_t\left(1_T\right) B(t,T)\sigma dW_t^Q. \end{cases} \tag{11}$$

At the same time, under the physical measure,

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

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

$$\nu_t^W = -\frac{\mu_t^P - Y_t}{B(t,T)\sigma}, \qquad \nu_t^Z = \frac{\mu_t^X - Y_t - \eta\phi\nu_t^W}{\eta\sqrt{1-\phi^2}}. \tag{12}$$

At any $t \in [0,T]$, the Radon-Nikodym derivative of Q with respect to P on $\mathcal{F}_t$ is

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

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

$$dM_{0,t} = -Y_t M_{0,t} dt - \nu_t^W M_{0,t} dW_t^P - \nu_t^Z M_{0,t} dZ_t^P.$$

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 consider an equally-weighted portfolio over six horizons: $N = 6$ and $\lambda^{(i)} = 1/N$ for all $i = 1, \ldots, 6$. We employ a monthly time grid and horizons $t_1, \ldots, t_6$ associated with six subsequent semesters. We set the target means equal to $h^{(i)} = 1.06$ for $i = 1, \ldots, 6$. 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 by investing in 6 equally weighted buy-and-hold sub-portfolios built at time 0. The cashflows obtained at the end of each semester from the liquidation of the related sub-portfolio are not re-invested.

We simulate both the classical and the risk-adjusted multi-period 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

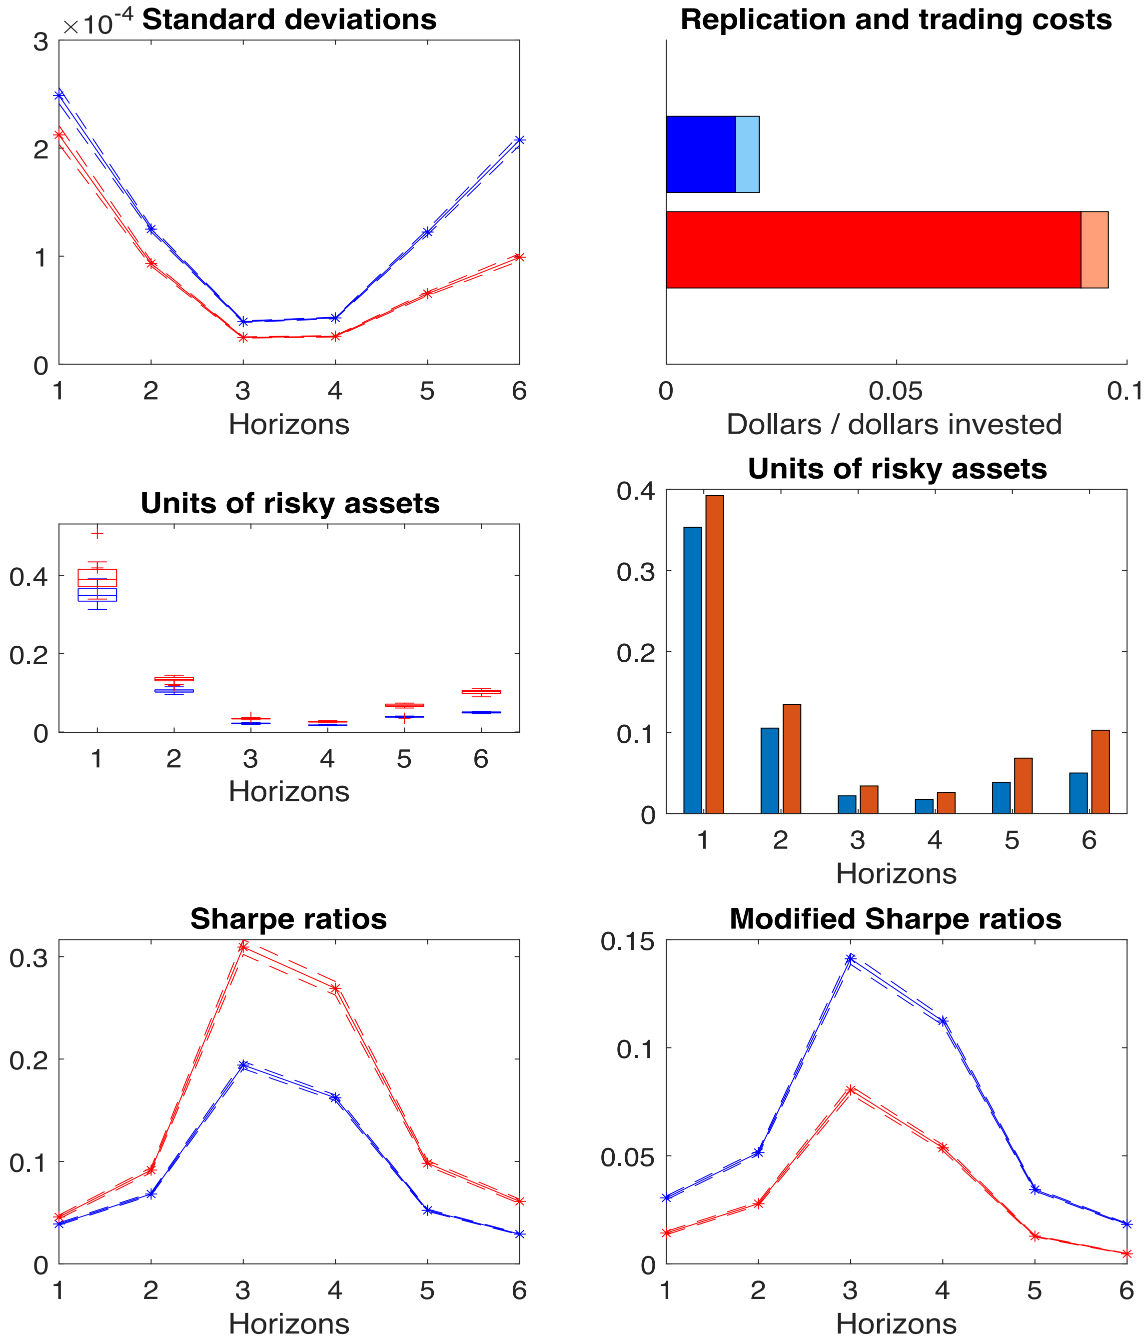

Figure 1: Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the problem of Subsection 5.2. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \dots, 6$. 90% confidence intervals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

sub-portfolio i with return process $u^{(i)}$ or $v^{(i)}$ and horizon t_i, for $i = 1, \ldots, 6$. 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 $\lambda^{(i)}$. Every simulated sub-portfolio matches perfectly the target means $h^{(i)}$ at the proper horizon for $i = 1, \ldots, 6$. As predicted by the theory, classical mean-variance sub-portfolios display lower standard deviations than our risk-adjusted strategies, whose advantage relies on a parsimonious implementation.

In our simulations the loadings of the risk-adjusted sub-portfolios are smaller than the ones of the classical mean-variance strategies, requiring to buy or sell fewer assets. We visualize this fact in the medium panels of Figure 5.2, where we plot the absolute values of $\alpha^{(i)}$ and $\widetilde{\alpha}^{(i)}$ at each horizon t_i. The graphs depict the units of risky assets - i.e. the ones associated with g_T and $M_{0,t_i}/\mathbb{E}[M_{0,t_i}^2]$ respectively - contained in each sub-portfolio. The exposure to the risky securities is higher at horizons near in time. However, at any horizon, the loadings in the risk-adjusted sub-portfolios are lower than the ones in the classical sub-portfolios (with slightly lower dispersion). Consequently, the implementation of the portfolio with return processes $v^{(i)}$ involves narrower long (or short) positions, both in g and in f, 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 risk-adjusted mean-variance frontier.

The commission shrinkage of the risk-adjusted 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 risk-adjusted solution is the best performing. The excess standard deviation of the risk-adjusted portfolio is fully compensated by its reduced transaction costs (in particular, replication costs), as captured by the modified Sharpe ratio.

5.3 A life annuity application

Still in the market of Subsection 5.1, we compare the risk-adjusted 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 the subscriber dies. The insurance company invests

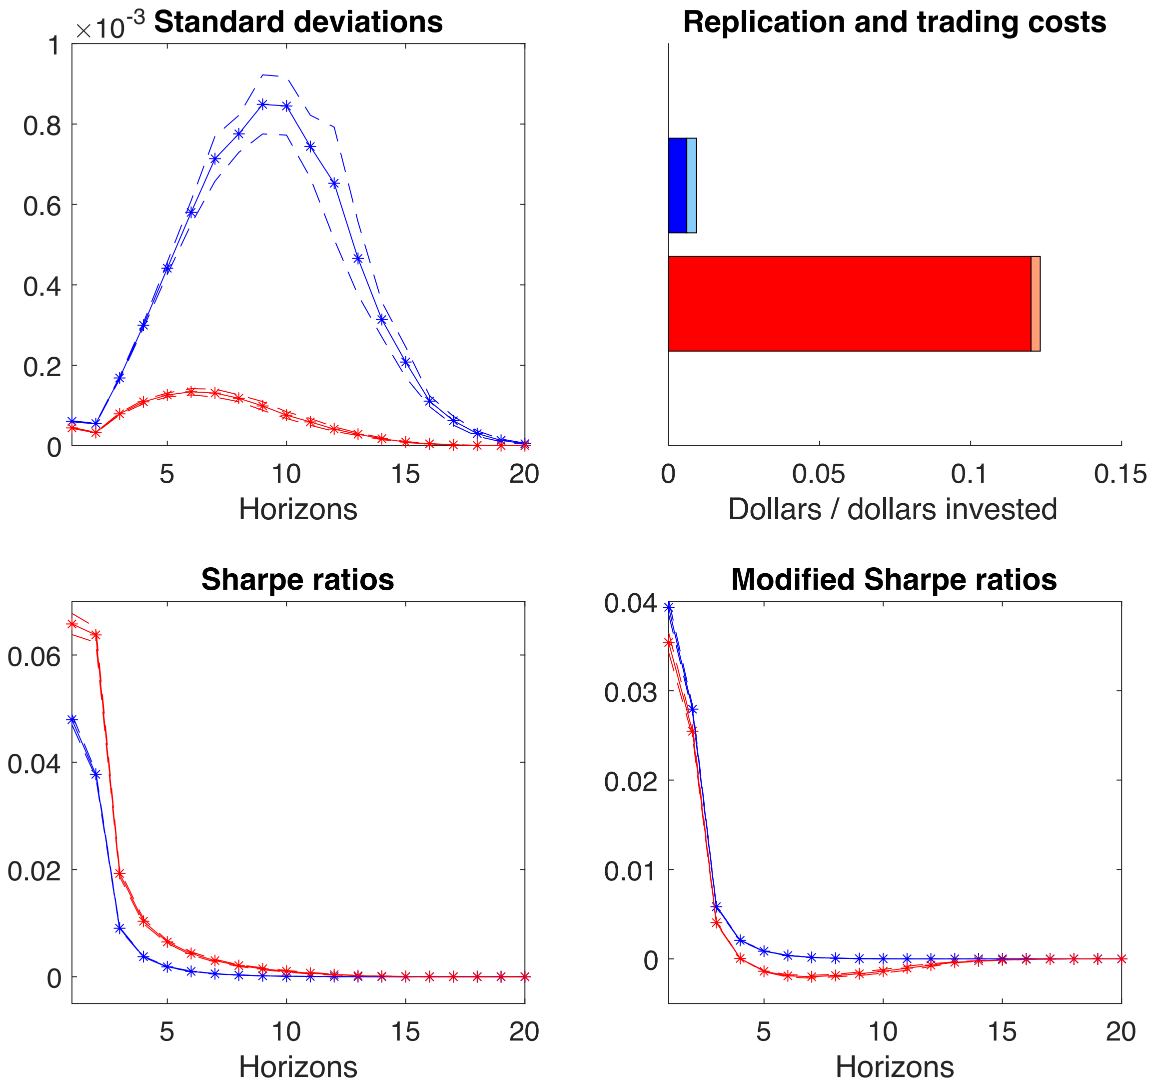

Figure 2: Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the life-annuity problem. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \ldots, 20$. 90% confidence intervals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

the received capital in N sub-portfolios with increasing horizons that allow to meet the future payments. For example, we can assume that each sub-portfolio has target return $h^{(i)} = 1.05$ for $i = 1, \dots, N$ with $N = 20$ years.

The random variable *time-until-death* captures the difference between the insured's age at death and the 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 sub-portfolio i depends on the proportion of survivors at the horizon-year t_i, i.e.

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

If the company aims at reducing the risk of each sub-portfolio, it can consider a (classical or risk-adjusted) mean-variance approach for each return process $u^{(i)}$ satisfying $\mathbb{E}[u_{t_i}^{(i)}] = 1.05$ 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 $\lambda^{(i)}$ 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 risk-adjusted sub-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 risk-adjusted 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 risk-adjusted strategy. Nevertheless, the introduction of the commissions reverses the conclusion: the classical mean-variance optimal portfolio turns out to have a lower (and sometimes negative) modified Sharpe ratio. This effect is mostly due to the number of payments in the life annuity contract, which requires the replication of many risky securities.

6 Mean-variance frontier and optimal consumption-investment

We provide a microeconomic foundation of the risk-adjusted mean-variance frontier of returns described by Corollary 6. 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 the return processes $g(s)$ and $f(s)$ in agreement with Theorem 9. Moreover, an analogue of time consistency property of risk-adjusted mean-variance returns can be retrieved in optimal investment policies.

6.1 Optimal consumption-investment problem

We consider the optimization problem of an agent that decides a consumption policy $c = \{c_\tau\}_{\tau\in[s,T]}$. The agent 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 the initial wealth by selecting a payoff stream (or wealth profile) with value $w = \{w_\tau\}_{\tau\in[s,T]}$ and, at any instant τ, the agent 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

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

where the process $b = \{b_\tau\}_{\tau\in[s,T]}$ defines a time-varying adapted bliss point. Moreover, the investor deflates the consumption c_τ by exploiting the pricing kernel $M_{s,\tau}$. This attitude reflects the use of returns discounted by the log-optimal portfolio in Sections 3 and 4. The intertemporal consumption-investment optimization problem to solve is

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

The related reduced form is

$$\max_w \ \mathbb{E}_s\left[\int_s^T U(i_\tau + w_\tau)\, d\tau\right] \qquad \text{sub } w_s = \mathbb{E}_s\left[\int_s^T M_{s,\tau} w_\tau d\tau\right]. \tag{13}$$

Proposition 10 *If in Problem* (13) *the income stream is null and the bliss point is*

$$b_\tau = \frac{\beta_s \pi_\tau(1_T)}{T-s} M_{s,\tau} \qquad \forall \tau \in [s,T]$$

with $\beta_s \in L^0(\mathcal{F}_s)$, then the optimal payoff stream w^* defines the risk-adjusted mean-variance return in $[s,T]$ given by

$$\frac{(T-s)w^*}{w_s} = \frac{\beta_s \pi_s(1_T)}{w_s} f(s) + \left(1 - \frac{\beta_s \pi_s(1_T)}{w_s}\right) g(s).$$

Proof of Proposition 10. The Lagrangian function is

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

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 times $\tau \in [s,T]$ and states $\omega \in \Omega$. The first-order condition implies that (at any time and in any state) $U'\left(i_\tau + w_\tau\right) - \lambda_s M_{s,\tau} = 0$. Therefore,

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

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,

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

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

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

Consequently, the optimal payoff stream w^* is associated with the return process defined, for all $\tau \in [s,T]$, by

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

which lies on the risk-adjusted mean-variance frontier in $[s,T]$ by Theorem 9. ■

6.2 Time consistency of optimal cashflows

Inspired by the time consistency of the risk-adjusted mean-variance frontier shown in Corollary 8, we investigate whether a similar feature is kept in the optimal consumption-investment problem. Specifically, once Problem (13) is solved by a payoff stream $w^* = \{w^*_\tau\}_{\tau\in[s,T]}$ on the time interval $[s,T]$, we assess whether the restriction of w^* is also optimal on the subperiod $[s,t]$ with $t \leqslant T$. In particular, we consider the problem

$$\max_{w} \ \mathbb{E}_s\left[\int_s^t U\left(i_\tau + w_\tau\right) d\tau\right] \qquad \text{sub } \tilde{w}_s = \mathbb{E}_s\left[\int_s^t M_{s,\tau} w_\tau d\tau\right], \tag{14}$$

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

Proposition 11 *Under the assumptions of Proposition 10, if w^* solves Problem* (13) *with initial wealth w_s, then it also solves Problem* (14) *with initial wealth*

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

Proof of Proposition 11. Following the same steps as in the proof of Proposition 10, 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, for any $\tau \in [s,t]$, the optimal payoff stream is

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

and it coincides with the one prescribed by Proposition 10. ■

The risk-adjusted 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 to obtain 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).

An analogous reasoning to Proposition 11 shows that w^* is optimal also on the terminal subperiod $[t, T]$, according to

$$\max_{w} \ \mathbb{E}_s \left[\int_t^T U\left(i_\tau + w_\tau\right) d\tau \right] \qquad \text{sub } \hat{w}_s = \mathbb{E}_s \left[\int_t^T M_{s,\tau} w_\tau d\tau \right], \tag{15}$$

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

Corollary 12 *Under the assumptions of Proposition 10, if w^* solves Problem* (13) *with initial wealth w_s, then it also solves Problem* (15) *with*

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

Although Problem (15) 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 property of w^* that we show requires, in fact, the same information set. This approach is in line with 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 period. Hence, our construction is consistent with a rational inattention approach, as described in Sims (2003) or Abel, Eberly, and Panageas (2013). Indeed, one can assume that our investor makes a decision at time s for the whole period $[s, T]$ because of 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.

7 Conclusions

We obtain a conditional orthogonal decomposition of asset return processes in the spirit of Hansen and Richard (1987) by employing the series of returns discounted by the log-optimal portfolio. The associated risk-adjusted mean-variance frontier features an important time consistency property, with practical advantages for multi-period 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.

S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23(8):2970–3016, 2010.

G. Bekaert and J. Liu. Conditioning information and variance bounds on pricing kernels. *Review of Financial Studies*, 17(2):339–378, 2004.

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, 3 edition, 2009.

N.L. Bower, H.U. Gerber, J.C. Hickman, D.A. Jones, and C.J. Nesbitt. *Actuarial mathematics.* The Society of Actuaries, Schaumburg, Illinois, 2 edition, 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*, 446(1):970–1017, 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. 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.

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.

C.A. Favero, F. Ortu, A. Tamoni, and H. Yang. Implications of return predictability for consumption dynamics and asset pricing. *Journal of Business & Economic Statistics*, 38 (3):527–541, 2020.

W.E. Ferson and A.F. Siegel. The efficient use of conditioning information in portfolios. *Journal of Finance*, 56(3):967–982, 2001.

W.E. Ferson and A.F. Siegel. Stochastic discount factor bounds with conditioning information. *Review of Financial Studies*, 16(2):567–595, 2003.

A.R. Gallant, L.P. Hansen, and G. Tauchen. Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics*, 45(1-2):141–179, 1990.

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 R. Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991.

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.

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.

J.B. Jr Long. The numeraire portfolio. *Journal of Financial Economics*, 26(1):29–69, 1990.

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.

I.W.R. Martin and C. Wagner. What is the expected return on a stock? *Journal of Finance*, 74(4):1887–1929, 2019.

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.

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.

F. Severino. Long-term risk with stochastic interest rates. $working\ paper\ http://ssrn.com/abstract = 3113718$, 2021.

W.F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442, 1964.

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.

P.M. Van Staden, D.-M. Dang, and P.A. Forsyth. Time-consistent mean–variance portfolio optimization: A numerical impulse control approach. *Insurance: Mathematics and Economics*, 83:9–28, 2018.

O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

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 Forward measure and numéraire changes

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

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

See Theorem 1 and Example 2 in Geman, El Karoui, and Rochet (1995). Moreover,

$$J_t = \mathbb{E}_t\left[L_{t,T} J_T\right] = e^{r_0^T T - r_t^T (T-t) - \int_0^t Y_\tau d\tau} \qquad \forall t \in [0,T]$$

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

$$G_t = \mathbb{E}_t\left[G_T\right] = \mathbb{E}_t\left[L_T J_T\right] = L_t J_t \qquad \forall t \in [0,T]$$

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

B Additional simulations: 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 horizons. Generalizations with a higher number of assets can be developed in a similar way.

In the system of equations (11) 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 T-bond, while χ gives the instantaneous correlation with the old stock:

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

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. (12) 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 B, where we represent (scaled) standard deviations, (scaled) Sharpe ratios and (scaled) modified Sharpe ratios across horizons, transaction costs and units of risky assets in each sub-portfolio, where risky assets coincide with the log-optimal portfolio (in the risk-adjusted approach) and the portfolio replicating the pricing kernel (in the classical frontier). As the modified Sharpe ratio shows, in this simulation the risk-adjusted approach outperforms the standard mean-variance optimization when replication and trading costs are taken into account.

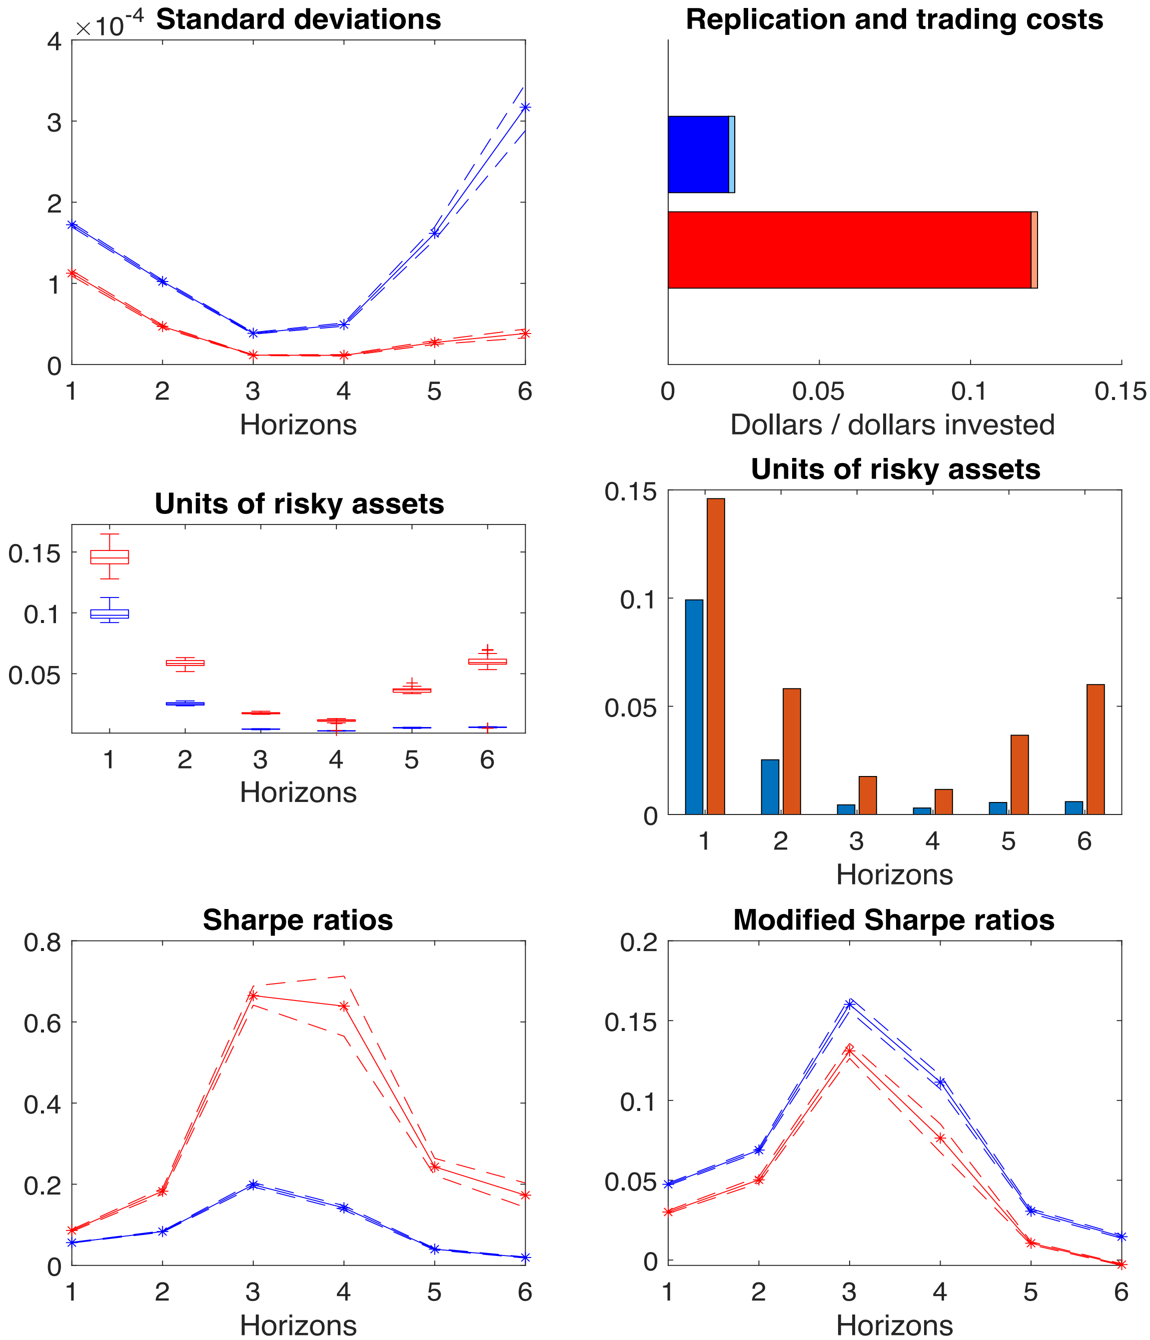

Figure 3: Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the problem of Appendix B. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \ldots, 6$. 90% confidence intervals for these variables are represented. The top-right panel represents the transaction costs of the risk-adjusted portfolio (blue for replication costs, light blue for trading costs) and of the classical mean-variance portfolio (red for replication costs, light red for trading costs). Medium panels contain the box-and-whisker plot at 25^{th} and 75^{th} percentiles and the bar plot of loadings $|\alpha^{(i)}|$ and $|\widetilde{\alpha}^{(i)}|$ at all horizons.