



**PERSISTENCE-BASED CAPITAL
ALLOCATION ALONG THE
FOMC CYCLE**

**FULVIO ORTU
PIETRO REGGIANI
FEDERICO SEVERINO**

The purpose of the **Working Papers** is to disseminate the results of research conducted by CIRANO research members in order to solicit exchanges and comments. These reports are written in the style of scientific publications. The ideas and opinions expressed in these documents are solely those of the authors.

Les cahiers de la série scientifique visent à rendre accessibles les résultats des recherches effectuées par des chercheurs membres du CIRANO afin de susciter échanges et commentaires. Ces cahiers sont rédigés dans le style des publications scientifiques et n'engagent que leurs auteurs.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the government of Quebec, and grants and research mandates obtained by its research teams.

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du gouvernement du Québec, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO Partners – Les partenaires du CIRANO

Corporate Partners – Partenaires corporatifs

Autorité des marchés financiers
Bank of Canada
Bell Canada
BMO Financial Group
Business Development Bank of Canada
Caisse de dépôt et placement du Québec
Desjardins Group
Énergir
Hydro-Québec
Innovation, Science and Economic Development Canada
Intact Financial Corporation
Manulife Canada
Ministère de l'Économie, de l'Innovation et de l'Énergie
Ministère des finances du Québec
National Bank of Canada
Power Corporation of Canada
PSP Investments
Ville de Montréal

Academic Partners – Partenaires universitaires

Concordia University
École de technologie supérieure
École nationale d'administration publique
HEC Montréal
McGill University
National Institute for Scientific Research
Polytechnique Montréal
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval

CIRANO collaborates with many centers and university research chairs; list available on its website. *Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.*

© February 2024. Fulvio Ortu, Pietro Reggiani and Federico Severino. All rights reserved. *Tous droits réservés.* Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source. *Reproduction partielle permise avec citation du document source, incluant la notice ©.*

The observations and viewpoints expressed in this publication are the sole responsibility of the authors; they do not represent the positions of CIRANO or its partners. *Les idées et les opinions émises dans cette publication sont sous l'unique responsabilité des auteurs et ne représentent pas les positions du CIRANO ou de ses partenaires.*

Persistence-based capital allocation along the FOMC cycle

Fulvio Ortu ^{*}, *Pietro Reggiani* [†], *Federico Severino* ^{‡§}

Abstract/Résumé

The Federal Reserve holds two main sets of monetary policy meetings, the “Federal Open Market Committee” (FOMC) and the “Board Meetings”, which gather with six week and two week cadence respectively. Cieslak, Morse, and Vissing-Jorgensen (2019) show that the cadence of these meetings is associated with cycles of corresponding frequencies in stock markets. These can be fruitfully exploited through a portfolio strategy that invests in the whole market at alternate weeks (the even-week strategy). This simple investment rule is based on the cycles identified empirically but, so far, lacks a theoretical foundation. In this paper, we provide a rigorous framework to detect cycles in the stock market, and to determine optimal portfolio choices which profit from such cycles. We use the filtering approach for stationary time series of Ortu, Severino, Tamoni, and Tebaldi (2020) to isolate uncorrelated components of stock returns that are precisely associated with two- and six-week cycles. Then, we replicate these components using tradable assets from the U.S. market, and design an optimal portfolio strategy that maximizes the investor’s wealth and outperforms the even-week strategy.

La Federal Reserve organise deux séries principales de réunions de politique monétaire, le "Federal Open Market Committee" (FOMC) et les "Board Meetings", qui se réunissent avec une cadence de six semaines et de deux semaines respectivement. Cieslak, Morse et Vissing-Jorgensen (2019) montrent que la cadence de ces réunions est associée à des cycles de fréquences correspondantes sur les marchés boursiers. Ceux-ci peuvent être exploités de manière fructueuse par le biais d'une stratégie de portefeuille qui investit dans l'ensemble du marché une semaine sur deux (la stratégie des semaines paires). Cette règle d'investissement simple est basée sur les cycles identifiés empiriquement mais, jusqu'à présent, elle n'a pas de fondement théorique. Dans cet article, nous fournissons un cadre rigoureux pour détecter les cycles sur le marché boursier et pour déterminer les choix de portefeuille optimaux qui profitent de ces cycles. Nous utilisons l'approche de filtrage des séries temporelles stationnaires d'Ortu,

^{*} University Bocconi, Milan (Italy), Department of Finance and IGIER

[†] New York University (US), Finance Department

[‡] University Laval, Quebec (Canada), Department of Finance, Insurance and Real Estate

[§] We thank Francesco Corielli, Marzia A. Cremona, Carlo A. Favero and participants at The Pennsylvania State University (2022), at 16th International Conference on Computational and Financial Econometrics (King’s College London, 2022), at 3rd World Conference on Business, Management, Finance, Economics and Marketing (Eurasia, Paris, 2023) and at Università Bocconi (2023) for useful comments. We thank Yingdong Lin for valuable research assistance. F. Severino acknowledges the financial support of iA Financial Group Chair in Insurance and Financial Services, as well as the computing capability of the Digital Research Alliance of Canada.

Severino, Tamoni et Tebaldi (2020) pour isoler les composantes non corrélées des rendements boursiers qui sont précisément associées aux cycles de deux et six semaines. Ensuite, nous reproduisons ces composantes en utilisant des actifs négociables du marché américain et concevons une stratégie de portefeuille optimale qui maximise la richesse de l'investisseur et surpasse la stratégie des semaines paires.

Keywords/Mots-clés: Fed cycles, Even-week strategy, Return persistence, Optimal portfolios / Cycles de la Fed, Stratégie de la semaine paire, Persistance des rendements, Portefeuilles optimaux

JEL Codes/Codes JEL: G11, E32

Pour citer ce document / To quote this document

Ortu, F., Reggiani, P., & Severino, F. (2024). Persistence-based capital allocation along the FOMC cycle (2024s-02, Cahiers scientifiques, CIRANO.) <https://doi.org/10.54932/TUHB8180>

1 Motivations and main results

The identification of trends and cycles in financial markets is of utmost important for the investors. The presence of periodicities in asset prices can be explained by the way in which markets interiorize the information coming from diverse sources as, for instance, political decisions, corporate news and macroeconomic forecasts. Cieslak et al. (2019) provide strong evidence of two- and six-week cycles in U.S. stock market returns and relate them to the occurrence of Federal Reserve (Fed) meetings. In their view, two- and six-week patterns can be fruitfully exploited by a portfolio strategy that keeps only a share of the market index at alternate weeks: the *even-week strategy*. This simple investment rule agrees with the cycles previously isolated but, so far, misses a theoretical foundation. In this paper we provide a rigorous framework for the detection of stock market cycles and the determination of optimal portfolio choices that profit from two- and six-week periodicities. Specifically, we use the techniques in Ortu et al. (2020) to decompose the time series of market returns into uncorrelated subseries related to shocks with increasing persistence (or decreasing frequencies). Two of these *persistent components* are associated with two- and six-week cycles. By means of a factor model, we replicate these components by using the returns of traded securities. Then, we use the replicated version of the components as input in a dynamic optimal portfolio problem. The arising persistence-based strategy largely outperforms the even-week strategy because it exploits optimally the two- and six-week cycles of stock returns.

It is useful to summarize the origins of the even-week strategy. Fed open market operations are directed by the “Federal Open Market Committee” (FOMC) that gathers eight times per year (roughly every six weeks). Such meetings are scheduled in advance and, since 1994, the FOMC announces its decisions and publishes accompanying statements right after the meetings. Such press release provides the public with information about monetary policy decisions, causing important reactions in stock markets. As extensively acknowledged in the literature, exceptional stock returns are documented around each FOMC meeting (see Subsection 1.1). Interestingly, Cieslak et al. (2019) observe that stock returns are significantly higher not only during the week of the FOMC, but also two and four weeks later. This biweekly timing corresponds to the occurrence of additional Fed meetings (the “Board of Governors Meetings” or “Discount Rate Meetings”) that would be beneficial for financial markets. Most of Board Meetings are about the discount rate (or primary credit rate), i.e. the interest rate charged to commercial banks on loans they receive from regional Fed lending facilities. However, these meetings are also important venues for staff briefings about market financial health and the support provided by the Fed. Although these meetings have no transcripts, information leaks from the Fed around the meeting times and makes its way into the stock market (see the “Fed put” in Cieslak and Vissing-Jorgensen, 2021).

As a result, FOMC and Board Meetings turn out to generate the aforementioned two- and six-week cycles.

In this paper, we consider an arbitrage-free market with n risky securities. An investor assumes autoregressive dynamics for excess returns and maximizes a CARA utility function on terminal wealth, subject to the usual constraint on the wealth path (see Problem (2) in Subsection 3.1). This problem is faced, for instance, by fund managers who invest their clients' money with a given time horizon, in order to maximize the final value of their portfolio. The investor decides the amounts to invest in each of the risky securities. The even-week strategy of Cieslak et al. (2019) prescribes investing only in the market index and so $n = 1$ in their case. Nevertheless, our investor considers other ways to exploit two- and six-week stock market cycles. One possibility could be to build, via a factor model, some synthetic securities that respond only to two- and six-week shocks, and trade them. In general, the resolution of such multi-period problems relies on dynamic programming (see e.g. Chapter 9 in Back, 2010). Since we do not model returns as being independent and identically distributed, computations get complicated quickly, but thanks to the choice of the exponential utility function we can still find a closed form solution for optimal portfolio weights by adapting the techniques in Chryssikou (1998).

To detect patterns with heterogeneous durations in the time series of asset returns, a suitable tool is the Extended Wold Decomposition (EWD henceforth) of Ortu et al. (2020). The latter provide a methodology to decompose a weakly stationary process into uncorrelated components associated with increasing persistence levels (e.g. weekly, monthly or yearly durations). Such subseries respond to shocks that impinge the market and last for specific time scales (e.g. weekly, monthly or yearly shocks). As we illustrate in Subsection 2.1, the null correlation between persistent components permits to quantify the importance of each time scale through a variance decomposition. Moreover, the construction is entirely in the time domain and the persistent components of market returns can be seen as asset returns themselves. This feature triggers their replication through traded securities or indices.

More precisely, starting from the daily time series r_t of (cumulated) excess returns on U.S. stock markets, we employ a suitable version of the EWD to determine the persistent components $r_t^{(2)}$ at time scale 2 and $r_t^{(3)}$ at time scale 3. The innovations on these time scales involve shocks of 9 and 27 days, respectively. As a result, the components $r_t^{(2)}$ and $r_t^{(3)}$ are able to capture the two- and six-week stock market cycles described by Cieslak et al. (2019), that correspond to roughly 10 and 30 working days. Later, the time series of $r_t^{(2)}$ and $r_t^{(3)}$ are replicated by projecting them on ten traded factors, and the replicates are used as asset returns in our optimal investment problem.

Our optimal portfolio recipe outperforms the even-week strategy because it properly exploits market cycles. However, this persistence-based strategy is more difficult to implement because it requires dynamic investments in several factors. The issue is, then, to analyze the trade-off between the two approaches in detail. Hence, in the final part of the paper we introduce robustness checks, as well as transaction and short-selling costs. The aim is to make the optimal investment strategy feasible for the investor in a real market.

1.1 Additional related literature

The impact of Fed meetings on stock markets is analyzed in Lucca and Moench (2015), who focus on FOMC meetings, which take place every six weeks. They identified a pre-FOMC announcement drift, i.e. abnormally high stock returns during the 24 hours preceding the release of monetary policy announcements from the FOMC. Evidence of abnormal stock market returns on FOMC meeting dates was already provided by Tori (2001) and Savor and Wilson (2013) renew the empirical evidence of significant excess returns on scheduled days in which macroeconomic news are supposed to be released. More recently, Ernst et al. (2019) dig into the reasons behind the origins of this too-much-return puzzle. The announcement drifts have a sizable impact on U.S. bond markets too, as documented by Brooks et al. (2019).

Wachter and Zhu (2022) introduce a rational investor model explaining the striking empirical findings about FOMC announcements. Boguth et al. (2019) provide some insights on the way in which the information from such announcements is incorporated into equity prices. Other explanations of FOMC-driven market anomalies arise from the resolution of uncertainty due to FOMC meetings (Hu et al., 2022), in particular tail uncertainty or downside risk (Beckmeyer et al., 2021). In addition, Laarits (2022) and Cocoma (2022) propose alternative theoretical mechanisms that could generate the pre-FOMC announcement drifts.

In general, beyond the two- and six-week cycles of Cieslak et al. (2019), the financial literature is rich of contexts where shocks with heterogeneous duration are present. Some examples are given by stock return volatility (Campbell and Hentschel, 1992), the Gross National Product (Cochrane, 1988) and the yield curve (Cieslak and Povala, 2015). To capture disturbances with heterogeneous persistence, we use the EWD. The methodology rests on a low-pass filter that makes averages of subsequent shocks and permits to isolate innovations with lower and lower frequencies. However, differently from existing methods (Müller and Watson, 2008), the EWD is developed entirely in the time domain and does not require the use of frequencies. The adaptation to multivariate time series is provided by Cerreia-Vioglio et al. (2023).

In spite of the widespread empirical evidence about market cycles, few papers construct portfolios that optimally respond to shocks with heterogeneous duration. Chaudhuri and Lo (2015), Crouzet et al. (2017) and Di Virgilio et al. (2019) provide some attempts. The EWD allows us to fill this gap by using the persistent components of stock returns as asset returns, after their replication via a factor model.

2 Persistence-based FED cycle detection

In this section we show how to elicit the persistent components of market returns associated with two- and six-week cycles.

2.1 Persistence-based decomposition of excess returns

We employ a decomposition of market excess returns that generalizes the persistence-based decomposition of Ortu et al. (2020). The latter moves from the fact that, by the Wold decomposition (Wold, 1938), any zero-mean weakly stationary time series $\mathbf{x} = \{x_t\}_t$ is an infinite moving average

$$x_t = \sum_{k=0}^{+\infty} \alpha_k \varepsilon_{t-k}, \quad (1)$$

where the fundamental innovations $\varepsilon = \{\varepsilon_t\}_t$ constitute a unit variance white noise and α_h are the impulse response functions. The EWD aggregates the fundamental innovations in a way to obtain an orthogonal decomposition of x_t into uncorrelated components $x_t^{(j)}$ associated with increasing time scales j :

$$x_t = \sum_{j=1}^{+\infty} x_t^{(j)}, \quad x_t^{(j)} = \sum_{k=0}^{+\infty} \beta_k^{(j)} \varepsilon_{t-k2^j}^{(j)}.$$

Each detail process $\varepsilon^{(j)} = \{\varepsilon_t^{(j)}\}_t$ is an $MA(2^j - 1)$ with respect to the fundamental innovations and $\beta_k^{(j)}$ is the scale-specific response associated with scale j and time-shift $k2^j$. The explicit expressions of $\beta_k^{(j)}$ and $\varepsilon_t^{(j)}$ exploit the impulse response functions and the fundamental innovations, respectively (Theorem 1 in Ortu et al., 2020).

The innovations of the *persistent component* $x_t^{(j)}$ evolve on a 2^j -step grid, capturing higher and higher persistence (or lower and lower frequencies) as the time scale j increases. This intuition is justified by spectral analysis considerations related to the use of the scaling operator as a low-pass filter.

Moreover, the orthogonality of the persistent components and the unit variance of the

detail processes $\varepsilon_t^{(j)}$ induce a variance decomposition of x_t :

$$\text{var}(x_t) = \sum_{j=1}^{+\infty} \text{var}(x_t^{(j)}) = \sum_{j=1}^{+\infty} \sum_{k=0}^{+\infty} (\beta_k^{(j)})^2.$$

As we outline in Appendix B, the EWD can be generalized to any base N different from 2. As we discuss in the next subsections, In our application it is convenient to choose $N = 3$. Indeed, when $N = 3$, an application of the scaling operator smooths the effects of innovations lasting up to 3 periods, leaving such residual disturbances at time scale 1. Another application of the operator will smooth the effects of innovations lasting up to 9 periods, leaving the residual disturbances (involving shocks from 3 to 9 periods) at time scale 2. In short, the persistent component at scale 1 will be associated with 3-day shocks, the persistent component at scale 2 will be associated with 9-day shocks, and so on. An orthogonal decomposition of x_t into the sum of uncorrelated persistent components $x_t^{(j)}$ still obtains, as well a variance decomposition across the persistent components. All formulas are in Appendix B for a general base N .

Empirically, the persistent components are easily estimated once the moving average representation in eq. (1) is known. Hence, the first step is the estimation of an autoregressive form for x_t . Then, it is easy to retrieve the corresponding moving average representation of x_t . After that, the impulse responses and the fundamental innovations permit to compute the scale-specific responses and the detail processes, respectively.

2.2 Sample construction

The FOMC began to regularly announce its decisions and issue companion statements after the meetings in 1994. In particular, the first meeting of the year occurred on February 3, 1994. Therefore, it is sensible to start our analysis at the beginning of 1994. Moreover, the totality of the ten factors that we will use in Subsection 3.3 to replicate the persistent components is defined until the end of December 2019 (the USD index that we employ is replaced by another index in 2020). Hence, we end our sample the week before the last FOMC meeting of 2019. Summing up, the sample considered to estimate the persistent components goes from January 3, 1994 to December 6, 2019. In addition, the choice of the terminal date permits to exclude the COVID-19 pandemic from the analysis. At the beginning of the pandemic, central bank interventions were unusually frequent, a fact that is not compatible with two- and six-week cycles.

Regarding the datasets used, we obtain the dates of Federal Reserve meetings from web scraping of the Fed website. As to market returns, we consider the first Fama-French factor and the risk-free rate from Kenneth R. French website, *Fama/French 5 Factors (2x3) [Daily]*


series. In this database, the risk-free rate is the one-month Treasury bill rate, updated at the beginning of each month. Moreover, the first factor (the market) comes from value-weighted returns of all NYSE, AMEX, and NASDAQ firms.

The market return time series contains missing data due, for instance, to stock market closure on some days. We do a simple imputation by attributing zero returns to these days. In doing so, and by excluding all week-ends, we are able to obtain a full Monday-Friday structure in the data, which permits to easily implement weekly portfolio strategies by investing on Monday. Then, following Cieslak et al. (2019), we compute the 5-day forward cumulated returns on each day in the sample, even when markets are closed. The resulting time series constitute the object of our analysis.


The estimation of the persistent components at scale 2 and 3 on the sample from January 3, 1994 to December 6, 2019 requires some initial data as input. Indeed, the innovations at increasing time scales involve longer moving averages of the original shocks. Therefore, these components can be estimated from June 21, 1994. Since the next FOMC meeting took place on July 5, 1994, when dealing with the optimal portfolio problem, we consider returns data from Monday July 4, 1994 to Friday December 6, 2019. Then, the investment period under scrutiny will span from January 31, 2000 to December 6, 2019 and data before January 31, 2000 will be employed to fit the models required by the optimization algorithm at that date. See Subsection 3.2 and the following.

2.3 Persistent components of market excess returns


We first apply the EWD in base 3 to the time series r_t of 5-day forward cumulated excess returns on U.S. stock markets in order to build the persistent components of returns $r_t^{(j)}$ for time scales $j = 1, \dots, 6$. The truncation at scale 6 is dictated by the sample size. Indeed, each single innovation at scale 6 involves 729 days and so the required amount of data for the procedure is large. As mentioned at the end of Subsection 2.1, the first step is to estimate an autoregressive form for r_t by the Bayesian Information Criterion (Schwarz, 1978), then its moving average representation and, finally, its persistent components. See Panel (a) in Fig. 1. According to the persistence-based variance decomposition, the components $r_t^{(2)}$ and $r_t^{(3)}$ turn out to explain the higher share of return variance (roughly 38% and 32% in the sample from January 3, 1994 to December 6, 2019). See Panel (b) in Fig. 1. Since these components are associated with shocks lasting 9 and 27 working days respectively, the analysis provides further support to the detection of two- and six-week cycles by Cieslak et al. (2019). We repeat the analysis on a rolling window of past 1250 days and we plot the time series of variance ratios in Panel (c) of Fig. 1. On average, scale 2 explains roughly 40% of the sample variance, while scale 3 explains 31.5%.



(a) Market and persistent components.



(b) Relative variance.



(c) Relative variance through rolling window.

Figure 1: Panel (a): realizations of 5-day forward cumulated market excess returns and their persistent components at time scales 2 and 3. Panel (b): relative variance explained by each persistent component from January 1994 to December 2019. Panel (c): relative variance explained by the 2nd and 3rd persistent components of 5-day forward cumulated market excess returns computed on a rolling window of 1250 past days, from January 1994 to December 2019.

3 Optimal capital allocation

We first present the general dynamic portfolio problem to be solved by a CARA investor. Then, we apply it to two different contexts in Subsections 3.2 and 3.3.

The first context is a counterfactual context, where we assume that there exists two assets that are traded and whose (5-day forward cumulated) returns coincide with the two components $r_t^{(2)}$ and $r_t^{(3)}$ obtained by the EWD of market returns. This context is counterfactual (indeed, 1 \$ invested in 2000 becomes roughly 47 \$ after 20 years) but it is a useful benchmark to ascertain that a lot can be gained from investing in the persistent components.

Then, we move to the second context, in which we replicate the two persistent components using ten traded factors. Here, 1 \$ invested in the portfolio strategy based on the replicated components gives (almost) 8 \$ after 20 years, as opposed to (almost) 4 \$ obtained by the even-week strategy.

Finally, in Subsection 3.4 we provide a distributional analysis via bootstrap.

3.1 Dynamic optimal portfolio problem

We consider a market with a constant risk-free rate r_f and n risky securities. A maturity T is fixed. At any discrete time t between 0 and T , the excess returns are collected in the vector r_t . An investor has a CARA utility function with absolute risk aversion parameter equal to γ and is endowed with an initial wealth W_0 . The investor faces a sequential investment decision at times $t = 0, 1, \dots, T - 1$ concerning the final wealth at time T . The excess returns on the risky assets are estimated through a VAR(p) process and, at any t , the amounts invested in the risky securities are contained in the n -dimensional vector u_t . The investor's optimal portfolio problem is

$$\begin{aligned} \max_{u_0, \dots, u_{T-1}} \quad & \mathbb{E} \left[-e^{-\gamma W_T} \right] \\ \text{sub} \quad & W_t = (1 + r_f)W_{t-1} + u'_{t-1}r_t \\ & r_t = \mu + B_1 r_{t-1} + \dots + B_p r_{t-p} + \varepsilon_t, \end{aligned} \tag{2}$$

where, B_i are the matrices of the returns sensitivities in the VAR(p), and $\varepsilon_t \sim \mathcal{N}(0, \Sigma)$ are independent normally distributed random vectors with zero mean and covariance matrix Σ . No intermediate consumption is allowed and no transaction cost or friction are taken into account.

Paired with the normality of returns, CARA preferences permit to obtain a closed-form solution to the investor problem. Such preferences are used, e.g., by Wang (1993) and an example of dynamic portfolio choice with CARA utility and ARMA(1,1) excess returns is

provided by Balvers and Mitchell (1997). In general, Problem (2) is solved by dynamic programming, as done in Section 5.1 of Chryssikou (1998) with VAR(1) excess returns. When excess returns are modeled via a VAR(p), we obtain the following proposition. In Appendix A, we provide a derivation for the special case with $p = 2$, by using a Bellman equation and backward induction.

Proposition 1 *Consider Problem (2). For any $k = 1, \dots, T$, the optimal portfolio choice at time $T - k$ is*

$$u_{T-k}^* = \frac{1}{\gamma(1+r_f)^{k-1}} \left(\Sigma^{-1} \mathbb{E}_{T-k} [r_{T-k+1}] - \sum_{i=1}^{\min\{k-1,p\}} B_i' \Sigma^{-1} E_{k,i} \right), \quad (3)$$

where, for any $i = 1, \dots, \min\{k-1, p\}$,

$$E_{k,i} = \mathbb{E}_{T-k} [r_{T-k+i+1} | r_{T-k+1} = \dots = r_{T-k+i} = 0].$$

As it is apparent from eq. (3), in the CARA-normal set-up the demand for risky assets is unaffected by initial wealth. The optimal amount of risky assets u_{T-k}^* depends on risk aversion, the risk-free rate, the time to maturity and the features of asset returns. As shown in Chryssikou (1998), when $p = 1$, the solution at $T - 1$ is discontinuous with respect to previous investment choices. This is related to the number of lags in the autoregressive process. When excess returns are modeled as a VAR(p) process, the same discrepancy holds between the chosen amounts in the $T - p$ last periods and the previous portfolio choices. This explains the extremes in the summation term in eq. (3). Moreover, as can be seen in the proof of the proposition, the CARA-normal setup can be reduced to an intertemporal mean-variance problem. See, e.g., Cochrane (2014).

Regarding our application, in Problem (2) we use the series of 5-day forward cumulated excess returns and a time unit of 5 days (Monday to Friday). The CARA investor makes the portfolio choice on Monday, based on the past 5-day forward cumulated returns, and holds it until the following Monday. For instance, time t may be Monday, February 4, 2019 and time $t+1$ Monday, February 11, 2019. In fact, we have built the persistent components at scales 2 and 3 from the series of 5-day forward cumulated excess returns and the investor considers the sub-sample of them consisting of all Mondays.

To make the comparison with the even-week strategy, it is important to ascertain that such a strategy is available to the investor. A week when an FOMC meeting occurs is a *week 0*; the week after is a *week 1* and so on, until the counter is reset to 0 at the week of the next FOMC meeting. The even-week strategy can be embedded in Problem (2) by considering only one risky return (the return on the market) and selecting the amount

$u_{T-k} = W_{T-k}$ on even weeks of the FOMC cycle, and $u_{T-k} = 0$ on odd weeks (where the entire wealth is invested in the risk-free asset). Hence, the even-week strategy is feasible for our investor who makes the portfolio choice on Monday and does not review the position until the following Monday.

Finally, to enhance the comparisons, we also consider the *buy-and-hold strategy* on the market. Also in this case, only the market return is present in Problem (2) and the amount invested in that is $u_{T-k} = W_{T-k}$ for all k .

In case the wealth associated with the persistence-based investment strategy turns negative (or null), we close the position and we suppose returns to be null until the end of the investment period. Moreover, if the initial amounts invested in the persistent components is negative (or null), i.e. a short position is optimal at the starting date, we do not make the comparison between the three strategies. Indeed, the even-week and the buy-and-hold strategies do not require short positions and a comparison with the persistence-based strategy would not be appropriate in this case.

A delicate point to take into account is the initial value of the even-week and the buy-and-hold strategies. Indeed, the optimal portfolio choice of Problem (2) is independent of the investor's initial wealth W_0 . From Proposition 1, we only find the optimal amount to invest in the risky securities available. When this investment is positive, we suppose that the investor does not trade the risk-free asset at time 0. Hence, the cost of the strategy is the sum of the amounts invested in the risky securities, i.e. the sum of the entries of the vector u_0^* . Then, we set W_0 equal to this value. To make a proper comparison, we consider the contemporaneous even-week strategy with initial cost W_0 , meaning that the investor invests W_0 in the market in the first week of investment, which is a *week 0*.

3.2 Optimal capital allocation when persistent components are traded

Due to the importance of the components $r_t^{(2)}$ and $r_t^{(3)}$ in subsamples illustrated in Panel (c) of Fig. 1, we put ourselves in a context where two risky securities with cumulated returns $r_t^{(2)}$ and $r_t^{(3)}$ are traded. Although counterfactual, this context provides a way to gauge the potential quantitative importance of these persistent components.

In particular, we consider a specification of Problem (2) where the agent is allowed to invest in exactly two risky assets with (5-day forward) cumulated returns $r_t^{(2)}$ and $r_t^{(3)}$. In this problem, the returns are supposed to follow a VAR(p) model. The order p of the autoregressive can be inferred from the Akaike Information Criterion (Akaike, 1974). We apply this criterion in the 250-week return window preceding the investment period. Every four weeks, we move the window and repeat the estimation. Moreover, Problem (2) exploits a constant risk-free rate. Therefore, we set this rate as the risk-free rate at the beginning




Figure 2: Optimal wealth path obtained by investing in $r_t^{(2)}$ and $r_t^{(3)}$, together with the wealth paths of the even-week and the buy-and-hold strategies. Initial wealth is normalized to 1 \$. Investment period from January 31, 2000 to December 6, 2019.

of each period of portfolio optimization (and so it is updated every four weeks).

The investment period that we consider starts from a week where the FOMC meeting takes place (a *week 0*). This makes the comparison with the even-week strategy well-posed. Fig. 2 displays the optimal wealth path obtained by investing in the components at scale 2 and 3, as well as the even-week and the buy-and-hold strategies, after normalizing the initial wealth to 1 \$. The return on the persistence-based strategy is astonishing: roughly 47 \$ against a 1 \$ initial investment. The even-week strategy outperforms the buy-and-hold one, but its terminal wealth is only close to 4 \$. Here, the persistent components are supposed to be the returns of securities traded in the market, hence the optimal wealth is expected to be upward biased in this analysis.

3.3 Optimal capital allocation when persistent components are replicated

We now put ourselves in a more realistic context where the persistent components at scale 2 and 3 can be replicated by ten market indices. Our CARA investor will use the replications of the persistent components to build the optimal portfolio by solving Problem (2).

In our replication exercise, beyond Fama-French five factors, we use the USD index, the AAA and BBB bond indices, a commodity and a volatility index. Such factors capture different risk exposures. Specifically, our analysis makes use of the following daily data.

- Fama and French (2015) five factors (MKT, SMB, HML, RMW, CMA), as well as the risk-free rate from Kenneth R. French website, *Fama/French 5 Factors (2x3) [Daily]* series.
- The USD index, i.e. the Nominal Major Currencies U.S. Dollar Index (Goods Only) from the FRED (Federal Reserve Economic Data) database of St Louis Fed, *DTWEXM* series.
- The AAA (and BBB) indices, i.e. the ICE BofA AAA (or BBB) U.S. Corporate Index Total Return from the FRED database of St Louis Fed, *BAMLCC0A1AAATRIV* (or *BAMLCC0A4BBBTRIV*) series.
- The GSCI index, i.e. the S&P GSCI Commodity Total Return (SPGSCITR) from Bloomberg Terminal.
- The VIX index from the CBOE website.

The last five factors come from the literature on hedge funds, in particular from Fung and Hsieh (1997) and Hasanhodzic and Lo (2007), who build upon the original asset-class factor model by Sharpe (1992).

We first provide some descriptive information about the ten factors and their persistent components. Consistent with the approach taken so far, we begin with cumulating the returns of each factor over 5 days forward. We then perform the persistence analysis. As Fig. 3 shows, scales 2 and 3 are the most relevant for all the ten factors. Together they explain from 60% to 75% of the variance of each index. This feature is in line with the variance decomposition of the market factor provided in Fig. 1(b). As a result, evidence of FED-induced cycles emerges here as well.

Table 1 contains the correlation matrix of the (cumulated) returns on the ten factors, as well as the correlation coefficients between the market and its components $r_t^{(2)}$ and $r_t^{(3)}$. Consistent with our optimal portfolio problem, we only consider the returns on Mondays. Regarding the ten factors, the highest correlations (in absolute value) are found between AAA and BBB and between MKT and VIX. Such correlation may reveal some level of multicollinearity, which we will address in Subsection 4.3. As expected, the sample correlation between the two components is almost null due to their orthogonality. Finally, Table 2 shows the coefficients of the linear regression of the market persistent components at scale 2 and 3 on the (cumulated) returns on the ten factors. MKT and RMW are significant predictors for both components at the levels indicated in Table 2. Other significant factors for $r_t^{(2)}$ are SMB, AAA, BBB and VIX. As to $r_t^{(3)}$, HML and USD constitute additional significant factors.

A further step is the replication of the persistent components $r_t^{(2)}$ and $r_t^{(3)}$ via the linear regressions

$$r_t^{(2)} = \sum_{i=1}^{10} b_i f_{i,t} + \varepsilon_t, \quad r_t^{(3)} = \sum_{i=1}^{10} c_i f_{i,t} + \eta_t \quad (4)$$

where each $f_{i,t}$ is the 5-day forward cumulated return of one of the ten factors and ε_t and η_t are the error terms. The regression coefficients b_i and c_i refer to the components at scale 2 or 3, respectively. Following Hasanhodzic and Lo (2007), we omit the constant terms in the regressions since the coefficients b_i (resp. c_i) are the weights in the portfolios replicating the persistent components. Accordingly, we impose the sums of the b_i 's and of the c_i 's to be both equal to 1. Therefore, we run the following two regressions:

$$r_t^{(2)} - f_{1,t} = \sum_{i=2}^{10} b_i (f_{i,t} - f_{1,t}) + \varepsilon_t, \quad r_t^{(3)} - f_{1,t} = \sum_{i=2}^{10} c_i (f_{i,t} - f_{1,t}) + \eta_t.$$

To avoid the look-ahead bias, such linear regressions are fitted in a time window of 250 weeks preceding the investment period. After obtaining the two linear clones $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$, we use them in Problem (2) instead of the persistent components at scale 2 and 3. Differently from before, the CARA investor trades an optimal amount of ten indices instead of two assets. The optimal amount invested in each index turns out to be $[b_i, c_i]u_{T-k}^*$. To make

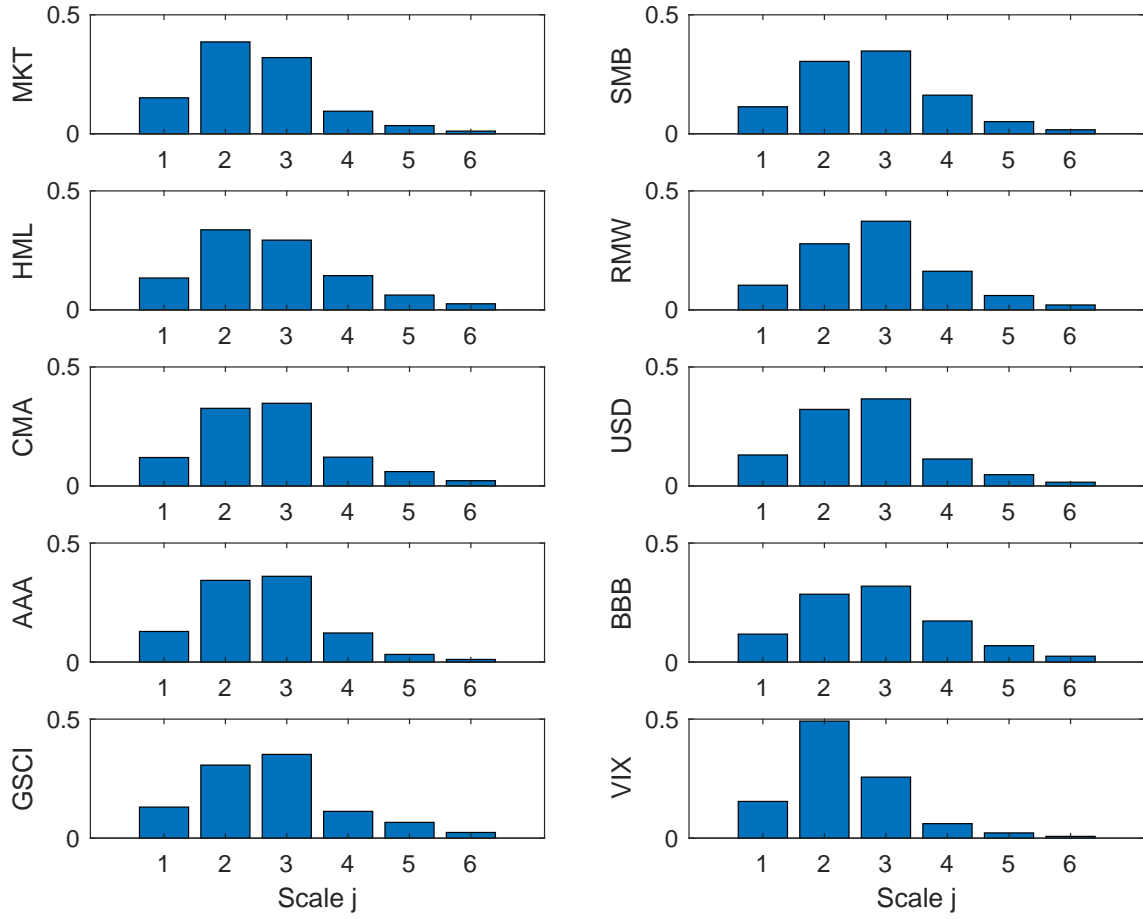


Figure 3: Relative variance explained by the persistent components of 5-day forward cumulated returns of the ten factors from January 1994 to December 2019.

Table 1: Correlation matrix of the (5-day forward cumulated) ten factors from July 4, 1994 to December 6, 2019 with weekly sampling (on Monday). Then, correlation between the market and the persistent components at scale 2 and 3.

	MKT	SMB	HML	RMW	CMA	USD	AAA	BBB	GSCI	VIX
MKT	1	0.149	-0.004	-0.400	-0.338	-0.143	-0.081	0.015	0.258	-0.705
SMB	0.149	1	-0.006	-0.358	-0.018	-0.075	-0.144	-0.088	0.127	-0.098
HML	-0.004	-0.006	1	0.168	0.490	-0.069	-0.033	-0.025	0.139	0.022
RMW	-0.400	-0.358	0.168	1	0.286	0.050	0.118	0.024	-0.114	0.199
CMA	-0.338	-0.018	0.490	0.286	1	0.002	-0.006	-0.017	-0.057	0.165
USD	-0.143	-0.075	-0.069	0.050	0.002	1	-0.122	-0.163	-0.308	0.038
AAA	-0.081	-0.144	-0.033	0.118	-0.006	-0.122	1	0.811	-0.074	0.122
BBB	0.015	-0.088	-0.025	0.024	-0.017	-0.163	0.811	1	0.031	0.025
GSCI	0.258	0.127	0.139	-0.114	-0.057	-0.308	-0.074	0.031	1	-0.186
VIX	-0.705	-0.098	0.022	0.199	0.165	0.038	0.122	0.025	-0.186	1

	MKT	$r_t^{(2)}$	$r_t^{(3)}$
MKT	1	0.635	0.562
$r_t^{(2)}$	0.635	1	-0.032
$r_t^{(3)}$	0.562	-0.032	1

Table 2: Coefficients of the linear regression of the market persistent components on the (5-day forward cumulated) ten factors from July 4, 1994 to December 6, 2019 with weekly sampling (on Monday):

$$r_t^{(2)} = \alpha_0^{(2)} + \sum_{i=1}^{10} \alpha_i^{(2)} f_{i,t} + \xi_t^{(2)}, \quad r_t^{(3)} = \alpha_0^{(3)} + \sum_{i=1}^{10} \alpha_i^{(3)} f_{i,t} + \xi_t^{(3)}.$$

Here, $\xi_t^{(2)}$ and $\xi_t^{(3)}$ denote the error terms. t -statistics are in parentheses and *, ** and *** denote a p-value lower than 0.10, 0.05 and 0.01 respectively.

	$r_t^{(2)}$	$r_t^{(3)}$
Intercept	0.001** (2.456)	0.000 (0.259)
MKT	0.347*** (16.820)	0.315*** (14.652)
SMB	-0.084*** (-3.428)	0.030 (1.161)
HML	0.008 (0.325)	0.045* (1.705)
RMW	0.057* (1.861)	-0.080** (-2.492)
CMA	-0.055 (-1.365)	-0.044 (-1.029)
USD	-0.042 (-1.192)	-0.063* (-1.733)
AAA	0.306*** (4.470)	0.051 (0.712)
BBB	-0.488*** (-6.619)	0.068 (0.879)
GSCI	0.006 (0.579)	-0.010 (-0.857)
VIX	-0.014*** (-4.518)	0.003 (0.843)
N	1327	1327
R^2	0.442	0.329

the estimation dynamic, we keep the 250-week rolling window and update the estimates every four weeks.

The left panel of Fig. 4 shows the optimal wealth paths for the persistence-based, the even-week and the buy-and-hold strategies. The replication of components $r_t^{(2)}$ and $r_t^{(3)}$ is responsible for lowering the return on the persistence-based strategy with respect to the one obtained in Fig. 2. For an investment of 1 \$, we observe a decrease of terminal wealth from (approximately) 47 \$ to a value close to 8 \$. However, the ranking between the three considered investment strategies keeps unchanged, starting from the second year from the start of the investment. The right panel of Fig. 4 depicts the fraction of wealth invested in each of the ten factors by the persistence-based strategy. Portfolio weights do not require extreme positions except for the crisis periods of 2001 and 2008-2009, during which they are more volatile.

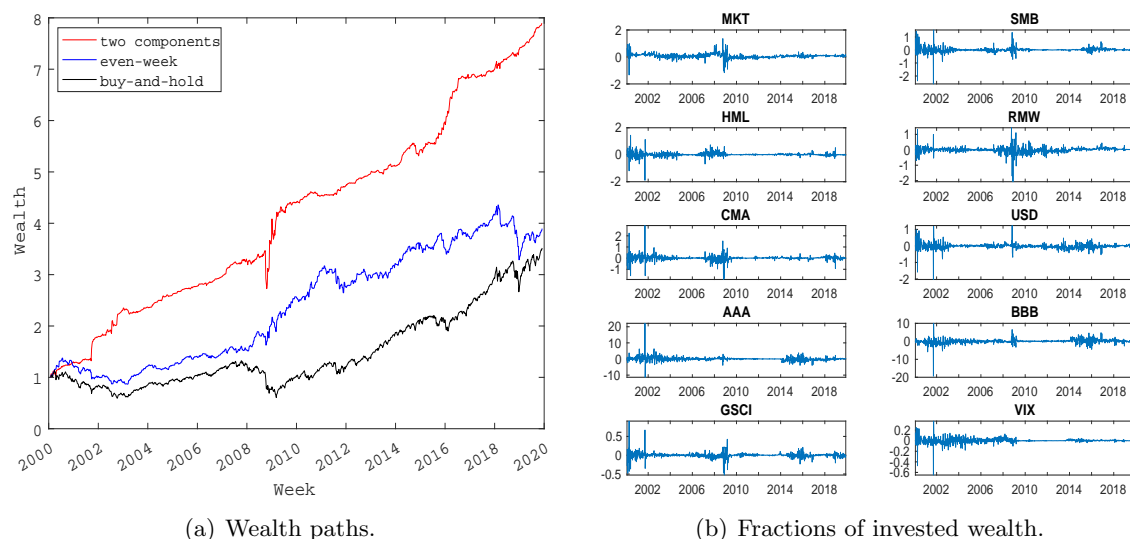


Figure 4: Panel (a): optimal wealth path obtained by investing in the replicated persistent components $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$, together with the wealth paths of the even-week and the buy-and-hold strategies. Initial wealth is normalized to 1 \$. Panel (b): fraction of wealth invested in the ten factors by the persistence-based strategy. Investment period from January 31, 2000 to December 6, 2019.

We finally compare our portfolio strategies in terms of Sharpe ratio, by considering the (weekly) mean excess return and the empirical standard deviation over the investment period. Results are collected in Table 3. The persistence-based strategy features the highest mean return, the lower volatility and the highest Sharpe ratio, both when the components are traded and when they are replicated by ten factors.

Table 3: Weekly mean excess returns, standard deviations and Sharpe ratios obtained by investing in the traded components $r_t^{(2)}$ and $r_t^{(3)}$, the replicated components $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$, the even-week and the buy-and-hold strategies. Investment period from January 31, 2000 to December 6, 2019.

	Mean	Standard deviation	Sharpe ratio
Traded components	0.0035	0.0138	0.2540
Replicated components	0.0018	0.0142	0.1263
Even-week	0.0012	0.0178	0.0656
Buy-and-hold	0.0012	0.0246	0.0494

3.4 Bootstrap

To provide a distributional view on the results of Fig. 2, we implement a bootstrap procedure (Efron and Tibshirani, 1986) in case the persistent components are the returns of traded securities. To implement the bootstrap, we first fit an $AR(77)$ model to the daily time series of market returns and the risk-free rate, from January 3, 1994 to December 6, 2019. The choice of 77 lags in the autoregressive model is based on the AIC of market returns. Then, we shuffle the residuals and we add them to the fitted values in order to obtain a bootstrap sample of daily returns. We use the same permutation of residuals for the risk-free rate. Then, we set to zero the returns on the days of market closure, in order to mimic the original data-set. After that, we compute 5-day forward cumulated returns and we estimate the persistent components at scales 2 and 3 by using a rolling window of 1250 past observations. We repeat this procedure 500 times. Both the persistence-based, the even-week and the buy-and-hold strategies yield zero (or negative) wealth in some bootstrap samples. We filter out such bootstrap samples and, in Fig. 5, we plot for each week the 5% and 95% percentiles of the distributions of the (normalized) wealth for our strategies, in logarithmic scale. The buy-and-hold strategy is omitted to improve the graph readability.

Although there is a small overlap of the bootstrap intervals of the two strategies at the beginning of the investment period, starting from 2002 there is no intersection between the two. Our persistence-based strategy consistently outperforms the even-week strategy from 2002 on. By considering a 1 \$ initial investment, at the end of 2019, the wealth obtained from the persistence-based strategy spans from (approximately) 26.5 \$ to 1333.0 \$, while the even-week strategy provides an approximate terminal wealth between 1.3 \$ and 7.6 \$. The bootstrap intervals may be extremely wide due to the variability introduced during the bootstrap sample creation over long time series. However, the ranking between the persistence-based and the even-week strategy is also maintained in this distributional




Figure 5: 5% and 95% bootstrap percentiles of the optimal wealth paths obtained by investing in $r_t^{(2)}$ and $r_t^{(3)}$, together with the 5% and 95% bootstrap percentiles of the wealth paths of the even-week strategy. Initial wealth is normalized to 1 \$. Investment period from January 31, 2000 to December 6, 2019. The y -axis is in logarithmic scale.

analysis.

4 Robustness

We now provide some robustness analyses to confirm the superior performance of the persistent-based strategy with respect to the even-week and the buy-and-hold strategies. In Subsection 4.1 the persistent components are supposed to be the returns of traded securities (as in Subsection 3.2), while in Subsections 4.2 and 4.3 the components are replicated by the ten factors of Subsection 3.3. The first is a benchmark framework where we can easily move the starting date of the investment period forward, consider a fixed investment term or move the terminal date backward. The other is a more suitable framework to study the impact of transaction and short-selling costs, as well as the potential multicollinearity issues arising in the factor model

4.1 Subsample robustness

We propose three kinds of robustness analysis in case $r_t^{(2)}$ and $r_t^{(3)}$ are the returns of traded securities. As reference period, we keep the investment window from January 31, 2000 to December 6, 2019.

We first move the starting date from one Monday to the next one, by considering only Mondays of even weeks, and we make the portfolio strategy comparison on every (shorter and shorter) investment period ending on December 6, 2019. In case the persistence-based strategy prescribes a negative initial investment in the persistent components, we do not make the comparison with the even-week and the buy-and-hold strategy. In Fig. 6 we represent the gross return of each strategy on a weekly basis and the average Sharpe ratios computed by using weekly returns. The end of the sample features some instability of the returns on the persistence-based strategy, which is much lower before 2018. This suggests that long investment periods provide stable weekly returns on such strategy. In any case, the performance of the persistent-based strategy is remarkable throughout the whole sample, in terms of both weekly returns and Sharpe ratios. In particular, the Sharpe ratios of the persistent-based strategy are always positive.

In the previous analysis, the investment horizon becomes shorter and shorter as long as the starting date is shifted forward. Hence, we now ask ourselves what would happen in case a fixed investment horizon is considered. We keep the same framework as the last analysis and consider a 3-year investment horizon for the strategy comparison. Fig. 7 traces Fig. 6 except for the last part of the sample, which is missing because the investment term goes beyond the used data. Both returns and Sharpe ratios of the persistence-based strategy are

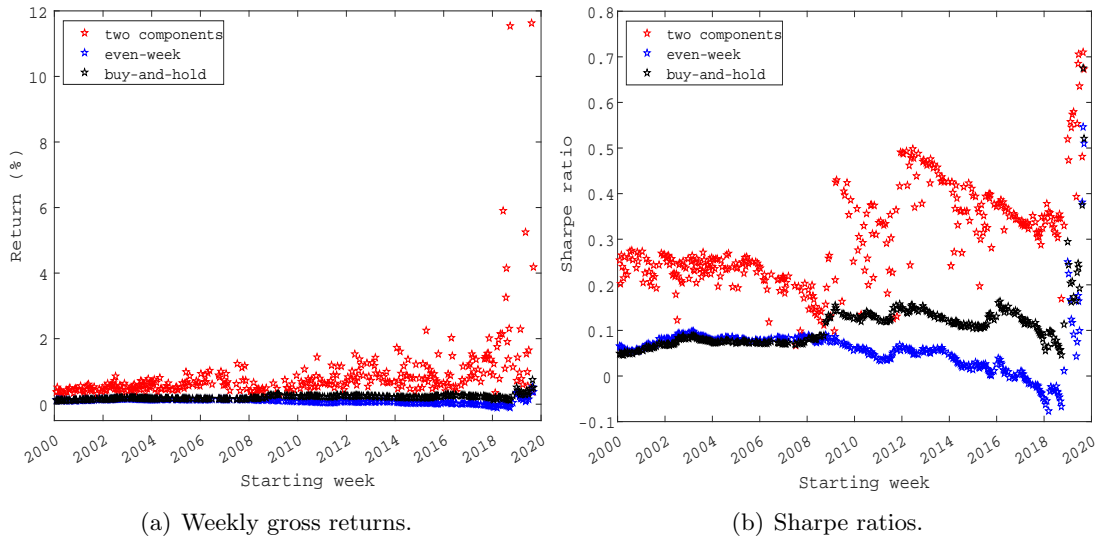


Figure 6: Investment strategies ending on December 6, 2019. The starting date is increasing over time, from Monday to Monday. Only Mondays of even weeks are considered. Panel (a): gross returns (in percentage points) of the three types of strategies converted on a weekly basis. Panel (b): average Sharpe ratios of the three types of strategies. The plotted values refer to the day in which the investment starts.

always positive. The returns of such strategy are variable over time and we can recognize some cycles in the corresponding series of Sharpe ratios. Overall, the persistence-based strategy features a prominent performance with respect to both the even-week and the buy-an-hold strategy throughout the whole sample. The extra return of the persistence-based strategy depends on the initial date of the 3-year investment.

Finally, we keep the starting date fixed on January 31, 2000 and we move the terminal date in the investment window from January 31, 2000 to December 6, 2019. The terminal date is moved from one Monday to the previous one. Similarly to before, we make the comparison of our three investment strategies on every (longer and longer) investment period starting on January 31, 2000. This creates a “term structure” of portfolio returns that we plot in the left panel of Fig. 8. The right panel contains the average Sharpe ratios, similarly to the previous figures. The ranking among the persistence-based, the even-week and the buy-and-hold strategy is consistent across maturities. In line with Fig. 6, short horizons feature the highest returns and Sharpe ratios, while long horizons are rather stable. Fig. 8 even suggests a long-term convergence to a weekly return of roughly 0.35% (and a Sharpe ratio of 0.25) over time. In short, the persistent-based strategy largely outperforms the other strategies under consideration, with a stable effect for long maturities.

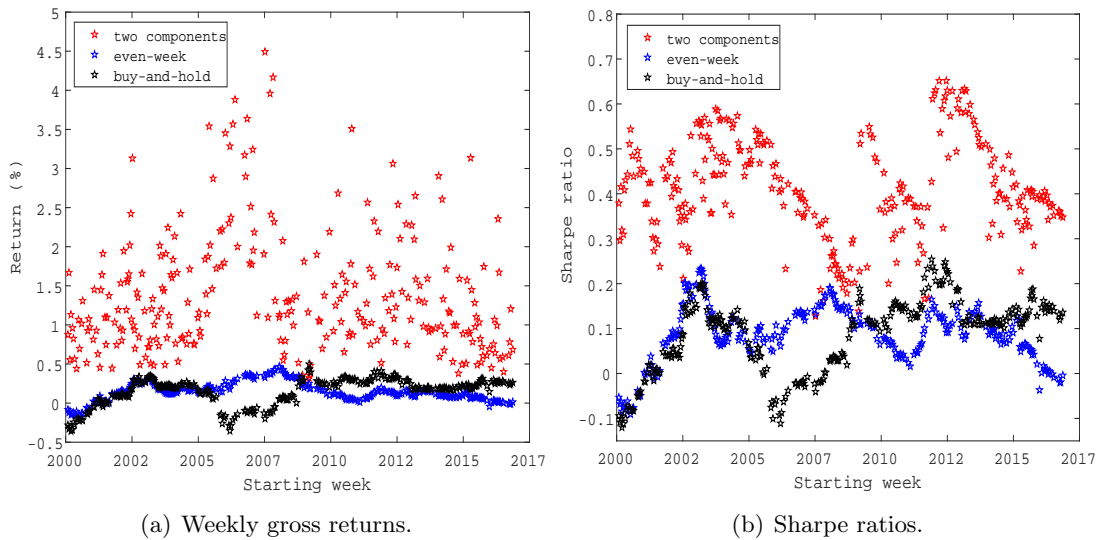


Figure 7: Investment strategies over a 3-year (156-week) term. The starting date is increasing over time, from Monday to Monday. Only Mondays of even weeks are considered. Panel (a): gross returns (in percentage points) of the three types of strategies converted on a weekly basis. Panel (b): average Sharpe ratios of the three types of strategies. The plotted values refer to the day in which the investment starts.

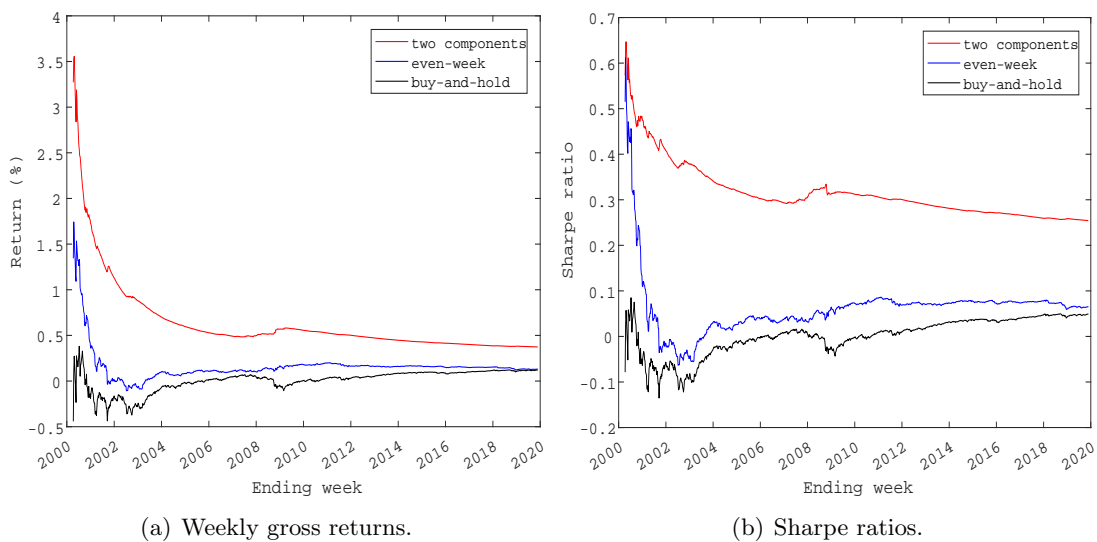


Figure 8: Investment strategies starting on January 31, 2000. The terminal date is decreasing over time, from Monday to Monday. Panel (a): gross returns (in percentage points) of the three types of strategies converted on a weekly basis. Panel (b): average Sharpe ratios of the three types of strategies. The plotted values refer to the day in which the investment ends.

4.2 Transaction and short-selling costs

We now consider the case in which the persistent components at scale 2 and 3 are replicated by the ten factors described in Subsection 3.3. In this case, the implementability of the persistence-based strategy is an important aspect to study. In particular, we focus on transaction costs and short-selling costs.

Transaction costs are supposed to be proportional to the amount invested in risky assets. According to the mutual fund expense ratios documented in Duvall (2020), a conservative approach is to consider a proportion of 2%. Since, in our optimization problem, the invested amounts are multiplied by the asset returns, we apply this percentage directly to market returns at the beginning. Then, we follow the usual steps.

As to short-selling costs, Weitzner (2023) documents (in Table 1D) average stock loan fees that amount to roughly 2.5% (annual rate). Hence, to be conservative, in our analysis we reduce the wealth associated with short positions by 3% (annual rate).

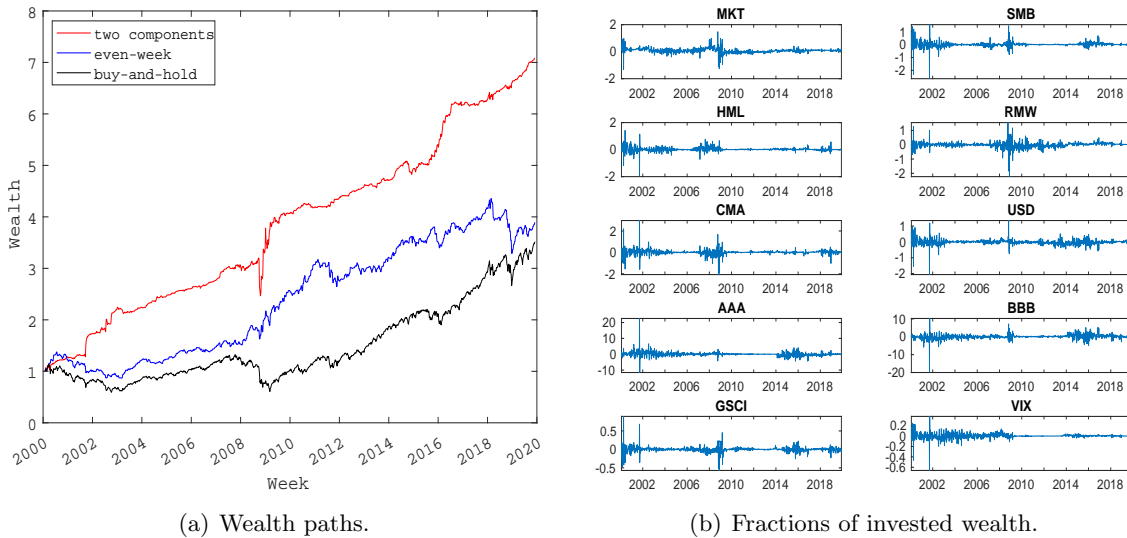


Figure 9: Panel (a): optimal wealth path obtained by investing in the replicated persistent components $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$ with transaction and short-selling costs, together with the wealth paths of the even-week and the buy-and-hold strategies. Initial wealth is normalized to 1 \$. Panel (b): fraction of wealth invested in the ten factors by the persistence-based strategy with transaction and short-selling costs. Investment period from January 31, 2000 to December 6, 2019.

We consider the investment window between January 31, 2000 and December 6, 2019, and we plot the results in Fig. 9, which traces the previous Fig. 4. The introduction of transaction and short-selling costs decreases the optimal wealth of the persistence-based

strategy, while the other strategies are not particularly affected. In fact, neither the even-week nor the buy-and-hold strategy requires short selling. For an investment of 1 \$, the terminal wealth of our persistence-based strategy is around 7 \$ instead of the 8 \$ obtained in Fig. 4. Nevertheless, the optimal investment in $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$ keeps outperforming the even-week and the buy-and-hold strategies, starting from the second year from the start of the investment. As to the portfolio weights required by the persistence-based strategy, there are some tiny adjustments with respect the weights displayed in the right panel of Fig. 4, due to the introduction of transaction and short-selling costs.

4.3 The multicollinearity issue

The predictors employed in our factor models can potentially suffer from multicollinearity. For instance, we already observed in Table 1 that AAA and BBB indices are rather correlated, as well as MKT and VIX. The literature on variable selection offers a plethora of methods to address this issues as the stepwise or the best subset selection. Here, we propose to use the LASSO, a penalized regression that permits to decrease the number of predictors (Tibshirani, 1996). We set the tuning parameter as the minimum of the mean square error. The structure of the exercise is the same as the one in Subsection 3.3: we only replace the multiple regression with the LASSO.

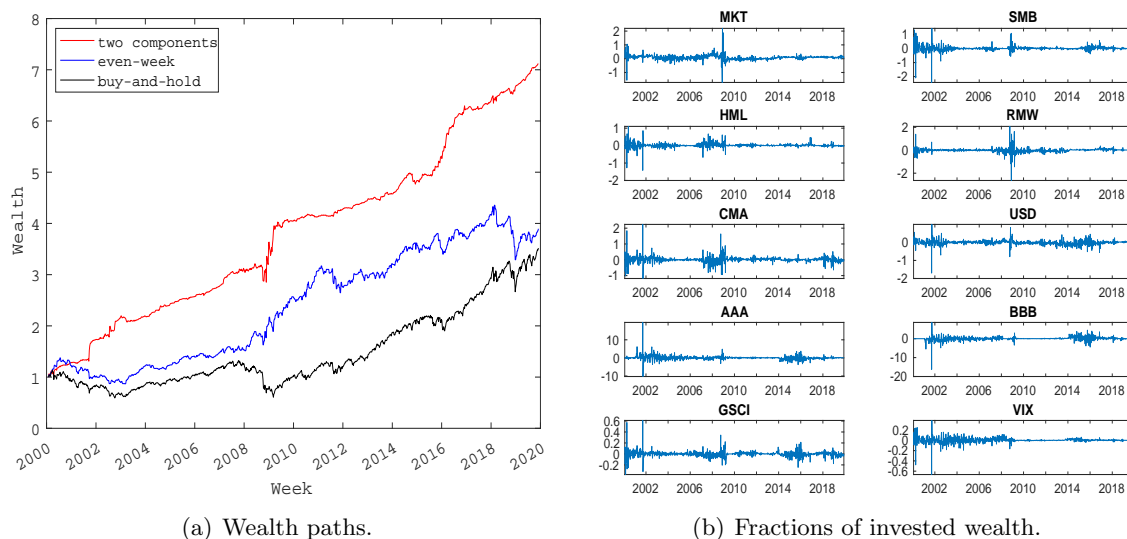


Figure 10: Panel (a): optimal wealth path obtained by investing in the replicated persistent components $\hat{r}_t^{(2)}$ and $\hat{r}_t^{(3)}$ obtained through LASSO, together with the wealth paths of the even-week and the buy-and-hold strategies. Initial wealth is normalized to 1 \$. Panel (b): fraction of wealth invested in the ten factors by the persistence-based strategy obtained through LASSO. Investment period from January 31, 2000 to December 6, 2019.

We fix the investment window from January 31, 2000 to December 6, 2019, and we plot the results in Fig. 10, which traces Fig. 4. We observe a reduction of terminal wealth for the persistence-based strategy from a value close to 8 \$ (in Fig. 4) to a value slightly higher than 7 \$, with the ranking of the three investment strategies not changing. Interestingly, the Sharpe ratio of the persistence-based strategy obtained via the LASSO exceeds the one obtained in the replication exercise of Subsection 3.3: 0.1298 versus 0.1263. This value also exceeds the Sharpe ratios of the even-week and the buy-and-hold strategies reported in Table 3. In addition, the factor models chosen by the LASSO in each rolling window can be more parsimonious than the ten-factor model. This fact increases the frequency of zero portfolio weights for some indices, making the implementation of the persistence-based strategy less demanding. For instance, the initial period requires no investment in BBB, as we can see in the right panel of Fig. 10.

5 Conclusions

The paper enriches the literature of portfolio management by providing a methodology to exploit the persistent patterns of stock returns induced by periodic meetings of the Fed. The contribution is both theoretical and empirical. First of all, we provide a renewed evidence of the presence of two- and six-week cycles in the stock market by analyzing the persistence of market returns. Then, we model such cycles so that they are usable in a capital allocation problem. The investment strategies based on the elicited persistent components turn out to be profitable ways of exploiting such market cycles.

A large number of factors is available today and the financial literature keeps creating them. Using an expression from Cochrane (2011), the so-called factor zoo counts more than 500 factors today (Harvey and Liu, 2019). This poses many issues such as multicollinearity and non-stationarity, which can be addressed by machine learning techniques, such as penalized regressions or other methods for variable selection and dimensionality reduction. Feng et al. (2020) put things in order by analyzing and proposing novel approaches. The replication of the persistent components of market returns could benefit from this literature. A better replication could improve the accuracy of optimal portfolio weights and generate higher wealth for the investor.

We expect our portfolio recipe to be beneficial for a large audience of investors that could improve the performance of their portfolios by properly exploiting market cycles induced by the Fed meetings.

References

- H. Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.
- K. Back. *Asset pricing and portfolio choice theory*. Oxford University Press, 2010.
- R.J. Balvers and D.W. Mitchell. Autocorrelated returns and optimal intertemporal portfolio choice. *Management Science*, 43(11):1537–1551, 1997.
- H. Beckmeyer, N. Branger, and T. Grünthaler. Fed tails: FOMC announcements and stock market uncertainty. *Working paper*, 2021.
- O. Boguth, V. Grégoire, and C. Martineau. Price pressure and efficiency on FOMC announcements. *Working paper*, 2019.
- J. Brooks, M. Katz, and H. Lustig. Post-FOMC announcement drift in US bond markets. Technical report, National Bureau of Economic Research, 2019.
- J.Y. Campbell and L. Hentschel. No news is good news: An asymmetric model of changing volatility in stock returns. *Journal of Financial Economics*, 31(3):281–318, 1992.
- S. Cerreia-Vioglio, F. Ortu, F. Severino, and C. Tebaldi. Multivariate Wold decompositions: a Hilbert A-module approach. *Decisions in Economics and Finance*, 46:45–96, 2023.
- S.E. Chaudhuri and A.W. Lo. Spectral analysis of stock-return volatility, correlation, and beta. In *2015 IEEE Signal Processing and Signal Processing Education Workshop (SP/SPE)*, pages 232–236, 2015.
- E. Chryssikou. *Multiperiod portfolio optimization in the presence of transaction costs*. PhD thesis, Massachusetts Institute of Technology, 1998.
- A. Cieslak and P. Povala. Expected returns in Treasury bonds. *The Review of Financial Studies*, 28(10):2859–2901, 2015.
- A. Cieslak and A. Vissing-Jorgensen. The economics of the Fed put. *The Review of Financial Studies*, 34(9):4045–4089, 2021.
- A. Cieslak, A. Morse, and A. Vissing-Jorgensen. Stock returns over the FOMC cycle. *The Journal of Finance*, 74(5):2201–2248, 2019.
- J.H. Cochrane. How big is the random walk in GNP? *Journal of Political Economy*, 96(5): 893–920, 1988.

- J.H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108, 2011.
- J.H. Cochrane. A mean-variance benchmark for intertemporal portfolio theory. *The Journal of Finance*, 69(1):1–49, 2014.
- P. Cocoma. Explaining the pre-announcement drift. *Working paper*, 2022.
- N. Crouzet, I. Dew-Becker, and C.G. Nathanson. A model of multi-frequency trade. *Northwestern University Working Paper*, 2017.
- D. Di Virgilio, F. Ortu, F. Severino, and C. Tebaldi. Optimal asset allocation with heterogeneous persistent shocks and myopic and intertemporal hedging demand. In Itzhak Venezia, editor, *Behavioral Finance: The Coming of Age*, pages 57–108. World Scientific, 2019.
- J. Duvall. Trends in the expenses and fees of funds, 2019. Technical Report 1, ICI Research Perspective, 2020.
- B. Efron and R. Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, pages 54–75, 1986.
- R. Ernst, T. Gilbert, and C.M. Hrdlicka. More than 100% of the equity premium: How much is really earned on macroeconomic announcement days? *Working paper*, 2019.
- E.F. Fama and K.R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015.
- G. Feng, S. Giglio, and D. Xiu. Taming the factor zoo: A test of new factors. *The Journal of Finance*, 75(3):1327–1370, 2020.
- W. Fung and D.A. Hsieh. Empirical characteristics of dynamic trading strategies: The case of hedge funds. *Review of Financial Studies*, 10(2):275–302, 1997.
- C.R. Harvey and Y. Liu. A census of the factor zoo. *Working paper*, 2019.
- J. Hasanhodzic and A.W. Lo. Can hedge-fund returns be replicated?: The linear case. *Journal of Investment Management*, 5(2):5–45, 2007.
- G.X. Hu, J. Pan, J. Wang, and H. Zhu. Premium for heightened uncertainty: Explaining pre-announcement market returns. *Journal of Financial Economics*, 145(3):909–936, 2022.

- T. Laarits. Pre-announcement risk. *Working paper*, 2022.
- D.O. Lucca and E. Moench. The pre-FOMC announcement drift. *The Journal of Finance*, 70(1):329–371, 2015.
- U.K. Müller and M.W. Watson. Testing models of low-frequency variability. *Econometrica*, 76(5):979–1016, 2008.
- F. Ortú, F. Severino, A. Tamoni, and C. Tebaldi. A persistence-based wold-type decomposition for stationary time series. *Quantitative Economics*, 11(1):203–230, 2020.
- P. Savor and M. Wilson. How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis*, 48(2):343–375, 2013.
- G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, pages 461–464, 1978.
- W.F. Sharpe. Asset allocation: Management style and performance measurement. *Journal of Portfolio Management*, 18(2):7–19, 1992.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- C.R. Tori. Federal Open Market Committee meetings and stock market performance. *Financial Services Review*, 10(1-4):163–171, 2001.
- J.A. Wachter and Y. Zhu. A model of two days: Discrete news and asset prices. *The Review of Financial Studies*, 35(5):2246–2307, 2022.
- J. Wang. A model of intertemporal asset prices under asymmetric information. *The Review of Economic Studies*, 60(2):249–282, 1993.
- G. Weitzner. The term structure of short selling costs. *Review of Finance*, pages 1–37, 2023.
- H. Wold. *A study in the analysis of stationary time series*. Almqvist & Wiksells Boktryckeri, 1938.

Appendix

A Dynamic CARA portfolio optimization

The general solution of Problem (2), presented in Proposition 1, is particularly lengthy. Therefore, here we suppose that the excess returns of the risky assets follow a VAR(2) process and we solve the problem

$$\begin{aligned} & \max_{u_0, \dots, u_{T-1}} \mathbb{E} [-e^{-\gamma W_T}] \\ & \text{sub} \quad W_t = (1 + r_f)W_{t-1} + u'_{t-1}r_t \\ & \quad \quad r_t = \mu + B_1r_{t-1} + B_2r_{t-2} + \varepsilon_t. \end{aligned} \quad (5)$$

We now provide a derivation for the special case of Proposition 1 with $p = 2$.

Proposition 2 *Consider Problem (5). For any $k = 1, \dots, T$, the optimal portfolio choice at time $T - k$ is*

$$u_{T-k}^* = \frac{1}{\gamma(1 + r_f)^{k-1}} \left(\Sigma^{-1} \mathbb{E}_{T-k} [r_{T-k+1}] - \sum_{i=1}^{\min\{k-1, 2\}} B_i' \Sigma^{-1} E_{k,i} \right), \quad (6)$$

where

$$\begin{aligned} E_{k,1} &= \mathbb{E}_{T-k} [r_{T-k+2} | r_{T-k+1} = 0] = \mu + B_2 r_{T-k}, \\ E_{k,2} &= \mathbb{E}_{T-k} [r_{T-k+3} | r_{T-k+1} = r_{T-k+2} = 0] = \mu. \end{aligned}$$

We propose an adaptation of the proofs of Theorem 5.1 and Proposition 5.2 of Chryssikou (1998) to the case $p = 2$. The original results are valid for $p = 1$ only. We provide the explicit derivation for the first steps of the algorithm corresponding to $k = 1, \dots, 4$. A formal proof can be obtained by induction.

Problem (5) can be stated through a Bellman equation and solved via backward induction. We denote by \mathbf{r}_t the realizations r_τ for $\tau = 0, \dots, t$. For $t = 0, 1, \dots, T - 1$, the value function at t is

$$V_t(W_t, \mathbf{r}_t) = \max_{u_0, \dots, u_t} \mathbb{E}_t [V_{t+1}(W_{t+1}, \mathbf{r}_{t+1})]$$

with the terminal condition $V_T(W_T, \mathbf{r}_T) = -e^{-\gamma W_T}$.

At time $T - 1$, we solve

$$\begin{aligned} V_{T-1}(W_{T-1}, \mathbf{r}_{T-1}) &= \max_{u_{T-1} \in \mathbb{R}^n} \mathbb{E}_{T-1} [-e^{-\gamma W_T}] \\ & \text{sub} \quad W_T = (1 + r_f)W_{T-1} + u'_{T-1}r_T \\ & \quad \quad r_T = \mu + B_1r_{T-1} + B_2r_{T-2} + \varepsilon_T. \end{aligned}$$

Since ε_T is normally distributed, by Proposition 5.1 in Chryssikou (1998), we rewrite the conditional expectation as

$$\begin{aligned}\mathbb{E}_{T-1}[-e^{-\gamma W_T}] &= \mathbb{E}_{T-1}[-\exp(-\gamma((1+r_f)W_{T-1} + u'_{T-1}r_T))] \\ &= -\exp(-\gamma(1+r_f)W_{T-1}) \mathbb{E}_{T-1}[\exp(-\gamma u'_{T-1}(\mu + B_1 r_{T-1} + B_2 r_{T-2} + \varepsilon_T))] \\ &= -\exp\left(-\gamma(1+r_f)W_{T-1} - \gamma u'_{T-1}(\mu + B_1 r_{T-1} + B_2 r_{T-2}) + \frac{1}{2}\gamma^2 u'_{T-1} \Sigma u_{T-1}\right).\end{aligned}$$

As a result, the maximization problem is equivalent to

$$\min_{u_{T-1} \in \mathbb{R}^n} -\gamma u'_{T-1}(\mu + B_1 r_{T-1} + B_2 r_{T-2}) + \frac{1}{2}\gamma^2 u'_{T-1} \Sigma u_{T-1}.$$

Since we obtain a convex optimization problem, first-order conditions are necessary and sufficient to find the minimum. The first-order condition is

$$-\gamma(\mu + B_1 r_{T-1} + B_2 r_{T-2}) + \gamma^2 \Sigma u_{T-1} = 0.$$

Thus, the optimal investment choice is

$$u_{T-1}^* = \frac{1}{\gamma} \Sigma^{-1}(\mu + B_1 r_{T-1} + B_2 r_{T-2}) = \frac{1}{\gamma} \Sigma^{-1} \mathbb{E}_{T-1}[r_T]$$

and it can be retrieved in eq. (6). By replacing u_{T-1}^* in the objective function, we get the value function at $T-1$:

$$V_{T-1}(W_{T-1}, \mathbf{r}_{\mathbf{T}-1}) = -\exp\left(-\gamma(1+r_f)W_{T-1} - \frac{1}{2} \mathbb{E}'_{T-1}[r_T] \Sigma^{-1} \mathbb{E}_{T-1}[r_T]\right).$$

We now move to time $T-2$ and consider the problem

$$\begin{aligned}V_{T-2}(W_{T-2}, \mathbf{r}_{\mathbf{T}-2}) &= \max_{u_{T-2} \in \mathbb{R}^n} \mathbb{E}_{T-2}[V_{T-1}(W_{T-1}, \mathbf{r}_{\mathbf{T}-1})] \\ \text{sub} \quad W_{T-1} &= (1+r_f)W_{T-2} + u'_{T-2} r_{T-1} \\ r_{T-1} &= \mu + B_1 r_{T-2} + B_2 r_{T-3} + \varepsilon_{T-1}.\end{aligned}$$

In the expression of $V_{T-1}(W_{T-1}, \mathbf{r}_{\mathbf{T}-1})$, r_{T-1} appears both in $-\gamma(1+r_f)W_{T-1}$ and in $\mathbb{E}'_{T-1}[r_T] \Sigma^{-1} \mathbb{E}_{T-1}[r_T]$. We take the expectation at $T-2$, use Proposition 5.1 in Chryssikou (1998) and compute the first-order condition. We obtain the optimal investment choice

$$\begin{aligned}u_{T-2}^* &= \frac{1}{\gamma(1+r_f)} (\Sigma^{-1}(\mu + B_1 r_{T-2} + B_2 r_{T-3}) - B'_1 \Sigma^{-1}(\mu + B_2 r_{T-2})) \\ &= \frac{1}{\gamma(1+r_f)} (\Sigma^{-1} \mathbb{E}_{T-2}[r_{T-1}] - B'_1 \Sigma^{-1} E_{2,1}).\end{aligned}$$

By defining $\Lambda_1 = (\Sigma^{-1} + B'_1 \Sigma^{-1} B_1)^{-1}$ and denoting with $|\cdot|$ the matrix determinant, the value function at $T-2$ turns out to be

$$\begin{aligned}V_{T-2}(W_{T-2}, \mathbf{r}_{\mathbf{T}-2}) &= -\sqrt{\frac{|\Lambda_1|}{|\Sigma|}} \exp(-\gamma(1+r_f)^2 W_{T-2}) \\ &\quad \exp\left(-\frac{1}{2} \mathbb{E}'_{T-2}[r_{T-1}] \Sigma^{-1} \mathbb{E}_{T-2}[r_{T-1}] - \frac{1}{2} E'_{2,1} \Sigma^{-1} E_{2,1}\right).\end{aligned}$$

Then, at $T - 3$, we face the problem

$$\begin{aligned} V_{T-3}(W_{T-3}, \mathbf{r}_{\mathbf{T}-3}) &= \max_{u_{T-3} \in \mathbb{R}^n} \mathbb{E}_{T-3} [V_{T-2}(W_{T-2}, \mathbf{r}_{\mathbf{T}-2})] \\ \text{sub} \quad W_{T-2} &= (1 + r_f)W_{T-3} + u'_{T-3}r_{T-2} \\ r_{T-2} &= \mu + B_1r_{T-3} + B_2r_{T-4} + \varepsilon_{T-2}. \end{aligned}$$

Similarly to before, we obtain the optimal investment choice

$$\begin{aligned} u_{T-3}^* &= \frac{1}{\gamma(1 + r_f)^2} (\Sigma^{-1}(\mu + B_1r_{T-3} + B_2r_{T-4}) - B'_1\Sigma^{-1}(\mu + B_2r_{T-3}) - B'_2\Sigma^{-1}\mu) \\ &= \frac{1}{\gamma(1 + r_f)^2} (\Sigma^{-1}\mathbb{E}_{T-3}[r_{T-2}] - B'_1\Sigma^{-1}E_{3,1} - B'_2\Sigma^{-1}E_{3,2}). \end{aligned}$$

By defining $\Lambda_2 = (\Sigma^{-1} + B'_1\Sigma^{-1}B_1 + B'_2\Sigma^{-1}B_2)^{-1}$, the value function at $T - 3$ is

$$\begin{aligned} V_{T-3}(W_{T-3}, \mathbf{r}_{\mathbf{T}-3}) &= -\sqrt{\frac{|\Lambda_1||\Lambda_2|}{|\Sigma|^2}} \exp(-\gamma(1 + r_f)^3W_{T-3}) \\ &\quad \exp\left(-\frac{1}{2}\mathbb{E}'_{T-3}[r_{T-2}]\Sigma^{-1}\mathbb{E}_{T-3}[r_{T-2}] - \frac{1}{2}E'_{3,1}\Sigma^{-1}E_{3,1} - \frac{1}{2}E'_{3,2}\Sigma^{-1}E_{3,2}\right). \end{aligned}$$

At $T - 4$ the problem is

$$\begin{aligned} V_{T-4}(W_{T-4}, \mathbf{r}_{\mathbf{T}-4}) &= \max_{u_{T-4} \in \mathbb{R}^n} \mathbb{E}_{T-4} [V_{T-3}(W_{T-3}, \mathbf{r}_{\mathbf{T}-3})] \\ \text{sub} \quad W_{T-3} &= (1 + r_f)W_{T-4} + u'_{T-4}r_{T-3} \\ r_{T-3} &= \mu + B_1r_{T-4} + B_2r_{T-5} + \varepsilon_{T-3}. \end{aligned}$$

As can be seen from the expressions of $V_{T-2}(W_{T-2}, \mathbf{r}_{\mathbf{T}-2})$ and $V_{T-3}(W_{T-3}, \mathbf{r}_{\mathbf{T}-3})$, the value functions at $T - 2$ and $T - 3$ have a similar form. Thus, we can write the optimal investment choice u_{T-4}^* similarly to u_{T-3}^* :

$$\begin{aligned} u_{T-4}^* &= \frac{1}{\gamma(1 + r_f)^3} (\Sigma^{-1}(\mu + B_1r_{T-4} + B_2r_{T-5}) - B'_1\Sigma^{-1}(\mu + B_2r_{T-4}) - B'_2\Sigma^{-1}\mu) \\ &= \frac{1}{\gamma(1 + r_f)^3} (\Sigma^{-1}\mathbb{E}_{T-4}[r_{T-3}] - B'_1\Sigma^{-1}E_{4,1} - B'_2\Sigma^{-1}E_{4,2}). \end{aligned}$$

Moreover, by defining $\Lambda_3 = \Lambda_2$, the value function at $T - 4$ is

$$\begin{aligned} V_{T-4}(W_{T-4}, \mathbf{r}_{\mathbf{T}-4}) &= -\sqrt{\frac{|\Lambda_1||\Lambda_2||\Lambda_3|}{|\Sigma|^3}} \exp(-\gamma(1 + r_f)^4W_{T-4}) \\ &\quad \exp\left(-\frac{1}{2}\mathbb{E}'_{T-4}[r_{T-3}]\Sigma^{-1}\mathbb{E}_{T-4}[r_{T-3}]\right) \\ &\quad \exp\left(-\frac{1}{2}E'_{4,1}\Sigma^{-1}E_{4,1} - \frac{1}{2}E'_{4,2}\Sigma^{-1}E_{4,2} - \frac{1}{2}\mu'\Sigma^{-1}\mu\right). \end{aligned}$$

By induction, it is straightforward to see that the solutions at the previous periods $T - k$ for $k \geq 5$ have a similar form. In general, the optimal portfolio choice at any $T - k$ is given by eq. (6) and the value function at time $T - k$ is

$$V_{T-k}(W_{T-k}, \mathbf{r}_{T-k}) = -\sqrt{\frac{\prod_{j=1}^{k-1} |\Lambda_j|}{|\Sigma|^{k-1}}} \exp\left(-\gamma(1+r_f)^k W_{T-k}\right) \\ \exp\left(-\frac{1}{2} \mathbb{E}'_{T-k}[r_{T-k+1}] \Sigma^{-1} \mathbb{E}_{T-k}[r_{T-k+1}] - \frac{1}{2} \sum_{i=1}^{\min\{k-1, 2\}} E'_{k,i} \Sigma^{-1} E_{k,i} - \frac{k-2-1}{2} \mu' \Sigma^{-1} \mu\right)$$

with

$$\Lambda_j = \left(\Sigma^{-1} + \sum_{i=1}^{\min\{j-1, 2\}} B_i' \Sigma^{-1} B_i \right)^{-1}.$$

B Extended Wold Decomposition in base N

The EWD of Ortu et al. (2020) exploits 2 as reference number. Indeed, their construction exploits a dyadic procedure inspired by the Discrete Haar Transform. However, it is possible to deduce alternative orthogonal decompositions by using any natural number N as base. A given zero-mean weakly stationary time series $\mathbf{x} = \{x_t\}_{t \in \mathbb{Z}}$ is decomposed into the sum of uncorrelated components whose persistence is now classified according to the time grid N^j at any scale j . The notation employed can be retrieved in Section A.1 of the online supplement of Ortu et al. (2020).

Let $\mathcal{H}_t(\varepsilon)$ be the Hilbert space spanned by the sequence of the fundamental innovations of x_t of eq. (1):

$$\mathcal{H}_t(\varepsilon) = \left\{ \sum_{k=0}^{+\infty} a_k \varepsilon_{t-k} : \sum_{k=0}^{+\infty} a_k^2 < +\infty \right\}.$$

Given $N \in \mathbb{N}$, define the scaling operator $\mathbf{R} : \mathcal{H}_t(\varepsilon) \rightarrow \mathcal{H}_t(\varepsilon)$ by

$$\mathbf{R} : \sum_{k=0}^{+\infty} a_k \varepsilon_{t-k} \mapsto \sum_{k=0}^{+\infty} \frac{a_k}{\sqrt{N}} \sum_{i=0}^{N-1} \varepsilon_{t-Nk-i} = \sum_{k=0}^{+\infty} \frac{a_{\lfloor \frac{k}{N} \rfloor}}{\sqrt{N}} \varepsilon_{t-k},$$

where $\lfloor \cdot \rfloor$ associates any real number c with the integer $\lfloor c \rfloor = \max\{n \in \mathbb{Z} : n \leq c\}$. The scaling operator is isometric on $\mathcal{H}_t(\varepsilon)$ and its adjoint is the operator $\mathbf{R}^* : \mathcal{H}_t(\varepsilon) \rightarrow \mathcal{H}_t(\varepsilon)$ such that

$$\mathbf{R}^* : \sum_{k=0}^{+\infty} a_k \varepsilon_{t-k} \mapsto \sum_{k=0}^{+\infty} \frac{1}{\sqrt{N}} \left(\sum_{i=0}^{N-1} a_{Nk+i} \right) \varepsilon_{t-k}.$$

The wandering subspace $\mathcal{L}_t^{\mathbf{R}} = \mathcal{H}_t(\varepsilon) \ominus \mathbf{R} \mathcal{H}_t(\varepsilon)$ is

$$\mathcal{L}_t^{\mathbf{R}} = \left\{ \sum_{k=0}^{+\infty} a_k \varepsilon_{t-k} \in \mathcal{H}_t(\varepsilon) : a_{Nk+N-1} = - \sum_{i=0}^{N-2} a_{Nk+i} \quad \forall k \in \mathbb{N}_0 \right\}.$$

In addition, by defining the innovations

$$\tilde{\varepsilon}_t^{(j)} = \frac{1}{\sqrt{N^{j-1}}} \sum_{k=0}^{N^{j-1}-1} \varepsilon_{t-k}$$

for all $j \in \mathbb{N}$ and $t \in \mathbb{Z}$, we have

$$\mathbf{R}^{j-1} \mathcal{L}_t^{\mathbf{R}} = \left\{ \sum_{k=0}^{+\infty} \left(\sum_{i=0}^{N-2} a_{Nk+i} \tilde{\varepsilon}_{t-N^j k - N^{j-1} i}^{(j)} - \left(\sum_{i=0}^{N-2} a_{Nk+i} \right) \tilde{\varepsilon}_{t-N^j k - N^{j-1} (N-1)}^{(j)} \right) \in \mathcal{H}_t(\varepsilon) : a_h \in \mathbb{R} \right\}.$$

Following the same steps as Theorem 1 in Ortú et al. (2020), we decompose x_t into the sum of the uncorrelated persistent components $x_t^{(j)}$:

$$x_t = \sum_{j=1}^{+\infty} x_t^{(j)}, \quad x_t^{(j)} = \sum_{k=0}^{+\infty} \sum_{q=0}^{N-1} \beta_{k,q}^{(j)} \tilde{\varepsilon}_{t-kN^j - qN^{j-1}}^{(j)}$$

with

$$\beta_{k,q}^{(j)} = \frac{1}{\sqrt{N^{j-1}}} \left(\sum_{i=0}^{N^{j-1}-1} \alpha_{N^j k + N^{j-1} q + i} - \frac{1}{N} \sum_{i=0}^{N^j-1} \alpha_{N^j k + i} \right).$$

As expected, the orthogonality of persistent components and the unit variance of the scale-specific innovations $\tilde{\varepsilon}_t^{(j)}$ allow for a variance decomposition of the original time series:

$$\text{var}(x_t) = \sum_{j=1}^{+\infty} \text{var}(x_t^{(j)}) = \sum_{j=1}^{+\infty} \sum_{k=0}^{+\infty} \sum_{q=0}^{N-1} \left(\beta_{k,q}^{(j)} \right)^2.$$