Dynamic Portfolio Management with Views at Multiple Horizons[1]

Attilio Meucci
attilio.meucci@symmys.com

Marco Nicolosi
marco.nicolosi@unipg.it

this revision: 16 April 2015
latest revision and code: http://www.symmys.com/node/831

Abstract

We introduce Dynamic Entropy Pooling, a quantitative technique to perform dynamic portfolio construction with discretionary, non-synchronous views. With Dynamic Entropy Pooling, the portfolio manager can embed in the allocation process signals with life spans ranging from minutes to years, calendar views, autocorrelation stress-testing, and the traditional views on expectations, correlations and volatilities.

After introducing the theoretical framework for Dynamic Entropy Pooling, we show how to solve the respective portfolio construction problem by means of dynamic programming with time-dependent coefficients. To understand the optimal exposures ensuing from Dynamic Entropy Pooling we analyze a variety of relevant sub-cases and we present two case-studies.

Fully documented code is available at http://www.symmys.com/node/831.

JEL Classification: C1, G11

Keywords: dynamic strategies, discretionary allocation, mean-reversion, multivariate Ornstein-Uhlenbeck, Kullback Leibler divergence, relative entropy, optimal policy, dynamic programming, calculus of variations.

[1]The authors are grateful to Marcello Colasante, Stefano Herzel, Federico Mosca, and Mustafa Pinar for their feedback, and to Tai-Ho Wang at the Department of Mathematics at Baruch College - CUNY for inviting Marco Nicolosi to NYC and work on this project with Attilio Meucci

1 Introduction

Portfolio construction is about blending market views or signals into optimized portfolios. Conceptually, there are three dimensions of complexity in constructing portfolios.

First, simple "myopic" sequences of one-period problems versus a more complex multi-period framework.

Second, systematic trading based on rules and algorithms that process market signals, versus the more complex discretionary trading, where trades follow from the bespoke subjective bearish/bullish views of analysts, strategists, or macroeconomists.

Third, simple mean-variance optimization versus true optimization that accounts for constraints and market impact.

We summarize the main approaches in the table below, which we proceed to comment.

	Discretionary	Multiperiod	Mkt impact
Grinold ('89)	×	×	×
Black-Litterman ('90)	✓	×	×
Entropy Pooling ('08)	✓	×	×
Davis-Lleo ('14)	×	✓	×
Garleanu-Pedersen ('13)	×	✓	✓
Dynamic Entropy Pooling	✓	✓	✓

Table 1: Methodologies for portfolio construction

The path to portfolio construction started decades ago, with the path-breaking work on systematic factors by [Rosenberg and Lanstein, 1985] and [Fama and French, 1993], which then became the systematic cross-sectional signals popularized in [Grinold, 1989] and later in [Grinold and Kahn, 1999], which in turn were adopted by systematic hedge funds worldwide, see e.g. [Asness et al., 2013]. In these early "Grinold" approaches, dynamic portfolios are built from a sequence of myopic one-period mean-variance optimizations, where the expected returns are set systematically as deterministic functions of securities characteristics. In essence, Grinold laid the foundation for portfolio construction with an effective systematic, one-period, no-market impact framework.

Next, portfolio construction has focused on processing more complex, discretionary views. The celebrated model by [Black and Litterman, 1990] is among the first approaches to use quantitative techniques to combine subjective views within a base-case risk model. In Black-Litterman, the distribution of the risk factors is assumed normal and the views are expressed on the expectation of linear combinations of the risk factors. The Entropy Pooling approach in [Meucci, 2008] extends the Black-Litterman methodology to allow for fully general market distributions and for views on general features of such distributions. Conceptually, a key difference between the Black-Litterman methodology and Entropy Pooling is that Black-Litterman expresses views on the *parameters* of the market under rigid (normal) parametric assumptions, whereas Entropy Pooling expresses views on *features* of the market, such as expectation or correlations, regardless of the market distribution.

Despite its generality, the original Entropy Pooling approach operates within a one-period framework: the subjective views or stress-tests refer to the distribution of the market at one specific investment horizon. However, portfolio managers operate with multiple horizons. Furthermore, the views often stem from signals with different spans: signals from microstructure analysis can last a few minutes; signals from technical analysis have a life of hours or days; econometric mean reversion materializes over days or months; and macroeconomic signals can be of the order of years.

True multi-period portfolio construction is tackled in [Grinold, 2007], [Grinold, 2010] and [Gârleanu and Pedersen, 2013], where systematic signals follow a VAR(1) process.

[Davis and Lleo, 2013] provide a theoretical framework for multi-period portfolio view-processing with discretionary views. However, the framework can be implemented in practice only with systematic views that follow a VAR(1) process. Furthermore, the authors do not show how to build their view-processing framework within a multi-period optimization.

The above history of portfolio construction is summarized in Table 1.

In this article we introduce Dynamic Entropy Pooling, a quantitative approach to blend views and stress tests at multiple horizons, and to embed such views and stress-tests into a multi-period optimization process. Dynamic Entropy Pooling allows the discretionary portfolio manager to implement a truly quantitative portfolio construction process. Portfolio managers can focus on the view generation process, by filtering the opinions of analysts, strategists, or macroeconomists, while leaving to Dynamic Entropy Pooling the burden to optimally reflect their views in a dynamic portfolio.

First of all, we model the market risk drivers (such as interest-rate term-structures, implied volatility surfaces, stock log-prices, etc.). Using the nomenclature of [Black and Litterman, 1990], we call the estimated process of the drivers the "prior" model. More precisely, in Dynamic Entropy Pooling we assume a multivariate Ornstein-Uhlenbeck process as the prior model. This process allows us to model mean-reversion, cointegration, explosive behaviors, as well as the standard arithmetic and geometric Brownian motion.

Second, we model views, scenarios, and stress-tests on the risk drivers, which are constraints on the multivariate stochastic process followed by the risk drivers. To preserve analytical tractability, in Dynamic Entropy Pooling we model the views as statements on expectations and cross-covariances of arbitrary linear combinations of the multivariate process at arbitrary times. This approach accommodates a number of practical applications: in particular, it is possible to process cross-assets calendar views and stress-testing on volatility propagation and cross-assets autocorrelations.

Third, we compute the posterior process for the dynamic drivers, which embeds the views in the prior model, using Entropy Pooling framework in [Meucci, 2008]. In particular, we adopt the analytical implementation of Entropy Pooling, which allows for computational speed even in large-dimensional markets with an arbitrary number of trading periods.

Fourth, we map the risk drivers into the portfolio profit and loss (P&L) over each trading period. To this purpose, we make the standard assumption that the P&L is linear in the increments of the risk drivers through exposures (durations for rates, deltas and vegas for implied volatilities, etc.). We highlight that the risk drivers can also include external variables [Chen et al., 1986], such as macroeconomic variables (CPI, GDP, etc.), which do not affect directly the portfolio P&L, but on which the manager can express views that indirectly affect the P&L via correlations and autocorrelations.

Fifth and last, we optimize a multi-period objective function. In Dynamic Entropy Pooling we consider as objective function the discounted expected stream of the future P&L's, penalized with a risk aversion term and with a market impact term, as in [Gârleanu and Pedersen, 2013]. As a matter of fact, our approach and results coincides with [Gârleanu and Pedersen, 2013] in the special case of no discretionary views and no constraints. Our work generalizes

[Gârleanu and Pedersen, 2013] in three directions: first and foremost, we include discretionary views in the portfolio construction process; second, we allow for arbitrary equality and inequality linear constraints; and third, we account for estimation risk.

This article is organized as follows. In Section 2 we lay in full generality the theoretical foundations of dynamic portfolio management in the presence of multi-period views. This entails adapting to the present dynamic, multi-period context the ten steps of quantitative portfolio management in [Meucci, 2011].

In Section 3 we introduce Dynamic Entropy Pooling: a tractable, yet flexible and general set of modeling assumptions, which make it possible to implement in practice the general approach to dynamic portfolio management in the presence of multi-period discretionary views and constraints, as laid out in the previous section.

In Section 4 we present the simplest solution of the Dynamic Entropy Pooling framework, namely when there is no market impact. In this case we obtain the optimal policy via a series of simple mean-variance optimizations, each reflecting a different set of multi-horizon views.

In Section 5 we include market impact in the Dynamic Entropy Pooling framework. To determine the optimal policy in this scenario, we solve analytically a sequence of Bellman equations with time-dependent coefficients, where the coefficients follow from a backward recursion.

In Section 6 we consider the most general case of Dynamic Entropy Pooling: market impact, constraints, and estimation updating. To maximize the multiperiod objective function in this context, we use a heuristic approach, solving for the optimal deterministic policy via calculus of variations, only preserving the first step of the policy, and then repeating the process at each step. As it turns out, this last solution to Dynamic Entropy Pooling nests the analytical solutions in the simpler frameworks of Sections 4 and 5. We also tested complex linear policies such as those in [Brandt et al., 2009] and [Moallemi and Sağlam, 2012], which can be computed via second-order cone programming. However, the repeated deterministic policy provides the same results as the repeated linear policy.

In Section 7 we show two simple low-dimensional case studies, to further our understanding of the Dynamic Entropy Pooling approach.

Section 8 concludes. A technical appendix contains results that can be skipped at first reading. Fully documented code is available at http://www.symmys.com/node/831.

2 Theory: a general framework for discretionary multi-period portfolio management

In this section we introduce in full generality a framework for multi-period portfolio management, which we summarize in the table below, adapting the general ten-step approach to portfolio management in [Meucci, 2011] to the present multi-period environment

Step	Computation
1-2-3	Market dynamics
$\tilde{2}$-$\tilde{3}$	Discretionary views
4-5-6	Profit-and-Loss
7	Ex-ante evaluation
8	Portfolio construction
9	Execution
10	Performance analysis

Table 2: Framework for discretionary multi-period portfolio management

To facilitate the understanding of the general framework, we illustrate each concept with an oversimplified example. In Section 3 we discuss Dynamic Entropy Pooling, which is a realistic, flexible model to implement the general framework laid out here.

2.1 Market dynamics (1: Quest for Invariance; 2: Estimation; 3: Projection)

The randomness in the market is driven by a set of $\bar{n}$ risk drivers whose values at time t are collected in the vector $\boldsymbol{X}_t \equiv (X_{1,t}, \ldots, X_{\bar{n},t})'$, such as interest-rate term-structures, implied volatilities surfaces, etc. The first ingredient is the process followed by the key risk-drivers, which we collect in a single vector

$$\boldsymbol{X}_{t \rightsquigarrow \bar{t}} | \boldsymbol{i}_t \equiv \begin{pmatrix} \boldsymbol{X}_t \\ \boldsymbol{X}_{t+1} \\ \vdots \\ \boldsymbol{X}_{\bar{t}} \end{pmatrix} | \boldsymbol{i}_t \sim f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}} | \boldsymbol{i}_t}, \tag{1}$$

where $\boldsymbol{i}_t$ is the information available at time t. The process of the risk drivers is monitored at discrete unit steps $t, t+1, \ldots, \bar{t}$, where the last point $\bar{t}$ is arbitrarily far in the future, and where the time unit corresponds to the rebalancing frequency for a given investment style.

Example 1 *Consider a trivial market of $\bar{n}$ stocks and a cash account with zero interest rates. The risk drivers $\mathbf{X}$ for the stocks can be set as the stock prices. Assume that the stocks are traded at most once a day, and thus time is measured in days. A simplistic baseline process could be a random walk with normal shocks with expected value $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\sigma}^2$ (multivariate Brownian motion). The process $\mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t$ follows the distribution:*

$$\begin{pmatrix} \boldsymbol{X}_t \\ \boldsymbol{X}_{t+1} \\ \vdots \\ \boldsymbol{X}_{\bar{t}} \end{pmatrix} | \boldsymbol{i}_t \sim N\left(\begin{pmatrix} \boldsymbol{x}_t \\ \boldsymbol{x}_t + \boldsymbol{\mu} \\ \vdots \\ \boldsymbol{x}_t + (\bar{t}-t)\boldsymbol{\mu} \end{pmatrix}, \begin{pmatrix} 0\boldsymbol{\sigma}^2 & \mathbf{0} & \cdot & \cdot \\ \mathbf{0} & 1\boldsymbol{\sigma}^2 & \mathbf{0} & \vdots \\ \vdots & \vdots & \vdots & \vdots \\ \cdot & \cdot & \cdot & (\bar{t}-t)\boldsymbol{\sigma}^2 \end{pmatrix} \right), \tag{2}$$

where the variance is a block diagonal matrix. Note that the first diagonal block is zero. This is just a compact notation to indicate that the marginal distribution of $\mathbf{X}_t$ is a Dirac delta about $\boldsymbol{x}_t$, as at time t the risk drivers $\mathbf{X}_t$ have already realized. And we can think of the delta distribution as the limit of a Gaussian density, as the variance approaches zero.

2.2 Exposures and P&L (4: Pricing; 5: Aggregation; 6: Attribution)

The second ingredient is $\boldsymbol{b}_t \equiv (b_{1,t}, \ldots, b_{\bar{n},t})'$, the $\bar{n}$-dimensional vector of exposures to the risk drivers at time t, and we collect the exposures at different times in a unique vector:

$$\boldsymbol{b}_{t \rightsquigarrow \bar{t}} \equiv \begin{pmatrix} \boldsymbol{b}_t \\ \boldsymbol{b}_{t+1} \\ \cdot \\ \boldsymbol{b}_{\bar{t}} \end{pmatrix}. \tag{3}$$

The exposures are control variables, which steer the portfolio profit and loss (P&L) through time. More precisely, the P&L (without market impact) Π_{s+1} over the unit period from time s to $s+1$, for any future time $s \geq t$, depends on the path of the risk drivers over the period and the exposures at the beginning of the period

$$\Pi_{s+1} = p\&l(\boldsymbol{X}_{s \rightsquigarrow s+1}, \boldsymbol{b}_s), \tag{4}$$

where $p\&l$ is a suitable deterministic function.

Example 2 *In our example the exposures* $\boldsymbol{b}_s$ *are the number of shares at time* s. *Then the generic P&L between time* s *and* $s+1$ *reads*

$$\Pi_{s+1} = \boldsymbol{b}_s' \Delta \boldsymbol{X}_{s+1}, \tag{5}$$

where $\Delta \boldsymbol{X}_{s+1} = \boldsymbol{X}_{s+1} - \boldsymbol{X}_s$ *is the increment on the trading interval; and where the capital necessary to rebalance the exposures* $\boldsymbol{b}$ *comes from a cash account, which does not contribute to the P&L.*

Given the sequence of the exposures, the distribution of the risk drivers determine the distribution of the path of the profits and losses $\boldsymbol{\Pi}_{t \rightsquigarrow \bar{t}}$, conditioned on current information:

$$(\boldsymbol{b}_{t \rightsquigarrow \bar{t}}, f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}) \mapsto \boldsymbol{\Pi}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t \equiv \begin{pmatrix} \Pi_{t+1} \\ \Pi_{t+2} \\ \cdot \\ \Pi_{\bar{t}} \end{pmatrix} |\boldsymbol{i}_t \sim f_{\boldsymbol{\Pi}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}. \tag{6}$$

Example 3 *In our example, the distribution of the profit and loss sequence, given the sequence of the exposures, as defined in (3), comes from the distribution of the process of the risk drivers (2):*

$$\boldsymbol{\Pi}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t \equiv \begin{pmatrix} \Pi_{t+1} \\ \Pi_{t+2} \\ \cdot \\ \Pi_{\bar{t}} \end{pmatrix} |\boldsymbol{i}_t \sim N\left(\begin{pmatrix} \boldsymbol{b}_t'\boldsymbol{\mu} \\ \boldsymbol{b}_{t+1}'\boldsymbol{\mu} \\ \cdot \\ \boldsymbol{b}_{\bar{t}-1}'\boldsymbol{\mu} \end{pmatrix}, \begin{pmatrix} \boldsymbol{b}_t'\boldsymbol{\sigma}^2\boldsymbol{b}_t & \boldsymbol{0} & \cdot & \cdot \\ \boldsymbol{0} & \boldsymbol{b}_{t+1}'\boldsymbol{\sigma}^2\boldsymbol{b}_{t+1} & & \\ \cdot & & \cdot & \\ \cdot & & & \boldsymbol{b}_{\bar{t}-1}'\boldsymbol{\sigma}^2\boldsymbol{b}_{\bar{t}-1} \end{pmatrix} \right). \tag{7}$$

2.3 Ex-ante Evaluation (7)

For each generic time t there is an index of satisfaction $\mathbb{S}_t$ that depends on the distribution of the future P&L $\boldsymbol{\Pi}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t$. Since the P&L depends on the exposures and the process of the risk drivers as in (6), then the index of satisfaction $\mathbb{S}_t$ must be a function $\mathcal{S}_t$ of the exposures and the distribution of the risk drivers

$$\mathbb{S}_t\{\boldsymbol{\Pi}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t\} = \mathcal{S}_t(\boldsymbol{b}_{t\rightsquigarrow\bar{t}}, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}). \tag{8}$$

Example 4 *In our example, let us choose a myopic (one-period) mean-variance trade-off:*

$$\mathbb{S}_t\{\Pi_{t+1}, \Pi_{t+2}, \ldots, \Pi_{\bar{t}}|\boldsymbol{i}_t\} = \mathbb{E}_t\{\Pi_{t+1}\} - \gamma\, \mathbb{V}_t\{\Pi_{t+1}\}, \tag{9}$$

where $\mathbb{E}_t\{\cdot\}$ *and* $\mathbb{V}_t\{\cdot\}$ *are respectively the expectation and the variance operator conditioned to the information* $\boldsymbol{i}_t$*, and* γ *is the risk aversion parameter. Then satisfaction depends on the number of shares and the stock process (2), as follows*

$$\mathcal{S}_t(\boldsymbol{b}_{t\rightsquigarrow\bar{t}}, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}) = \boldsymbol{b}_t\boldsymbol{\mu} - \gamma \boldsymbol{b}_t'\boldsymbol{\sigma}^2\boldsymbol{b}_t. \tag{10}$$

2.4 Portfolio Construction (8)

When constructing a portfolio at the generic time t, we process the information available at time t and decide on a set of exposures that we deem optimal.

More precisely, at time t, the future exposures $\{\boldsymbol{b}_s\}_{s\geq t}$ must be specified as the output of a policy. A policy is a set of functions calibrated at time t, which we denote by $p_s(\boldsymbol{i}_s)$, which map the information $\boldsymbol{i}_s$ available at time s into a vector of exposures $\boldsymbol{b}_s$, for any $s \geq t$

$$\boldsymbol{b}_s = p_s(\boldsymbol{i}_s), \text{ for any } s \geq t. \tag{11}$$

Therefore, the portfolio is fully specified by a policy:

$$p_{t\rightsquigarrow\bar{t}} : \begin{pmatrix} \boldsymbol{i}_t \\ \boldsymbol{i}_{t+1} \\ \cdot \\ \boldsymbol{i}_{\bar{t}} \end{pmatrix} \mapsto \begin{pmatrix} p_t(\boldsymbol{i}_t) \\ p_{t+1}(\boldsymbol{i}_{t+1}) \\ \cdot \\ p_{\bar{t}}(\boldsymbol{i}_{\bar{t}}) \end{pmatrix}. \tag{12}$$

Then, with minor abuse of notation

$$\mathbb{S}_t\{\boldsymbol{\Pi}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t\} = \mathcal{S}_t(p_{t\rightsquigarrow\bar{t}}, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}), \tag{13}$$

where the function $\mathcal{S}_t$ depends on the information $\boldsymbol{i}_t$. This setting is fully general: it includes one-period optimization, as well as multi-period targets and dynamic rebalancing (see e.g. [Merton, 1992]).

Example 5 *In our example, where the satisfaction is given by the one-period mean-variance trade-off (9), only the first term $p_t(\boldsymbol{i_t})$ of the policy function plays a role. Therefore from (10):*

$$\mathcal{S}_t(p_{t\rightsquigarrow\bar{t}}, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}) = [p_t(\boldsymbol{i}_t)]'\boldsymbol{\mu} - \gamma[p_t(\boldsymbol{i}_t)]'\boldsymbol{\sigma}^2[p_t(\boldsymbol{i}_t)]. \tag{14}$$

The optimization framework yields the optimal policy calibrated at time t

$$p^*_{t\rightsquigarrow\bar{t}} = \text{argmax}_{p_{t\rightsquigarrow\bar{t}}\in\mathcal{C}_t}\, \mathcal{S}_t(p_{t\rightsquigarrow\bar{t}}, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}), \tag{15}$$

where $\mathcal{C}_t$ is a set of potential additional constraints.

Example 6 *In our example, leaving the problem unconstrained, we obtain*

$$p^*_{t\rightsquigarrow\bar{t}} = \begin{pmatrix} \frac{1}{2\gamma}(\boldsymbol{\sigma}^2)^{-1}\boldsymbol{\mu} \\ p_{t+1}(\boldsymbol{i}_{t+1}) \\ \vdots \end{pmatrix} = \text{argmax}_{p_{t\rightsquigarrow\bar{t}}}\{p'_t(\boldsymbol{i}_t)\boldsymbol{\mu} - \gamma p'_t(\boldsymbol{i}_t)\boldsymbol{\sigma}^2 p_t(\boldsymbol{i}_t)\}, \tag{16}$$

where we emphasize how the subsequent policy functions $p_{t+1}, p_{t+2}...$ are left undetermined, as the objective function is a myopic one-period function. Then the optimal allocation for the next period is the standard mean-variance trade-off

$$t \rightarrow \boldsymbol{b}^*_t = \frac{1}{2\gamma}(\boldsymbol{\sigma}^2)^{-1}\boldsymbol{\mu}. \tag{17}$$

2.5 Further steps (9: Execution; 10: Performance Analysis)

We decide the strategy to achieve the next-step optimal exposure $\boldsymbol{b}^*_t$ in such a way to minimize market impact. Then, we evaluate ex-post the performance achieved by our process. These last two steps are beyond the scope of the present article.

2.6 Embedding the views ($\tilde{2}$: Estimation; $\tilde{3}$: Projection)

Views are statements on the potential outcomes of the risk drivers. For instance, the classical views a-la Black-Litterman are statements on expectations of linear combinations of the process. Similarly, one can express statements on correlations, volatilities, tails. As sketched in the process map in Table 2, in order to embed views in the allocation process, we need to revisit the very beginning of the process, namely the formulation of the process for the risk drivers, Steps 2-3.

More formally, the views are a set of statements $\mathcal{V}$ on the distribution of the process $f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}$ followed by the risk drivers $\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t$. We denote the distribution of a process that satisfies the views as follows

$$f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t} \in \mathcal{V}_t. \tag{18}$$

The above notation highlights that the views act as a constraint on the set of processes that the risk drivers can follow. Notice that since we are considering the whole process for the drivers $\boldsymbol{X}_{t\rightarrow\bar{t}}|\boldsymbol{i}_t$, the views include calender statements on different features of the process at different times, such as expectations on growth of a driver, or stress-testing of autocorrelations.

Example 7 *Suppose that we have a view that the return of the n-th stock over the next period is twice the forecast in the base-case process (2)*

$$\mathbb{E}_t\{\Delta X_{n,t+1}\} = 2\mu_n. \tag{19}$$

We easily obtain a process that satisfies the views by shifting all the future expectations of the n-th stock in the original Brownian motion (2), as follows:

$$\mathbb{E}_t\{X_{n,s}\} = x_{n,t} + (s - t + 1)\mu_n \qquad \textit{for any} \qquad s = t+1, \ldots, \bar{t}. \tag{20}$$

To embed the views in the investment process, we start from a base-case risk model, i.e. a model for the joint distribution of the process of the risk drivers $f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}$, as in (1). In the terminology of [Black and Litterman, 1990], this is the "prior" distribution, which in the original Black-Litterman framework is set in terms of CAPM-like equilibrium arguments.

Then, we need a mechanism that replaces the process for the risk drivers with a new process that satisfies the views

$$f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t} \notin \mathcal{V}_t \rightarrow \bar{f}_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t} \in \mathcal{V}_t. \tag{21}$$

Then, simply by following all the subsequent steps in Table 2, we automatically ensure an allocation for the next period consistent with the views

$$\bar{f}_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t} \in \mathcal{V}_t \overset{(\text{Table } 2)}{\longrightarrow} b_t^*. \tag{22}$$

Example 8 *In our example, let us suppose that the revised process (20) is the most natural perturbation of the original Brownian motion. Then we obtain the revised optimal allocation*

$$\boldsymbol{b}_t^* = \frac{1}{\gamma}(\boldsymbol{\sigma}^2)^{-1}\begin{pmatrix}\mu_1\\ \vdots\\ 2\mu_n\\ \vdots\\ \mu_{\bar{n}}\end{pmatrix}. \tag{23}$$

As expected, the optimal allocation provides a larger exposure to the n-th stock, given the bullish view on it (19).

3 Practice: Dynamic Entropy Pooling

In this section we present Dynamic Entropy Pooling, a practical implementation of all the steps for dynamic portfolio management and for views processing, which was laid out in full generality in Section 2. Unlike in the toy examples in Section 2, Dynamic Entropy Pooling is both flexible and able to realistically model a wide variety of real-life investment scenarios.

3.1 Market dynamics (1: Quest for Invariance; 2: Estimation; 3: Projection)

The first ingredient is a prior model for the evolution of the risk drivers, as prescribed by the general theoretical framework (1). The univariate Brownian motion in the toy example (2) is not sufficient to model general market dynamics. Instead, we model the $\bar{n}$ risk drivers $\boldsymbol{X}_t$ as a multivariate Ornstein-Uhlenbeck process

$$d\boldsymbol{X}_t = (-\boldsymbol{\theta}\boldsymbol{X}_t + \boldsymbol{\mu})dt + \boldsymbol{\sigma} d\boldsymbol{W}_t, \tag{24}$$

following [Meucci, 2009], where the reader can find all the proofs of the statements to follow.

In this expression $\boldsymbol{\theta}$ is the $\bar{n} \times \bar{n}$ transition matrix, namely a square matrix that defines the deterministic portion of the evolution of the process; $\boldsymbol{\mu}$ is a fully generic $\bar{n} \times 1$ vector, which represents the unconditional expectation when this is defined; $\boldsymbol{\sigma}$ is the $\bar{n} \times \bar{n}$ scatter generator, namely a full-rank square matrix that induces the dispersion of the process; $\boldsymbol{W}_t$ is an $\bar{n} \times 1$ vector of independent Brownian motions.

Note that $\boldsymbol{\theta}$ is fully generic, and that the process (24) is mean reverting only if the eigenvalues of $\boldsymbol{\theta}$ have positive real part. In this case $\boldsymbol{\theta}^{-1}\boldsymbol{\mu}$ represents the long term expectations of the factors. If $\boldsymbol{\theta}$ has eigenvalues whose real part is null or negative, then it means that a combination of factors has an explosive behavior, respectively like a random-walk or exponentially explosive. If $\boldsymbol{\theta}$ has some eigenvalues whose real part is positive, then there exists cointegration.

Note that $\boldsymbol{\sigma}$ is only required to be full rank. Furthermore, we assume without loss of generality that $\boldsymbol{\sigma}$ is also symmetric. Indeed, if we start with a fully arbitrary $\boldsymbol{\sigma}$, we can then replace it with the Riccati root of the matrix $\boldsymbol{\sigma}\boldsymbol{\sigma}'$, i.e. the matrix $\boldsymbol{z}$ that solves the two equations $\boldsymbol{z}^2 = \boldsymbol{\sigma}\boldsymbol{\sigma}'$ and $\boldsymbol{z} = \boldsymbol{z}'$.

To estimate the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ of the multivariate Ornstein-Uhlenbeck process (24) we rely on standard econometric analysis. In the process, we superimpose a generalized risk-parity condition, similar in spirit to the equilibrium condition in [Black and Litterman, 1990], to ensure that we do not obtain unwieldy corner solutions for the optimal portfolios

$$\mu_n \propto \sqrt{[\boldsymbol{\sigma}^2]_{n,n}}, \qquad n = 1, \ldots, \bar{n}. \tag{25}$$

Then the prior distribution $f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$ followed by the market process is jointly normal:

$$\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t \sim N(\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}}, \boldsymbol{\sigma}^2_{t \rightsquigarrow \bar{t}}). \tag{26}$$

For the expectation in (26) we have:

$$\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}} \equiv \begin{pmatrix} e^{-0\boldsymbol{\theta}}\boldsymbol{x}_t + \left(\mathbb{I}_{\bar{n}} - e^{-0\boldsymbol{\theta}}\right)\boldsymbol{\theta}^{-1}\boldsymbol{\mu} \\ e^{-1\boldsymbol{\theta}}\boldsymbol{x}_t + \left(\mathbb{I}_{\bar{n}} - e^{-1\boldsymbol{\theta}}\right)\boldsymbol{\theta}^{-1}\boldsymbol{\mu} \\ \vdots \\ e^{-(\bar{t}-t)\boldsymbol{\theta}}\boldsymbol{x}_t + \left(\mathbb{I}_{\bar{n}} - e^{-(\bar{t}-t)\boldsymbol{\theta}}\right)\boldsymbol{\theta}^{-1}\boldsymbol{\mu} \end{pmatrix}, \tag{27}$$

where $\mathbb{I}_{\bar{n}}$ is the $\bar{n} \times \bar{n}$ identity matrix. If $\boldsymbol{\theta}$ is singular, $\boldsymbol{\theta}^{-1}$ means $\lim_{\epsilon \to 0} \boldsymbol{\theta}_{\epsilon}^{-1}$,where $\boldsymbol{\theta}_{\epsilon}$ is any invertible perturbation of $\boldsymbol{\theta}$, see Appendix A.1.

As for the covariance in (26), we first define the square matrix $\boldsymbol{\sigma}_\tau^2$, for a generic $\tau \geq 0$ in terms of its stacked columns

$$vec(\boldsymbol{\sigma}_\tau^2) \equiv (\boldsymbol{\theta} \oplus \boldsymbol{\theta})^{-1} \left(\mathbb{I}_{\bar{n}^2} - e^{-(\boldsymbol{\theta} \oplus \boldsymbol{\theta})\tau}\right) vec(\boldsymbol{\sigma}^2), \tag{28}$$

where $\oplus$ is the Kronecker sum, and $\mathbb{I}_{\bar{n}^2}$ is the $\bar{n}^2 \times \bar{n}^2$ identity matrix. Notice that $\boldsymbol{\theta} \oplus \boldsymbol{\theta}$ is invertible if and only if $\boldsymbol{\theta}$ is invertible. If $\boldsymbol{\theta}$ is not invertible then $(\boldsymbol{\theta} \oplus \boldsymbol{\theta})^{-1} \equiv \lim_{\epsilon \to 0} (\boldsymbol{\theta}_\epsilon \oplus \boldsymbol{\theta}_\epsilon)^{-1}$, where $\boldsymbol{\theta}_\epsilon$ is an invertible perturbation of $\boldsymbol{\theta}$, see Appendix A.1 for more details. Then

$$\boldsymbol{\sigma}^2_{t \rightsquigarrow \bar{t}} \equiv \begin{pmatrix} \boldsymbol{\sigma}_0^2 & \boldsymbol{\sigma}_0^2 e^{-\boldsymbol{\theta}'} & \boldsymbol{\sigma}_0^2 e^{-2\boldsymbol{\theta}'} & \cdot & \boldsymbol{\sigma}_0^2 e^{-(\bar{t}-t)\theta'} \\ e^{-\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & \boldsymbol{\sigma}_1^2 & \boldsymbol{\sigma}_1^2 e^{-\boldsymbol{\theta}'} & \cdot & \boldsymbol{\sigma}_1^2 e^{-(\bar{t}-t-1)\boldsymbol{\theta}'} \\ e^{-2\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & e^{-\boldsymbol{\theta}} \boldsymbol{\sigma}_1^2 & \boldsymbol{\sigma}_2^2 & & \\ \cdot & & & \cdot & \\ e^{-(\bar{t}-t)\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & & & & \boldsymbol{\sigma}^2_{\bar{t}-t} \end{pmatrix}. \tag{29}$$

As in the example (2), the zero matrix $\boldsymbol{\sigma}_0^2$ reflects the fact that at time t the risk drivers have already realized $\mathbf{X}_t = \boldsymbol{x}_t$, and therefore the distribution of $\mathbf{X}_t$ conditioned on information at time t, is normal with variance that tends to zero. This is a compact notation for degenerate deterministic variables, which we treat as normal with a small variance that we set to zero at the end of the process.

The multivariate Ornstein-Uhlenbeck process generalizes the toy example (2) in several directions: it includes the Brownian motion as a special case, it is suitable to model diffusion (random walk), mean-reversion and cointegration. Furthermore, through suitable transformations of the risk drivers, it models a variety of additional dynamics, such as for instance the multivariate geometric Brownian motion.

3.2 Embedding the views ($\tilde{2}$: Estimation - $\tilde{3}$: Projection)

Now we need a mechanism to embed the views in the allocation process, as prescribed by the general theoretical framework (21).

In the toy example (20) we used a heuristic to embed the views in the market process. In this section, we apply the Entropy Pooling approach in [Meucci, 2008] to stochastic processes, viewed as (large dimensional) random variables.

We recall that relative entropy minimization is widely applied in physics and statistics ([Cover and Thomas, 2006]), that it generalizes Bayesian updating ([Caticha and Giffin, 2006]), and that it has already found other applications to economics and finance ([Avellaneda, 1999], [Hansen and Sargent, 2007], [Breuer and Csiszar, 2013], [Glasserman and Xu, 2014]).

More precisely, starting from a prior distribution $f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$ for the process of the risk drivers $\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t$ conditioned on currently available information, the posterior distribution is a new distribution for the whole process $\bar{f}_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t} \neq f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$ that should be as close as possible to the prior model $f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$, and at the same time satisfy the views, without adding any unnecessary additional structure.

Let us define the (pseudo) distance $\mathcal{E}$ between a generic distribution $g_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$ for the process of the risk drivers and the prior distribution $f_{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}|\boldsymbol{i}_t}$ in terms of the relative entropy

$$\mathcal{E}(g, f) \equiv \int g(\boldsymbol{x}_t, \ldots, \boldsymbol{x}_{\bar{t}}) \ln \frac{g(\boldsymbol{x}_t, \ldots, \boldsymbol{x}_{\bar{t}})}{f(\boldsymbol{x}_t, \ldots, \boldsymbol{x}_{\bar{t}})} d\boldsymbol{x}_t \cdots \boldsymbol{x}_{\bar{t}}, \tag{30}$$

where for ease of notation we wrote the pdf's as f instead of $f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}$. Then the posterior distribution is defined as the closest distribution to the prior $f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t}$ that yet satisfies the views $\mathcal{V}_t$

$$\bar{f}_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t} \equiv \underset{g\in\mathcal{V}_t}{\operatorname{argmin}}\, \mathcal{E}\{(g, f_{\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t})\}. \tag{31}$$

In Dynamic Entropy Pooling, for the prior, we use the multivariate Ornstein-Uhlenbeck process (24).

For the views, we consider statements on expectations and covariances of arbitrary linear combinations of the process at arbitrary times

$$\mathcal{V}_t : \begin{cases} \mathbb{E}^g_t\{\boldsymbol{v}_{\mu,t}\boldsymbol{X}_{t\rightsquigarrow\bar{t}}\} \equiv \boldsymbol{\mu}_{view;t} \\ \mathbb{C}v^g_t\{\boldsymbol{v}_{\sigma,t}\boldsymbol{X}_{t\rightsquigarrow\bar{t}}\} \equiv \boldsymbol{\sigma}^2_{view;t}. \end{cases}, \tag{32}$$

where $\mathbb{C}v^g_t\{\cdot\}$ is the conditional covariance and $\mathbb{E}^g_t\{\cdot\}$ is the conditional expectation with respect the yet-to-be defined distribution g.

In this expression, the "view" matrices $\boldsymbol{v}_{\mu,t}$ and $\boldsymbol{v}_{\sigma,t}$ define arbitrary linear combinations of the process at the monitoring times for the views, and the conformable vector $\boldsymbol{\mu}_{view;t}$ and square, positive definite matrix $\boldsymbol{\sigma}^2_{view;t}$ define the extent of the views/stress tests as stated at time t. If $\bar{n}_\mu$ is the number of views on expectation and $\bar{n}_\sigma$ is the number of views on covariance, then $\boldsymbol{v}_{\mu,t}$ is a $\bar{n}_\mu \times \bar{n}(\bar{t}-t+1)$ matrix, $\boldsymbol{\mu}_{view;t}$ is a $\bar{n}_\mu \times 1$ vector, $\boldsymbol{v}_{\sigma,t}$ is a $\bar{n}_\sigma \times \bar{n}(\bar{t}-t+1)$ matrix, and $\boldsymbol{\sigma}^2_{view;t}$ is a $\bar{n}_\sigma \times \bar{n}_\sigma$ symmetric and positive definite matrix. For instance, a single view that the spread between second-period first entry and next-period second entry is 1%, i.e. $\mathbb{E}^g_t\{X_{1,t+2}\} - \mathbb{E}^g_t\{X_{2,t+1}\} = 0.01$ corresponds to a one-row $\boldsymbol{v}_{\mu,t} = (0,\ldots 0; 0, -1, 0, \ldots, 0; 1, 0, \ldots, 0; 0, \ldots)$ and to a view scalar $\mu_{view;t}$ =0.01.

Note that the views can be, and in general are, updated as time goes by. Therefore all the matrices defining the views are in general time dependent. Also, the views expire as time passes. Then, the number of views $\bar{n}_\mu$ and $\bar{n}_\sigma$ are time dependent but we did not indicate such a dependence explicitly for ease of notation.

The full-confidence posterior (31) of Dynamic Entropy Pooling then follows as in [Meucci, 2008]:

$$\boldsymbol{X}_{t\rightsquigarrow\bar{t}}|\boldsymbol{i}_t \sim N(\bar{\boldsymbol{\mu}}_{t\rightsquigarrow\bar{t}}, \bar{\boldsymbol{\sigma}}^2_{t\rightsquigarrow\bar{t}}), \tag{33}$$

where we now proceed to re-write the analytical expressions for $\bar{\boldsymbol{\mu}}_{t\rightsquigarrow\bar{t}}$ and $\bar{\boldsymbol{\sigma}}^2_{t\rightsquigarrow\bar{t}}$ in [Meucci, 2008] in terms of complementary projectors.

More precisely, let us define the $\bar{n}(\bar{t}-t+1) \times \bar{n}_\mu$ pseudo-inverse matrix of $\boldsymbol{v}_{\mu,t}$[2]:

$$\boldsymbol{v}^+_{\mu,t} \equiv \boldsymbol{\sigma}^2_{t\rightsquigarrow\bar{t}}\boldsymbol{v}'_{\mu,t}(\boldsymbol{v}_{\mu,t}\boldsymbol{\sigma}^2_{t\rightsquigarrow\bar{t}}\boldsymbol{v}'_{\mu,t})^{-1}, \tag{34}$$

and the two complementary projectors

$$\mathbb{P}_{\mu,t} \equiv (\mathbb{I}_{\bar{n}(\bar{t}-t+1)} - \boldsymbol{v}^+_{\mu,t}\boldsymbol{v}_{\mu,t}), \tag{35}$$

$$\mathbb{P}^\perp_{\mu,t} \equiv \boldsymbol{v}^+_{\mu,t}\boldsymbol{v}_{\mu,t}, \tag{36}$$

then the Dynamic Entropy Pooling posterior expected path in (33) reads:

$$\bar{\boldsymbol{\mu}}_{t\rightsquigarrow\bar{t}} \equiv \mathbb{P}_{\mu,t}\boldsymbol{\mu}_{t\rightsquigarrow\bar{t}} + \mathbb{P}^\perp_{\mu,t}(\boldsymbol{v}^+_{\mu,t}\boldsymbol{\mu}_{view;t}). \tag{37}$$

[2] This is a pseudo-inverse as $\boldsymbol{v}_{\mu,t}\boldsymbol{v}^+_{\mu,t} = \mathbb{I}_{\bar{n}_\mu}$, but $\boldsymbol{v}^+_{\mu,t}\boldsymbol{v}_{\mu,t} \neq \mathbb{I}_{\bar{n}(\bar{t}-t+1)}$.

The posterior expectation (37) is the sum of i) the projected prior expectation $\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}}$ defined in (27); ii) the complementary projection of $\boldsymbol{v}^{+}_{\mu,t}\boldsymbol{\mu}_{view;t}$, where the matrix $\boldsymbol{v}^{+}_{\mu,t}$ percolates the effect of the vector of views $\boldsymbol{\mu}_{view;t}$ across all the $\bar{n}(\bar{t}-t+1)$ entries of the path expectations $\mathbb{E}_t\{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}\}$.

It is easy to see that the posterior expected path (37) satisfies the view (32), as $\boldsymbol{v}_{\mu,t}\mathbb{P}_{\mu,t} = \boldsymbol{0}$ and $\boldsymbol{v}_{\mu,t}\mathbb{P}^{\perp}_{\mu,t} = \boldsymbol{v}_{\mu,t}$.

Also the Dynamic Entropy Pooling posterior covariance in (33) can be written in terms of complementary projectors

$$\mathbb{P}_{\sigma,t} \equiv \mathbb{I}_{\bar{n}(\bar{t}-t+1)} - \boldsymbol{v}^{+}_{\sigma,t}\boldsymbol{v}_{\sigma,t} \ , \tag{38}$$

$$\mathbb{P}^{\perp}_{\sigma,t} \equiv \boldsymbol{v}^{+}_{\sigma,t}\boldsymbol{v}_{\sigma,t} \ , \tag{39}$$

and reads

$$\bar{\boldsymbol{\sigma}}^{2}_{t \rightsquigarrow \bar{t}} \equiv \mathbb{P}_{\sigma,t}\boldsymbol{\sigma}^{2}_{t \rightsquigarrow \bar{t}}\mathbb{P}'_{\sigma,t} + \mathbb{P}^{\perp}_{\sigma,t}(\boldsymbol{v}^{+}_{\sigma,t}\boldsymbol{\sigma}^{2}_{view;t}(\boldsymbol{v}^{+}_{\sigma,t})')(\mathbb{P}^{\perp}_{\sigma,t})', \tag{40}$$

where $\boldsymbol{v}^{+}_{\sigma,t}$ is the $\bar{n}(\bar{t}-t+1) \times \bar{n}_{\sigma}$ pseudo-inverse matrix of $\boldsymbol{v}_{\sigma,t}$

$$\boldsymbol{v}^{+}_{\sigma,t} \equiv \boldsymbol{\sigma}^{2}_{t \rightsquigarrow \bar{t}}\boldsymbol{v}'_{\sigma,t}(\boldsymbol{v}_{\sigma,t}\boldsymbol{\sigma}^{2}_{t \rightsquigarrow \bar{t}}\boldsymbol{v}'_{\sigma,t})^{-1}. \tag{41}$$

The matrix $\boldsymbol{v}^{+}_{\sigma,t}$ percolates the $\bar{n}_{\sigma} \times \bar{n}_{\sigma}$ matrix of views $\boldsymbol{\sigma}^{2}_{view;t}$ across all the $\bar{n}(\bar{t}-t+1) \times \bar{n}(\bar{t}-t+1)$ entries of the path covariance $\mathbb{C}v_t\{\boldsymbol{X}_{t \rightsquigarrow \bar{t}}\}$. As in the case of the posterior expected values, the properties of the projectors $\mathbb{P}_{\sigma,t}$ and $\mathbb{P}^{\perp}_{\sigma,t}$ ensure that the posterior covariance matrix defined in Equation (40) satisfies the view (32).

We can add one last step: the posterior (33) follows by assuming full confidence in the views. If the confidence is less than full, we mix the prior (26) and the full-confidence posterior (33). To generalize the simple mixture to a multi-manager context, where confidence can be linked to the track record, and to a multi-confidence framework, see [Meucci, 2008].

Note that we resorted to Dynamic Entropy Pooling, rather than attempting to adapt [Black and Litterman, 1990], because it is not possible to adapt the original BL methodology to the present multi-period context in a straightforward manner. Indeed Black-Litterman sets views on the market *parameters*, which in this case are the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ of the multivariate Ornstein-Uhlenbeck process (24). Hence, first of all, we would need to express statements on complex functions of $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$; second, we would have to set prior distributions on such parameters that are analytically tractable. Neither step appears straightforward to us.

3.3 Exposures and P&L (4: Pricing; 5: Aggregation; 6: Attribution)

Next, we need to model how the P&L depends on a set of control variables, as prescribed by the general theoretical framework (4). To this purpose, we model the P&L generated by a portfolio over a generic time step as the sum of two terms: a "constant-exposure" P&L Π, and a "market impact" term MI that ensues from rebalancing the exposures.

The constant-exposure P&L is assumed linear in the exposures and in the increments of the risk drivers

$$\Pi_{t+1} = \boldsymbol{b}'_t \Delta \boldsymbol{X}_{t+1}. \tag{42}$$

The form given in Equation (42) is quite general and accurately models a variety of practical situations.

For instance, suppose that the n-th position is an equity share, or an index. Then the risk-driver is the log-value $X_{n,t} = \ln V_{n,t}$ and the P&L of a position with $h_{n,t}$ shares is $\Pi_{n,t+1} = h_{n,t}(V_{n,t+1} - V_{n,t})$. Then the portfolio P&L becomes

$$\Pi_{t+1} = \sum_n h_{n,t} V_{n,t} \times (\frac{V_{n,t+1}}{V_{n,t}} - 1) \approx \sum_n b_{n,t} \Delta X_{n,t+1}, \tag{43}$$

where $b_{n,t} = h_{n,t} V_{n,t}$ is the money exposure to the equity, and where we used the approximation $V_{n,t+1}/V_{n,t} - 1 \approx \ln(V_{n,t+1}/V_{n,t})$.

More in general, if the manger invests in equity shares, we can express the P&L in terms of a linear factor model such as market, value, momentum, size in [Carhart, 1997], or more general "alpha" ("style") or "beta" ("risk") factors [Fung and Hsieh, 1997]:

$$\Pi_{t+1} = \sum_k b_{k,t}^{style} \Delta X_{k,t+1}^{style}. \tag{44}$$

Similarly, suppose that the n-th position is a fixed-income instrument, such as a government bond futures. Let us denote by $Y_{k,t}$ the key-rates on the relevant term structure, such as the par yield curve, or the swap curve. Let us denote by $dv01_{n,k,t}$ the dollar-sensitivity per unit notional of the n-th instrument to the k-th key rate at time t. Then at first order

$$\Pi_{n,t+1} \approx -\sum_k dv01_{n,k,t} \Delta Y_{k,t+1}, \tag{45}$$

In order to account for low-rate regimes and high-rate regimes, we can define as risk drivers the shadow rates, i.e. the inverse-call transform of the rates proposed in [Meucci and Loregian, 2013]

$$X_{t,k} \equiv c^{-1}(Y_{t,k}). \tag{46}$$

Then the rates satisfy $\Delta Y_{t+1,k} \approx \varphi_{k,t} \Delta X_{t+1,k}$, where $\varphi_{k,t} \equiv \frac{dc}{dx}|_{x=X_{t,k}}$. Hence the P&L due to a set of government bond futures can be written as in (42) as follows

$$\Pi_{t+1} \approx \sum_k \underbrace{(-\sum_n h_{n,t} dv01_{n,k,t} \varphi_{k,t})}_{b_{k,t}} \Delta X_{k,t+1}. \tag{47}$$

Similarly, consider a portfolio of stock options, then its risk drivers are the log-value of the respective underlyings, as in (43), and the respective implied volatilities. Denoting the value of the underlying by $V_t = e^{X_t}$ and the implied volatility by Σ^{impl}, the P&L of the option reads at first order

$$\Pi_{n,t+1} \approx \delta_{n,t} V_{n,t} \Delta X_{n,t+1} + \upsilon_{n,t} \Delta \Sigma_{n,t+1}^{impl}, \tag{48}$$

where $\delta_{n,t}$ and $\upsilon_{n,t}$ are the delta and vega of the option and $\Sigma_{n,t+1}^{impl}$ is the implied volatility.

Then for a portfolio with $h_{n,t}$ holdings in the n-th option, the P&L is of the form (42)

$$\Pi_{t+1} \approx \sum_n \underbrace{h_{n,t} \delta_{n,t} V_{n,t}}_{b_{n,t}^{\delta}} \Delta X_{n,t+1} + \sum_n \underbrace{h_{n,t} \upsilon_{n,t}}_{b_{n,t}^{\sigma}} \Delta \Sigma_{n,t+1}^{impl}. \tag{49}$$

In reality, in a large portfolio of stock options, the underlyings are summarized by a style/risk factor model as in (44), and the VIX index is used as a proxy for the implied volatilities. Then

$$\Pi_{t+1} \approx \sum_k b_{k,t}^{style} \Delta X_{k,t+1}^{style} + b_t^{VIX} \Delta VIX_{t+1}. \tag{50}$$

It is important to note that the set of $\bar{n}$ risk drivers $\boldsymbol{X}_t$ may include external factors, such as inflation or other macro economic variables, that are not directly investable, but that have an indirect statistical effect on the $\bar{k}$ exposures $\boldsymbol{b}_t$ via the correlations of the investable risk drivers. To this purpose, we assume that the investable risk drivers are the first $\bar{k}$ among the $\bar{n}$ in $\boldsymbol{X}_t$. Then we generalize the linear expression for the P&L (42) to include a $\bar{k} \times \bar{n}$ matrix $\boldsymbol{\omega} \equiv (\mathbb{I}_{\bar{k}}|\mathbf{0}_{\bar{k}\times(\bar{n}-\bar{k})})$, as follows

$$\Pi_{t+1} = \boldsymbol{b}_t'\boldsymbol{\omega}\Delta\boldsymbol{X}_{t+1}. \tag{51}$$

As far as the market impact term is concerned, standard models for the market impact of transactions assume a super-linear relation with the variation of the holdings (see [Almgren and Chriss, 2000], [Almgren et al., 2005] and [Gatheral, 2010]). For computational purposes we use a quadratic dependence, similar to [Gârleanu and Pedersen, 2013] or [Grinold, 2006].

$$MI_t = a^2 + (\boldsymbol{b}_t - \boldsymbol{b}_{t^-})'\boldsymbol{c}^2(\boldsymbol{b}_t - \boldsymbol{b}_{t^-}) = a^2 + \Delta\boldsymbol{b}_t'\boldsymbol{c}^2\Delta\boldsymbol{b}_t, \tag{52}$$

where a^2 is the average cost of maintaining constant exposures; $\boldsymbol{c}^2$ is a symmetric positive definite matrix; and where we used the shift operator $\Delta\boldsymbol{b}_t \equiv \boldsymbol{b}_t - \boldsymbol{b}_{t-1}$ and the observation that, in our discrete-time model, $\boldsymbol{b}_{t^-} = \boldsymbol{b}_{t-1}$. Furthermore, in the practical implementation, we set $\boldsymbol{c}^2$ proportional to the variance of the risk-drivers returns as in [Gârleanu and Pedersen, 2013], although the specification of $\boldsymbol{c}^2$ is completely flexible.

Note that a portfolio with fully general exposure rebalancing $\boldsymbol{b}_t \to \boldsymbol{b}_{t+1}$ appears to violate the self-financing condition. However, self-financing is ensured, as long as we assume access to a pool of cash as part of the assets under management, which does not affect the P&L computation. Such cash is always present in fund management in practice.

3.4 Ex-ante Evaluation (7)

At this point we need an index of satisfaction to evaluate a potential investment, as prescribed by the general theoretical framework (8).

A myopic single-period mean-variance trade-off as in the toy example (9) is not satisfactory because the portfolio manager cares about all the stream of future performance.

Similar to [Gârleanu and Pedersen, 2013], in Dynamic Entropy Pooling we model satisfaction as the discounted expected stream of the constant-exposure P&L, penalized with a quadratic term that accounts for risk aversion and with a further term that accounts for market impact

$$\overline{\mathbb{S}}_t^{(\gamma,\eta)} \equiv \overline{\mathbb{E}}_t\{ \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \left(\overline{\mathbb{E}}_s\{\Pi_{s+1}\} - \frac{\gamma}{2}\overline{\mathbb{V}}_s\{\Pi_{s+1}\} - \frac{\eta}{2}\overline{\mathbb{E}}_s\{MI_s\}\right) \}. \tag{53}$$

In the above expression, λ is the discount rate, γ is the risk aversion penalty parameter and η is the market impact penalty parameter.

The index of satisfaction (53) is the practical objective function of a manager who is compensated on the basis of the performances on each period. A different approach, used mostly in the consumption/investment problem, would be to maximize the expected utility of the final wealth, as in [Karatzas et al., 1987].

To express the satisfaction (53) as a function of the exposures as in the general framework (8), we express the P&L as a linear function of the exposures (51). Leaving aside constant terms, which are irrelevant, we obtain

$$\overline{\mathbb{S}}_t^{(\gamma,\eta)} = \overline{\mathbb{E}}_t\{ \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \left(\boldsymbol{B}_s' \boldsymbol{\omega} \overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} \right. \\ \left. - \frac{\gamma}{2} \boldsymbol{B}_s' \boldsymbol{\omega} \overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\} \boldsymbol{\omega}' \boldsymbol{B}_s - \frac{\eta}{2} \Delta \boldsymbol{B}_s' \boldsymbol{c}^2 \Delta \boldsymbol{B}_s \right) \}. \tag{54}$$

Note that at the current time t the future exposures $\boldsymbol{B}_s$ for any time $s > t$ can be stochastic, if they are decided at time s, which is why we used the upper case notation. Also note that the satisfaction (54) depends on $\boldsymbol{b}_{t-1}$, the current "legacy" portfolio before trading at time t, due to the first market impact term in the sum.

The upper-bar notation in the index of satisfaction (53) and (54) emphasizes that we use the posterior distribution of the market (33), which embeds the discretionary views (32). If there are no views, the market follows the prior equilibrium (multivariate Ornstein-Uhlenbeck) process (26).

3.5 Portfolio construction (8)

To construct the Dynamic Entropy Pooling portfolio, we seek to maximize the index of satisfaction (54) via an optimal policy, as in the general theoretical framework (11)

$$\{\boldsymbol{b}_s^* = p_s^*(\boldsymbol{i}_s)\}_{s\geq t}, \text{ where} \\ \{p_s^*\}_{s\geq t} = \text{argmax}_{\{p_s\}_{s\geq t}\in\mathcal{C}} \overline{\mathbb{E}}_t\{ \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \left(p_s(\boldsymbol{I}_s)' \boldsymbol{\omega}\, \overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} \right. \\ \left. - \frac{\gamma}{2} p_s(\boldsymbol{I}_s)' \boldsymbol{\omega} \overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\} \boldsymbol{\omega}' p_s(\boldsymbol{I}_s) - \frac{\eta}{2} \Delta p_s(\boldsymbol{I}_s)' \boldsymbol{c}^2 \Delta p_s(\boldsymbol{I}_s) \right) \} \tag{55}$$

where $p \in \mathcal{C}$ denotes that the policy can be required to satisfy additional constraints and where $\boldsymbol{I}_s$ is the information set that which will become available at time $s \geq t$ (and thus it consists of random variables, this is why we used the capital letter notation).

To solve the general portfolio construction problem (55), we consider increasingly complex situations, where the more complex includes the simpler as a special case. We purposely show this progression, instead of jumping directly to solving the most complex scenario, because the intermediate steps and the respective techniques we use to solve them, are of practical relevance and allow for a deeper understanding of the problem at hand.

We summarize below the progression of the cases we consider, their definition, how we solve such problems, and in which section

Views	Constr.	Estim	How	Where
✓	×	×	Mean-variance	Section 4
×	×	×	Dyn prog (analyt)	GP (2013)
✓	×	×	Dyn prog (semi-analyt)	Section 5
✓	✓	✓	Calc of variations (QP)	Section 6

(56)

3.6 Further steps (9: Execution; 10: Performance Analysis)

Once we have computed the optimal future path of exposures (55), and in particular the optimal current exposure $\boldsymbol{b}_t^*$, we need to implement this exposure by buying or selling specific instruments, with holdings $\boldsymbol{h}_t^*$.

As discussed in Section 3.3, the exposures are linearly related to the holdings through an equation of the kind

$$\boldsymbol{b}_t = \boldsymbol{d}_t \boldsymbol{h}_t, \tag{57}$$

where $\boldsymbol{d}_t$ is a matrix that depends on the information available at time t. In general, $\boldsymbol{d}_t$ is not full-rank matrix, because the number of assets considered for investment is greater than the number of the risk drivers.

The strategy to achieve the optimal exposures $\boldsymbol{b}_t^*$ at time t is decided by minimizing the transaction costs at that point in time. Then optimal holdings read

$$\boldsymbol{h}_t^* = \underset{\boldsymbol{b}_t^* = \boldsymbol{d}_t \boldsymbol{h}_t}{\text{argmin}}\, TC(\boldsymbol{h}_t). \tag{58}$$

To summarize, at each trading time t, we compute the optimal exposure $\boldsymbol{b}_t^*$, we update $\boldsymbol{d}_t$ on the basis of the current information, and we solve (58) in order to determine the optimal holdings $\boldsymbol{h}_t^*$. Once the allocation has been implemented, we can evaluate the performance ex-post via standard performance attribution techniques.

4 No market impact, no constraints, no estimation risk: mean-variance

In this section we discuss a practical solution for the Dynamic Entropy Pooling framework introduced in Section 3. Such solution is viable for liquid products such as futures. The solution is analytical, and thus it allows us to understand in depth the role played by the views at different horizons in the Dynamic Entropy Pooling framework.

Accordingly, in the general portfolio construction problem (55) we set the market impact penalty to zero, and we do not constrain the problem

$$\eta \equiv 0, \qquad \mathcal{C} \equiv \emptyset. \tag{59}$$

We also assume that the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ of the multivariate Ornstein-Uhlenbeck process (24) have been properly estimated once and for all.

Then, the general Dynamic Entropy Pooling portfolio construction problem (55) becomes

$$\begin{aligned} &\{\boldsymbol{b}_s^* = p_s^*(\boldsymbol{i}_s)\}_{s\geq t}, \text{ where} \\ \{p_s^*\}_{s\geq t} &= \text{argmax}_{\{p_s\}_{s\geq t}}\, \overline{\mathbb{E}}_t\{ \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \left(p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\, \overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} \right. \\ &\quad \left. -\tfrac{\gamma}{2} p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\}\boldsymbol{\omega}' p_s(\boldsymbol{I}_s)\right) \} \end{aligned} \tag{60}$$

The objective function (60) is simply the expected sum of disentangled discounted mean-variance trade-offs

$$\begin{aligned} &\max_{\{p_s\}_{s\geq t}} \overline{\mathbb{E}}_t\{\textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\, \overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} - \tfrac{\gamma}{2} p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\}\boldsymbol{\omega}' p_s(\boldsymbol{I}_s)\} \\ &= \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \overline{\mathbb{E}}_t\{\max_{p_s}[p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\, \overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} - \tfrac{\gamma}{2} p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\}\boldsymbol{\omega}' p_s(\boldsymbol{I}_s)]\}. \end{aligned} \tag{61}$$

Therefore optimizing the discounted sum (60) is equivalent to optimizing for all s's the myopic (= one-period) mean-variance problems

$$\boldsymbol{b}_s^* = p_s^*(\boldsymbol{i}_s) \equiv \underset{\boldsymbol{b}}{\text{argmax}}\{\boldsymbol{b}'\boldsymbol{\omega}\overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}'_{s+1}\} - \tfrac{\gamma}{2}\boldsymbol{b}'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}'_{s+1}\}\boldsymbol{\omega}'\boldsymbol{b}\}, \quad s = t, t+1, \ldots. \tag{62}$$

The solution of the mean-variance problem (62) is readily obtained analytically

$$\boldsymbol{b}_s^* = \tfrac{1}{\gamma}(\boldsymbol{\omega}\bar{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\}\boldsymbol{\omega}')^{-1}\boldsymbol{\omega}\bar{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\}, \quad s = t, t+1, \ldots, \tag{63}$$

where the upper bar notation emphasizes that we are using the Dynamic Entropy Pooling posterior process in the expectations and covariances.

4.1 No views

First, let us look at the optimal exposures when there are no views, i.e. when the market $\boldsymbol{X}$ follows the prior equilibrium (multivariate Ornstein-Uhlenbeck) process (26).

The expected returns according to the prior distribution follow from (27) and read

$$\mathbb{E}_s\{\Delta \boldsymbol{X}_{s+1}\} = \left(\mathbb{I}_{\bar{n}} - e^{-\boldsymbol{\theta}}\right)\left(\boldsymbol{\theta}^{-1}\boldsymbol{\mu} - \boldsymbol{x}_s\right), \quad s = t, t+1, \ldots. \tag{64}$$

The covariance matrix according to the prior distribution follows from (28)-(29) and reads

$$\mathbb{C}v_s\{\Delta \boldsymbol{X}_{s+1}\} = \boldsymbol{\sigma}_1^2, \quad s = t, t+1, \ldots. \tag{65}$$

Then the optimal solution (63) reads

$$\boldsymbol{b}_s^* = \tfrac{1}{\gamma}(\boldsymbol{\omega}\boldsymbol{\sigma}_1^2\boldsymbol{\omega}')^{-1}\boldsymbol{\omega}(\mathbb{I}_{\bar{n}} - e^{-\boldsymbol{\theta}})\left(\boldsymbol{\theta}^{-1}\boldsymbol{\mu} - \boldsymbol{x}_s\right), \quad s = t, t+1, \ldots. \tag{66}$$

The optimal solution (66) is driven by the dislocation between the value $\boldsymbol{x}_s$ of the risk drivers observed at time s and $\boldsymbol{\theta}^{-1}\boldsymbol{\mu}$, which is the vector of the long term expected levels (if the risk drivers are mean reverting).

4.2 With views

Here we switch on the views (32), i.e. we assume that the market $\boldsymbol{X}$ follows the posterior distribution (33). Then to compute the optimal policy $\{\boldsymbol{b}_s^* = p_s^*(\boldsymbol{i}_s)\}_{s\geq t}$ in (63) we need the posterior expectations and the posterior covariances.

Let us denote by $\bar{t}$ the last time involving any of the views. The Dynamic Entropy Pooling posterior expectations for the next-step increment at the generic time s before $\bar{t}$ read

$$\begin{aligned}\overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} \equiv (\mathbb{P}_{\mu,s})_{s+1,\cdot}\, & \underbrace{\begin{pmatrix} \mathbb{I}_{\bar{n}} - e^{-0\boldsymbol{\theta}} \\ \mathbb{I}_{\bar{n}} - e^{-1\boldsymbol{\theta}} \\ \vdots \\ \mathbb{I}_{\bar{n}} - e^{-(\bar{t}-s)\boldsymbol{\theta}} \end{pmatrix}(\boldsymbol{\theta}^{-1}\boldsymbol{\mu} - \boldsymbol{x}_s)}_{\Delta\boldsymbol{\mu}_{s\rightsquigarrow\bar{t}}^{LongTerm}} \\ + (\mathbb{P}_{\mu,s}^{\perp})_{s+1,\cdot}\, & \underbrace{\left(\boldsymbol{v}_{\mu,s}^{+}\boldsymbol{\mu}_{view;s} - \begin{pmatrix}\boldsymbol{x}_s \\ \vdots \\ \boldsymbol{x}_s\end{pmatrix}\right)}_{\Delta\boldsymbol{\mu}_{s\rightsquigarrow\bar{t}}^{ViewMean}} \quad s \leq \bar{t}\end{aligned} \tag{67}$$

where $(\mathbb{P}_{\mu,s})_{s+1,\cdot}$ and $(\mathbb{P}_{\mu,s}^{\perp})_{s+1,\cdot}$ are the sub-matrices corresponding to the rows relative to time $s+1$ in the projector matrices $\mathbb{P}_{\mu,s}$ and $\mathbb{P}_{\mu,s}^{\perp}$, defined in (35)-(36).

The posterior expected return is the sum of the projection of a term that is proportional to the difference of the value $\boldsymbol{x}_s$ of the risk drivers observed at time s from the long term expected levels $\boldsymbol{\theta}^{-1}\boldsymbol{\mu}$, and of a term that is the complementary projection of the difference between the levels of the view $\boldsymbol{\mu}_{view;s}$, "percolated" all over the times and dimensions trough the matrix $\boldsymbol{v}^{+}_{\mu,s}$, and the value $\boldsymbol{x}_s$ of the risk drivers observed at time s.

As for the Dynamic Entropy Pooling posterior covariance matrix, rearranging terms from (40) we obtain for the generic time s before $\bar{t}$

$$\overline{\mathbb{C}}v_s\{\Delta\boldsymbol{X}_{s+1}\} \equiv \bar{\boldsymbol{\sigma}}^2_s = \boldsymbol{\sigma}^2_1 + [\boldsymbol{v}^{+}_{\sigma,s}(\boldsymbol{\sigma}^2_{view;s} - \boldsymbol{v}_{\sigma,s}\boldsymbol{\sigma}^2_{s\rightsquigarrow\bar{t}}\boldsymbol{v}'_{\sigma,s})(\boldsymbol{v}^{+}_{\sigma,s})']_{s+1,s+1}, \quad s \leq \bar{t}, \tag{68}$$

where $[\cdot]_{s+1,s+1}$ is the $(s+1, s+1)$ block sub-matrix.

Substituting the expectations (67) and the covariance matrix (68) in the optimal exposures (63), we obtain the Dynamic Entropy Pooling optimal solution

$$\begin{aligned} \boldsymbol{b}^*_s = & \underbrace{\tfrac{1}{\gamma}(\boldsymbol{\omega}\bar{\boldsymbol{\sigma}}^2_s\boldsymbol{\omega}')^{-1}\boldsymbol{\omega}(\mathbb{P}_{\mu,s})_{s+1,\cdot}\Delta\boldsymbol{\mu}^{LongTerm}_{s\rightsquigarrow\bar{t}}}_{\boldsymbol{b}^{LongTerm}_s} \\ & + \underbrace{\tfrac{1}{\gamma}(\boldsymbol{\omega}\bar{\boldsymbol{\sigma}}^2_s\boldsymbol{\omega}')^{-1}\boldsymbol{\omega}(\mathbb{P}^{\perp}_{\mu,s})_{s+1,\cdot}\Delta\boldsymbol{\mu}^{ViewMean}_{s\rightsquigarrow\bar{t}}}_{\boldsymbol{b}^{ViewMean}_s} \quad s \leq \bar{t}. \end{aligned} \tag{69}$$

Notice that when there are no views, $\boldsymbol{v}_{\mu,s} = \boldsymbol{v}_{\sigma,s} = \emptyset$, and $\mathbb{P}_{\mu,s} = \mathbb{I}_{\bar{n}(\bar{t}-s+1)}$, $\mathbb{P}^{\perp}_{\mu,s} = \boldsymbol{0}$, $\bar{\boldsymbol{\sigma}}^2_s = \boldsymbol{\sigma}^2_1$, and thus the optimal exposures based on the posterior (69) become the optimal exposures based on the prior (66). Hence, at the generic time $s > \bar{t}$ the optimal policy is the prior policy (66).

5 Market impact, no constraints, no estimation risk: dynamic programming

In this section we add market impact to the Dynamic Entropy Pooling framework discussed in Section 4, thereby providing a viable practical solution for illiquid markets. As it turns out, the solution is an analytical recursion.

Accordingly, in the general portfolio construction problem (55) we leave the market impact term, though we do not constrain the problem

$$\eta \neq 0, \qquad \mathcal{C} \equiv \emptyset. \tag{70}$$

Furthermore, as in the previous Section 4, we assume that the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ of the multivariate Ornstein-Uhlenbeck process (24) have been properly estimated once and for all.

Then, the general Dynamic Entropy Pooling portfolio construction problem (55) reads

$$\begin{aligned} & \{\boldsymbol{b}^*_s = p^*_s(\boldsymbol{i}_s)\}_{s\geq t}, \text{ where} \\ \{p^*_s\}_{s\geq t} = & \operatorname{argmax}_{\{p_s\}_{s\geq t}} \overline{\mathbb{E}}_t\{ \textstyle\sum_{s=t}^{\infty} e^{-\lambda(s-t)} \big(p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\,\overline{\mathbb{E}}_s\{\Delta\boldsymbol{X}_{s+1}\} \\ & -\tfrac{\gamma}{2}p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta\boldsymbol{X}_{s+1}\}\boldsymbol{\omega}'p_s(\boldsymbol{I}_s) - \tfrac{\eta}{2}\Delta p_s(\boldsymbol{I}_s)'\boldsymbol{c}^2\Delta p_s(\boldsymbol{I}_s)\big)\,\}. \end{aligned} \tag{71}$$

We solve the problem (71) by dynamic programming. The value function for the generic period $s \geq t$ is a function $v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s)$ of the legacy exposures $\boldsymbol{b}_{s-1}$ and the value $\boldsymbol{x}_s$ of the risk drivers observed at time s. The value function satisfies the Bellman equation

$$\begin{aligned} v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s) = & \max_{\boldsymbol{b}}\{\boldsymbol{b}'\boldsymbol{\omega}\overline{\mathbb{E}}_s\{\Delta\boldsymbol{X}_{s+1}\} - \tfrac{\gamma}{2}\boldsymbol{b}'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta\boldsymbol{X}_{s+1}\}\boldsymbol{\omega}'\boldsymbol{b} \\ & - \tfrac{\eta}{2}(\boldsymbol{b}-\boldsymbol{b}_{s-1})'\boldsymbol{c}^2(\boldsymbol{b}-\boldsymbol{b}_{s-1}) + e^{-\lambda}\overline{\mathbb{E}}_s\{v_{s+1}(\boldsymbol{b}, \boldsymbol{X}_{s+1})\}\}, \quad s = t, t+1, \ldots \end{aligned} \tag{72}$$

where the current information set $\boldsymbol{i}_s \supset \{\boldsymbol{b}_{s-1}, \boldsymbol{x}_s\}$ consists of deterministic variables at time s, and thus they are in lower-case notation.

The conditional one-period covariance is a deterministic function $\overline{\mathbb{C}}v_s\{\Delta \boldsymbol{X}_{s+1}\} = \bar{\boldsymbol{\sigma}}^2_s$, see (68). The conditional one-period expected returns are affine in the risk drivers

$$\overline{\mathbb{E}}_s\{\Delta \boldsymbol{X}_{s+1}\} = \boldsymbol{\alpha}_s + \boldsymbol{\beta}_s \boldsymbol{x}_s, \tag{73}$$

with coefficients that follow from (67)

$$\boldsymbol{\alpha}_s \equiv (\mathbb{P}_{\mu,s})_{s+1,\cdot} \begin{pmatrix} \mathbb{I}_{\bar{n}} - e^{-0\boldsymbol{\theta}} \\ \mathbb{I}_{\bar{n}} - e^{-1\boldsymbol{\theta}} \\ \vdots \\ \mathbb{I}_{\bar{n}} - e^{-(\bar{t}-s)\boldsymbol{\theta}} \end{pmatrix} \boldsymbol{\theta}^{-1}\boldsymbol{\mu} + (\mathbb{P}^{\perp}_{\mu,s})_{s+1,\cdot}\boldsymbol{v}^{+}_{\mu,s}\boldsymbol{\mu}_{view;s} \tag{74}$$

$$\boldsymbol{\beta}_s \equiv (\mathbb{P}_{\mu,s})_{s+1,\cdot} \begin{pmatrix} e^{-0\boldsymbol{\theta}} \\ e^{-1\boldsymbol{\theta}} \\ \vdots \\ e^{-(\bar{t}-s)\boldsymbol{\theta}} \end{pmatrix} - \mathbb{I}_{\bar{n}}. \tag{75}$$

Substituting the conditional expectations (73) and covariance $\bar{\boldsymbol{\sigma}}^2_s$ in the Bellman equation (72) we obtain

$$\begin{aligned} v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s) = \max_{\boldsymbol{b}}\{&\boldsymbol{b}'\boldsymbol{\omega}(\boldsymbol{\alpha}_s + \boldsymbol{\beta}_s\boldsymbol{x}_s) - \tfrac{\gamma}{2}\boldsymbol{b}'\boldsymbol{\omega}\bar{\boldsymbol{\sigma}}^2_s\boldsymbol{\omega}'\boldsymbol{b} - \tfrac{\eta}{2}(\boldsymbol{b} - \boldsymbol{b}_{s-1})'\boldsymbol{c}^2(\boldsymbol{b} - \boldsymbol{b}_{s-1}) \\ &+ e^{-\lambda}\overline{\mathbb{E}}_s\{v_{s+1}(\boldsymbol{b}, \boldsymbol{X}_{s+1})\}\}. \end{aligned} \tag{76}$$

We test a value function quadratic in $\boldsymbol{b}_s$ and $\boldsymbol{x}_{s+1}$ with time-depending coefficients $\boldsymbol{\psi}_s \equiv \{\boldsymbol{\psi}_{bb,s}, \boldsymbol{\psi}_{bx,s}, \boldsymbol{\psi}_{xx,s}, \boldsymbol{\psi}_{b,s}, \boldsymbol{\psi}_{x,s}, \psi_{0,s}\}$, as follows

$$v_{s+1}(\boldsymbol{b}_s, \boldsymbol{x}_{s+1}) = -\tfrac{1}{2}\boldsymbol{b}'_s\boldsymbol{\psi}_{bb,s}\boldsymbol{b}_s + \boldsymbol{b}'_s\boldsymbol{\psi}_{bx,s}\boldsymbol{x}_{s+1} + \tfrac{1}{2}\boldsymbol{x}'_{s+1}\boldsymbol{\psi}_{xx,s}\boldsymbol{x}_{s+1} + \boldsymbol{\psi}'_{b,s}\boldsymbol{b}_s + \boldsymbol{\psi}'_{x,s}\boldsymbol{x}_{s+1} + \psi_{0,s}. \tag{77}$$

To determine the coefficients $\boldsymbol{\psi}_s$, we substitute the quadratic ansatz (77) in the Bellman equation (76) and we impose that it is satisfied for each $\boldsymbol{x}_s$ and $\boldsymbol{b}_{s-1}$. The result is a recursive definition

$$\boldsymbol{\psi}_{s-1} = g_s(\boldsymbol{\psi}_s), \tag{78}$$

where the explicit expression for the recursion function g_s is provided in (117)-(122) in the appendix.

Then, as we show in Appendix A.2, the optimal Dynamic Entropy Pooling policy is determined by the coefficients $\boldsymbol{\psi}_s$ of the value function

$$\begin{aligned} \boldsymbol{b}^*_s = (\gamma\boldsymbol{\omega}\bar{\boldsymbol{\sigma}}^2_s\boldsymbol{\omega}' + \eta\boldsymbol{c}^2 + e^{-\lambda}\boldsymbol{\psi}_{bb,s})^{-1}[\;&\underbrace{\eta\boldsymbol{c}^2\boldsymbol{b}_{s-1}}_{\text{legacy exposures}} \\ &+ \underbrace{(\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}}))\boldsymbol{x}_s}_{\text{current risk drivers}} + \underbrace{(\boldsymbol{\omega} + e^{-\lambda}\boldsymbol{\psi}_{bx,s})\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s}}_{(\star)\text{ future views}}], \quad s = t, t+1, \ldots \end{aligned} \tag{79}$$

The optimal solution (79) at time s has three additive contributions. The first contribution is linear in the legacy exposure $\boldsymbol{b}_{s-1}$. The second contribution is linear in the risk drivers' current values $\boldsymbol{x}_s$. The third contribution is linear in the future views. More precisely, applying the recursion (78) to $\boldsymbol{\psi}_{b,s}$ ((120) in the appendix), the last contribution in the exposures (79) can be written as

$$(\star) = \underbrace{\textstyle\sum_{t=s}^{\bar{t}} e^{-(t-s)\lambda}\boldsymbol{d}_{s,t}(\boldsymbol{\omega} + e^{-\lambda}\boldsymbol{\psi}_{bx,t})\boldsymbol{\alpha}_t}_{\text{future views } \boldsymbol{\mu}_{view;t}\ (74)} + \underbrace{e^{-(\bar{t}-s+1)\lambda}\boldsymbol{d}_{s,\bar{t}}\boldsymbol{\psi}_{b,\bar{t}}}_{\text{no views}}, \tag{80}$$

where $\bar{t}$ is far enough in the future that it is past the last view; and where $\boldsymbol{d}_{s,t} \equiv \prod_{u=s+1}^{t}(\eta \boldsymbol{c}^2 \boldsymbol{q}_u^{-1})$ for $t > s$ and $\boldsymbol{d}_{s,s} \equiv \mathbb{I}_{\bar{k}}$.

The first term in the sum (80) is proportional to $\boldsymbol{\alpha}_s$, that is, from (74), it is proportional to the extent of the views $\boldsymbol{\mu}_{view;s}$ at the current time s. More in general, the generic term for $t \geq s$ in the sum (80) is proportional to $\boldsymbol{\alpha}_t$, that is, it is proportional to the views $\boldsymbol{\mu}_{view;t}$ that will be still alive at the future time $t \geq s$. However the influence of the future views is dampened by an exponential decaying factor.

To summarize, starting from a suitable initialization $\boldsymbol{\psi}_{\bar{t}}$, we compute the coefficients $\boldsymbol{\psi}_s$ of the value function (77) and thus the optimal policy (79) at all times $s \leq \bar{t}$. We discuss below how to properly initialize $\boldsymbol{\psi}_{\bar{t}}$.

It is simple to check that in the case of no market impact, replacing $\eta = 0$ in the expressions (117)-(122) of the coefficients $\boldsymbol{\psi}_s$, and then in the optimal exposure (79), we obtain, as expected, the optimal mean-variance allocation (69).

5.1 No views

First, let us look at the optimal exposures when there are no views, i.e. when the market $\boldsymbol{X}$ follows the prior equilibrium (multivariate Ornstein-Uhlenbeck) process (26).

When the market follows the prior process the recursion function (78) is time independent, or $g_s = g$, because all the coefficients that define g_s are time independent $\boldsymbol{\alpha}_s = \boldsymbol{\alpha}$, $\boldsymbol{\beta}_s = \boldsymbol{\beta}$ from (74) and (75), and $\bar{\boldsymbol{\sigma}}_s^2 = \boldsymbol{\sigma}_1^2$ from (68). Hence, the value function (77) is also time-independent. Therefore we can drop the time subscript from the value function coefficients $\boldsymbol{\psi}_s = \boldsymbol{\psi}$. The recursion (78) then becomes an implicit equation that *defines* the coefficients $\boldsymbol{\psi}$

$$\boldsymbol{\psi} \equiv g(\boldsymbol{\psi}). \tag{81}$$

We show in Appendix A.2 the explicit analytical expression of the values $\boldsymbol{\psi} \equiv \{\boldsymbol{\psi}_{bb}, \boldsymbol{\psi}_{bx}, \boldsymbol{\psi}_{xx}, \boldsymbol{\psi}_{b}, \boldsymbol{\psi}_{x}, \psi_0\}$ that solve equation (81).

Then the optimal solution (79) becomes the same as [Gârleanu and Pedersen, 2013]. Indeed, they also consider the infinite horizon objective (71) and a VAR(1) dynamics for the markets.

5.2 Views, market impact

Here we switch on the views (32), i.e. we assume that the market $\boldsymbol{X}$ follows the posterior distribution (33). Then to compute the optimal exposure $\boldsymbol{b}_s^*$ in (63) we need the posterior expectations and the posterior covariances.

When the views are switched on, we must initialize $\boldsymbol{\psi}_{\bar{t}}$ at a future time $\bar{t}$. Accordingly, we choose $\bar{t}$ as the last horizon that involves any view. After $\bar{t}$, there are no more active views, and thus the optimal exposures must be the same as in the prior case discussed in Section 5.1. Hence at time $\bar{t}$ we set the coefficients $\boldsymbol{\psi}_{\bar{t}}$ to match the prior, time independent coefficients (81), or

$$\boldsymbol{\psi}_{\bar{t}} \equiv \boldsymbol{\psi}. \tag{82}$$

With the boundary condition (82) we can now compute analytically all the coefficients $\boldsymbol{\psi}_s$ of the value function (77) for $s \leq \bar{t}$ using the recursion (78). Finally, with the coefficients $\boldsymbol{\psi}_s$, we compute the optimal policy (79).

6 Market impact, constraints, estimation risk: calculus of variations

In this section we consider the most general Dynamic Entropy Pooling scenario, adding constraints to the framework in Section 5, and accounting for the issue of estimation risk. As it turns out, the solution is numerical, and yet very efficient. Indeed, the optimal exposures follow from solving a sequence of quadratic programs.

First, in view of our numerical approach, we modify the general portfolio infinite-horizon construction problem (55) to include a large, fixed, yet finite, number $\bar{\tau}$ of relevant future steps at any point in time

$$\begin{aligned} &\{\boldsymbol{b}^*_s = p^*_s(\boldsymbol{i}_s)\}_{s\geq t}, \text{ where} \\ \{p^*_s\}_{s\geq t} = &\operatorname{argmax}_{\{p_s\}_{s\geq t}\in\mathcal{C}}\, \overline{\mathbb{E}}_t\{\textstyle\sum_{s=t}^{t+\bar{\tau}} e^{-\lambda(s-t)}\left(p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{E}}_s\{\Delta\boldsymbol{X}_{s+1}\}\right. \\ &\left.-\tfrac{\gamma}{2}p_s(\boldsymbol{I}_s)'\boldsymbol{\omega}\overline{\mathbb{C}}v_s\{\Delta\boldsymbol{X}_{s+1}\}\boldsymbol{\omega}'p_s(\boldsymbol{I}_s) - \tfrac{\eta}{2}\Delta p_s(\boldsymbol{I}_s)'\boldsymbol{c}^2\Delta p_s(\boldsymbol{I}_s)\right)\}. \end{aligned} \tag{83}$$

In practice the number of future time steps $\bar{\tau}$ is such that the discount satisfies $e^{-\lambda\bar{\tau}} \approx 10^{-2}$.

Second, we search for a "deterministic" policy, i.e. a set of decisions $\{\boldsymbol{b}_{t,t}, \boldsymbol{b}_{t,t+1}, \ldots, \boldsymbol{b}_{t,t+t+\bar{\tau}}\}$ that disregards any information. In the above notation we emphasize that the policy is computed at time t.

Third, we include linear equality or inequality constraints. Let us stack the path of exposure into one vector as in (3):

$$\boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} \equiv \begin{pmatrix} \boldsymbol{b}_{t,t} \\ \boldsymbol{b}_{t,t+1} \\ \cdot \\ \boldsymbol{b}_{t,t+\bar{\tau}} \end{pmatrix}. \tag{84}$$

Then, we formulate linear constraints as follows

$$\boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} \in \mathcal{C}_t \Leftrightarrow \begin{cases} \boldsymbol{m}_t \boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} = \boldsymbol{u}_t \\ \tilde{\boldsymbol{m}}_t \boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} \leq \tilde{\boldsymbol{u}}_t \end{cases}, \tag{85}$$

where $\boldsymbol{m}_t$ and $\tilde{\boldsymbol{m}}_t$ are conformable matrices while $\boldsymbol{u}_t$ and $\tilde{\boldsymbol{u}}_t$ are conformable vectors.

By including linear constraints, we can now assume exposures $\boldsymbol{b}_{t,s}$ to all the risk drivers $\boldsymbol{X}_s$, regardless weather the drivers are investable. Indeed if a driver is not investable, we simply constrain the respective exposure to zero. Hence, we can simplify the problem, dropping "$\boldsymbol{\omega}$" from the objective (83).

Fourth and last, we assume that the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ of the multivariate Ornstein-Uhlenbeck process (24) are constantly re-estimated, and similarly for the market impact parameters $\boldsymbol{c}$, and thus we denote them by $(\hat{\boldsymbol{\theta}}_t, \hat{\boldsymbol{\mu}}_t, \hat{\boldsymbol{\sigma}}_t, \hat{\boldsymbol{c}}_t)$.

With the above changes, the Dynamic Entropy Pooling portfolio construction problem (83) at the generic time t becomes

$$\begin{aligned} \boldsymbol{b}^*_{t,t\rightsquigarrow t+\bar{\tau}} = &\operatorname{argmax}_{\boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}}\in\mathcal{C}_t}\, \overline{\mathbb{E}}_t\{\textstyle\sum_{s=t}^{t+\bar{\tau}} e^{-\lambda(s-t)}(\boldsymbol{b}'_{t,s}\overline{\mathbb{E}}_s\{\Delta\boldsymbol{X}_{s+1}|\hat{\boldsymbol{\theta}}_t, \hat{\boldsymbol{\mu}}_t, \hat{\boldsymbol{\sigma}}_t\} \\ &- \tfrac{\gamma}{2}\boldsymbol{b}'_{t,s}\overline{\mathbb{C}}v_s\{\Delta\boldsymbol{X}_{s+1}|\hat{\boldsymbol{\theta}}_t, \hat{\boldsymbol{\mu}}_t, \hat{\boldsymbol{\sigma}}_t\}\boldsymbol{b}_{t,s} - \tfrac{\eta}{2}\Delta\boldsymbol{b}'_{t,s}\hat{\boldsymbol{c}}^2_t\Delta\boldsymbol{b}_{t,s})\}. \end{aligned} \tag{86}$$

The optimization problem (86) is an instance of quadratic programming

$$\boldsymbol{b}^*_{t,t\rightsquigarrow t+\bar{\tau}} = \underset{\boldsymbol{b}_{t\rightsquigarrow t+\bar{\tau}}}{\operatorname{argmin}}\{\boldsymbol{b}'_{t,t\rightsquigarrow t+\bar{\tau}}\boldsymbol{q}_t\boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} - \boldsymbol{b}'_{t,t\rightsquigarrow t+\bar{\tau}}\boldsymbol{l}_t\} \tag{87}$$

$$\text{such that } \begin{cases} \boldsymbol{m}_t \boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} = \boldsymbol{u}_t \\ \tilde{\boldsymbol{m}}_t \boldsymbol{b}_{t,t\rightsquigarrow t+\bar{\tau}} \leq \tilde{\boldsymbol{u}}_t \end{cases}$$

for a suitable positive definite matrix $\boldsymbol{q}_t$ and vector $\boldsymbol{l}_t$.

In full generality, we can solve for the optimal path (87) numerically, using off-the-shelf packages, such as CVX for MATLAB (see [Grant and Boyd, 2014]). If we only consider equality constraints in (87), we can solve the problem analytically

$$\boldsymbol{b}^*_{t,t\rightsquigarrow t+\bar{\tau}} = \tfrac{1}{2}\mathbb{P}_{m,t}(\boldsymbol{q}_t^{-1}\boldsymbol{l}_t) + \mathbb{P}^{\perp}_{m,t}(\boldsymbol{m}_t^{+}\boldsymbol{u}_t), \tag{88}$$

where $\boldsymbol{m}_t^{+}$ is the pseudo-inverse of $\boldsymbol{m}_t$

$$\boldsymbol{m}_t^{+} \equiv \boldsymbol{q}_t^{-1}\boldsymbol{m}_t'(\boldsymbol{m}_t\boldsymbol{q}_t^{-1}\boldsymbol{m}_t')^{-1}; \tag{89}$$

$\mathbb{P}_{m,t} = \mathbb{I} - \boldsymbol{m}_t^{+}\boldsymbol{m}_t$ and $\mathbb{P}^{\perp}_{m,t} = \boldsymbol{m}_t^{+}\boldsymbol{m}_t$ are two complementary projectors; and $\mathbb{I}$ is the identity matrix of conformable dimension. The analytical solution (88) is useful in practice in a variety of situations, most notably to set to zero the exposures to some macroeconomic risk drivers on which the user can have views and that yet are not directly tradable, as explained in Section 3.3.

Since the construction problem (87) can be solved easily and efficiently, we repeat the process at each time step, with fresh estimates $(\hat{\boldsymbol{\theta}}_t, \hat{\boldsymbol{\mu}}_t, \hat{\boldsymbol{\sigma}}_t, \hat{\boldsymbol{c}}_t)$ that account for possible regime shifts. As a result, the optimal path of exposures becomes

$$\{\boldsymbol{b}^*_s \equiv \boldsymbol{b}^*_{s,s}\}_{s\geq t}. \tag{90}$$

The solution (90) is the most general and flexible among the ones proposed so far. Indeed, extensive numerical analysis showed that when the parameters $(\boldsymbol{\theta}, \boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{c})$ are fixed, the rolling cutoff horizon $\bar{\tau}$ is large enough, and there are no constraints other than possibly setting to zero some exposures, the general optimal path of exposures (90) matches perfectly the unconstrained solution with market impact and discretionary views (79). Hence (90) is the optimal solution in the presence of views with no constraints, and at the same it can be generalized to constrained problems.

7 Case studies

In this section we apply the optimal Dynamic Entropy Pooling path of exposures, derived and analyzed in Sections 4-5-6, to two low-dimensional case studies in fixed-income portfolio management.

7.1 One investable risk driver, one view

We consider the case of a single risk driver X_t corresponding to an investable asset (therefore $\boldsymbol{\omega} = 1$ in Equation (51)) and a single view on its expected value at a given time t^*. We emphasize that the time of the view t^* is a fixed date (such as "September 13, 1999"), which does not change as time t goes by. Similarly, the view quantification μ_{view} is fixed as time passes. Hence, at a generic time $s \geq t$ the view is

$$\overline{\mathbb{E}}_s\{X_{t^*}\} = \mu_{view}, \qquad s = t, t+1, \ldots. \tag{91}$$

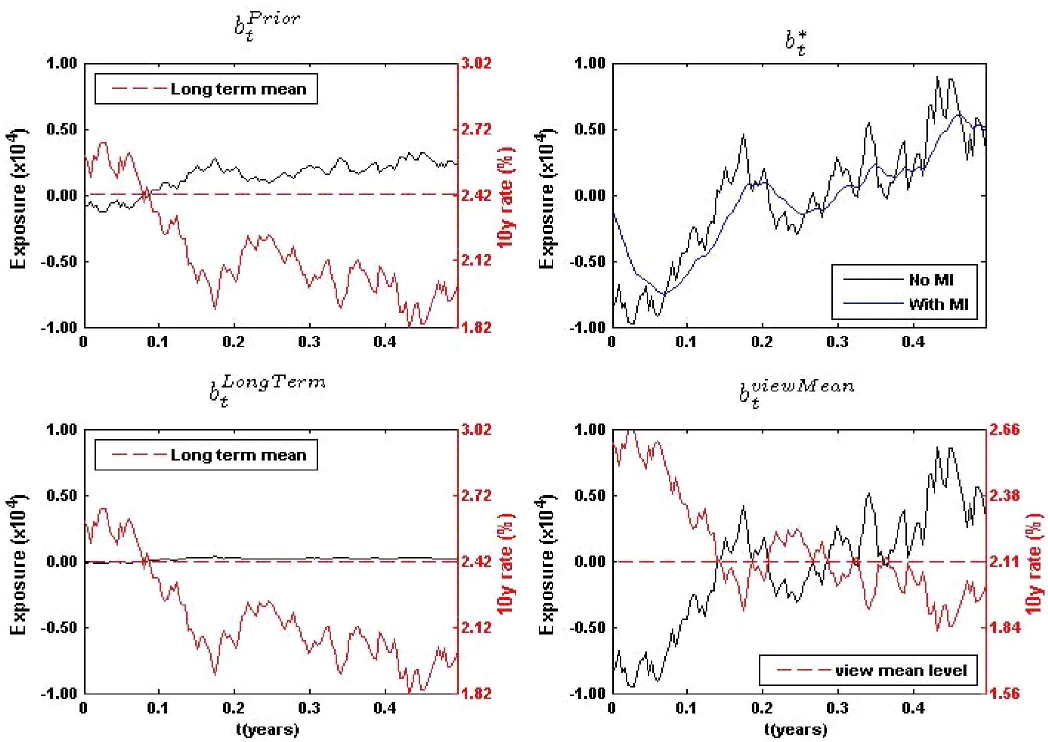

Figure 1: The top-left plot shows the optimal prior solution (93), as well as the 10y rate with its long term expectation on the right axis. The top-right plot shows the Dynamic Entropy Pooling posterior solution (94), both without accounting for market impact and accounting for market impact (smoother line). The bottom-left plot shows the long-term deviation contribution $b_t^{LongTerm}$ in (94), and the bottom-right plot shows the view contribution $b_t^{ViewMean}$ in (94), as well as the view μ_{view}.

The matrix $\boldsymbol{v}_{\mu,t}$ that qualifies the view at the generic time t as in (32) is a $(\bar{t}-t+1)$-dimensional row, where the end point $\bar{t}$ is an arbitrary time in the future, such that $\bar{t} > t^*$

$$\boldsymbol{v}_{\mu,t} = (0, 0, \ldots, \underset{\substack{\uparrow \\ (t^*-t+1)\text{-th}}}{1}, \ldots, 0, 0). \tag{92}$$

The pseudo-inverse matrix $\boldsymbol{v}_{\mu,t}^{+}$ is given by the (t^*-t+1)–th column of (29) divided by $\sigma^2_{t^*-t}$.

When the market impact is neglected, the prior solution is (66), which in our context reads

$$b_s^* = \frac{2\theta}{\gamma\sigma^2}\frac{1}{1+e^{-\theta}}\left(\frac{\mu}{\theta} - x_s\right), \qquad s = t, t+1, \ldots \bar{t}. \tag{93}$$

Similarly, when the market impact is neglected, the Dynamic Entropy Pooling posterior solution that reflects the views follows from the general formulation in (69), which, in terms of the model parameters θ, μ, σ^2, reads

$$b_s^* = \underbrace{\tfrac{2\theta}{\gamma\sigma^2}\tfrac{1}{1+e^{-\theta}}(1-\tfrac{1+e^{\theta}}{1+e^{\theta(t^*-s)}})(\tfrac{\mu}{\theta}-x_s)}_{b_s^{LongTerm}} + \underbrace{\tfrac{2\theta}{\gamma\sigma^2}\tfrac{e^{\theta}}{e^{\theta(t^*-s)}-e^{-\theta(t^*-s)}}(\mu_{view}-x_s)}_{b_s^{ViewMean}}, \quad s = t, t+1, \ldots, \bar{t}. \tag{94}$$

We comment on the result in the case of a mean reverting process, corresponding to $\theta > 0$.

The first contribution $b_s^{LongTerm}$ is proportional to the distance of the current value of the risk driver x_s from its long term expectation μ/θ. If the current value of the risk driver is below (above) its long term mean, it is expected to increase (decrease) and therefore the corresponding contribution to the exposure is positive (negative). Such term does not depend on the kind of view (bullish or bearish for instance) but only on the time at which the view refers.

The second contribution $b_s^{ViewMean}$ accounts for the view, depending on the difference between the view μ_{view} and the current level of the risk driver x_s. If such a difference is positive (negative) the view is bullish (bearish) and the contribution is positive (negative).

For small values of the mean reversion parameter θ, that is, when the process is almost a Brownian motion, the long-term component is small and the optimal posterior solution is well approximated by the view term

$$\theta \approx 0 \Rightarrow b_s^* \approx b_s^{ViewMean} \approx \tfrac{1}{\gamma\sigma^2(t^*-s)}(\mu_{view} - x_s). \tag{95}$$

Also, when the view is very far in the future, the effect of the view vanishes and the posterior solution is well approximated by the prior (93)

$$t^* \to \infty \Rightarrow b_s^* \approx \tfrac{2\theta}{\gamma\sigma^2}\tfrac{1}{1+e^{-\theta}}(\tfrac{\mu}{\theta} - x_s). \tag{96}$$

To compute the fully-fledged optimal path in the presence of market impact, we simply adapt the general solution (79) to the present case.

Example 9 *We consider a fixed-income portfolio where the risk driver is the 10y government rate. To better model low-rate environments we use the 10y shadow rate obtained from the inverse-call transform in [Meucci and Loregian, 2013]. We calibrate the Ornstein-Uhlenbeck process for the shadow rate using 10 years of daily observations. We assume daily rebalancing and we compute the optimal path exposure for a period of 6 months.*
At time $t = 0$ the observed value of the rate is x_0 =2.61%. The estimated parameters (on a daily basis) are μ =0.0302×10^{-3}, θ =1.2469×10^{-3}, σ^2 =0.2295×10^{-6}. The long term expectation is $\theta^{-1}\mu$ =2.42%.
We model a view that the expected value of the 10-year rate will be $\mu_{view} \equiv x_0 -$50 basis points at t^* =1 year from today. We set the risk aversion parameter as $\gamma \equiv 10^{-2}$ and the trading aversion parameter as $\eta \equiv$0.5. We set the intertemporal discount parameter λ in such a way that the half life of the discount factor is 1 month.
Figure 1 displays the results for a simulated path of the 10y rate.

7.2 Two risk drivers (investable/non-investable), two views

In this case study, we consider $\bar{n} = 2$ risk drivers $\boldsymbol{X}_t \equiv \binom{X_{1,t}}{X_{2,t}}$, where $X_{1,t}$ is investable (such as an interest rate, which is investable via futures, or via fixed-income indices and ETF's), and $X_{2,t}$ is not investable (such as an equity index for a fixed-income portfolio manager, or a macroeconomic index, such as inflation). In this case $\boldsymbol{\omega} = (1, 0)$ in (51).

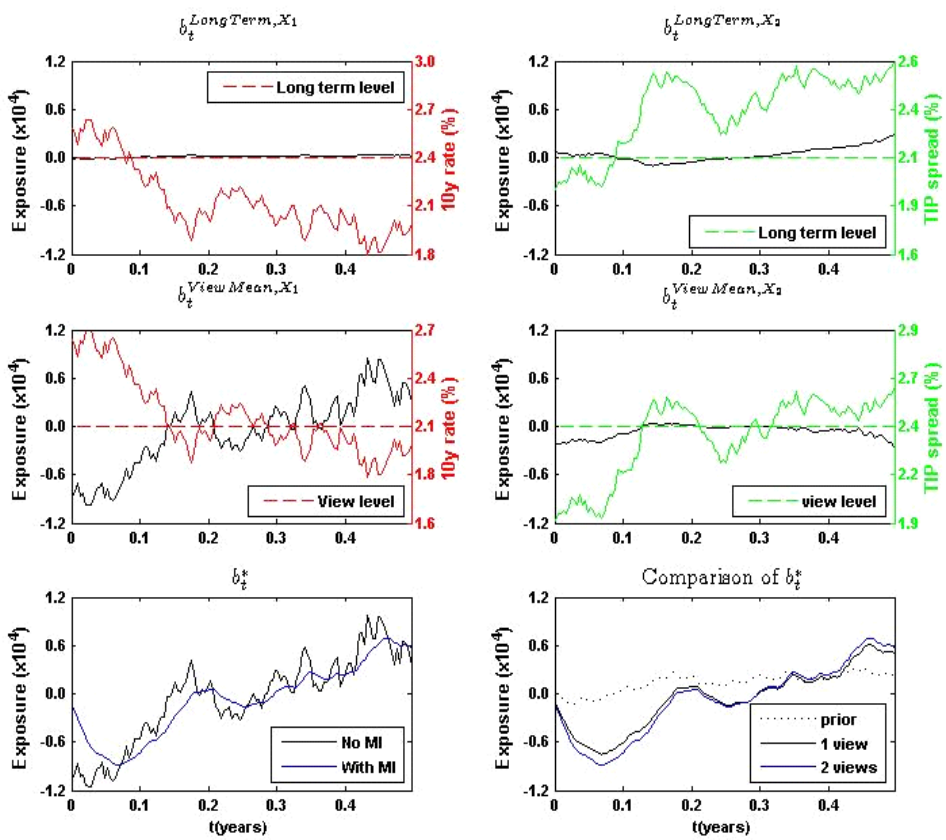

Figure 2: Decomposition of the Dynamic Entropy Pooling posterior exposure to the 10 year rate (101) for a simulated joint path of 10 year rate and 5 year TIP spread (top and middle plots). Also, Dynamic Entropy Pooling posterior solution with and without accounting for market impact (bottom-left). Comparison of Dynamic Entropy Pooling posterior solution with the prior solution, and with the posterior obtained when only the view on the rate is considered (bottom-right).

As in the previous case study we state a view on X_1 at time t^*. Furthermore we consider a second view on X_2 at a prior time $t^{**} < t^*$, as follows

$$\overline{\mathbb{E}}_s\{X_{1,t^*}\} = \mu_{view;1}, \quad \overline{\mathbb{E}}_s\{X_{2,t^{**}}\} = \mu_{view;2}, \qquad s = t, t+1, \ldots \tag{97}$$

We emphasize that the times of views t^{**} and t^* and the views $\mu_{view;1}$ and $\mu_{view;2}$ are fixed as time passes.

Then the matrix $\boldsymbol{v}_{\mu,t}$ that qualifies the views as in (32) is a $2\times 2(\bar{t}-t+1)$ matrix that in this context reads

$$\boldsymbol{v}_{\mu,t} = \left(\begin{array}{ccccccccc} 0 & \ldots & 0 & \ldots & 0 & \ldots & 1 & \ldots & 0 \\ 0 & \ldots & \underset{2(t^{**}-t+1)\text{-th}}{\underset{\uparrow}{1}} & \ldots & 0 & \ldots & \underset{(2(t^{*}-t+1)-1)\text{-th}}{\underset{\uparrow}{0}} & \ldots & 0 \end{array} \right), \tag{98}$$

where the end point $\bar{t}$ is an arbitrary time in the future, such that $\bar{t} > t^{*}$.

We specialize on the situation where the parameter $\boldsymbol{\theta}$ of the Ornstein-Uhlenbeck process satisfies

$$\boldsymbol{\theta} = \left(\begin{smallmatrix} \theta_{1,1}>0 & \theta_{1,2}=0 \\ \theta_{2,1}=0 & \theta_{2,2}>0 \end{smallmatrix} \right). \tag{99}$$

Hence, the processes of $X_{1,t}$ and $X_{2,t}$ are mean-reverting around their long term means $\frac{\mu_1}{\theta_{1,1}}$ and $\frac{\mu_2}{\theta_{2,2}}$ respectively.

If we disregard market impact, the prior solution for the exposure to the first investable risk driver follows from (66) and reads as in the univariate case (93), which we report here, adapted to the new bivariate notation

$$b^{*}_{1,s} = \frac{2\theta_{1,1}}{\gamma\sigma^2_{1,1}}\frac{1}{1+e^{-\theta_{1,1}}}\left(\frac{\mu_1}{\theta_{1,1}} - x_{1,s}\right), \qquad s = t, t+1, \ldots \bar{t}, \tag{100}$$

where $\sigma^2_{1,1}$ is the $(1,1)$ component of $\boldsymbol{\sigma}^2$.

The Dynamic Entropy Pooling posterior exposure in absence of market impact (69) can be expressed in terms of four contributions, due to the long-term and view effect of the two separate drivers

$$b^{*}_{1,s} = \underbrace{b^{LongTerm,x_1}_{1,s} + b^{LongTerm,x_2}_{1,s}}_{b^{LongTerm}_{1,s}} + \underbrace{b^{ViewMean,x_1}_{1,s} + b^{ViewMean,x_2}_{1,s}}_{b^{ViewMean}_{1,s}}, \quad s = t, t+1, \ldots, \bar{t} \tag{101}$$

where

$$b^{LongTerm,x_1}_{1,s} \equiv \frac{2\theta_{1,1}}{\gamma\sigma^2_{1,1}}\frac{1}{1+e^{-\theta_{1,1}}}\left(1-\delta\frac{1-e^{-\theta_{1,1}(t^{*}-s)}}{1-e^{-\theta_{1,1}}}\right)\left(\frac{\mu_1}{\theta_{1,1}} - x_{1,s}\right) \tag{102}$$

$$b^{LongTerm,x_2}_{1,s} \equiv -\varrho\frac{2\theta_{1,1}}{\gamma\sigma^2_{1,1}}\frac{1}{1-e^{-2\theta_{1,1}}}\left(1-e^{-\theta_{2,2}(t^{**}-s)}\right)\left(\frac{\mu_2}{\theta_{2,2}} - x_{2,s}\right) \tag{103}$$

$$b^{ViewMean,x_1}_{1,s} \equiv \delta\frac{2\theta_{1,1}}{\gamma\sigma^2_{1,1}}\frac{1}{1-e^{-2\theta_{1,1}}}(\mu_{view;1} - x_{1,s}), \tag{104}$$

$$b^{ViewMean,x_2}_{1,s} \equiv \varrho\frac{2\theta_{1,1}}{\gamma\sigma^2_{1,1}}\frac{1}{1-e^{-2\theta_{1,1}}}(\mu_{view;2} - x_{2,s}). \tag{105}$$

Notice how the second driver, which may or may not be investable, influences the exposure to the first risk driver.

The coefficients δ and ϱ in (102)-(105) depend on the model parameters as well as the times involved in the problem and are defined as

$$\delta(s, t^{**}, t^{*}; \boldsymbol{\theta}, \boldsymbol{\sigma}) \equiv (\tilde{\sigma}^2_{3,7}\tilde{\sigma}^2_{6,6} - \tilde{\sigma}^2_{3,6}\tilde{\sigma}^2_{6,7})/(\tilde{\sigma}^2_{6,6}\tilde{\sigma}^2_{7,7} - (\tilde{\sigma}^2_{6,7})^2) \tag{106}$$

$$\varrho(s, t^{**}, t^{*}; \boldsymbol{\theta}, \boldsymbol{\sigma}) \equiv (\tilde{\sigma}^2_{3,6}\tilde{\sigma}^2_{7,7} - \tilde{\sigma}^2_{3,7}\tilde{\sigma}^2_{6,7})/(\tilde{\sigma}^2_{6,6}\tilde{\sigma}^2_{7,7} - (\tilde{\sigma}^2_{6,7})^2), \tag{107}$$

where $\tilde{\sigma}^2_{i,j}$ is a short notation to indicate the (i, j)-th entry of the 8×8 sub-covariance matrix of (29) corresponding to the four times $\{s, s+1, t^{**}, t^*\}$ for the two risk drivers.

To compute the fully-fledged optimal path in the presence of market impact, we simply adapt the general solution (79) to the present case.

Example 10 *We consider the same framework as in Example 9, where the 10 years government shadow rate is the investable risk driver X_1, and where we add expected inflation, as measured by the 5 year TIP spread, as a non-investable risk driver X_2.*
The 5 years expected inflation at time t =0 is $x_{2,0}$ =1.93%. The estimated parameters of the bivariate Ornstein-Uhlenbeck process are

$$\boldsymbol{\mu} = 10^{-3} \begin{pmatrix} 0.0302 \\ 0.1894 \end{pmatrix} \qquad \boldsymbol{\theta} = 10^{-3} \begin{pmatrix} 1.2469 & 0 \\ 0 & 8.9692 \end{pmatrix} \qquad \boldsymbol{\sigma}^2 = 10^{-6} \begin{pmatrix} 0.2295 & 0.0594 \\ 0.0594 & 0.1262 \end{pmatrix}. \tag{108}$$

Note that no equilibrium condition (25) has been imposed, as the TIP spread is an external risk driver. The long term level of the TIP spread is $\mu_2/\theta_{2,2}$ =2.11%. All the other parameters are set as in Example 9. The additional view is that the expected value of the TIP spread will be $\mu_{view;2} = x_{2,0}$+50 basis points at t^{**} =0.75 years from today. Figure 2 displays the results.

8 Conclusions

We presented the Dynamic Entropy Pooling, a quantitative approach to discretionary portfolio management, which allows the manager to process views and stress testing at multiple horizons.

To preserve tractability and yet flexibility, we assume as the prior model for the risk drivers a multivariate (not necessarily mean-reverting) Ornstein-Uhlenbeck process, and we model the views as statements on expectations and covariances of arbitrary linear combinations of the process, and at arbitrary times.

The ensuing optimal Dynamic Entropy Pooling policy is computed via dynamic programming with time-varying coefficients, for unconstrained problems, or via a sequence of quadratic programs, for more general constrained problems.

References

[Almgren and Chriss, 2000] Almgren, R. and Chriss, N. (2000). Optimal execution of portfolio transactions. *Journal of Risk*, 3(Winter):5–39.

[Almgren et al., 2005] Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005). Direct estimation of equity market impact. *Risk Magazine*, 18(7):57–62.

[Asness et al., 2013] Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013). Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985.

[Avellaneda, 1999] Avellaneda, M. (1999). Minimum-entropy calibration of asset-pricing models. *International Journal of Theoretical and Applied Finance*, 1:447–472.

[Black and Litterman, 1990] Black, F. and Litterman, R. (1990). Asset allocation: combining investor views with market equilibrium. *Goldman Sachs Fixed Income Research*.

[Brandt et al., 2009] Brandt, M., Santa-Clara, P., and Valkanov, R. (2009). Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9):3411–3447.

[Breuer and Csiszar, 2013] Breuer, T. and Csiszar, I. (2013). Systematic stress tests with entropic plausibility constraints. *Journal of banking and finance*, 37:1552–1559.

[Carhart, 1997] Carhart, M. (1997). On persistence in mutual fund performance. *Journal of Finance*, 52:57–82.

[Caticha and Giffin, 2006] Caticha, A. and Giffin, A. (2006). Updating probabilities. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering, AIP Conf. Proc.*

[Chen et al., 1986] Chen, N., Roll, R., and Ross, S. (1986). Economic forces and the stock market. *The Journal of Business*, 59:383–403.

[Cover and Thomas, 2006] Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley, 2nd edition.

[Davis and Lleo, 2013] Davis, M. and Lleo, S. (2013). Black-litterman in continuous time: the case for filtering. *Quantitative Finance Letters*, 1(1):10.1080/21649502.2013.803794.

[Fama and French, 1993] Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

[Fung and Hsieh, 1997] Fung, W. and Hsieh, D. A. (1997). Empirical characterizations of dynamic trading strategies: The case of hedge funds. *Review of Financial Studies*, 10:275–302.

[Gârleanu and Pedersen, 2013] Gârleanu, N. and Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance*, 68(6):2309–2340.

[Gatheral, 2010] Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759.

[Glasserman and Xu, 2014] Glasserman, P. and Xu, X. (2014). Robust risk measurement and model risk. *Quantitative Finance*, 1:29–58.

[Grant and Boyd, 2014] Grant, M. and Boyd, S. (2014). CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx.

[Grinold, 1989] Grinold, R. (1989). The fundamental law of active management. *Journal of Portfolio Management*, 15(3):30–37.

[Grinold, 2006] Grinold, R. (2006). A dynamic model of portfolio management. *Journal of Investment Management*, 4:5–22.

[Grinold, 2007] Grinold, R. (2007). Dynamic portfolio analysis. *Journal of Portfolio Management*, 34:12–26.

[Grinold, 2010] Grinold, R. (2010). Signal weighting. *The Journal of Portfolio Management*, 36:24–34.

[Grinold and Kahn, 1999] Grinold, R. C. and Kahn, R. (1999). *Active Portfolio Management. A Quantitative Approach for Producing Superior Returns and Controlling Risk*. McGraw-Hill, 2nd edition.

[Hansen and Sargent, 2007] Hansen, L. P. and Sargent, T. J. (2007). *Robustness*. Princeton University Press.

[Karatzas et al., 1987] Karatzas, I., Lehoczky, J. P., and Shreve, S. E. (1987). Optimal portfolio and consumption decisions for a small investor on a finite horizon. *SIAM Journal of Control and Optimization*, 27:1157–1186.

[Merton, 1992] Merton, R. C. (1992). *Continuous-Time Finance*. Blackwell.

[Meucci, 2008] Meucci, A. (2008). Fully Flexible Views: Theory and practice. http://symmys.com/node/158. Risk, 21(10), 97-102.

[Meucci, 2009] Meucci, A. (2009). Review of statistical arbitrage, cointegration, and multivariate Ornstein-Uhlenbeck. http://symmys.com/node/132. Working Paper.

[Meucci, 2011] Meucci, A. (2011). The Prayer: Ten-step checklist for advanced risk and portfolio management. http://symmys.com/node/63. Garp Risk Professional.

[Meucci and Loregian, 2013] Meucci, A. and Loregian, A. (2013). Neither "normal" not "lognormal": Modeling interest rates across all regimes. http://symmys.com/node/601. Working Paper.

[Moallemi and Sağlam, 2012] Moallemi, C. C. and Sağlam, M. (2012). Dynamic portfolio choice with linear rebalancing rules. http://ssrn.com/abstract=2011605. Working paper.

[Rosenberg and Lanstein, 1985] Rosenberg, B.and Reid, K. and Lanstein, R. (1985). Persuasive evidence of market inefficiency. *Journal of Portfolio Management*, 11:9–17.

A Appendix

In this appendix we discuss technical results that can be skipped at first reading.

A.1 Limit inversion

We consider here the case of a MVOU process (24) with a matrix $\boldsymbol{\theta}$ that has some null eigenvalues. In this case, the matrices $\boldsymbol{\theta}$ and $\boldsymbol{\theta} \oplus \boldsymbol{\theta}$ are not invertible. Nevertheless the mean and the variance matrix as defined in (27) and (28) are still well defined. In fact, both the expressions (27) and (28) depend on a term that is $\left(\mathbb{I} - e^{-\boldsymbol{\alpha} t}\right) \boldsymbol{\alpha}^{-1}$, where $\boldsymbol{\alpha} = \boldsymbol{\theta}$ in (27), $\boldsymbol{\alpha} = \boldsymbol{\theta} \oplus \boldsymbol{\theta}$ in (28) and $\mathbb{I}$ is the identity matrix with conformable dimension. If $\boldsymbol{\alpha}$ is a $\bar{n} \times \bar{n}$ matrix, such an expression has to be computed as

$$\left(\mathbb{I} - e^{-\boldsymbol{\alpha} t}\right) \boldsymbol{\alpha}^{-1} \equiv \lim_{\rho_{\bar{\imath}} \to 0} \boldsymbol{v} Diag\left(\tfrac{1-e^{-\rho_1 t}}{\rho_1}, \ldots, \tfrac{1-e^{-\rho_{\bar{n}} t}}{\rho_{\bar{n}}}\right) \boldsymbol{v}^{-1}, \tag{109}$$

where $\rho_1, \ldots \rho_{\bar{n}}$ and $\boldsymbol{v}$ are respectively the eigenvalues and the matrix of the eigenvectors of $\boldsymbol{\alpha}$, in such a way that $\boldsymbol{\alpha} = \boldsymbol{v} Diag(\rho_1, \ldots, \rho_{\bar{n}}) \boldsymbol{v}^{-1}$, and $\bar{\imath}$ labels the null eigenvalues.

A.2 The general Bellman equation

We first compute the expectation of the value function $v_{s+1}(\boldsymbol{b}_s, \boldsymbol{X}_{s+1})$ in (76)

$$\begin{aligned} \overline{\mathbb{E}}\{v_{s+1}(\boldsymbol{b}_s, \boldsymbol{X}_{s+1})|\boldsymbol{i}_s\} = & -\tfrac{1}{2}\boldsymbol{b}_s'\boldsymbol{\psi}_{bb,s}\boldsymbol{b}_s + \boldsymbol{b}_s'\boldsymbol{\psi}_{bx,s}\left(\boldsymbol{\alpha}_s + (\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s\right) + \frac{1}{2}\mathrm{tr}(\boldsymbol{\psi}_{xx,s}\bar{\boldsymbol{\sigma}}_s^2) \\ & + \frac{1}{2}\left(\boldsymbol{\alpha}_s + (\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s\right)'\boldsymbol{\psi}_{xx,s}\left(\boldsymbol{\alpha}_s + (\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s\right) \\ & + \boldsymbol{\psi}_{b,s}'\boldsymbol{b}_s + \boldsymbol{\psi}_{x,s}'\left(\boldsymbol{\alpha}_s + (\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s\right) + \psi_{0,s}. \end{aligned} \tag{110}$$

Then the Bellman equation (76) reads

$$v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s) = \max_{\boldsymbol{b}_s}\{-\tfrac{1}{2}\boldsymbol{b}_s'\boldsymbol{q}_s\boldsymbol{b}_s + \boldsymbol{b}_s'\boldsymbol{l}_s + d_s\}, \tag{111}$$

where

$$\boldsymbol{q}_s \equiv \gamma\boldsymbol{\omega}\bar{\boldsymbol{\sigma}}_s^2\boldsymbol{\omega}' + \eta\boldsymbol{c}^2 + e^{-\lambda}\boldsymbol{\psi}_{bb,s} \tag{112}$$

$$\boldsymbol{l}_s \equiv \boldsymbol{\omega}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s} + [\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})]\boldsymbol{x}_s + \eta\boldsymbol{c}^2\boldsymbol{b}_{s-1} \tag{113}$$

$$\begin{aligned} d_s \equiv & -\tfrac{\eta}{2}\boldsymbol{b}_{s-1}'\boldsymbol{c}^2\boldsymbol{b}_{s-1} \\ & + e^{-\lambda}\tfrac{1}{2}\boldsymbol{x}_s'(\boldsymbol{\beta}_s' + \mathbb{I}_{\bar{n}})\boldsymbol{\psi}_{xx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s + e^{-\lambda}(\boldsymbol{\alpha}_s'\boldsymbol{\psi}_{xx,s} + \boldsymbol{\psi}_{x,s}')(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})\boldsymbol{x}_s \\ & + \tfrac{1}{2}e^{-\lambda}\mathrm{tr}(\boldsymbol{\psi}_{xx,s}\bar{\boldsymbol{\sigma}}_s^2) + \tfrac{1}{2}e^{-\lambda}\boldsymbol{\alpha}_s'\boldsymbol{\psi}_{xx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{x,s}'\boldsymbol{\alpha}_s + e^{-\lambda}\psi_{0,s} \end{aligned} \tag{114}$$

The optimal solution of the right hand side of Equation (111) is

$$\boldsymbol{b}_s^* = \boldsymbol{q}_s^{-1}\boldsymbol{l}_s. \tag{115}$$

that substituted in (111) gives

$$v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s) = \tfrac{1}{2}\boldsymbol{l}_s'\boldsymbol{q}_s^{-1}\boldsymbol{l}_s + d_s\ . \tag{116}$$

Using the quadratic expression for $v_s(\boldsymbol{b}_{s-1}, \boldsymbol{x}_s)$, and the definition of $\boldsymbol{q}_s$, $\boldsymbol{l}_s$ and d_s given in Equations (112), (113) and (114) respectively, we impose that the above equation holds for any $\boldsymbol{x}_s$ and $\boldsymbol{b}_{s-1}$. This gives the following set of equations for the coefficients of the value function

$$\boldsymbol{\psi}_{bb,s-1} = \eta\boldsymbol{c}^2 - \eta\boldsymbol{c}^2\boldsymbol{q}_s^{-1}\eta\boldsymbol{c}^2 \tag{117}$$

$$\boldsymbol{\psi}_{bx,s-1} = \eta\boldsymbol{c}^2\boldsymbol{q}_s^{-1}(\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})) \tag{118}$$

$$\begin{aligned} \boldsymbol{\psi}_{xx,s-1} = & (\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}}))'\boldsymbol{q}_s^{-1}(\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})) \\ & + e^{-\lambda}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})'\boldsymbol{\psi}_{xx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}}) \end{aligned} \tag{119}$$

$$\boldsymbol{\psi}_{b,s-1} = \eta\boldsymbol{c}^2\boldsymbol{q}_s^{-1}(\boldsymbol{\omega}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s}) \tag{120}$$

$$\begin{aligned} \boldsymbol{\psi}_{x,s-1} = & (\boldsymbol{\omega}\boldsymbol{\beta}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}}))'\boldsymbol{q}_s^{-1}(\boldsymbol{\omega}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s}) \\ & + e^{-\lambda}(\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})'(\boldsymbol{\psi}_{xx,s}'\boldsymbol{\alpha}_s + \boldsymbol{\psi}_{x,s}) \end{aligned} \tag{121}$$

$$\begin{aligned} \psi_{0,s-1} = & \tfrac{1}{2}(\boldsymbol{\omega}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s})'\boldsymbol{q}_s^{-1}(\boldsymbol{\omega}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{bx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{b,s}) \\ & + \tfrac{1}{2}e^{-\lambda}\mathrm{tr}(\boldsymbol{\psi}_{xx,s}\bar{\boldsymbol{\sigma}}_s^2) + \tfrac{1}{2}e^{-\lambda}\boldsymbol{\alpha}_s'\boldsymbol{\psi}_{xx,s}\boldsymbol{\alpha}_s + e^{-\lambda}\boldsymbol{\psi}_{x,s}'\boldsymbol{\alpha}_s + e^{-\lambda}\psi_{0,s}. \end{aligned} \tag{122}$$

A.3 The value function coefficients for the prior case

As explained in Section 5.1, when the market follows the prior model (26), we can drop the time subscript from the value function coefficients. Therefore equations (117)-(122) can be solved analytically.

First we solve Equation (117), that using the definition (112) reads

$$\boldsymbol{\psi}_{bb} = \eta \boldsymbol{c}^2 - \eta \boldsymbol{c}^2 \left(\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}' + \eta \boldsymbol{c}^2 + e^{-\lambda} \boldsymbol{\psi}_{bb}\right)^{-1} \eta \boldsymbol{c}^2. \tag{123}$$

Then using the Woodbury identity

$$(\boldsymbol{a}^{-1} + \boldsymbol{c}^{-1})^{-1} = \boldsymbol{a} - \boldsymbol{a}(\boldsymbol{c} + \boldsymbol{a})^{-1} \boldsymbol{a} \tag{124}$$

to manipulate the right hand side of (123), where $\boldsymbol{a} = (\eta \boldsymbol{c}^2)$ and $\boldsymbol{c} = \left(\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}' + e^{-\lambda} \boldsymbol{\psi}_{bb}\right)$, Equation (123) becomes

$$\boldsymbol{\psi}_{bb}^{-1} = (\eta \boldsymbol{c}^2)^{-1} + \left(\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}' + e^{-\lambda} \boldsymbol{\psi}_{bb}\right)^{-1}. \tag{125}$$

Multiplying (125) first by $\boldsymbol{\psi}_{bb}$ on the left and then by $\left(\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}' + e^{-\lambda} \boldsymbol{\psi}_{bb}\right)$ on the right, we obtain

$$\boldsymbol{\psi}_{bb} (\eta \boldsymbol{c}^2)^{-1} \boldsymbol{\psi}_{bb} + \boldsymbol{\psi}_{bb} \left(e^{\lambda} (\eta \boldsymbol{c}^2)^{-1} (\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}') + (e^{\lambda} - 1) \mathbb{I}_{\bar{k}}\right) - e^{\lambda} \gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}' = \boldsymbol{0}\,. \tag{126}$$

Redefining

$$\widehat{\boldsymbol{\psi}}_{bb} \equiv (\eta \boldsymbol{c}^2)^{-\frac{1}{2}} \boldsymbol{\psi}_{bb} (\eta \boldsymbol{c}^2)^{-\frac{1}{2}} \tag{127}$$

$$\hat{\boldsymbol{\sigma}}^2 \equiv e^{\lambda} (\eta \boldsymbol{c}^2)^{-\frac{1}{2}} (\gamma \boldsymbol{\omega} \boldsymbol{\sigma}_1^2 \boldsymbol{\omega}') (\eta \boldsymbol{c}^2)^{-\frac{1}{2}} \tag{128}$$

Equation (126) becomes

$$\widehat{\boldsymbol{\psi}}_{bb}^2 + \widehat{\boldsymbol{\psi}}_{bb} \left(\hat{\boldsymbol{\sigma}}^2 + \mathbb{I}_{\bar{k}} (e^{\lambda} - 1)\right) - \hat{\boldsymbol{\sigma}}^2 = \boldsymbol{0}\,, \tag{129}$$

whose solution is

$$\widehat{\boldsymbol{\psi}}_{bb} = \left(\tfrac{1}{4} (\hat{\boldsymbol{\sigma}}^2 + \mathbb{I}_{\bar{k}} (e^{\lambda} - 1))^2 + \hat{\boldsymbol{\sigma}}^2\right)^{\frac{1}{2}} - \tfrac{1}{2} (\hat{\boldsymbol{\sigma}}^2 + \mathbb{I}_{\bar{k}} (e^{\lambda} - 1))\,, \tag{130}$$

where the positive determination is the only one for which $\widehat{\boldsymbol{\psi}}_{bb}$ is positive definite. Finally, $\boldsymbol{\psi}_{bb}$ is obtained by inverting (127).

Equation (118) can be solved by vectorizing it and then using the identity

$$vec(\boldsymbol{abc}) = (\boldsymbol{c}' \otimes \boldsymbol{a}) vec(\boldsymbol{b}) \tag{131}$$

for any matrices $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{c}$ with the conformable dimensions. In vectorized form, (118) reads

$$vec(\boldsymbol{\psi}_{bx}) = vec(\eta \boldsymbol{c}^2 \boldsymbol{q}^{-1} \boldsymbol{\omega} \boldsymbol{\beta}) + e^{-\lambda} ((\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}}) \otimes \eta \boldsymbol{c}^2 \boldsymbol{q}^{-1}) vec(\boldsymbol{\psi}_{bx})\,, \tag{132}$$

which yields to

$$vec(\boldsymbol{\psi}_{bx}) = (\mathbb{I}_{\bar{k}\bar{n}} - e^{-\lambda} ((\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}}) \otimes \eta \boldsymbol{c}^2 \boldsymbol{q}^{-1}))^{-1} vec(\eta \boldsymbol{c}^2 \boldsymbol{q}^{-1} \boldsymbol{\omega} \boldsymbol{\beta})\,. \tag{133}$$

Also Equation (119) is solved by vectorizing it and then using the identity (131). The result is

$$vec(\boldsymbol{\psi}_{xx}) = (\mathbb{I}_{\bar{n}^2} - e^{-\lambda}(\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}}) \otimes (\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}}))^{-1} \times vec\left((\boldsymbol{\omega\beta} + e^{-\lambda}\boldsymbol{\psi}_{bx}(\boldsymbol{\beta} + \mathbb{I}_{\bar{n}}))'\boldsymbol{q}^{-1}(\boldsymbol{\omega\beta} + e^{-\lambda}\boldsymbol{\psi}_{bx}(\boldsymbol{\beta} + \mathbb{I}_{\bar{n}}))\right). \tag{134}$$

Equation (120) is easily solved by collecting the terms in $\boldsymbol{\psi}_b$

$$(\mathbb{I}_{\bar{k}} - e^{-\lambda}\eta\boldsymbol{c}^2\boldsymbol{q}^{-1})\boldsymbol{\psi}_b = \eta\boldsymbol{c}^2\boldsymbol{q}^{-1}\left(\boldsymbol{\omega\alpha} + e^{-\lambda}\boldsymbol{\psi}_{bx}\boldsymbol{\alpha}\right). \tag{135}$$

Hence the solution is

$$\boldsymbol{\psi}_b = (\boldsymbol{q}(\eta\boldsymbol{c}^2)^{-1} - e^{-\lambda}\mathbb{I}_{\bar{k}})^{-1}\left(\boldsymbol{\omega} + e^{-\lambda}\boldsymbol{\psi}_{bx}\right)\boldsymbol{\alpha} \tag{136}$$

The solution of Equation (121) is

$$\boldsymbol{\psi}_x = (\mathbb{I}_{\bar{n}} - e^{-\lambda}(\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}}))^{-1} \times [(\boldsymbol{\omega\beta} + e^{-\lambda}\boldsymbol{\psi}_{bx}(\boldsymbol{\beta} + \mathbb{I}_{\bar{n}}))'\boldsymbol{q}^{-1}(\boldsymbol{\omega\alpha} + e^{-\lambda}\boldsymbol{\psi}_{bx}\boldsymbol{\alpha} + e^{-\lambda}\boldsymbol{\psi}_b) + e^{-\lambda}(\boldsymbol{\beta}' + \mathbb{I}_{\bar{n}})\boldsymbol{\psi}'_{xx}\boldsymbol{\alpha}]. \tag{137}$$

Finally the solution of (122) is

$$\boldsymbol{\psi}_0 = \tfrac{1}{1-e^{-\lambda}}\Big(\tfrac{1}{2}(\boldsymbol{\omega\alpha} + e^{-\lambda}\boldsymbol{\psi}_{bx}\boldsymbol{\alpha} + e^{-\lambda}\boldsymbol{\psi}_b)'\boldsymbol{q}^{-1}(\boldsymbol{\omega\alpha} + e^{-\lambda}\boldsymbol{\psi}_{bx}\boldsymbol{\alpha} + e^{-\lambda}\boldsymbol{\psi}_b) + \tfrac{1}{2}e^{-\lambda}\mathrm{tr}(\boldsymbol{\psi}_{xx}\bar{\boldsymbol{\sigma}}^2) + \tfrac{1}{2}e^{-\lambda}\boldsymbol{\alpha}'\boldsymbol{\psi}_{xx}\boldsymbol{\alpha} + e^{-\lambda}\boldsymbol{\psi}'_x\boldsymbol{\alpha}\Big). \tag{138}$$