

On horizon-consistent mean-variance portfolio allocation

Simone Cerreia-Vioglio^{1†}, Fulvio Ortu^{2†}, Francesco
Rotondi^{3,4†} and Federico Severino^{5*†}

¹Department of Decision Sciences and IGIER, Università
Bocconi, Via Roentgen 1, Milano, 20136, MI, Italy.

²Department of Finance and IGIER, Università Bocconi, Via
Roentgen 1, Milano, 20136, MI, Italy.

³Department of Mathematics, Università degli Studi di Padova,
Via Trieste 63, Padova, 35121, PD, Italy.

⁴Department of Finance, Università Bocconi, Via Roentgen 1,
Milano, 20136, MI, Italy.

⁵Department of Finance, Insurance and Real Estate, Université
Laval, 2325 Rue de la Terrasse, Québec, G1V0A6, QC, Canada.

*Corresponding author(s). E-mail(s):

federico.severino@fsa.ulaval.ca;

Contributing authors: simone.cerreia@unibocconi.it;

fulvio.ortu@unibocconi.it; francesco.rotondi@math.unipd.it;

†These authors contributed equally to this work. This version of the article has been accepted for publication on *Annals of Operations Research* on May 23, 2022, after peer review but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <http://dx.doi.org/10.1007/s10479-022-04798-x>. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>.

Abstract

We analyze the problem of constructing multiple buy-and-hold mean-variance portfolios over increasing investment horizons in continuous-time arbitrage-free stochastic interest rate markets. The orthogonal

approach to the 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 horizon 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.

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

JEL Classification: G11 , G12

1 Motivations and main results

The mean-variance approach to portfolio optimization first formalized by the seminal work of Markowitz (1952) is a cornerstone in finance theory. In the standard formulation of the problem, an investor at time zero has to build a buy-and-hold portfolio to be liquidated at a given fixed investment horizon. When setting up this portfolio, the investor sets the desired portfolio's expected return and tries to minimize its variance. For any possible expected return, the optimal minimum variance portfolio lies on the so-called mean-variance frontier. However, if the investor at time zero needs to set up many optimal (in this mean-variance sense) buy-and-hold portfolios to be liquidated at different investment horizons, the standard one-period approach is of no help and the investor should solve many separate problems dealing with one investment horizon per time. Working on this limitation, in this paper we propose a novel approach to *multi-horizon* mean-variance portfolio allocation. More precisely, we formally describe a new way to solve a static multi-horizon portfolio allocation problem as a whole, rather than as a set of separate problems. Despite our approach turns out to be slightly suboptimal in a frictionless financial market, it does prove to be competitive when realistic trading and replication costs are accounted for.

This multi-horizon portfolio allocation problem is of utmost importance for insurance companies, pension funds and any financial intermediary managing long-term multiple cashflows, such as annuities. For example, consider

an investor who wants to meet N expected return targets at N subsequent horizons by investing in N buy-and-hold portfolios.¹ These portfolios have to attain the targets while displaying the minimum possible variance each (the proper formalization of such an example is in Section 5). According to the standard one-period mean-variance approach, the investor should solve N different problems that would lead to N optimal portfolios. In general, the components of these N optimal portfolios are going to be completely different. On the contrary, following our approach, the investor will still need to build N different buy-and-hold portfolios, but the components of all of these portfolios will be always the same. When transaction costs (and, in particular, replication ones) are taken into account, our solution leads to substantial savings.

Our approach is based on two fundamental building blocks of modern finance theory: the orthogonal characterization of the mean-variance frontier, as first derived by Hansen and Richard (1987), and the properties of the log-optimal portfolio when used as a numéraire, as derived by Long (1990). In their celebrated paper, Hansen and Richard (1987) solve the standard one-period mean-variance allocation problem providing an orthogonal decomposition of the set of all attainable portfolio returns. Using this decomposition, they describe the returns of the portfolios on the mean-variance frontier as linear combinations of only two fundamental ones: the return associated to the only traded stochastic discount factor and the one associated with the risk-free security. However, when taking a multi-horizon perspective, this approach suffers from the same limitation described above: returns that lie on the mean-variance frontier at T generally do not exhibit this desirable property at any intermediate date t before T . As a consequence, frontiers at different horizons are generated by different portfolios.

To tackle this issue, we propose an orthogonal decomposition of returns expressed in units of the log-optimal portfolio. This portfolio, which is the one that maximizes the expected value of the terminal wealth of a log-utility investor, was first introduced by Kelly (1956), while its properties if used as a numéraire were formalized by Long (1990). Using this portfolio as a numéraire, we obtain a new mean-variance frontier, that we call *risk-adjusted* frontier, which is spanned by the same two securities (a risky one and a riskless one) irrespective of the time horizon. Considering again the multi-horizon problem above, according to our risk-adjusted approach, the investor has to replicate only one fundamental security as all the N optimal portfolios involve different units of the same assets. As a result, after incorporating transaction costs in the analysis, risk-adjusted mean-variance portfolios can display a higher Sharpe ratio than classical mean-variance portfolios. Numerical examples of the magnitude of these savings are given in Section 5, in the contexts of fixed-income markets and life annuities.

¹In a multi-horizon portfolio allocation problem, the restriction to buy-and-hold portfolios only might seem too strict. However, when transaction costs are taken into account, rebalancing strategies might become quite expensive and possibly suboptimal, like in the case of replication of a simple derivative in Soner et al. (1995).

To give a snapshot of our construction, we consider a continuous-time arbitrage-free market with finite horizon T , stochastic interest rates and 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 et al. (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 2 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

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

in the spirit of Hansen and Richard (1987). Here $g(s)$ is the so-called *log-optimal return*, namely the return of the log-optimal portfolio; $e(s)$ is the *mean excess return*, namely the difference between the zero-coupon T -bond return and the log-optimal one; $n(s)$ is an additional zero-price return that represents idiosyncratic risk 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 Theorem 3). As risk-adjusted mean variance returns can be represented as linear combinations of just two out of three components of the orthogonal decomposition derived above, a Two-fund Separation Theorem holds (Theorem 5) 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 $t \leq T$ in an analogous way and the resulting risk-adjusted mean variance frontier at t is always spanned by the same returns $g(s)$ and $f(s)$, namely the returns on the very same portfolios (the log-optimal portfolio and the pure discount T -bond). 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 *horizon 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 horizon consistency property holds for our mean-variance returns: returns on the risk-adjusted mean-variance frontier at horizon T are risk-adjusted mean-variance returns at horizon t , too (Proposition 4). 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 horizon-consistent. 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 et al. (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 horizon consistency property that allows to generate optimal returns via the same two securities across a sequence of maturities. The literature mainly concentrates on time consistency of portfolio or consumption choices, which 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 et al. (2018) provide a detailed summary of the two literature streams – one related to precommitment, the other to time consistency. The mean-variance theory proposed here is peculiar in this respect. Indeed, it shares some aspects of both streams: the horizon consistency transfers the logic of time consistency to the investment horizon 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 maturities.

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 et al. (1990), among the others. See, e.g. the review in Favero et al. (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 $B_t = e^{\int_0^t Y_\tau d\tau}$ at any time t . Pure discount bonds with any possible maturity and face value equal to 1 are traded. Let $\{\pi_t(1_T)\}_{t \in [0, T]}$ denote the price of a pure discount T -bond at time 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 .

Additional risky securities, with adapted price processes, can be traded in the market. While completeness of the financial market is not required here, we assume that the *log-optimal portfolio* (or *growth optimal portfolio* or *numéraire portfolio*) is traded in the market. This self-financing portfolio, first introduced by Kelly (1956), is the one that maximizes the expected utility on the terminal wealth of a log-utility investor with initial wealth equal to 1 (see also Chap. 20 of Björk, 2009, for the formal definition and the properties of this portfolio in a continuous time framework). Let $N = \{N_t\}_{t \in [0, T]}$ be the price process of the log-optimal portfolio at any t . This portfolio has a very important property: prices of traded securities expressed in units of the log-optimal portfolio are martingales under P (see Long, 1990, and Sect. 26.9 in Björk, 2009). Under this assumption, the market is free of arbitrage opportunities. We refer to Subsection 2.1.1 for the explicit construction of the log-optimal portfolio via self-financing trading in a generic financial market.

Along with the log-optimal portfolio, another standard choice for the numéraire is the money market account B . Using this as numéraire, we obtain a *risk-neutral measure*, under which prices of traded securities denominated in

units of the money market account B are martingales. Notice that, since we are not assuming market completeness, there could be infinitely many risk-neutral measures.

There is yet another standard way to denominate prices. As a third different numéraire, we consider the pure discount T -bond, with price $\pi_t(1_T)$. Using this as numéraire, we obtain the *forward measure* with horizon T (or T -forward measure). See Geman et al. (1995).

It is now important to identify one precise risk-neutral (and T -forward) measure and link it to the original physical measure. As a consequence of the first order conditions of the optimal investment problem underlying the log-optimal portfolio, the inverse of its value process is a stochastic discount factor for the market. We denote this stochastic discount factor by $M = \{M_t\}_{t \in [0, T]}$ and we set $M_{t, T} = M_T/M_t$. By construction, clearly $M_{0, t} = N_t^{-1}$.

Among the possibly infinitely many risk-neutral measures, we label by Q the only one whose Radon-Nikodym derivative w.r.t. P on \mathcal{F}_T , $L_T = dQ/dP$, satisfies $L_t = B_t/N_t$, or, equivalently, $M_t = e^{-\int_0^t Y_\tau d\tau} L_t$, with $L_t = \mathbb{E}_t[L_T]$.² 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 . Moreover, we define $L_{t, T} = L_T/L_t$ at any time $t \in [0, T]$.

We also select in the same way a precise T -forward measure. In particular, we label by F the only forward measure whose Radon-Nikodym derivative with respect to P on \mathcal{F}_T , $G_T = dF/dP$, satisfies $G_t = e^{r_0^T T - r_t^T (T-t)}/N_t$, or, equivalently, $M_t = e^{r_t^T (T-t) - r_0^T T} G_t$, with $G_t = \mathbb{E}_t[G_T]$. Notice that, since $e^{-\int_0^T Y_\tau d\tau} L_T$ belongs to $L^2(\mathcal{F}_T)$, G_T belongs to $L^2(\mathcal{F}_T)$. We also set $G_{t, T} = G_T/G_t$. Further details on these changes of measure are provided in App. A.

As this precise forward measure F will play a key role in the paper, we point out that, under F , the pricing kernel in any interval $[s, t]$ with $s \leq t \leq T$ becomes

$$M_{s, t} = e^{r_t^T (T-t) - r_s^T (T-s)} G_{s, t}. \quad (1)$$

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

$$\pi_t(h_T) = \mathbb{E}_t[M_{t, T} h_T] = e^{-r_t^T (T-t)} \mathbb{E}_t^F[h_T]. \quad (2)$$

We find it useful to summarize the different numéraires and the related probability measures we introduced so far in Table 1.

2.1.1 The log-optimal portfolio construction

Here we provide the recipe to construct the self-financing strategy whose value is the log-optimal portfolio $N_t = M_{0, t}^{-1}$ in a rather general setting.³ We

²Here, \mathbb{E}_t denotes the conditional expectation with respect to \mathcal{F}_t under the measure P .

³We point out that, here and throughout the paper, we adopt the somehow generic notation M_t for the only stochastic discount factor whose value process equals the inverse of the value process of the log-optimal portfolio.

8 *On horizon-consistent mean-variance portfolio allocation*

Numéraire	Symbol	Equivalent Martingale Measure	Radon-Nikodym derivative w.r.t. P on \mathcal{F}_t	Equivalent expression for M_t
Log-optimal portfolio	N_t	P	-	-
Money market account	B_t	Q	L_t	$e^{-\int_0^t Y_\tau d\tau} L_t$
T -zero coupon bond	$\pi_t(1_T)$	F	G_t	$e^{r_t^T(T-t) - r_0^T T} G_t$

Table 1 Summary of the numéraires and the related Equivalent Martingale Measures involved in the paper.

assume that the pricing kernel $M_{0,t}$ is the continuous Itô semimartingale with dynamics

$$\frac{dM_{0,t}}{M_{0,t}} = -Y_t dt - \boldsymbol{\nu}'_t d\mathbf{W}_t^P,$$

where $\boldsymbol{\nu} = \{\boldsymbol{\nu}_t\}_{t \in [0, T]}$ is a d -dimensional adapted process with entries $\nu^{(i)} = \{\nu_t^{(i)}\}_{t \in [0, T]}$ such that $\int_0^T \mathbb{E}[(\nu_t^{(i)})^2] dt < +\infty$ for $i = 1, \dots, d$, and $\mathbf{W}^P = \{\mathbf{W}_t^P\}_{t \in [0, T]}$ is a d -dimensional independent Wiener process. The vector $\boldsymbol{\nu}$ represents, as usual, the market price of risk.

By applying the multidimensional Itô's formula (Theorem 4.16 in Björk, 2009) to the function $\varphi(t, M_{0,t}) = M_{0,t}^{-1} = N_t$, we get

$$\frac{dN_t}{N_t} = [Y_t + \boldsymbol{\nu}'_t \boldsymbol{\nu}_t] dt + \boldsymbol{\nu}'_t d\mathbf{W}_t^P. \quad (3)$$

To construct the log-optimal portfolio one needs to rewrite the previous expression in terms of the infinitesimal price increments of the traded securities. For instance, from the money market account dynamics $dB_t = Y_t B_t dt$, it is immediate to retrieve $dt = dB_t / (Y_t B_t)$. If d risky securities with values $\mathbf{X} = \{\mathbf{X}_t\}_{t \in [0, T]}$ are traded, we can find the adapted processes $\boldsymbol{\theta}^B = \{\theta_t^B\}_{t \in [0, T]}$ and (the d -dimensional) $\boldsymbol{\theta} = \{\boldsymbol{\theta}_t\}_{t \in [0, T]}$ such that

$$dN_t = \theta_t^B dB_t + \boldsymbol{\theta}'_t d\mathbf{X}_t.$$

This equation represents the dynamics of the self-financing portfolio that provides the log-optimal portfolio by investing in θ_t^B units of the money market account and in $\boldsymbol{\theta}_t$ units of the risky assets. An example with explicit formulas can be found at the end of Subsection 5.1.

2.2 The Hilbert modules H_s^t and linear pricing functionals

In 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 et al. (2016) show that $L_s^1(\mathcal{F}_t)$ is an L^0 -module with the multiplicative decomposition $L_s^1(\mathcal{F}_t) = L^0(\mathcal{F}_s)L^1(\mathcal{F}_t)$.⁴

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 $\mathbb{E}_s[|\hat{z}(\tau)|] \in L^0(\mathcal{F}_s)$ for all $\tau \in [s, t]$. See, e.g., Chap. VII, Sect. 1 of Shiryaev (1996). For any $t \in [s, T]$ we define the space

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

H_s^t contains the price processes discounted by the log-optimal portfolio with the appropriate square-integrability condition.⁵ 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 $H_s^{t_1}$.⁶

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

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

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 management. Moreover, the inner product structure delivers a natural notion of orthogonality: two processes \hat{z}, \hat{v} in H_s^t are orthogonal when $\langle \hat{z}, \hat{v} \rangle_s^t = \mathbb{E}_s[\hat{z}_t \hat{v}_t] = 0$. Our inner product mimics the structure of Hansen and Richard (1987) and Gallant et al. (1990), who define a conditional asset pricing framework under P . Here, we will apply such an approach to the martingale processes induced by discounted prices in a risk neutral framework.

Importantly, H_s^t is a selfdual pre-Hilbert module or, more simply, a Hilbert module (see Proposition 9 in App. B). The selfduality property allows for an inner product representation of any L^0 -linear and bounded functional on H_s^t (see Definition 2 in Cerreia-Vioglio et al., 2017). This is true, in particular, for linear pricing functionals, consistently with the asset pricing literature: see, e.g. Harrison and Kreps (1979), Ross (1978) and Hansen and Richard (1987).

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 $\hat{h} = \{\hat{h}_\tau\}_{\tau \in [s, t]}$ defined by $\hat{h}_\tau = M_{s,\tau} \pi_\tau(h_t)$. Such process

⁴Clearly, $L_s^1(\mathcal{F}_t)$ 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 Chap. II, Sect. 7 of Shiryaev (1996).

⁵ 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).

⁶Indeed, the conditional expectation of $\hat{z}_{t_1}^2$ is always defined as an extended real random variable and $\mathbb{E}_s[\hat{z}_{t_1}^2] \leq \mathbb{E}_s[\hat{z}_{t_2}^2]$.

belongs to H_s^t and $\hat{h}_s = \pi_s(h_t) = \mathbb{E}_s[M_{s,t}h_t]$. Hence, the no-arbitrage price of eq. (2) induces the L^0 -valued functional $\Pi_s : H_s^t \rightarrow L^0(\mathcal{F}_s)$ such that

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

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

$$\Pi_s(\hat{h}) = \left\langle \hat{g}^t(s), \hat{h} \right\rangle_s^t \quad \text{with} \quad \hat{g}_\tau^t(s) = 1 \quad \forall \tau \in [s, t] \quad (4)$$

for any $\hat{h} \in H_s^t$. The constant conditional martingale $\hat{g}^t(s)$ clearly belongs to H_s^t . This process plays a fundamental role in our return decomposition and its financial meaning is related to the log-optimal portfolio (see 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 $\langle \cdot, \cdot \rangle_s^t$ and, as a consequence, we retrieve a decomposition of returns. As illustrated in Sect. 3.3 of Cerreia-Vioglio et al. (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 1, we determine in Corollary 2 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}_τ -measurable random variable $u_\tau(s)$ satisfying

$$\mathbb{E}_s[M_{s,\tau}u_\tau(s)] = 1 \quad \forall \tau \in [s, T] \quad \text{and} \quad \mathbb{E}_s[M_{s,T}^2u_T^2] \in L^0(\mathcal{F}_s). \quad (5)$$

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 $\mathbb{E}_s[M_{s,T}^2h_T^2] \in L^0(\mathcal{F}_s)$. 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 (5) 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

$$\hat{u}_\tau^T(s) = M_{s,\tau} u_\tau(s) \quad \forall \tau \in [s, T] \quad \text{and} \quad \hat{u}_s^T(s) = 1. \quad (6)$$

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 $t \leq T$. Indeed, we define $\hat{u}^t(s)$ in H_s^t as the restriction of $\hat{u}^T(s)$ on $[s, t]$:

$$\hat{u}_\tau^t(s) = M_{s,\tau} u_\tau(s) \quad \forall \tau \in [s, t] \quad \text{and} \quad \hat{u}_s^t(s) = 1. \quad (7)$$

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

$$f_\tau(s) = \frac{\pi_\tau(1_T)}{\pi_s(1_T)}, \quad \hat{f}_\tau^T(s) = G_{s,\tau} \quad \forall \tau \in [s, T], \quad (8)$$

where $G_{s,\tau}$ is defined in Subsection 2.1. Indeed, by using the relation in (6) and the expression of the pricing kernel in eq. (1), for any τ in $[s, T]$ we find

$$\hat{f}_\tau^T(s) = M_{s,\tau} \frac{\pi_\tau(1_T)}{\pi_s(1_T)} = e^{r_\tau^T(T-\tau) - r_s^T(T-s)} G_{s,\tau} \frac{\pi_\tau(1_T)}{\pi_s(1_T)} = G_{s,\tau}.$$

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 $M_{s,T}$. 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

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

Indeed, the no-arbitrage price at time τ associated with $M_{s,T}$ is $\pi_\tau(M_{s,T}) = \mathbb{E}_\tau[M_{\tau,T} M_{s,T}]$, while the price at time s is $\pi_s(M_{s,T}) = \mathbb{E}_s[M_{s,T}^2]$. By taking the ratio of no-arbitrage prices and by using eq. (1), we find the return

$$\begin{aligned} u_\tau(s) &= \frac{\mathbb{E}_\tau[M_{\tau,T} M_{s,T}]}{\mathbb{E}_s[M_{s,T}^2]} = \frac{e^{-r_\tau^T(T-\tau)} e^{-r_s^T(T-s)} \mathbb{E}_\tau[G_T^2]}{G_\tau G_s} \frac{G_s^2}{e^{-2r_s^T(T-s)} \mathbb{E}_s[G_T^2]} \\ &= \frac{e^{-r_\tau^T(T-\tau)} \mathbb{E}_\tau[G_T^2]}{G_\tau} \frac{G_s}{e^{-r_s^T(T-s)} \mathbb{E}_s[G_T^2]} = \frac{\mathbb{E}_\tau[G_T^2]}{M_{s,\tau} \mathbb{E}_s[G_T^2]}. \end{aligned}$$

The conditional martingale $\hat{u}^T(s)$ associated with this return comes from the relation in (6).

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)

$$\hat{H}_s^t = \{\hat{l}^t(s) \in H_s^t : \mathbb{E}_s[M_{s,t} \hat{l}^t(s)] = 0\}$$

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

where $\iota(s)$ and $\hat{\iota}^t(s)$ are linked by the relation in (7) and $\hat{g}^t(s)$ is defined in eq. (4). Precisely, the process $\hat{g}^t(s)$ in H_s^t and the associated return process $g(s)$ are respectively defined by

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

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

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

3.3 The mean excess return $e(s)$

Fix again $t \in [s, T]$. From the definition of $\hat{f}^T(s)$ in eq. (8), 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 $\mathbb{E}_s[(\hat{f}_t^t(s) - \hat{e}_t^t(s))\hat{\iota}_t^t(s)] = 0$ for all $\hat{\iota}^t(s)$ 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

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

from the definition of $\hat{e}^t(s)$. Similarly to before, from the relation in (7), we define $e(s)$ by

$$e_\tau(s) = f_\tau(s) - g_\tau(s) \quad \forall \tau \in [s, t]. \quad (9)$$

Hence, $e(s)$ embodies the meaning of *mean excess return* because it is the difference between the zero-coupon T -bond return and the log-optimal return.

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 1 (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 $\hat{n}^t(s) \in \hat{H}_s^t$ 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 1 We first show that

$$\mathbb{E}_s \left[\left(\hat{e}_t^t(s) \right)^2 \right] = \mathbb{E}_s \left[G_{s,t} \hat{e}_t^t(s) \right] = \text{var}_s \left(G_{s,t} \right). \quad (10)$$

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

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

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

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

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 \hat{H}_s^t because both $\hat{u}^t(s) - \hat{g}^t(s)$ and $\hat{e}^t(s)$ are in \hat{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 , $\mathbb{E}_s[\hat{e}_t^t(s) \hat{n}_t^t(s)] = \mathbb{E}_s[G_{s,t} \hat{n}_t^t(s)] = 0$. \square

A straightforward application of Theorem 1 delivers an orthogonal decomposition of asset returns in the time window $[s, t]$.

Corollary 2 (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*

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

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 2 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 (7). Then, the result follows directly from Theorem 1. In addition, $\mathbb{E}_s [M_{s,t} n_t(s)] = 0$ because $\hat{n}^t(s)$ belongs to \mathring{H}_s^t . \square

The proof of Theorem 1 exploits the definition of the projection coefficient ω_s in $L^0(\mathcal{F}_s)$, that turns out to be

$$\omega_s = \frac{\mathbb{E}_s^F [M_{s,t} u_t(s)] - 1}{\text{var}_s(G_{s,t})}. \quad (11)$$

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

Let's now go back to the original one-period mean variance allocation problem. If the investor fixes the expected return on the desired portfolio (under the physical measure) and looks for the portfolio displaying the lowest possible variance (again, under the physical measure), the solution to the problem is unique. However, as we already pointed out in Section 1, this approach is of little help in a multi-horizon framework.

The solution we propose instead starts from the very same required portfolio expected return (under the physical measure), but aims at reducing as much as possible the variance (under the physical measure) of the portfolio returns denominated in units of the log-optimal portfolio. Despite this alternative solution being suboptimal, we will show in this section how our solution enjoys a desirable horizon-consistency property that allows for substantial savings when transactions costs are accounted for.

From the point of view of the formal derivation, however, our solution requires an intermediate step. Indeed, it is not possible to directly characterize the set of returns with a given expected value (under P) that also display the lowest possible variance when denominated in units of the log-optimal portfolio

(again, under P). Therefore, we must first set up a solvable parallel mean-variance allocation problem, where the constraint on the expected return is expressed under the T -forward measure. Then we will map back the solution to this parallel problem to the original framework.

This parallel problem starts from the definition of risk-adjusted mean-variance returns. Then, we show how to decompose the whole processes of returns discounted by the log-optimal portfolio to obtain the horizon consistency property that we describe in Subsection 4.1. Afterwards, we illustrate a useful Two-fund Separation Theorem.

Definition 1 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 $\text{var}_s(M_{s,t}u_t(s))$ 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 (7) is a *conditional mean-variance martingale* in $[s, t]$. Such conditional martingale minimizes $\text{var}_s(\hat{u}_t^t(s))$ for the given $\mathbb{E}_s^F[\hat{u}_t^t(s)]$.

The expected returns of Definition 1 can also be written under the physical measure: $\mathbb{E}_s^F[M_{s,t}u_t(s)] = \mathbb{E}_s[G_{s,t}M_{s,t}u_t(s)]$ and $\mathbb{E}_s^F[\hat{u}_t^t(s)] = \mathbb{E}_s[G_{s,t}\hat{u}_t^t(s)]$, and this is how we link this parallel risk-adjusted problem to the original one. However, in order to be able to prove the following theorem, we cannot provide Definition 1 neither using the same measure (either P or F) nor using the same random variable (either $M_{s,t}u_t(s)$ or $G_{s,t}M_{s,t}u_t(s)$) as, in either way, we would not be able to derive an orthogonal decomposition of returns. Therefore, we state Definition 1 in terms of risk-adjusted expected values to move closer to the expression of portfolio weights in eq. (11). Nevertheless, once the risk-adjusted mean-variance frontiers are built, the investor can map risk-adjusted expected returns to physical ones and select portfolios starting from their expected returns under P . See eq. (14) and the comments below, as well as the applications in Section 5.

Theorem 3 (Risk-adjusted mean-variance returns) *Let $t \in [s, T]$. Consider return processes $u(s)$ in $[s, t]$ such that $\mathbb{E}_s^F[M_{s,t}u_t(s)] = k_s$ for some $k_s \in L^0(\mathcal{F}_s)$. Among them, the return process that minimizes $\text{var}_s(M_{s,t}u_t(s))$ is*

$$u(s) = g(s) + \omega_s e(s) \quad \text{with} \quad \omega_s = \frac{k_s - 1}{\text{var}_s(G_{s,t})}.$$

Conversely, if $u(s)$ is a return process in $[s, t]$ such that $u(s) = g(s) + \omega_s e(s)$ for some $\omega_s \in L^0(\mathcal{F}_s)$, then $u(s)$ is a risk-adjusted mean-variance return in $[s, t]$.

Proof of Theorem 3 The proof relies on the fact that return processes $u(s)$ in $[s, t]$ can be associated with conditional martingales in $\hat{u}_t^t(s) \in H_s^t$ via the relation in (7). We also have that $\mathbb{E}_s^F[\hat{u}_t^t(s)] = \mathbb{E}_s^F[M_{s,t}u_t(s)]$ and $\text{var}_s(\hat{u}_t^t(s)) = \text{var}_s(M_{s,t}u_t(s))$.

Given the return processes $u(s)$ in $[s, t]$ such that $\mathbb{E}_s^F[M_{s,t}u_t(s)] = k_s$ for some $k_s \in L^0(\mathcal{F}_s)$, we consider the conditional martingales $\hat{u}_t^t(s) \in H_s^t$ with $\hat{u}_t^t(s) = 1$

such that $\mathbb{E}_s^F[\hat{u}_t^t(s)] = k_s$ and we show that, among them, the conditional martingale that minimizes $\text{var}_s(\hat{u}_t^t(s))$ is

$$\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s) \quad \text{with} \quad \omega_s = \frac{k_s - 1}{\text{var}_s(G_{s,t})}. \quad (12)$$

This immediately implies that the required return process that minimizes $\text{var}_s(M_{s,t}u_t(s))$ is $u(s) = g(s) + \omega_s e(s)$ with the same weight ω_s , as in the theorem statement.

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

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

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

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

By Theorem 1, $\mathbb{E}_s \left[\left(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s) \right) \hat{n}_t^t(s) \right] = 0$ and so

$$\begin{aligned} \mathbb{E}_s \left[\left(\hat{u}_t^t(s) \right)^2 \right] &= \mathbb{E}_s \left[\left(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s) \right)^2 \right] + \mathbb{E}_s \left[\left(\hat{n}_t^t(s) \right)^2 \right] \\ &\geq \mathbb{E}_s \left[\left(\hat{g}_t^t(s) + \omega_s \hat{e}_t^t(s) \right)^2 \right]. \end{aligned} \quad (13)$$

Therefore, $\text{var}_s(\hat{u}_t^t(s))$ is minimized by the conditional martingale with $\hat{n}^t(s) = 0$. This proves the required result in (12).

Conversely, suppose that $u(s)$ is a return process in $[s, t]$ such that $u(s) = g(s) + \omega_s e(s)$ for some $\omega_s \in L^0(\mathcal{F}_s)$ and consider the conditional martingale $\hat{u}^t(s) \in H_s^t$ defined by $\hat{u}^t(s) = \hat{g}^t(s) + \omega_s \hat{e}^t(s)$. Then, $\hat{u}_s^t(s) = \hat{g}_s^t(s) + \omega_s \hat{e}_s^t(s) = 1 + 0 = 1$ and, by the definition of $\hat{g}^t(s)$ and eq. (10),

$$\begin{aligned} \mathbb{E}_s^F \left[\hat{u}_t^t(s) \right] &= \mathbb{E}_s^F \left[\hat{g}_t^t(s) \right] + \omega_s \mathbb{E}_s^F \left[\hat{e}_t^t(s) \right] \\ &= \mathbb{E}_s \left[G_{s,t} \right] + \omega_s \mathbb{E}_s \left[G_{s,t} \hat{e}_t^t(s) \right] \\ &= 1 + \omega_s \text{var}_s(G_{s,t}). \end{aligned}$$

By Theorem 1, any other conditional martingale $\hat{v}^t(s) \in H_s^t$ with $\hat{v}_s^t(s) = 1$ and $\mathbb{E}_s^F[\hat{v}_t^t(s)] = 1 + \omega_s \text{var}_s(G_{s,t})$ satisfies

$$\hat{v}^t(s) = \hat{g}^t(s) + \bar{\omega}_s \hat{e}^t(s) + \hat{n}^t(s), \quad \bar{\omega}_s = \frac{1 + \omega_s \text{var}_s(G_{s,t}) - 1}{\text{var}_s(G_{s,t})} = \omega_s$$

for some $\hat{n}^t(s) \in \hat{H}_s^t$. Hence, $\hat{v}^t(s) = \hat{u}^t(s) + \hat{n}^t(s)$ and, as noted in the relation (13),

$$\mathbb{E}_s \left[\left(\hat{v}_t^t(s) \right)^2 \right] = \mathbb{E}_s \left[\left(\hat{u}_t^t(s) \right)^2 \right] + \mathbb{E}_s \left[\left(\hat{n}_t^t(s) \right)^2 \right] \geq \mathbb{E}_s \left[\left(\hat{u}_t^t(s) \right)^2 \right].$$

As a result, $\hat{u}^t(s)$ is a conditional mean-variance martingale. Therefore, $u(s)$ is a risk-adjusted mean-variance return in $[s, t]$. \square

As an example, consider the zero-coupon T -bond return process $f(s)$ in $[s, T]$. By eq. (9), such return process satisfies $f(s) = g(s) + e(s)$ and so, by Theorem 3, $f(s)$ minimizes the conditional variance of any $M_{s,T}u_T(s)$ with $\mathbb{E}_s^F[M_{s,T}u_T(s)] = \pi_s(1_T)\mathbb{E}[G_{s,T}^2]$.

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 $\mathbb{E}_s[u_t(s)] = \tilde{h}_s$, then the weight ω_s is univocally determined by

$$\omega_s = \frac{\tilde{h}_s - \mathbb{E}_s[g_t(s)]}{\mathbb{E}_s[e_t(s)]} = \frac{\tilde{h}_s - \mathbb{E}_s[g_t(s)]}{\mathbb{E}_s[f_t(s)] - \mathbb{E}_s[g_t(s)]}. \quad (14)$$

As far as concrete applications are concerned, the actual implementation of a portfolio with a given expected return under the physical measure and that minimizes the variance in units of the log-optimal portfolio requires a two-step procedure. First, the investors need to derive the risk-adjusted mean variance frontier, as done in Theorem 3. This delivers a set of optimal (in the sense of Definition 1) risk-adjusted expected returns and variances. Then, the investors need to find the risk-adjusted expected return that matches the expected return they are interested in under the physical measure. Once this physical expected return is mapped into a risk-adjusted one, the first step of the procedure delivers the desired optimal allocation. By construction, the variance under P of this portfolio will be larger than the one derived following the standard mean-variance approach. However, we now show that our optimal risk-adjusted portfolios enjoy a horizon-consistency property that helps in saving transaction costs.

4.1 Horizon consistency

A fundamental property of our approach to mean-variance portfolio analysis is horizon 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 $t \leq T$. This feature is ultimately due to the fact that the decomposition of Corollary 2 involves the whole return processes in the time range $[s, T]$ and so there is a mechanical overlap with the decompositions built at shorter horizons.

From Theorem 3, the risk-adjusted mean-variance frontiers with different horizons (e.g. t and T) are generated by the same two return processes $g(s)$ and $e(s)$. However, it is important to note that such frontiers are generally different because the returns $g_t(s)$ and $g_T(s)$ usually have different first and second moments. The same is true for $e_t(s)$ and $e_T(s)$. Hence, a security or a buy-and-hold portfolio can belong to all the risk-adjusted mean-variance frontiers while featuring variable expected return and variance depending on the considered horizon.

To establish the horizon consistency of risk-adjusted mean-variance returns, for simplicity, we express the result by using the time indices t and T , but the result clearly holds for any $t_1, t_2 \in [s, T]$ with $t_1 \leq t_2$.

Proposition 4 (Risk-adjusted mean-variance returns horizon 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 Proposition 4 Let $u(s)$ be a risk-adjusted mean-variance return in $[s, T]$. By Theorem 3, $u(s) = g(s) + \omega_s e(s)$ for some $\omega_s \in L^0(\mathcal{F}_s)$. Such decomposition holds algebraically at any time in $[s, t]$. By Theorem 3 again, $u(s)$ 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 at the same time our six-month and one-year horizon risk-adjusted mean-variance frontiers, based on the information available today. Proposition 4 ensures that risk-adjusted mean-variance returns on the yearly frontier lie also on the six-month one. 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 optimal risk-adjusted mean-variance returns are generated always by the same two return processes $g(s)$ and $e(s)$, regardless the horizon. This means that, when standing at time zero and building two different risk-adjusted mean-variance buy-and-hold portfolios (one with a six-months maturity, one with a one-year maturity), we just need to invest in two return processes, $g(s)$ and $e(s)$, for both portfolios. As pointed out before, we can select the two risk-adjusted mean-variance portfolios in such a way that their returns match our target expected returns under the physical measure, which is the starting point of all the applications of our technique in Section 5.

However, the mean excess return $e(s)$ is defined in Subsection 3.3 from purely theoretical reasons. To build the risk-adjusted mean-variance frontiers, one must assess whether $e(s)$ is the return of a traded security (or a portfolio) in the market. Luckily, this question can be easily answered by observing that $e(s)$ is the difference between the zero-coupon T -bond return and the log-optimal return (see eq. (9)). Such returns are attainable in the market (see Subsection 2.1.1 for the log-optimal portfolio) and so the risk-adjusted mean-variance frontiers can be implemented by using a traded security and a feasible portfolio. This logic leads to a Two-fund Separation Theorem (Theorem 5 below), where the risk-adjusted frontiers are expressed in terms of $g(s)$ and $f(s)$. Theorem 5 establishes in our setting the celebrated result by Merton (1972), making the implementation of our frontiers immediate.

Theorem 5 (Two-fund Separation) *Given $t \in [s, T]$, $u(s)$ is a risk-adjusted mean-variance return in $[s, t]$ if and only if*

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

where $\alpha_s \in L^0(\mathcal{F}_s)$, $\alpha_s = 1 - \omega_s$ and ω_s is obtained from Theorem 3.

Proof of Theorem 5 Suppose that $u(s)$ is a risk-adjusted mean-variance return in $[s, t]$. Then, Theorem 3 guarantees that $u(s) = g(s) + \omega_s e(s)$ for some $\omega_s \in L^0(\mathcal{F}_s)$. By eq. (9), $e(s) = f(s) - g(s)$ and the desired result obtains. \square

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

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, \dots, t_N with $0 < t_1 < t_2 < \dots < t_N = T$. 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

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

with $h^{(i)} \in \mathbb{R}$ given, for $i = 1, \dots, N$. By construction, 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):⁷ at each date t_i

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

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

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

with $\tilde{\alpha}^{(i)} = 1 - \pi_0(1_{t_i})\tilde{w}^{(i)}$.⁸

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.

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 log-optimal portfolio), for any number N of horizons involved. 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 Theorem 3, for any $i = 1, \dots, N$, we consider a sub-portfolio with return process

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

where $\omega^{(i)}$ is chosen so that the expectation of $v_{t_i}^{(i)}$ meets the target $h^{(i)}$ as in eq. (14). By Theorem 5, 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 maturity t_i , as a consequence of horizon consistency. We finally compare the performance of the two families of sub-portfolios with returns $u^{(i)}$

⁷It is useful to remember that, in the notation of Hansen and Richard (1987), $r^* = M_{0,t_i}/\mathbb{E}[M_{0,t_i}^2]$ and $z^* = 1 - e^{-r_0^{t_i}(t_i-0)}M_{0,t_i}/\mathbb{E}[M_{0,t_i}^2]$.

⁸Indeed, such pure discount bonds belong to the frontier because

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

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 of 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 $c(|\tilde{\alpha}^{(i)}| + |1 - \tilde{\alpha}^{(i)}|)$ 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 $M_{0,t_i}/\mathbb{E}[M_{0,t_i}^2]$ at any horizon t_i entails a positive *fixed* cost C for any (possibly linearly independent) security.⁹ 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 maturities t_1, \dots, t_N , each horizon-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 implementation costs, respectively,

$$C + c \left(\left| \tilde{\alpha}^{(i)} \right| + \left| 1 - \tilde{\alpha}^{(i)} \right| \right) \quad \text{and} \quad \lambda^{(i)}C + c \left(\left| \alpha^{(i)} \right| + \left| 1 - \alpha^{(i)} \right| \right). \quad (15)$$

Accordingly, the overall implementation costs of the two portfolios are:

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

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 implementation costs of eq. (15):

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

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.

⁹To keep the terminology parsimonious, we use the terms *replication costs* even when we have the explicit formulas to build the portfolio under consideration, as it is for the log-optimal portfolio (Subsections 2.1.1 and 5.1). In this case, we face construction or approximations costs.

5.1 Reference market

As in App. 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

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

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 App. C.

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

$$\begin{cases} dX_t = X_t Y_t dt + \eta X_t \left[\phi dW_t^Q + \sqrt{1 - \phi^2} dZ_t^Q \right] \\ d\pi_t(1_T) = \pi_t(1_T) Y_t dt - \pi_t(1_T) B(t, T) \sigma dW_t^Q. \end{cases} \quad (16)$$

At the same time, under the physical measure,

$$\begin{cases} dX_t = X_t \mu_t^X dt + \eta X_t \left[\phi dW_t^P + \sqrt{1 - \phi^2} dZ_t^P \right] \\ d\pi_t(1_T) = \pi_t(1_T) \mu_t^P dt - \pi_t(1_T) B(t, T) \sigma dW_t^P, \end{cases} \quad (17)$$

where μ^X and μ^π 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,

$$\begin{bmatrix} \eta\phi & \eta\sqrt{1 - \phi^2} \\ -B(t, T)\sigma & 0 \end{bmatrix} \begin{bmatrix} \nu_t^W \\ \nu_t^Z \end{bmatrix} = \begin{bmatrix} \mu_t^X - Y_t \\ \mu_t^\pi - Y_t \end{bmatrix} \quad (18)$$

so that

$$\nu_t^W = -\frac{\mu_t^\pi - Y_t}{B(t, T)\sigma}, \quad \nu_t^Z = \frac{\mu_t^X - Y_t - \eta\phi\nu_t^W}{\eta\sqrt{1-\phi^2}}. \quad (19)$$

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

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

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 $\mu_t^\pi = (1 - \xi B(t, T)\sigma) Y_t$ 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

$$\frac{dM_{0,t}}{M_{0,t}} = -Y_t dt - \nu_t^W dW_t^P - \nu_t^Z dZ_t^P. \quad (20)$$

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 $\mu_t^X = Y_t + 0.05$.

As we described in Subsection 2.1.1, the log-optimal portfolio can be constructed via a self-financing strategy whose dynamics can be derived from the application of Itô's formula to eq. (20). Indeed, the price process of the log-optimal portfolio satisfies $N_t = M_{0,t}^{-1}$ at any time t . Similarly to eq. (3), we find

$$\frac{dN_t}{N_t} = \left[Y_t + (\nu_t^W)^2 + (\nu_t^Z)^2 \right] dt + \nu_t^W dW_t^P + \nu_t^Z dZ_t^P.$$

The latter equation can be expressed in terms of the infinitesimal price variations of the traded securities by recalling that $dt = dB_t/(Y_t B_t)$ and by inverting the linear system in (17):

$$\begin{cases} dZ_t^P = \frac{1}{\eta X_t \sqrt{1-\phi^2}} dX_t + \frac{\phi}{\pi_t(1_T)B(t, T)\sigma \sqrt{1-\phi^2}} d\pi_t(1_T) \\ \quad - \left(\mu_t^X + \frac{\eta\phi\mu_t^\pi}{B(t, T)\sigma} \right) \frac{1}{\eta\sqrt{1-\phi^2}Y_t B_t} dB_t \\ dW_t^P = -\frac{1}{\pi_t(1_T)B(t, T)\sigma} d\pi_t(1_T) + \frac{\mu_t^\pi}{Y_t B_t B(t, T)\sigma} dB_t. \end{cases}$$

We obtain

$$dN_t = \theta_t^B dB_t + \theta_t^\pi d\pi_t(1_T) + \theta_t^X dX_t,$$

where θ_t^B , θ_t^π and θ_t^X are the units of assets in the self-financing strategy with values N_t . Specifically,

$$\begin{aligned} \theta_t^B &= \frac{1}{B_t} (N_t - \theta_t^\pi \pi_t(1_T) - \theta_t^X X_t) \\ \theta_t^\pi &= \frac{1}{\pi_t(1_T)B(t, T)\sigma} \left(-\nu_t^W + \frac{\nu_t^Z \phi}{\sqrt{1-\phi^2}} \right) N_t \\ \theta_t^X &= \frac{\nu_t^Z}{\eta X_t \sqrt{1-\phi^2}} N_t. \end{aligned}$$

One can also observe that the asset units θ_t^π and θ_t^X are the solutions of the linear system

$$[\theta_t^X X_t, \quad \theta_t^\pi \pi_t(1_T)] = [\mu_t^X - Y_t, \quad \mu_t^\pi - Y_t] (\Sigma \Sigma')^{-1} N_t,$$

where Σ denotes the matrix in (18). This property is in line with the traditional approach for the log-optimal portfolio construction illustrated in Chapter 15 of Luenberger (1997) and in Chap. 20 of Björk (2009) with constant interest rates.

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 $i = 1, \dots, 6$. 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 $h^{(i)} = 1.06$ 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 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 sub-portfolio i with return process $u^{(i)}$ or $v^{(i)}$ 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 $\lambda^{(i)}$. Every simulated sub-portfolio matches perfectly the target means $h^{(i)}$ at the proper horizon for $i = 1, \dots, 6$. 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 $\tilde{\alpha}^{(i)}$ 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

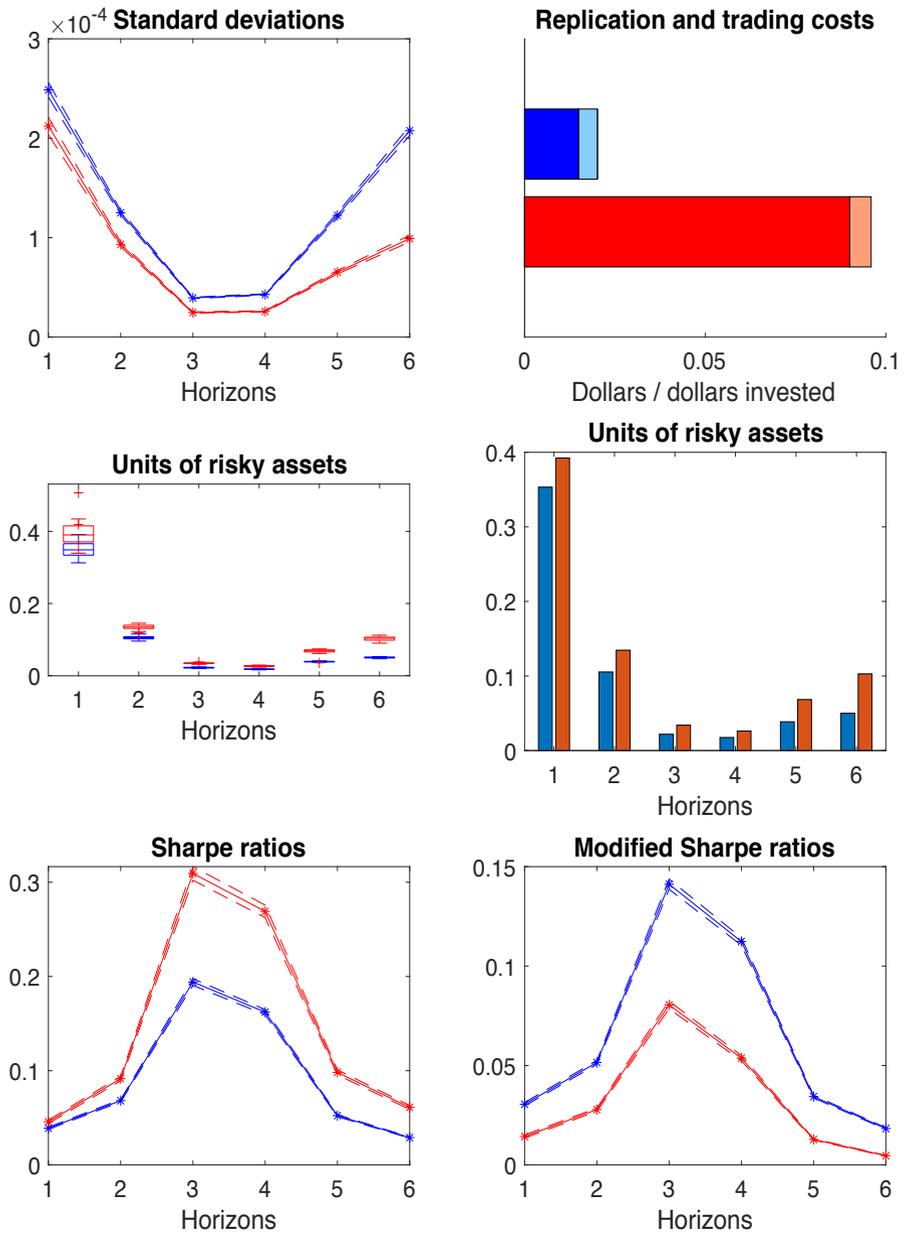


Fig. 1 Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the problem of Subsection 5.2. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \dots, 6$. 90% confidence intervals 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 25th and 75th percentiles and the bar plot of loadings $|\alpha^{(i)}|$ and $|\tilde{\alpha}^{(i)}|$ at all horizons.

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 reduction of the implementation costs 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. Chap. 5 in Bower et al., 1997). The annuity provides yearly payments to each subscriber until the subscriber dies. The insurance company invests 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

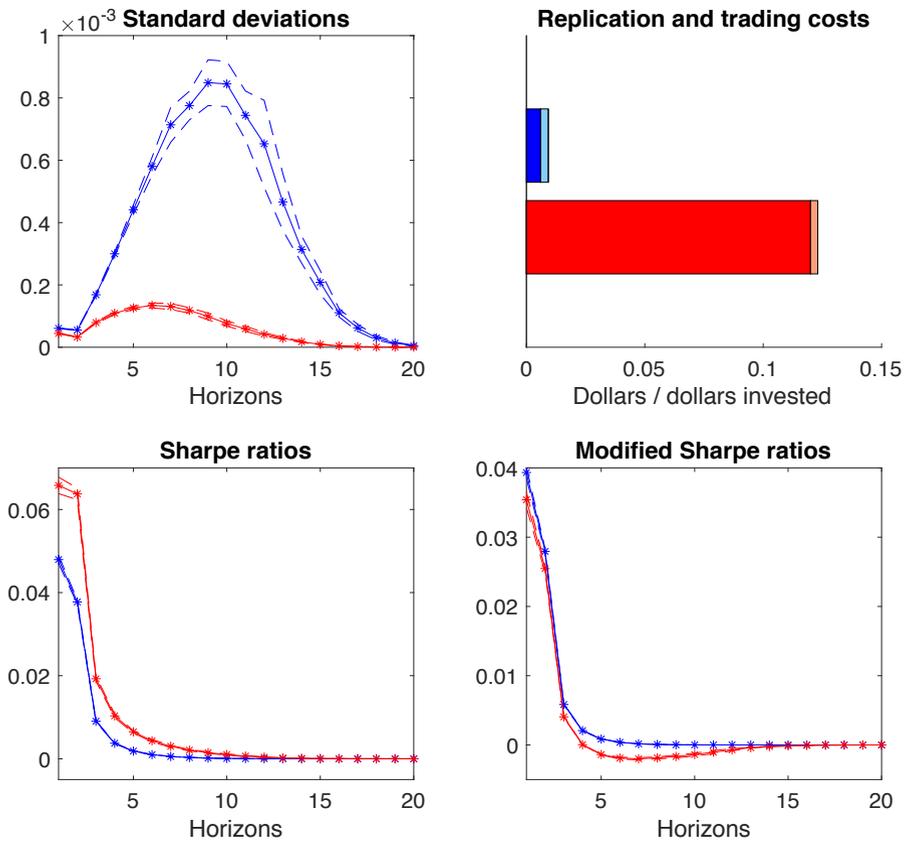


Fig. 2 Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the life-annuity problem. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \dots, 20$. 90% confidence intervals 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).

approach comes from the replication of one risky payoff instead of the $N = 20$ payoffs required by the classical mean-variance optimal strategies. The reduction in the implementation costs 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 implementation costs 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 Theorem 3. 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 5. Moreover, an analogue of horizon consistency of risk-adjusted mean-variance returns can be retrieved in optimal investment policies.

In order to simplify the statement of the problem and reduce technicalities, on top of the assumptions made in Subsection 2.1, we assume here that markets are complete. Therefore, M_t , the stochastic discount factor associated to the inverse of the log-optimal portfolio value process, is now the only stochastic discount factor in the market.

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

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

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

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

The related reduced form is

$$\max_w \mathbb{E}_s \left[\int_s^T U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } w_s = \mathbb{E}_s \left[\int_s^T M_{s,\tau} w_\tau d\tau \right]. \quad (21)$$

Proposition 6 *If in Problem (21) the income stream is null and the bliss point is*

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

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 6 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 times $\tau \in [s, T]$ 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,

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

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

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

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

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

which lies on the risk-adjusted mean-variance frontier in $[s, T]$ by Theorem 5. \square

6.2 Horizon consistency of optimal cashflows

Inspired by the horizon consistency of the risk-adjusted mean-variance frontier shown in Proposition 4, we investigate whether a similar feature is kept in the optimal consumption-investment problem. Specifically, once Problem (21) 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 $t \leq T$. In particular, we consider the problem

$$\max_w \mathbb{E}_s \left[\int_s^t U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } \tilde{w}_s = \mathbb{E}_s \left[\int_s^t M_{s,\tau} w_\tau d\tau \right], \quad (22)$$

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

Proposition 7 *Under the assumptions of Proposition 6, if w^* solves Problem (21) with initial wealth w_s , then it also solves Problem (22) with initial wealth*

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

Proof of Proposition 7 Following the same steps as in the proof of Proposition 6, 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

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

and it coincides with the one prescribed by Proposition 6. \square

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 7 shows that w^* is optimal also on the terminal subperiod $[t, T]$, according to

$$\max_w \mathbb{E}_s \left[\int_t^T U(i_\tau + w_\tau) d\tau \right] \quad \text{sub } \hat{w}_s = \mathbb{E}_s \left[\int_t^T M_{s,\tau} w_\tau d\tau \right], \quad (23)$$

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

Corollary 8 *Under the assumptions of Proposition 6, if w^* solves Problem (21) with initial wealth w_s , then it also solves Problem (23) with*

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

Although Problem (23) 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 horizon consistency 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 et al. (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 horizon consistency property, with practical advantages for multi-horizon 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.

Acknowledgments. We thank Anna Battauz, Marie-Claude Beaulieu, Claudio Fontana, Fabio Girardi, Massimo Marinacci and participants at AMASES XLIII Annual Meeting (2019), Goethe Universität Frankfurt (2021), 10th General AMaMeF Conference (2021), 37th AFFI International Conference (2021), CORS Annual Conference (2021), 29th EFMA Annual Conference (2021), Econometric Society European Meeting (2021) and New Zealand Finance Meeting (2021) for useful comments.

Declarations

Funding. 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).

Conflict of interest. The authors declare that they do not have any financial or non-financial conflicts of interest.

Availability of data and material. Not applicable.

Code availability. Codes are available upon request.

Authors' contribution All authors contributed equally to this work.

Appendix 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 et al. (1995). Moreover,

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

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

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

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

Appendix B The Hilbert modules H_s^t

Proposition 9 H_s^t is a selfdual pre-Hilbert module on $L^0(\mathcal{F}_s)$.

Proof of Proposition 9 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 $\hat{z} \in H_s^t$, $\mathbb{E}_s[\hat{z}_t^2] \in L_s^0(\mathcal{F}_s)$. 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) $\langle \hat{z}, \hat{z} \rangle_s^t = \mathbb{E}_s[\hat{z}_t^2] \geq 0$ 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 et al. (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

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

Since the selfduality of a pre-Hilbert L^0 -module is equivalent to the d_H -completeness (see Theorem 5 in Cerreia-Vioglio et al., 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 et al. (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 $n, m > N_\varepsilon$,

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

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 $(\mathbb{E}_s[|\hat{z}_\tau|])^2 \leq (\mathbb{E}_s[|\hat{z}_t|])^2 \leq \mathbb{E}_s[\hat{z}_t^2]$, 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

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

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

Appendix C 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 (16) 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:

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

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. (19) 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, 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.

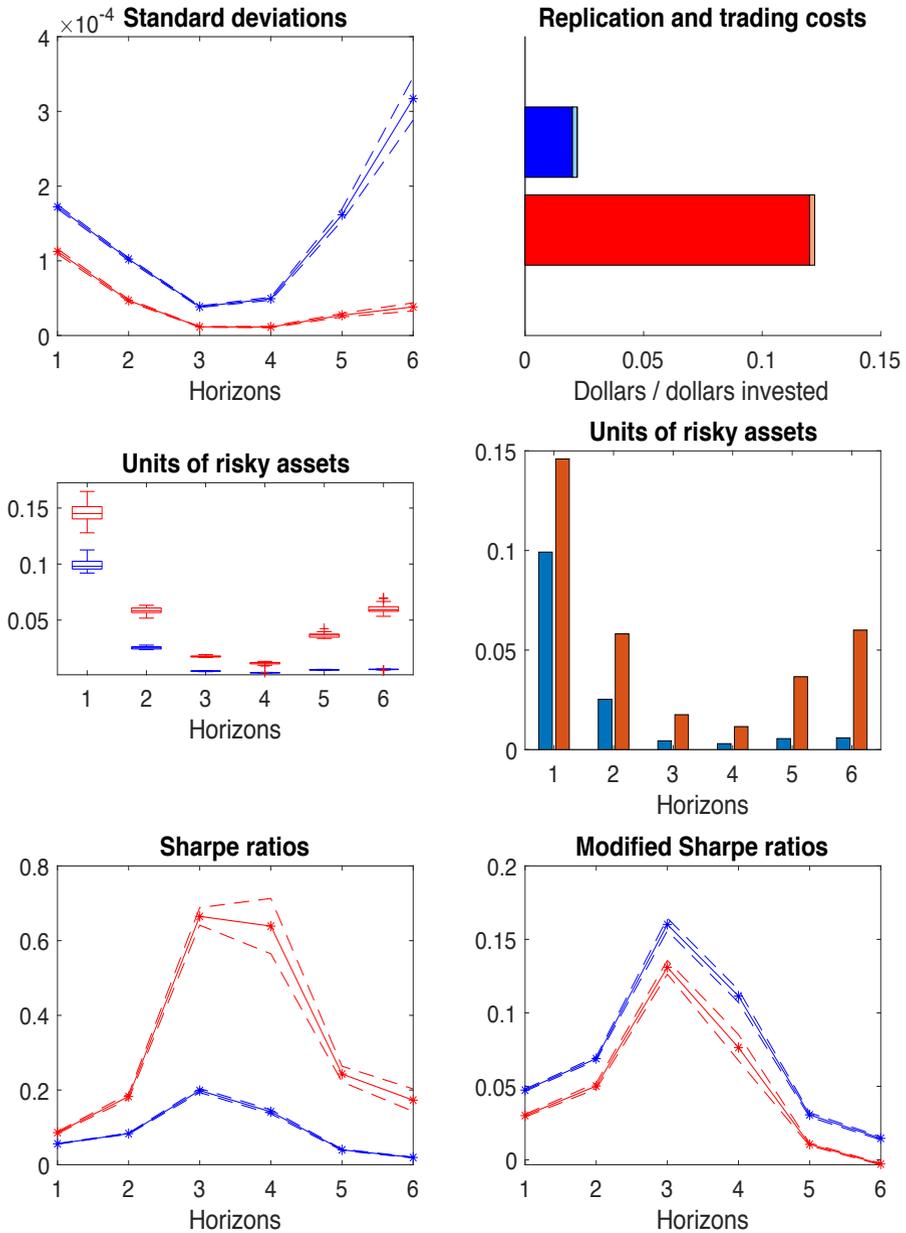


Fig. C1 Red (resp. blue) lines, bars and boxes refer to the classical (resp. risk-adjusted) mean-variance solution for the problem of App. C. Standard deviations, Sharpe ratios and modified Sharpe ratios are scaled by the weights $\lambda^{(i)}$ for all $i = 1, \dots, 6$. 90% confidence intervals for 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 25th and 75th percentiles and the bar plot of loadings $|\alpha^{(i)}|$ and $|\tilde{\alpha}^{(i)}|$ at all horizons.

References

- Abel, A., J. Eberly, and S. Panageas. 2013. Optimal inattention to the stock market with information costs and transactions costs. *Econometrica* 81(4): 1455–1481 .
- Basak, S. and G. Chabakauri. 2010. Dynamic mean-variance asset allocation. *Review of Financial Studies* 23(8): 2970–3016 .
- Bekaert, G. and J. Liu. 2004. Conditioning information and variance bounds on pricing kernels. *Review of Financial Studies* 17(2): 339–378 .
- Björk, T. 2009. *Arbitrage theory in continuous time* (3 ed.). Oxford University Press.
- Bower, N., H. Gerber, J. Hickman, D. Jones, and C. Nesbitt. 1997. *Actuarial mathematics* (2 ed.). The Society of Actuaries, Schaumburg, Illinois.
- Brigo, D. and F. Mercurio. 2006. *Interest rate models-theory and practice: with smile, inflation and credit* (2 ed.). Springer-Verlag, Berlin Heidelberg.
- Cerreia-Vioglio, S., M. Kupper, F. Maccheroni, M. Marinacci, and N. Vogelpoth. 2016. Conditional L_p -spaces and the duality of modules over f -algebras. *Journal of Mathematical Analysis and Applications* 444(2): 1045–1070 .
- Cerreia-Vioglio, S., F. Maccheroni, and M. Marinacci. 2017. Hilbert A-modules. *Journal of Mathematical Analysis and Applications* 446(1): 970–1017 .
- Cerreia-Vioglio, S., F. Maccheroni, and M. Marinacci. 2019. Orthogonal decompositions in Hilbert A-modules. *Journal of Mathematical Analysis and Applications* 470(2): 846–875 .
- Cochrane, J. 2014. A mean-variance benchmark for intertemporal portfolio theory. *Journal of Finance* 69(1): 1–49 .
- Czichowsky, C. 2013. Time-consistent mean-variance portfolio selection in discrete and continuous time. *Finance and Stochastics* 17(2): 227–271 .
- Favero, C., F. Ortu, A. Tamoni, and H. Yang. 2020. Implications of return predictability for consumption dynamics and asset pricing. *Journal of Business & Economic Statistics* 38(3): 527–541 .
- Ferson, W. and A. Siegel. 2001. The efficient use of conditioning information in portfolios. *Journal of Finance* 56(3): 967–982 .

- Ferson, W. and A. Siegel. 2003. Stochastic discount factor bounds with conditioning information. *Review of Financial Studies* 16(2): 567–595 .
- Gallant, A., L. Hansen, and G. Tauchen. 1990. Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics* 45(1-2): 141–179 .
- Geman, H., N. El Karoui, and J.C. Rochet. 1995. Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied Probability* 32(02): 443–458 .
- Hansen, L. and R. Jagannathan. 1991. Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99(2): 225–262 .
- Hansen, L. and S. Richard. 1987. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* 55(3): 587–613 .
- Harrison, J. and D. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20(3): 381–408 .
- Irle, A. and J. Sass. 2006. Optimal portfolio policies under fixed and proportional transaction costs. *Advances in Applied Probability* 38(4): 916–942 .
- Kelly, J.J. 1956. A new interpretation of information rate. *Bell System Technical Journal* 35(4): 917–926 .
- Leippold, M., F. Trojani, and P. Vanini. 2004. A geometric approach to multiperiod mean variance optimization of assets and liabilities. *Journal of Economic Dynamics and Control* 28(6): 1079–1113 .
- Li, D. and W.L. Ng. 2000. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance* 10(3): 387–406 .
- Long, J.J. 1990. The numeraire portfolio. *Journal of Financial Economics* 26(1): 29–69 .
- Luenberger, D. 1997. *Investment science*. Oxford university press.
- Marinacci, M. and F. Severino. 2018. Weak time-derivatives and no-arbitrage pricing. *Finance and Stochastics* 22(4): 1007–1036 .
- Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7(1): 77–91 .
- Martin, I. and C. Wagner. 2019. What is the expected return on a stock? *Journal of Finance* 74(4): 1887–1929 .

- Merton, R. 1972. An analytic derivation of the efficient portfolio frontier. *Journal of Financial and Quantitative Analysis* 7(04): 1851–1872 .
- Mossin, J. 1968. Optimal multiperiod portfolio policies. *Journal of Business* 41(2): 215–229 .
- Ross, S. 1978. A simple approach to the valuation of risky streams. *Journal of Business* 51(3): 453–475 .
- Severino, F. 2021. Long-term risk with stochastic interest rates. *working paper* <http://ssrn.com/abstract=3113718> .
- Sharpe, W. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19(3): 425–442 .
- Shiryayev, A. 1996. *Probability* (2 ed.). Springer-Verlag, New York.
- Sims, C. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50(3): 665–690 .
- Soner, H.M., S.E. Shreve, and J. Cvitanic. 1995. There is no nontrivial hedging portfolio for option pricing with transaction costs. *Annals of Applied Probability* 5(2): 327–355 .
- Strotz, R. 1955. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3): 165–180 .
- Tobin, J. 1958. Liquidity preference as behavior towards risk. *Review of Economic Studies* 25(2): 65–86 .
- Van Staden, P., D.M. Dang, and P. Forsyth. 2018. Time-consistent mean-variance portfolio optimization: A numerical impulse control approach. *Insurance: Mathematics and Economics* 83: 9–28 .
- Vasicek, O. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5(2): 177–188 .
- Zhou, X. and D. Li. 2000. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization* 42(1): 19–33 .