FACTOR INVESTING WITH BLACK-LITTERMAN-BAYES: INCORPORATING FACTOR VIEWS AND PRIORS IN PORTFOLIO CONSTRUCTION

PETTER N. KOLM[a] AND GORDON RITTER[b]

Abstract. The authors propose a general framework referred to as Black-Litterman-Bayes (BLB) for constructing optimal portfolios for factor-based investing. In the spirit of the classical Black-Litterman model, the framework allows for the incorporation of investor views and different priors on factor risk premia, including data-driven and benchmark priors. Computationally efficient closed-form formulas are provided for the (posterior) expected returns and return covariance matrix that result from integrating factor views into an APT multi-factor model. In a step-by-step procedure, the authors show how to build the prior and incorporate the factor views, demonstrating in a realistic empirical example, using a number of well-known cross-sectional U.S. equity factors, that the BLB approach can add value to mean-variance optimal multi-factor risk premia portfolios.

Keywords: Factor investing; Investment analysis; Bayesian statistics; Black-Litterman; Portfolio optimization; Portfolio theory; Risk premia

Key Findings

- The authors propose a general framework referred to as Black-Litterman-Bayes (BLB) for constructing optimal portfolios for factor-based investing.
- The framework allows for the incorporation of investor views and different priors on factor risk premia, including data-driven and benchmark priors.
- The authors provide computationally efficient closed-form formulas for the (posterior) expected returns and return covariance matrix.

(a) Courant Institute of Mathematical Sciences, New York University, 251 Mercer St., NY 10012, USA. Email: *petter.kolm@nyu.edu*
(b) Courant Institute of Mathematical Sciences, NYU Tandon School of Engineering, Baruch College, Columbia University and General Partner/CIO at Ritter Alpha, LP., *ritter@post.harvard.edu*
Date: November 10, 2020.

- In a realistic empirical example, using a number of well-known cross-sectional U.S. equity factors, they demonstrate that the BLB approach can add value to mean-variance optimal multi-factor risk premia portfolios.

1. Introduction

Risk premia and smart beta investing have attracted considerable attention from institutional and individual investors for its simplicity, transparency and low cost. Johnson (2018) estimated that at the end of 2018 there were almost 1,500 smart beta products traded on the public exchanges alone, amounting to about \$800 billion of assets under management worldwide. These funds represent a hybrid between active and passive investing, often deploying rules-based investment strategies and portfolio construction methodologies. Nevertheless, the main goal of these investment vehicles is to outperform the market just like that of active approaches. For this purpose, the majority of these funds rely on combining well-known *risk premia* into a portfolio, perhaps by a simple equally weighted approach or more sophisticated portfolio optimization schemes. Numerous academic studies provide support for this form of *factor investing*, confirming the presence of factor premia in different asset classes, domestically and internationally, and across various themes such as betting-against beta, carry, low volatility, momentum, quality, return seasonality, size and value (see, for example, Fama and French (1993), Carhart (1997), Subrahmanyam (2010), Moskowitz, Ooi, and Pedersen (2012), Asness, Moskowitz, and Pedersen (2013), Ang (2014), Frazzini and Pedersen (2014), Keloharju, Linnainmaa, and Nyberg (2016), Koijen, Moskowitz, Pedersen, and Vrugt (2018), and Baltussen, Swinkels, and Van Vliet (2019)).

While a number of studies have addressed portfolio construction for risk premia (see, for example, Bass, Gladstone, and Ang (2017), Bergeron, Kritzman, and Sivitsky (2018), Dopfel and Lester (2018), Bender, Le Sun, and Thomas (2018), and Aliaga-Diaz, Renzi-Ricci, Daga, and Ahluwalia (2020)), most of the work in this area has focused on factor models from the frequentist's perspective.

First introduced in two Goldman Sachs Fixed Income Research Notes in the early nineties (Black and Litterman, 1990; Black and Litterman, 1991b) and later published in academic journals (Black and Litterman, 1991a; Black and Litterman, 1992), perhaps the most well-known Bayesian portfolio construction methodology is the *Black-Litterman* (BL) model. The purpose of BL is to improve the inputs in portfolio construction. A key insight of Black and Litterman is that the unreasonable and unstable portfolio holdings frequently encountered in classical mean-variance optimization (MVO) is the result of feeding noisy and inconsistent risk and return forecasts into an optimizer. They realized that, risk and return forecasts have to be consistent with one another to be effective in optimization. Their solution resulted in the BL model, which combines the market implied views with the investors views, producing consistent risk and return forecasts (see, for example, Litterman and He (1999), Satchell and Scowcroft (2000), Fabozzi, Focardi, and Kolm (2006), Fabozzi, Kolm, Pachamanova, and Focardi (2007), Meucci (2009), and Fabozzi, Focardi, and Kolm (2010) for detailed expositions and discussions of the BL model). The BL method itself is often described as "being Bayesian," but the original authors do not elaborate directly on its connections with Bayesian statistics. Kolm and Ritter (2017) introduce the *Black-Litterman-Bayes* (BLB) model, a most general Bayesian portfolio optimization procedure. They elucidate the direct link between BL and Bayesian regression, highlighting its relationship to Bayesian networks (Pearl, 2014). In addition, they provide a detailed Bayesian treatment of views on factor risk premia in the context of multi-factor models. Specifically, they show that the Arbitrage Pricing Theory (APT) of (Ross, 1976), today a cornerstone in the practice of quantitative investing, has a natural Bayesian extension. They derive the associated BL optimal factor portfolio and discuss two common types of priors in these settings, benchmark-optimal and data-driven priors.

Several related studies explore BL and more general Bayesian approaches to portfolio construction including Jones, Lim, and Zangari (2007), Zhou (2009), Avramov and Zhou (2010), Figelman (2017), Rubesam and Hwang (2019), and Melas, Nagy, Kumar, and Zangari (2019).

In this article we explore the BLB model for factor portfolios and factor views, one of the models introduced in Kolm and Ritter (2017). In particular, we take as a premise that some asset managers and hedge funds employ systematic investment processes in which the principal goal is to optimally harvest various factor premia. In other words, these managers base their investment decisions on the view that certain factor risk premia are likely to deliver risk-adjusted returns that are attractive or cannot be obtained through the CAPM market portfolio alone.

In the original BL framework all views were on portfolios of securities. Factors are not portfolios; rather they are unobservable random variables which represent sources of risk that are common to all securities trading in the market. Hence the BL model is not directly applicable to views about factor performance. However, an extension of the same method used by Black and Litterman allows a portfolio manager to express views on factors directly.

In this article, we make several contributions. First, we show how to incorporate views on factor risk premia into APT models. Second, we provide concrete closed-form formulas for the BLB expected returns and covariances that result from integrating an investor's factor views into an APT model. Third, with a realistic empirical example, we demonstrate how the incorporation of factor views can add value to multi-factor risk premia portfolios. Specifically, we describe a step-by-step procedure of how to build the prior and incorporate the factor views, demonstrating that, in the event the views are prescient, the BLB portfolio strategy outperforms an analogous mean-variance optimization (MVO) strategy where the latter uses the views to build expected returns, but does not benefit from the BLB prior or likelihood.

2. The Black-Litterman Model

The BL model blends market implied expected returns with investor's views, taking into account the investor's confidence about these inputs (Black and Litterman, 1991a; Black and Litterman, 1992). In other words, the so-called *BL expected returns* are a confidence weighted average of the market implied expected returns and the investor's views.

We consider a market with n securities whose returns $\mathbf{r}_t$ are multivariate normally distributed with an expected return vector, $\boldsymbol{\mu}$, and return covariance matrix, $\boldsymbol{\Sigma}$, that is

$$\mathbf{r}_t \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})\,. \tag{1}$$

A mean-variance investor will choose their portfolio holdings, $\mathbf{h} = (h_1, \ldots, h_n)'$, by solving the mean-variance optimization (MVO) problem

$$\max_{\mathbf{h}} \mu_p - \frac{\lambda}{2}\sigma_p^2\,, \tag{2}$$

where $\mu_p := \mathbb{E}[\mathbf{h}'\mathbf{r}_t] = \mathbf{h}'\boldsymbol{\mu}$ is the expected portfolio return, $\sigma_p^2 := \mathbb{V}[\mathbf{h}'\mathbf{r}_t] = \mathbf{h}'\boldsymbol{\Sigma}\mathbf{h}$ is the portfolio variance and λ is the investor's risk aversion parameter that determines their trade-off between expected portfolio return and risk. The portfolio holdings $\mathbf{h}$ are in units of dollars, or whatever currency the investor is using. By solving (2) we obtain the optimal holdings

$$\mathbf{h}^* = \lambda^{-1}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}\,.$$

In the BL model an investor can have their own views. A view might be that "the German equity market will outperform a capitalization-weighted basket of the rest of the European equity markets by 5%" (Litterman and He, 1999). To express this view we let $\mathbf{p} = (p_1, \ldots, p_n)'$ denote the portfolio which is long one unit of the DAX index, and short a one-unit basket which holds each of the other major European indices (UKX, CAC40, AEX, etc.) in proportion to their respective aggregate market capitalizations, so that $\sum_i p_i = 0$. Denoting the expected return of this portfolio by $q = 0.05$, we can express the investor's view as $\mathbb{E}\left[\mathbf{p}'\mathbf{r}\right] = q$ where $\mathbf{r}$ is the vector of security returns over the considered time horizon. If the investor has k such views

$$\mathbb{E}\left[\mathbf{p}_i'\mathbf{r}\right] = q_i, \quad \text{for} \quad i = 1, \ldots, k\,, \tag{3}$$

the portfolios $\mathbf{p}_i$ are more conveniently arranged as rows of the $k \times n$ matrix, $\mathbf{P}$. Then the investor's views can be compactly expressed as

$$\mathbb{E}[\mathbf{P}\mathbf{r}] = \mathbf{q}\,, \tag{4}$$

where $\mathbf{q} = (q_1, \ldots, q_k)'$. However, the expectations (4) alone do not convey how confident an investor may be in their views. For this purpose, in the BL model the investor specifies a level of confidence in their views via the $k \times k$ matrix $\boldsymbol{\Omega}$ such that

$$\mathbf{q} = \mathbf{P}\boldsymbol{\mu} + \boldsymbol{\varepsilon}_{\mathbf{q}}, \quad \boldsymbol{\varepsilon}_{\mathbf{q}} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})\,. \tag{5}$$

In other words, an investor's views are represented as the (a) mean and (b) covariance matrix of linear combinations of the expected returns $\boldsymbol{\mu}$. This information is partial and indirect because the views are statements about portfolio returns rather than on individual security returns directly. Furthermore, this information is noisy, with the noise $\boldsymbol{\varepsilon}_{\mathbf{q}}$ assumed to be multivariate normally distributed, because the future is uncertain.

Black and Litterman were motivated by the principle that, in the absence of any sort of investor views and constraints, the mean-variance optimal portfolio holdings should be the the global CAPM equilibrium portfolio,

$\mathbf{h}_{\text{eq}}$. Consequently, in the absence of any views, the investor's model for security returns is given by (1) with

$$\boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\pi}, \mathbf{C}) \,, \tag{6}$$

where $\boldsymbol{\pi}$ is the global CAPM expected return in excess of the risk-free rate and the inverse of the covariance matrix $\mathbf{C}$ represents the amount of confidence the investor has in their estimate of $\boldsymbol{\pi}$. The distributional assumption (6) is referred to as the (CAPM) *prior*.

Given the security return distribution (1), the prior (6) and the investor's views (5), the BL model yields the (posterior) expected returns

$$\boldsymbol{\mu}_{\text{BL}} := (\mathbf{P}'\boldsymbol{\Omega}^{-1}\mathbf{P} + \mathbf{C}^{-1})^{-1}(\mathbf{P}'\boldsymbol{\Omega}^{-1}\mathbf{q} + \mathbf{C}^{-1}\boldsymbol{\pi}) \tag{7}$$

and (posterior) return covariance matrix

$$\boldsymbol{\Sigma}_{\text{BL}} := \mathbb{V}[\boldsymbol{\mu}_{\text{BL}}] = (\mathbf{P}'\boldsymbol{\Omega}^{-1}\mathbf{P} + \mathbf{C}^{-1})^{-1} \,, \tag{8}$$

frequently referred to as the *Black-Litterman expected returns* and *covariance matrix*, respectively. Kolm and Ritter (2017) provides a full Bayesian derivation of the BL model.

3. Black-Litterman for Factor Investing

3.1. **Linear Factor Models.** Taking *Arbitrage Pricing Theory* (APT) (Ross, 1976; Roll and Ross, 1980) as a starting point, we consider a market where the return in excess of the risk-free rate of the i-th security on day t is given by

$$r_{t,i} = x_{t,i,1}f_{t,1} + x_{t,i,2}f_{t,2} + \cdots + x_{t,i,k}f_{t,k} + \varepsilon_{t,i} \,, \tag{9}$$

where $f_{t,j}$, $j = 1, \ldots, k$, are k *factors*, $x_{t,i,j}$ denote the j-th *factor loading* of the i-th security, and $\varepsilon_{t,i}$ is the *residual return* (or *idiosyncratic return*) of the i-th security. We assume the residual returns $\varepsilon_{t,i}$ are white noise processes that are mutually uncorrelated with each other and across time, as well as uncorrelated with all factors, that is

$$\begin{aligned} \mathbb{E}[\varepsilon_{t,i}] &= 0 \,, && \text{for all } i \text{ and } t \,, \\ \operatorname{cov}[f_{t,j}, \varepsilon_{s,i}] &= 0 \,, && \text{for all } j, i, t \text{ and } s \,, \\ \operatorname{cov}[\varepsilon_{t,i}, \varepsilon_{s,j}] &= \begin{cases} \sigma_{t,i}^2, & \text{if } i = j \text{ and } t = s \,, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Additionally, we assume for simplicity that the residual returns are normally distributed

$$\varepsilon_{t,i} \sim \mathcal{N}(0, \sigma_{t,i}^2)\,. \tag{10}$$

Typically, factor loadings $x_{t,i,j}$ are known, exogenously-given and non-stochastic. Some examples of factors include size and value. For a size factor the loadings are typically some nonlinear transformation of the market capitalization of each stock, which are known and exogenous. For a value factor, the loadings are frequently taken to be some measure of book value or earnings of each company, divided by their market capitalization. The value factor is also often passed through a non-linear transformation to normalize it cross-sectionally. The factors $f_{t,j}$, $j = 1, \ldots, k$ are *unobservable* (or *hidden*) random variables which collectively account for all co-movement of security returns. As they cannot be observed directly, they must be obtained via statistical inference. In contrast to the residuals, in general different factors are not necessarily independent.

Using matrix/vector notation, we write (9)–(10) as

$$\mathbf{r}_t = \mathbf{X}_t \mathbf{f}_t + \boldsymbol{\varepsilon}_t\,, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{D}_t)\,, \tag{11}$$

where $\mathbf{r}_t := (r_{t,1}, \ldots, r_{t,n})'$ is an n-dimensional random vector of cross-sectional security returns in excess of the risk-free rate, $\mathbf{X}_t$ is the $n \times k$ matrix of factor loadings known before time t, $\mathbf{f}_t$ is the k-dimensional random vector process of factors, $\boldsymbol{\varepsilon}_t$ is the n-dimensional random vector process of residual return, and

$$\mathbf{D}_t := \mathrm{diag}(\sigma_{t,1}^2, \ldots, \sigma_{t,n}^2), \quad \text{where all } \sigma_{t,j} > 0, \tag{12}$$

is the (strictly positive) $n \times n$ diagonal matrix of residual variances.

We note that $\mathbf{r}_t$ represents close-to-close total returns from $t-1$ to t, when t denotes a day, or when more precision is required, the exact time of the equity market close on day t. The subscript "t" has the same meaning for the other random processes.

We assume the $\mathbf{f}_t$-process has finite first and second moments such that

$$\mathbb{E}[\mathbf{f}_t] = \boldsymbol{\mu}_{\mathbf{f}}\,, \text{ and } \mathbb{V}[\mathbf{f}_t] = \mathbf{F}_t\,, \tag{13}$$

where the elements of $\boldsymbol{\mu}_{\mathbf{f}}$ are referred to as the *factor risk premia*. From (11) and (13) we immediately obtain that the expected security returns and their covariance matrix are given by

$$\mathbb{E}[\mathbf{r}_t] = \mathbf{X}_t \boldsymbol{\mu}_{\mathbf{f}} \,, \tag{14}$$

$$\mathbb{V}[\mathbf{r}_t] = \mathbf{D}_t + \mathbf{X}_t \mathbf{F}_t \mathbf{X}_t' =: \boldsymbol{\Sigma}_t \,, \tag{15}$$

where $\mathbf{X}_t'$ denotes the transpose of $\mathbf{X}_t$. We note that (15) and the positivity assumption in (12) together imply that the security-level inverse $\boldsymbol{\Sigma}_t^{-1}$ exists. Classical MVO uses the security-level covariance matrix and expected returns. As is well-known, the historical sample covariance matrix is a poor estimator and should not be used for MVO. For instance, if we consider a universe that is similar to the Russell 3000, then one has $n = 3000$ securities and $n(n+1)/2 \approx 4,500,000$ independent elements in the covariance matrix, which is too many independent parameters to estimate from daily security returns. In this situation, APT models are especially useful in that they provide the covariance matrix estimator (15) which, unlike the sample covariance matrix, leads to stable, well-diversified optimal portfolios (Ritter, 2016).

We state without proof that the inverse of (15) can be obtained by the *Woodbury matrix inversion lemma*, given by

$$\boldsymbol{\Sigma}_t^{-1} = \mathbf{D}_t^{-1} - \mathbf{Z}_t[\mathbf{F}_t^{-1} + \mathbf{X}_t'\mathbf{Z}_t]^{-1}\mathbf{Z}_t' \,, \tag{16}$$

where $\mathbf{Z}_t := \mathbf{D}_t^{-1}\mathbf{X}_t$. The key to making computations (16) run quickly is to systematically avoid the bottleneck of creating and manipulating a full $n \times n$ matrix (as opposed to $n \times k$ or $k \times k$ matrices where k is much smaller than n). For example in the Markowitz case one should compute

$$\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t = \mathbf{D}_t^{-1}\boldsymbol{\mu}_t - \mathbf{Z}_t[\mathbf{F}_t^{-1} + \mathbf{X}_t'\mathbf{Z}_t]^{-1}\mathbf{Z}_t'\boldsymbol{\mu}_t \,, \tag{17}$$

observing that the right-hand side can be calculated *without* having to build a full $n \times n$ matrix.

3.2. **Black-Litterman for Factor Models.** Leveraging the general BLB approach in Kolm and Ritter (2017), we describe a BL-type model for factor models of the kind in (9)–(10) above. For simplicity we treat $\mathbf{F}_t$ as known. Therefore, once $\mathbf{X}_t, \mathbf{D}_t$ are given, then $\boldsymbol{\Sigma}_t$ is also known via (15), just as in the original BL model. It follows that the source of uncertainty is in the expected factor premia $\boldsymbol{\mu}_{\mathbf{f}}$. Next, we discuss two kinds of priors for factor premia.

3.2.1. *Data-Driven Priors.* If the $\mathbf{f}_t$-process is stationary, such that $\boldsymbol{\mu}_{\mathbf{f}}$ and $\mathbf{F}$ are approximately constant over time, then we can obtain a prior, $\boldsymbol{\pi}_{\mathbf{f}}$, from a (Bayesian) time-series model of the factor returns $\mathbf{f}_t$. One possible choice is the mean (or rolling mean) of a time-series of the OLS estimates

$$\widehat{\mathbf{f}}_t = (\mathbf{X}_t'\mathbf{X}_t)^{-1}\mathbf{X}_t'\mathbf{r}_{t+1}\,. \tag{18}$$

More sophisticated approaches such as that of hierarchical or mixed-effects models can also be used (Gelman, Carlin, Stern, and Rubin, 2003).

An advantage of data-driven priors is that they do not require a benchmark portfolio (such as the capitalization-weighted benchmark in the original BL model). Consequently, data-driven priors may be preferred in absolute return strategies where the effective benchmark is cash. See below for an empirical example which uses a data-driven prior in a hypothetical absolute-return strategy.

3.2.2. *Benchmark Priors.* If there is a benchmark portfolio, $\mathbf{h}_B$, then close in spirit to the original BL model is to find a benchmark-optimal prior. For a general discussion of benchmark-optimal priors in this setting, we refer the reader to Kolm and Ritter (2017). For instance, we could use the prior

$$\boldsymbol{\pi}_{\mathbf{f}} \sim \mathcal{N}(\boldsymbol{\xi}, \mathbf{V}) \tag{19}$$

where $\boldsymbol{\xi}$ is a k-dimensional parameter vector and $\mathbf{V}$ is a $k \times k$ matrix that is symmetric and positive definite.

Kolm and Ritter (2017) show that by combining this prior with the factor model (9)–(10), one obtains the *a priori* expected returns and return covariance matrix

$$\mathbb{E}_\pi[\mathbf{r}] = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Sigma}^{-1}\mathbf{X}\mathbf{H}^{-1}\mathbf{X}'\boldsymbol{\Sigma}^{-1})^{-1}\boldsymbol{\Sigma}^{-1}\mathbf{X}\mathbf{H}^{-1}\mathbf{V}^{-1}\boldsymbol{\xi}\,, \tag{20}$$

$$\mathbb{V}_\pi[\mathbf{r}] = (\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Sigma}^{-1}\mathbf{X}\mathbf{H}^{-1}\mathbf{X}'\boldsymbol{\Sigma}^{-1})^{-1}\,, \quad \text{with} \tag{21}$$

$$\mathbf{H} := \mathbf{V}^{-1} + \mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X}\,, \text{ and}$$

where all quantities are assumed to be sampled at the same time t. However, for readability purposes, in this section we omit the explicit time index, allowing the time-dependence to be implicit. Hence, the *a priori* mean-variance optimal portfolio is given by

$$\mathbf{h}_{\text{prior}} = (\lambda\mathbb{V}_\pi[\mathbf{r}])^{-1} \cdot \mathbb{E}_\pi[\mathbf{r}] = \lambda^{-1}\boldsymbol{\Sigma}^{-1}\mathbf{X}\mathbf{H}^{-1}\mathbf{V}^{-1}\boldsymbol{\xi}\,. \tag{22}$$

From (22) it is clear that unlike the original BL, any arbitrary benchmark portfolio *cannot* be realized as an *a priori* optimal portfolio. Those that can are necessarily of the form $\lambda^{-1}\mathbf{\Sigma}^{-1}\mathbf{\Pi}$, where $\mathbf{\Pi}$ is some linear combination of the columns of $\mathbf{X}$. These are precisely the portfolios which are optimal with respect to a set of individual risk premia that are in the factor model. As not every possible portfolio is realizable as *a priori* optimal, then the market portfolio may also not be. Nevertheless, if (a) the market is in a CAPM equilibrium and (b) one of the columns of $\mathbf{X}$ contains the CAPM betas, then the individual risk premia will be proportional to that column of $\mathbf{X}$ and the market portfolio is realizable as *a priori* optimal.

3.3. **Factor Views.** We are free to express views on factor risk premia in the same way we would express views in the original BL model. In particular, we can express views on any (linear) portfolios of risk premia. However, it less likely for portfolio managers to have views on linear combination of factors. Perhaps the most common and parsimonious situation is that we have a view on each factor risk premium that is independent of our views on other factors. Using the value and momentum factors as an example, a portfolio manager might have two independent views: (a) a view on the value premium, and (b) a view on the momentum premium. We can express views of this kind as

$$\mathbf{q} = \boldsymbol{\mu}_{\mathbf{f}} + \boldsymbol{\varepsilon}_{\mathbf{q}}, \quad \boldsymbol{\varepsilon}_{\mathbf{q}} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Omega}), \tag{23}$$

where $\mathbf{\Omega} = \mathrm{diag}(\omega_1^2, \ldots, \omega_k^2)$ is a diagonal $k \times k$ matrix with all $\omega_j > 0$.

3.4. **The Black-Litterman Expected Returns and Covariance Matrix for Factor Portfolios.** By combining the factor model (9)–(10), the factor prior (19) and the investors factor views (23), Kolm and Ritter (2017) show that the (posterior) expected returns and return covariance matrix are given by

$$\boldsymbol{\mu}_{\text{BL-fac}} = \mathbf{\Sigma}_{\text{BL-fac}} \cdot \mathbf{\Sigma}^{-1}\mathbf{X}(\widetilde{\mathbf{V}}^{-1} + \mathbf{X}'\mathbf{\Sigma}^{-1}\mathbf{X})^{-1}\widetilde{\mathbf{V}}^{-1}\widetilde{\boldsymbol{\xi}}, \tag{24}$$

$$\mathbf{\Sigma}_{\text{BL-fac}} = \left(\mathbf{\Sigma}^{-1} + \mathbf{\Sigma}^{-1}\mathbf{X}(\widetilde{\mathbf{V}}^{-1} + \mathbf{X}'\mathbf{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{\Sigma}^{-1}\right)^{-1}, \tag{25}$$

where all quantities are sampled at time t, and $\widetilde{\mathbf{V}} := (\mathbf{V}^{-1} + \mathbf{\Omega}^{-1})^{-1}$ and $\widetilde{\boldsymbol{\xi}} := \widetilde{\mathbf{V}}(\mathbf{V}^{-1}\boldsymbol{\xi} + \mathbf{\Omega}^{-1}\mathbf{q})$ are the posterior hyperparameters. We emphasize that $\mathbf{\Sigma}$ is the covariance matrix of security returns from the factor model (15). Consequently, from the expected returns (24) and covariance matrix (25), we obtain the mean-variance optimal portfolio

$$\mathbf{h}^* = \lambda^{-1}\mathbf{\Sigma}^{-1}\mathbf{\Pi}, \tag{26}$$

where

$$\mathbf{\Pi} := \mathbf{X}\widetilde{\boldsymbol{\mu}}_{\mathbf{f}}\,, \tag{27}$$

$$\widetilde{\boldsymbol{\mu}}_{\mathbf{f}} := (\mathbf{V}^{-1} + \mathbf{\Omega}^{-1} + \mathbf{X}'\mathbf{\Sigma}^{-1}\mathbf{X})^{-1}(\mathbf{V}^{-1}\boldsymbol{\xi} + \mathbf{\Omega}^{-1}\mathbf{q})\,. \tag{28}$$

The formulas (26)–(28) represent the solution to BLB optimization in the context of factor models. We observe that the security-level risk premia $\mathbf{\Pi} = \mathbf{X}\widetilde{\boldsymbol{\mu}}_{\mathbf{f}}$ are linear combinations of the factors which form the columns of $\mathbf{X}$. Intuitively, we can think of $\widetilde{\boldsymbol{\mu}}_{\mathbf{f}}$ as a set of factor risk premia "adjusted" to take account of the views, where these adjustments give more weight to factors which have high prior mean-variance ratios, ξ_i/V_{ii}, and/or high expected return-uncertainty ratios, q_i/ω_i^2.

3.5. **Discussion.** The original work of Black and Litterman was done under normality assumptions for the return distributions. This is mathematically convenient because the normal prior is a conjugate prior for the normal likelihood, so the posterior is also normal with all means and covariances obtainable in closed form. Nevertheless, it is of clear interest to investigate the applicability of similar methods beyond the normal distribution.

Chamberlain (1983) showed that under elliptically distributed return distributions, expected utility – for *any* concave utility function – is a function of *only* the portfolio's return and variance. Elliptical distributions is a large family of distributions, including many fat-tailed distributions such as the student's t-distributions. In the context of BLB optimization, this implies that we are free to use any form for the prior and the views – including models where the unknown parameter(s) need not simply be the mean return – as long as the final posterior distribution $p(\mathbf{r}_{t+1} \,|\, \mathbf{q})$ of the cross-sectional security return vector is a member of the elliptical family, then any risk-averse expected-utility maximizer will end up maximizing the posterior expected portfolio return minus a constant multiple of the posterior portfolio variance. Of course, with a different (e.g. non-normal) model the investor need not arrive at precisely the same expected return and variance formulas as did Black and Litterman. Additionally, if the prior is not conjugate to whatever likelihood is chosen, then computing the posterior moments may be difficult or even impossible.

One of the early successes of the Black-Litterman model was that it provided a form of regularization of the fundamentally-unstable inversion of the historical sample covariance matrix. However, the Black-Litterman technique is not primarily meant as a regularization technique; it has that property as a by-product because its likelihood and prior tend to pull the portfolios towards known, reasonable portfolios.

Importantly, with common usages of APT, regularization is strictly speaking no longer necessary, because the resulting covariance matrix is already stably invertible. For example, when the mean-variance portfolio is computed via (17), the phenomenon of "corner portfolios" does not occur. The resulting portfolios are typically well-diversified, as long as the factor model itself is reasonable.

We stress that even with a stable covariance matrix, the BL procedure is still tremendously valuable – perhaps even more so, since even the classical BL model does still use $\mathbf{\Sigma}$ as an input, while being less sensitive to its instabilities. However, as compared to MVO, the BL model gives the portfolio manager additional degrees of freedom to control the portfolio, in the form of ω_i which quantifies the uncertainty in the i-th view. A key difference is that the ω_i-parameters do not appear in the classical Markowitz formulation.

4. An Empirical Case Study

In this section we provide an empirical case study of the BLB model for factor models and investigate whether our methodology provide reasonable empirical results. Specifically, we construct an instance of BLB model (26)–(28) suitable for quantitatively-oriented institutional portfolio managers interested in systematically trading single-name U.S. equities in a factor-based fashion.

Our universe of stocks is chosen with certain liquidity thresholds in place. In particular, stocks must be priced above five dollars, bid-offer spreads must be less than hundred basis points, and must trade at least 10K shares and one million dollars of volume per day. We restrict attention to only common stock, and only the most liquid share class of each company. This typically results in a universe of 1800-2000 stocks, but in the period around the 2008 financial crisis the universe size briefly dropped to around 1400 stocks.

Our factor model includes a GICS-based industry classification with 58 industries which are either single GICS sub-industries, or several closely-related sub-industries merged together. Our model also includes a set of continuous numerical factors known colloquially as "style factors" because they are thought to represent investment styles such as value, contrarian, or dividend capture.

Roughly speaking, *factor timing* is defined as attempting to predict which factors will have superior factor returns over a forthcoming time period, with the goal of monetizing those views in the event they materialize. Any such predictions may be naturally translated into views on factors or combinations of factors.

In the following example, we simulate a hypothetical portfolio manager who possesses definite skill at factor timing. Then we illustrate how to use the factor-based BLB to construct a sequence of portfolios which successfully monetizes the portfolio manager's factor-timing views.

The views are expressed through (a) the k-dimensional vector $\mathbf{q}$ representing the expected returns on the respective factors, and (b) the $k \times k$ matrix $\mathbf{\Omega}$ which represent the uncertainties in those views. Specifically, we take $\mathbf{\Omega}$ to be diagonal, where $\Omega_{ii} = \omega_i^2$ and ω_i can be thought of as the width of a confidence interval around q_i, for each $i = 1, \ldots, k$. There are many possible uses of the concept of "uncertainty in a view" in the investment-management industry.

In order to simulate a manager who is skilled at factor timing, we assume that at the very beginning of our historical data, this portfolio manager made a single call regarding the expected returns and uncertainties on a handful of factors, and has no views on the remaining factors. In particular, the portfolio manager chooses $\mathbf{q}$ and $\mathbf{\Omega_f}$ at the beginning of 2007, and remains steadfast in those views until the present time. With the benefit of hindsight, we simulate the performance of this hypothetical portfolio manager's strategy, where some of their factor bets turned out to be correct. We strongly stress that this study is meant to simulate what happens *in the event the portfolio manager made the right call at the beginning*, as a way of illustrating the efficacy of the BLB portfolio construction with factors. In this article, we neither simulate nor provide guidance on the mechanism of factor timing itself.

Specifically, we assume that the portfolio manager decided, at the beginning of the sample period (January 2007), has decided to harvest the size (the small-cap effect) and liquidity premia. In addition, the portfolio manager also decided to assign negative premia to firms with high short interest, and a positive premium to a quality factor based on accounting metrics. The directions of these four factor bets align with well-known results from behavioural finance. In Figure 1, we depict the cumulative factor return performance these four factors over the 2007-2020 sample period. Further, we assume that the portfolio manager was prescient in deciding to take no explicit view on the well-known and much-studied value and momentum factors, neither of which performed well in the U.S. equity market over the sample period of 2007-2020. We refer to this portfolio manager's views as the *prescient views*, because they benefit from having taken no explicit view on value and momentum (within the US equity market) over the period.

FIGURE 1. Cumulative Factor Return Performance.

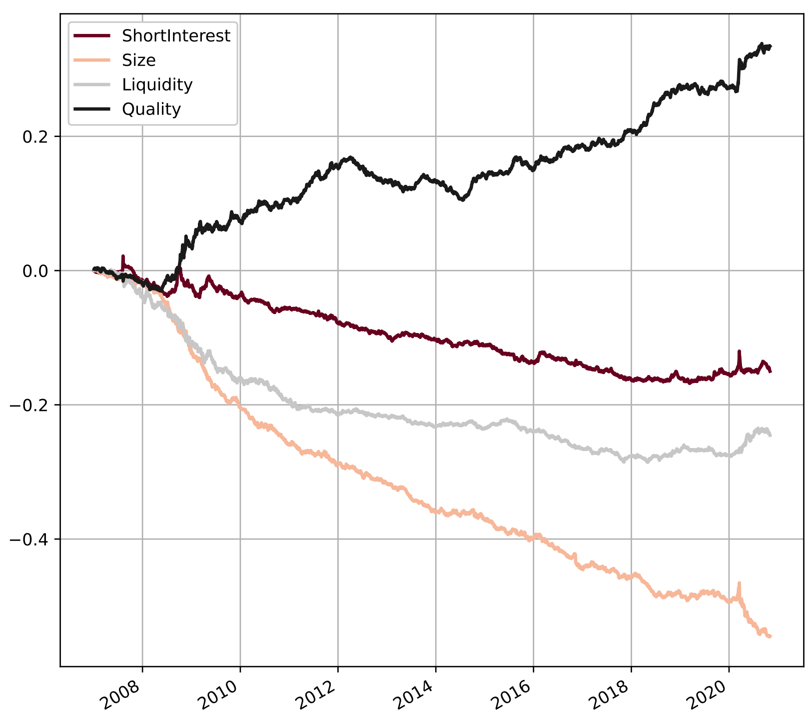

Note: Cumulative factor return performance for the four factors which we selected for inclusion in the prescient views, over the sample period, 2007-2020. The y-axis is the cumulative sum of $\widehat{\mathbf{f}}_t$ *for all dates* t *between 2007-01-01 and 2020-11-01.*

The procedure for our empirical study is as follows:

(1) For each day t, construct the matrix of factor loadings $\mathbf{X}_t$, also called the *design matrix*, by loading the data as a pandas data frame, transforming the loadings, and applying the function `patsy.dmatrix` with the appropriate model formula. The transformation entails reshaping the raw loading (e.g. market cap) to be Gaussian and centered within the estimation universe.

(2) For each day t, estimate the factor returns $f_{t,j}$ using ordinary least-squares as $\widehat{\mathbf{f}}_t = (\mathbf{X}_t'\mathbf{X}_t)^{-1}\mathbf{X}_t'\mathbf{r}_{t+1}$.

(3) Choose an in-sample period for the factor covariance matrix. We choose 2007-2015 to be in-sample with respect to the factor covariance, $\mathbf{F}$.

(4) Estimate the factor covariance $\mathbf{F}$ as the sample covariance of $\widehat{\mathbf{f}}_t$ with t restricted to the in-sample period. We separate the calculation of the factor covariance into volatility estimation, and correlation estimation. The factor covariance is given by $\mathbf{F} = \mathbf{SRS}$ where $\mathbf{S}$ is a diagonal matrix containing the factor volatilities, and $\mathbf{R}$ is the factor correlation matrix. Note that $\mathbf{R}$ was estimated in Python by the method `numpy.corrcoef` applied to the coefficients $\widehat{\mathbf{f}}_t$ estimated in a previous step. We mitigate the effect of estimation error in $\mathbf{F}^{-1}$ by applying a lower bound of 0.01 to the factor volatilities $\mathbf{S}$.

(5) For a market-neutral study, we set our prior to $\boldsymbol{\xi} = \mathbf{0}$, and $\mathbf{V} = \mathbf{SR}_s\mathbf{S}$ where $\mathbf{R}_s$ denotes the Ledoit-Wolf shrinkage estimator for the correlation matrix (see, Ledoit and Wolf (2004)). Under this prior, an all-cash benchmark portfolio is *a priori* optimal. Practitioners may wish to choose a prior under which a different benchmark, such as the S&P 500, is optimal; details of this procedure are in the discussion of eqn. (29).

(6) We set $\mathbf{q}$ and $\boldsymbol{\Omega}$ to be consistent with the prescient views outlined above. Specifically, we specify one view per factor. For all industry and style factors not explicitly represented in the prescient views, we set $q_i = 0$ and $\omega_i = 0.05$, representing no view on their expected return with a high level of uncertainty. For the factors involved in the prescient views, we set q_i equal to negative three basis points for size, negative one basis point for liquidity, negative one-half basis point for short interest, and positive one basis point for quality. We also set a premium of one basis point for the intercept factor.

(7) For each t, produce forward-looking estimates of $\mathbf{D}_t$ using historical residual vols blended with implied vols.
(8) For each t, compute $\mathbf{Z}_t := \mathbf{D}_t^{-1}\mathbf{X}_t$ without expanding $\mathbf{D}_t^{-1}$ as a full $n \times n$ matrix, by multiplying each i-th row of $\mathbf{X}_t$ by D_{ii}^{-1}. Note that $\mathbf{Z}_t$ has the same dimensions as $\mathbf{X}_t$, both are $n \times k$ matrices. Hence as discussed above, $\mathbf{Z}_t$ is easier to compute with than a $n \times n$ matrix.
(9) For each t, define the intermediate calculation result $\mathbf{K}_t := \mathbf{\Sigma}_t^{-1}\mathbf{X}_t$ and compute

$$\begin{aligned}
\mathbf{K}_t &= \mathbf{Z}_t - \mathbf{Z}_t[\mathbf{F}^{-1} + \mathbf{X}_t'\mathbf{Z}_t]^{-1}\mathbf{Z}_t'\mathbf{X}_t \\
\widetilde{\boldsymbol{\mu}}_{\mathbf{f}} &= (\mathbf{V}^{-1} + \mathbf{\Omega}^{-1} + \mathbf{X}_t'\mathbf{K}_t)^{-1}(\mathbf{V}^{-1}\boldsymbol{\xi} + \mathbf{\Omega}^{-1}\mathbf{q}) \\
\mathbf{h}_t^* &= \lambda^{-1}\mathbf{K}_t\widetilde{\boldsymbol{\mu}}_{\mathbf{f}}
\end{aligned}$$

(10) Finally, we compute the portfolio return $r_{p,t+1} = (\mathbf{h}_t^*)'\mathbf{r}_{t+1}$.

Although the BLB method for factor models as presented here is not much more complicated than a standard MVO approach to the same problem, we feel that any added complexity should be justified through added performance. Therefore to gauge the value-add by the portfolio construction technique, we also include results for an MVO baseline which is proportional to $\mathbf{\Sigma}_t^{-1}\boldsymbol{\mu}_t$ where $\boldsymbol{\mu}_t = \mathbf{X}_t\mathbf{q}$. Note that the baseline includes $\mathbf{q}$ and hence also benefits from the prescient views, but does not benefit from the additional Bayesian portfolio construction of BLB. Of course for the baseline computation, we use (16); hence our method does not entail the (notoriously unstable) security-level sample covariance matrix.

In Figure 2, we show a comparison between the BLB and MVO approaches. We have scaled the two portfolios so that they realize the same volatility over the sample period. The BLB portfolio sequence achieves a Sharpe ratio of 1.16 while the MVO portfolio is about 0.9. There are several reasons for the difference, but perhaps the most important difference is the presence of $\mathbf{\Omega}$ in the BLB portfolio. Also, the BLB portfolio is higher gross market value but achieves lower exposure to factors on which the manager has no views, in part because of the shrinkage towards $\mathbf{V}^{-1}\boldsymbol{\xi}$. In any case, this illustrates that the BLB method provides the manager with additional degrees of freedom which have the potential to reduce risk relative to MVO optimized portfolios also benefiting from the same prescient factor bets.

The average turnover of the two strategies we consider is comparable; each turns over an average of about 10-15% of gross market value per day. Such high turnover is unsurprising as the portfolio formation process is unaware of costs. The turnover is mostly driven by changes in the factor exposures themselves, and changes in the residual variance forecasts. This is likely too much turnover to be directly implementable, but the t-cost aware optimization that we discuss in the conclusions section would be implementable.

Figure 2. Comparison of BLB and MVO Simulations.

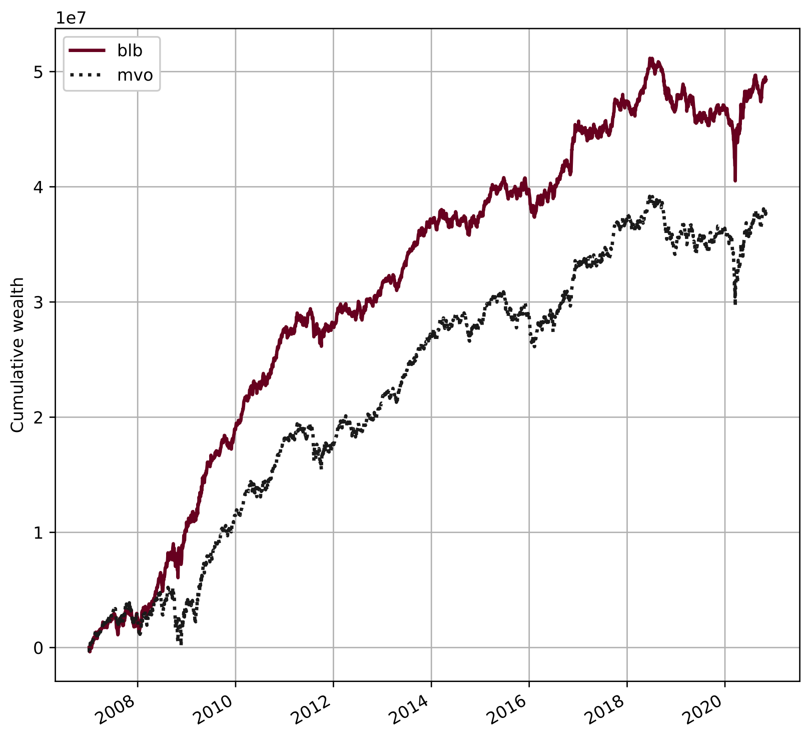

Note: Cumulative wealth from simulations of the MVO (dashes) and BLB (solid) multi-factor portfolios with prescient factor views and data-driven priors. The portfolios are scaled such that they realize the same volatility over the sample period, 2007-2020.

Although we shall not do so here, it is possible to run the above procedure with a benchmark which is approximately the market-capitalization weighted benchmark with holdings vector $\mathbf{h}_\mathrm{B}$. This is achieved by choosing a normal prior $\mathcal{N}(\boldsymbol{\xi}, \mathbf{V})$ having the property that the benchmark portfolio is *a priori* optimal, meaning that the benchmark is optimal under the expected returns and covariance matrix described by (20) and (21).

Mathematically, referring to (22), this amounts to solving for $\boldsymbol{\xi}$ in the equation

$$\lambda \boldsymbol{\Sigma}_t \mathbf{h}_\mathrm{B} = \mathbf{X}_t \mathbf{H}^{-1} \mathbf{V}^{-1} \boldsymbol{\xi}\,. \tag{29}$$

Solving for $\boldsymbol{\xi}$ in this way might not be possible exactly, in which case one may solve for it in the least-squares sense.

5. Conclusions

In this article, we proposed a general framework referred to as Black-Litterman-Bayes (BLB) for constructing optimal portfolios for factor-based investing. In the spirit of the classical Black-Litterman model, this framework allows for the incorporation of investor views and different priors on factor risk premia, including data-driven and benchmark priors. We provided computationally efficient closed-form formulas for the (posterior) expected returns and return covariance matrix that result from integrating factor views into an APT multi-factor model.

Finally, in a step-by-step procedure, we showed how to build the prior and incorporate the factor views, demonstrating in a realistic empirical example using a number of well-known cross-sectional U.S. equity factors, that the BLB approach can add value to mean-variance optimal multi-factor risk premia portfolios.

In closing, we make a few observations regarding applying these results in transaction-cost-aware MVO. The original BL model and BLB extension with factor views that we have presented here, share in common that the portfolio construction step is a standard mean-variance optimization. The BL and BLB methods are based on adjusting the posterior security return density (where only the first two posterior moments are needed, in the case of elliptical distributions). In particular, we denote by $\mathbf{h}_0$ the starting portfolio and $c(\mathbf{h}_0, \mathbf{h})$ the anticipated cost (predicted before trading begins) to trade from $\mathbf{h}_0 \to \mathbf{h}$. For institutional asset managers, $c(\mathbf{h}_0, \mathbf{h})$ is typically dominated by slippage costs due temporary and permanent impact. Nevertheless, the optimization problem for BLB with costs is

$$\max_{\mathbf{h}} \; \mathbf{h}' \boldsymbol{\mu}_{\text{BL-fac}} - \frac{\lambda}{2} \mathbf{h}' \boldsymbol{\Sigma}_{\text{BL-fac}} \mathbf{h} - c(\mathbf{h}_0, \mathbf{h})$$

where notation refers to eqns. (24) and (25).

References

Aliaga-Diaz, Roger, Giulio Renzi-Ricci, Ankul Daga, and Harshdeep Ahluwalia (2020). "Portfolio Optimization with Active, Passive, and Factors: Removing the Ad Hoc Step". In: *The Journal of Portfolio Management* 46.4, pp. 39–51.

Ang, Andrew (2014). *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen (2013). "Value And Momentum Everywhere". In: *The Journal of Finance* 68.3, pp. 929–985.

Avramov, Doron and Guofu Zhou (2010). "Bayesian Portfolio Analysis". In: *Annu. Rev. Financ. Econ.* 2.1, pp. 25–47.

Baltussen, Guido, Laurens Swinkels, and Pim Van Vliet (2019). *Global Factor Premiums*. Tech. rep.

Bass, Robert, Scott Gladstone, and Andrew Ang (2017). "Total Portfolio Factor, Not Just Asset, Allocation". In: *The Journal of Portfolio Management* 43.5, pp. 38–53.

Bender, Jennifer, Jerry Le Sun, and Ric Thomas (2018). "Asset Allocation vs. Factor Allocation—Can We Build a Unified Method?" In: *The Journal of Portfolio Management* 45.2, pp. 9–22.

Bergeron, Alain, Mark Kritzman, and Gleb Sivitsky (2018). "Asset Allocation and Factor Investing: An Integrated Approach". In: *The Journal of Portfolio Management* 44.4, pp. 32–38.

Black, Fischer and Robert Litterman (1992). "Global Portfolio Optimization". In: *Financial Analysts Journal*, pp. 28–43.

Black, Fischer and Robert B. Litterman (1990). *Asset Allocation: Combining Investor Views with Market Equilibrium*. Tech. rep.

– (1991a). "Asset Allocation: Combining Investor Views with Market Equilibrium". In: *The Journal of Fixed Income* 1.2, pp. 7–18.

– (1991b). *Global Asset Allocation with Equities, Bonds and Currencies*. Tech. rep.

Carhart, Mark M. (1997). "On Persistence in Mutual Fund Performance". In: *The Journal of Finance* 52.1, pp. 57–82.

Chamberlain, Gary (1983). "A Characterization of the Distributions That Imply Mean—Variance Utility Functions". In: *Journal of Economic Theory* 29.1, pp. 185–201.

Dopfel, Frederick E. and Ashley Lester (2018). "Optimal Blending of Smart Beta and Multifactor Portfolios". In: *The Journal of Portfolio Management* 44.4, pp. 93–105.

Fabozzi, Frank J., Sergio M. Focardi, and Petter N. Kolm (2006). "Incorporating Trading Strategies in the Black-Litterman Framework". In: *The Journal of Trading* 1.2, pp. 28–37.

– (2010). *Quantitative Equity Investing: Techniques and Strategies*. John Wiley & Sons.

Fabozzi, Frank J., Petter N. Kolm, Dessislava A. Pachamanova, and Sergio M. Focardi (2007). *Robust Portfolio Optimization and Management*. John Wiley & Sons.

Fama, Eugene F. and Kenneth R. French (1993). "Common Risk Factors in the Returns On Stocks And Bonds". In: *Journal of Financial Economics* 33.1, pp. 3–56.

Figelman, Ilya (2017). "Black–Litterman with a Factor Structure Applied to Multi-Asset Portfolios". In: *The Journal of Portfolio Management* 44.2, pp. 136–155.

Frazzini, Andrea and Lasse Heje Pedersen (2014). "Betting Against Beta". In: *Journal of Financial Economics* 111.1, pp. 1–25.

Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin (2003). *Bayesian Data Analysis*. 2nd. Taylor & Francis.

Johnson, Ben (2018). *Strategic Beta ETF Market Still Expanding*. URL: `https://www.morningstar.co.uk/uk/news/182188/strategic-beta-etf-market-still-expanding.aspx`.

Jones, Robert C., Terence Lim, and Peter J. Zangari (2007). "The Black-Litterman Model for Structured Equity Portfolios". In: *The Journal of Portfolio Management* 33.2, pp. 24–33.

Keloharju, Matti, Juhani T. Linnainmaa, and Peter Nyberg (2016). "Return Seasonalities". In: *The Journal of Finance* 71.4, pp. 1557–1590.

Koijen, Ralph S.J., Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt (2018). "Carry". In: *Journal of Financial Economics* 127.2, pp. 197–225.

Kolm, Petter N. and Gordon Ritter (2017). "On the Bayesian Interpretation of Black-Litterman". In: *European Journal of Operational Research* 258.2, pp. 564–572.

Ledoit, Olivier and Michael Wolf (2004). "Honey, I Shrunk The Sample Covariance Matrix". In: *The Journal of Portfolio Management* 30.4, pp. 110–119.

Litterman, Robert and Guangliang He (1999). *The Intuition Behind Black-Litterman Model Portfolios*. Tech. rep.

Melas, Dimitris, Zoltán Nagy, Navneet Kumar, and Peter Zangari (2019). "Integrating Factors in Market Indexes and Active Portfolios". In: *The Journal of Portfolio Management* 45.6, pp. 16–29.

Meucci, Attilio (2009). *Risk And Asset Allocation*. Springer Science & Business Media.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen (2012). "Time Series Momentum". In: *Journal of Financial Economics* 104.2, pp. 228–250.

Pearl, Judea (2014). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Ritter, Gordon (2016). "Stable Linear-Time Optimization in Arbitrage Pricing Theory Models". In: *Risk Magazine*.

Roll, Richard and Stephen A. Ross (1980). "An Empirical Investigation of the Arbitrage Pricing Theory". In: *The Journal of Finance* 35.5, pp. 1073–1103.

Ross, Stephen A (1976). "The Arbitrage Theory of Capital Asset Pricing". In: *Journal of Economic Theory* 13.3, pp. 341–360.

Rubesam, Alexandre and Soosung Hwang (2019). *Do Smart Beta ETFs Capture Factor Premiums? A Bayesian Perspective*. Tech. rep. IÉSEG Working Paper Series 2018-ACF-04.

Satchell, Stephen and Alan Scowcroft (2000). "A Demystification of the Black–Litterman Model: Managing Quantitative and Traditional Portfolio Construction". In: *Journal of Asset Management* 1.2, pp. 138–150.

Subrahmanyam, Avanidhar (2010). "The Cross-Section of Expected Stock Returns: What Have We Learnt From the Past Twenty-Five Years of Research?" In: *European Financial Management* 16.1, pp. 27–42.

Zhou, Guofu (2009). "Beyond Black–Litterman: Letting the Data Speak". In: *The Journal of Portfolio Management* 36.1, pp. 36–45.