Portfolio Tilts Using Views on Macroeconomic Regimes

Redouane Elkamhi, Jacky S. H. Lee, and Marco Salerno

Redouane Elkamhi
is an associate professor of finance in the Rotman School of Management at the University of Toronto in Toronto, Ontario, Canada.
redouane.elkamhi@rotman.utoronto.ca

Jacky S. H. Lee
is the senior managing director of the Total Portfolio Group at Healthcare of Ontario Pension Plan Trust Fund in Toronto, Ontario, Canada.
jlee5@hoopp.com

Marco Salerno
is a principal of the Total Portfolio Group at Healthcare of Ontario Pension Plan Trust Fund in Toronto, Ontario, Canada.
msalerno@hoopp.com

KEY FINDINGS

- We provide a methodology to tilt a portfolio according to investors' views on the likelihood of economic regimes rather than expected returns.
- We compare the stability of our methodology with the Black–Litterman model for different levels of errors in expected returns and likelihood of economic regimes. We find that our methodology is more robust to input errors.
- Our methodology ensures consistency across asset classes and leads to less extreme asset weights, which is a desirable feature in practical applications.

ABSTRACT

Long-term investors tilt their portfolios given their views on the evolving investment landscape. In the literature, portfolio tilting is often implemented with methodologies that use investors' views on point estimates of conditional assets' expected returns. These conditional return expectations are notoriously difficult to estimate, and using them often results in unstable portfolio weights when existing methodologies are applied. We avoid such shortcomings by providing a methodology that incorporates views on the likelihood of economic regimes (e.g., growth and inflation surprises) instead. Using data on equities, bonds, and commodities, we show—both in simulation and empirically—that our approach generates stable portfolio weights and outperformance that is minimally affected by forecast errors.

Strategic asset allocation refers to the set of decisions that aligns a portfolio's asset mix (e.g., weights of bonds, equities, and commodities) with the investors' long-term investment goals and objectives.[1] Strategic asset allocation is the key driver of portfolio returns for long-term institutional investors, such as pension and sovereign wealth funds. These investors often thrive to build portfolios that are well diversified. However, occasionally their views on the investment landscape change substantially that necessitate material portfolio shifts. How to translate investors' views on economic covariates to asset allocation decisions is the central question addressed by this article.

There exist several methodologies that allow investors to tilt their portfolios. Black and Litterman (1992) (BL hereafter) was the first methodology of its kind that

[1] Differently from tactical asset allocation, strategic asset allocation does not change frequently (e.g., once or twice a year maximum) because the goal is not to chase trends in the market but rather to position the portfolio to achieve the desired goal.

sparked a myriad of extensions (e.g., Jones, Lim, and Zangari 2007; Qian 2011).[2] Notwithstanding these advancements, existing methodologies require investors to express their views as point estimates of expected returns or covariances, or both.[3] Although covariances can be estimated reasonably well, it is known that estimating expected returns is subject to large estimation errors (Merton 1980), which lead to unstable portfolio weights (Best and Grauer 1991; Britten-Jones 1999; DeMiguel, Garlappi, and Uppal 2009). Even more importantly, views expressed as point estimates of expected returns across asset classes could imply inconsistent risk premiums when analyzed by asset pricing models. Such inconsistencies and estimation biases represent the Achilles' heel of most existing methodologies.

Our main contribution to the literature is to develop a novel methodology to incorporate investors' views on the likelihood of economic regimes, instead of point estimates of returns or covariances, to enhance the investment decision process. Our approach may be more aligned with what chief investment officers and asset managers consider for asset allocation, which often are related to the degree that asset prices have reflected the current or projected macroeconomic environments. Our work contributes to the literature on regime-based asset allocation (e.g., Sheikh and Sun 2012; Nystrup et al. 2015, 2017; Nystrup, Madsen, and Lindström 2018; Schmieder and Kollár 2020; Zheng, Xu, and Zhang 2021), but we differentiate from previous studies in two major ways: (1) our framework requires investors to provide their views in terms of probabilities of economic regimes, rather than point estimates of assets' expected returns and covariances; (2) our framework separates investors' views from the methodologies for portfolio construction.

Although it is necessary that investors need ways to come up with views on economic regimes that they define, such discussion is not the focus of this article. Our aim is to provide a framework for tilting portfolios given views on the likelihood of economic regimes, similar to how BL provided a framework for tilting portfolios given views on assets' expected returns.

The benefits of our approach can be summarized as follows. First, although our methodology requires the calculation of historical returns and covariances, different from existing methodologies our approach does not necessitate investors to provide forecasts on expected returns or covariances. Second, our framework greatly reduces or eliminates the curse of dimensionality because the number of economic regimes is much smaller than the number of assets for which investors can have views on. Last, by separating investors' views from portfolio construction, our methodology allows investors to use their preferred portfolio construction rule. The latter feature is a notable improvement in flexibility over existing methodologies, which are tied to a specific allocation rule. For instance, BL and its extensions are tied to mean–variance optimization.

Our methodology consists of four steps: Investors (1) define the economic regimes based on their preferred market or macroeconomic variables and determine their likelihood (prior probabilities), which can be done using historical estimates; (2) form their views on the probabilities of being in each regime in the future; (3) compute the posterior probabilities of economic regimes by combining their views with the prior probabilities; and (4) calculate the conditional historical moments required by their preferred allocation strategy (e.g., mean–variance, risk parity, etc.) and combine them using the posterior probabilities across economic regimes. For our methodology to work well, the characteristics of asset returns across those economic regimes need to be considerably different.

[2] More recently, the literature has extended further to provide ways to allow investors to incorporate views using forecasts of factor returns and covariances (e.g., Figelman 2017; Kolm and Ritter 2020).

[3] Throughout this article we refer to views, estimates, and forecasts of expected returns interchangeably.

We start our analysis by comparing the theoretical performance of our framework with the BL model using Monte Carlo simulations. Our goal with this controlled experiment is to analyze how portfolio performance is affected by forecast errors on expected returns (and consequently the BL model) versus forecast errors on the likelihood of economic regimes (and consequently our approach). To focus on the effects of forecast errors, we let investors have correct forecasts on average across simulations, but the forecasts made in each simulation are subject to forecast errors with respect to the true values. We demonstrate that our approach outperforms BL on two fronts. First, our model produces higher out-of-sample portfolio Sharpe ratios (SRs) on average. Second, the distribution of out-of-sample SRs for our model has significantly less variance. In other words, for our model the information ratio of the out-of-sample portfolio SRs is higher.

Our findings show that even small forecast errors in BL can lead to large variation in portfolio weights. Our approach leads to relatively more stable portfolio weights, which contributes to the higher average out-of-sample SR. For both models, we compute the levels of estimation error associated with different out-of-sample portfolio SRs. We find that, for a given portfolio SR, the BL model requires much more precise views in order to perform as well as our framework.[4] These simulation results suggest that, if investors have prescient views, using our methodology should lead to a more robust out-of-sample performance compared to BL. This is our key analytical result.

Next, we provide a simple example to demonstrate a practical application of our methodology. We define economic regimes based on two fundamental macroeconomic variables: inflation and real gross domestic product (GDP) growth because they are well-documented systemic factors that affect asset returns (see Campbell and Viceira 2001; Bansal and Yaron 2004; Bansal and Shaliastovich 2013; Kung 2015; Boons et al. 2020). Critically, we highlight that (a) our choice of variables is meant as an example only and may not span all relevant future economic regimes, and (b) investors, depending on their beliefs and utilities, may define their regimes differently. We consider four different regimes based on rising and falling inflation and economic growth combinations.[5] To apply our framework, investors form views on the probabilities of these four economic regimes. Our framework endogenously translates such views to changes in the portfolio weights without imposing an additional structure to link assets and macroeconomic factors.

We use data on US equities, US fixed income, and a commodity index to conduct our empirical example. Our demonstration shows that, when investors have prescient views on the probability of economic regimes one-year ahead, a mean–variance portfolio with views outperforms that without views. Our example shows that, if investors have prescient views on economic regimes (but not expected returns), they can improve the SR of their portfolio by applying our methodology.

A FRAMEWORK BASED ON THE LIKELIHOOD OF REGIMES

Key Differences between Our Framework and the Existing Literature

Our framework differentiates from the existing literature (e.g., Black and Litterman 1992; Jones, Lim, and Zangari 2007; Qian 2011; Figelman 2017; Kolm and Ritter 2020)

[4] As an example, for a BL view that has an uncertainty of 1/264 of the asset covariances, our model achieves the same average portfolio SR with a 10% error in a view for the regime probabilities.

[5] We define our regimes based on surprises with respect to the expected value. For example, *rising growth and rising inflation* is defined as the regime in which both inflation and economic growth are higher than expected.

because investors are not asked to forecast expected returns or covariances in order to embed their views in portfolio construction. We develop a methodology that allows investors to incorporate their views on the likelihood of economic regimes. For ease of exposition, we choose to use economic growth and inflation as an example throughout this article.[6]

For a given set of macroeconomic factors (e.g., economic growth and inflation) investors can have views on the likelihood of observing a particular regime (e.g., rising/falling economic growth or inflation or any combination of them). In other words, we shift the paradigm from forecasting expected returns to forecasting likelihood of regimes. Our article is relevant for all those investors that cannot (or are not confident to) forecast point estimates of expected returns but are confident in their ability to express the likelihood of certain economic regimes to happen.

There is another potential reason why our approach might be preferred by some investors. While using the BL model, investors could provide views that are inadvertently not consistent with each other. For example, private equity, public equity, and infrastructure are three different asset classes that are all positively exposed to economic growth. Investors might have views on one asset class (e.g., private equity) that imply a higher economic growth forecast, whereas their view on a different asset class (e.g., infrastructure) implies a lower economic growth forecast. This inconsistency could lead to unrealistic portfolio weights (e.g., a large positive private equity weight versus a low or negative infrastructure weight). As we explain in detail later in the article, using our methodology investors would avoid such concerns because the relations between assets classes (e.g., private equity and infrastructure) are consistent because they are anchored to their historical estimates.

This article presents one example of how these regimes can be defined but various alternative definitions can be employed. Investors can apply a different model in splitting the world into regimes. Our framework would work as long as (1) assets exhibit different characteristics in the various regimes historically (e.g., higher returns in one regime versus the other), (2) these characteristics are believed to persist into the future, and (3) investors have the ability to forecast the likelihood of these economic regimes. The following sections describe how the prior likelihood of regimes, investors' views, and the posterior likelihood of regimes are computed. Finally, we describe how these likelihoods translate into asset allocation.

Prior Probabilities on Regime Likelihood

We elect to use surprises of inflation and real GDP growth against their market expectations to define our macroeconomic regimes. Our choice for this example is guided by evidence that both inflation and real GDP growth are main drivers of asset returns (see, e.g., Campbell and Viceira 1999; Campbell and Vuolteenaho 2004). In addition, it is well-known that investors consider both the levels and future possible trajectories of real GDP growth and inflation in their investment decisions, as is clear from the amount of time that financial news and commentators spend on discussing these two topics daily. In other words, asset returns have different characteristics in regimes characterized by different inflation and growth surprises.

Let x_t and π_t denote the actual real GDP growth rate and inflation rate, respectively. We define $\hat{x}_t \equiv \mathbb{E}(x_t | I(t-1))$ and $\hat{\pi}_t \equiv \mathbb{E}(\pi_t | I(t-1))$ the expectations about x_t and

[6] Economic growth and inflation are not the only variables that can be used to define economic regimes. Our framework can be applied to any set of variables that allow investors to discriminate returns across different economic regimes and for which investors are confident in their forecasting abilities. For example, if economic regimes defined on the Industrial Production Index discriminate returns well (i.e., there is variation in average returns across regimes), then this variable could be used.

π_t of an investor conditional on the information available until time $t-1$. We define $S_{x,t} \equiv x_t - \hat{x}_t$ as the surprise for the real GDP growth rate at time t; similarly, define $S_{\pi,t} \equiv \pi_t - \hat{\pi}_t$ as the surprise for the inflation rate at time t. For ease of notation, we drop the subscript t for the remaining of this section.

We assume the prior probability distribution of the surprises $s = [S_\pi\ S_x]'$ is a bivariate normal

$$\mathbb{P}(s) = \mathcal{N}(\mu, \Sigma) \tag{1}$$

where μ is the vector of prior means of surprises and Σ is their variance-covariance matrix. Based on the distribution defined in Equation 1, we can calculate the probability of being in each of the four economic regimes based on the direction (e.g., positive or negative) of the surprises. For example, the probability of being in the regime in which both inflation and GDP surprises are negative is

$$\Pr(S_\pi < 0\ \ \&\ \ S_x < 0) = F_s(0,0)$$

where $F_s(\cdot)$ is the cumulative density function of the bivariate normal with mean μ and variance-covariance matrix Σ. The probabilities of being in the other three economic regimes are

$$\begin{aligned}
\Pr(S_\pi \geq 0\,\&\, S_x < 0) &= \Pr(S_x < 0) - \Pr(S_\pi < 0\,\&\, S_x < 0) = F_s(\infty, 0) - F_s(0,0)\\
\Pr(S_\pi < 0\,\&\, S_x \geq 0) &= \Pr(S_\pi < 0) - \Pr(S_\pi < 0\,\&\, S_x < 0) = F_s(0, \infty) - F_s(0,0)\\
\Pr(S_\pi \geq 0\,\&\, S_x \geq 0) &= 1 - \Pr(S_\pi < 0) - \Pr(S_\pi < 0\,\&\, S_x < 0) = 1 - F_s(0,\infty) - F_s(0,0)
\end{aligned}$$

When we demonstrate our framework in the section "Empirical Examples," we define the prior distribution for these economic surprises based on historical estimates. More details on how we defined those surprises are provided in that section.

Investors' Views on Regime Likelihood

We allow investors to express views on probabilities of being these regimes and find a distribution that is consistent with those probabilities. Let the investors' views on probabilities be denoted as follows

$$\begin{aligned}
\Pr(S_\pi \geq 0\,\&\, S_x \geq 0|v) &\equiv Q_{HGHI}\\
\Pr(S_\pi \geq 0\,\&\, S_x < 0|v) &\equiv Q_{LGHI}\\
\Pr(S_\pi < 0\,\&\, S_x < 0|v) &\equiv Q_{LGLI}\\
\Pr(S_\pi < 0\,\&\, S_x \geq 0|v) &\equiv Q_{HGLI}
\end{aligned}$$

where $\Pr(S_\pi \geq 0\ \&\ S_x \geq 0|v)$ is the probability of being in the regime with positive inflation and growth surprises according to investor's views v. We assume that the distribution associated with the investors' views is

$$\mathbb{P}(v) = \mathcal{N}(\mu_v, \Omega(\rho)) \tag{2}$$

$$\Omega(\rho) = \sqrt{diag(\Sigma)}\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\sqrt{diag(\Sigma)} \tag{3}$$

where $\mu_v = [\mu_\pi \ \mu_x]'$ is the mean of surprises associated with the investors' views and $\Omega(\rho)$ is the variance-covariance matrix as a function of ρ, which is the correlation between the surprises. μ_v and ρ need to be calibrated.[7]

Let $F_v(\cdot)$ be the cumulative density function of the bivariate normal with mean μ_v and variance-covariance matrix $\Omega(\rho)$. Given μ_v and ρ, the probabilities associated with the distribution in Equation 2 can be calculated as follows:

$$\begin{aligned}
&\Pr(S_\pi < 0 \ \& \ S_x < 0) = F_v(0,0)\\
&\Pr(S_\pi \geq 0 \ \& \ S_x < 0) = \Pr(S_x < 0) - \Pr(S_\pi < 0 \ \& \ S_x < 0) = F_v(\infty, 0) - F_v(0,0)\\
&\Pr(S_\pi < 0 \ \& \ S_x \geq 0) = \Pr(S_\pi < 0) - \Pr(S_\pi < 0 \ \& \ S_x < 0) = F_v(0, \infty) - F_v(0,0)\\
&\Pr(S_\pi \geq 0 \ \& \ S_x \geq 0) = 1 - \Pr(S_\pi < 0) - \Pr(S_\pi < 0 \ \& \ S_x < 0) = 1 - F_v(0, \infty) - F_v(0, 0)
\end{aligned}$$

We find the parameters μ_v and ρ such that the probabilities associated with the distribution in Equation 2 are the same as the investors' views. Formally, we solve the following problem using standard numerical procedures:

$$\begin{aligned}
\min_{\mu_v, \rho} [Q_{HGHI} - (1 - F_v(0, \infty) - F_v(0,0))]^2 + [Q_{LGHI} - (F_v(\infty, 0) - F_v(0,0))]^2 \\
+ [Q_{LGLI} - F_v(0,0)]^2 + [Q_{HGLI} - (F_v(0, \infty) - F_v(0,0))]^2
\end{aligned} \tag{4}$$

Posterior Distribution on Regime Likelihood

The prior distribution is denoted by $\mathbb{P}(s)$ (see Equation 1) while investors' views imply a conditional distribution $\mathbb{P}(v)$ (see Equation 2). To combine the investor views and the prior distribution, we consider a mixture of the two distributions. Specifically, we calculate the posterior distribution as

$$\mathbb{P}(s|v) = (1 - \gamma)\mathbb{P}(s) + \gamma\mathbb{P}(v) \tag{5}$$

where γ is a parameter that allows investors to determine how confident they are in their views. The confidence parameter γ varies from zero (least confident) to 1 (the most confident). We elect to use the posterior distribution specified in Equation 5 because of its clear interpretability: When investors are very confident in their views, their posterior should be closer to the distribution consistent with their views.[8]

Exhibit 1 provides a visualization of our framework. Panel A shows the prior distribution of growth and inflation surprises that, in this example, places equal weights to the four economic regimes. In Panel B, we assume that investors have views about the probabilities of the four economic regimes that differ from the prior distribution.

We estimate the distribution that is consistent with investors' views as explained earlier and we plot in Panel B of Exhibit 1. Using Equation 5, we combine the distribution implied by investors' views and the prior distribution. Panels C and D show the resulting distributions for two different levels of investors' confidence, $\gamma = 0.2$ and $\gamma = 0.5$, respectively. The effect of γ is intuitive. When investors are more confident in their views (i.e., higher gamma), the posterior probabilities are closer to investors' views. Furthermore, the two panels show that a mixture of normal distributions can generate nonnormal distributions.

---

[7] Equation 3 implies that $\Omega(\cdot)$ can have a different correlation coefficient compared to the prior distribution from Equation 1, but it keeps the same volatilities.

[8] We would like to point out that the prior, conditional, and posterior distributions defined in Equations 1, 2, and 5 have been assumed to be normal in this article but different distributional assumptions are possible. That is, our framework is flexible and can be applied to various definitions of economic regimes, priors, and conditional distributions.

EXHIBIT 1

Visualization of Our Framework for Views on Economic Regimes

NOTES: This exhibit provides a visualization of our framework that proposes views on economic regimes (i.e., growth and inflation). Panel A plots the distribution of growth and inflation surprises that investors can estimate using historical information (the prior distribution). In Panel B, we plot the distribution that is consistent with investors' views on probabilities of economic regimes (the views distribution). Panels C and D show the resulting distributions that combine the prior and views distributions for two different levels of investors' confidence, $\gamma = 0.2$ and $\gamma = 0.5$, respectively.

Portfolio Allocation Given Posterior Probabilities

The final step of our approach is to design the optimal portfolio taking into account the posterior probabilities. In our approach, changes in the regime probabilities lead to changes in the optimal portfolio allocation as asset returns have different characteristics across regimes. By using a regime-switching model, our framework allows for any type of allocation rule; that is, investors can use their preferred allocation rule (e.g., risk-parity, mean–variance, etc.) to build a portfolio. As noted, this offers more flexibility than traditional methodologies (e.g., BL) which are built around mean-variance.

To simplify, we discuss our methodology using mean–variance as an example, but our framework can be extended to other allocation rules. Assume that there are

S economic regimes. Investors know the historical conditional expected returns μ_s and covariance matrices Σ_s for all $s \in S$ as well as the probability that each of these regimes happen p_s. Define w_s be the vector of optimal weights if regime s were to happen with probability 1, and let $w \equiv \Sigma_{s \in S}\ w_s p_s$. We can define a portfolio on the efficient frontier as follows:

$$w^* = \arg\min_{\{w_s\}_{s\in S}} \frac{1}{2} \sum_{s \in S} p_s w_s' \Sigma_s w_s \tag{6}$$

subject to the following two constraints:

$$\sum_{s \in S} p_s w_s' \mu_s = \bar{\mu} \quad \text{and} \quad \sum_{s \in S} p_s w_s' I_N = 1 \tag{7}$$

where $\bar{\mu}$ is the required return and N is the number of assets. Solving Equation 7 yields the following tangency portfolio (i.e., highest SR)

$$w_u^* = \frac{\Sigma_u^{-1} \mu_u}{|\mathbf{1}' \Sigma_u^{-1} \mu_u|} \tag{8}$$

where $\mu_u = \Sigma_{s \in S}\ p_s \mu_s$, $\Sigma_u = \Sigma_{s \in S}\ p_s \Sigma_s$, and $\mathbf{1}$ is a vector of ones with the same dimension as μ_u.

Equation 8 is the standard solution for the optimal mean–variance portfolio using the probability-weighted expected returns μ_u and covariances Σ_u. It follows that applying this methodology to other strategies (e.g., mean–variance with short sales constraints, maximum diversification, risk-parity, etc.) boils down to calculating the required probability-weighted statistics before applying optimization. That is, for mean–variance we required the calculation of both expected returns and covariances, for risk-parity investors should calculate only the covariances, and so on.

Our approach relies on investors knowing μ_s (conditional historical returns in our framework), Σ_s (covariance matrices) and p_s (probability of regime s) for all economic regimes. In the section "Empirical Examples," we used historical estimates for these variables. It is important that the differences of these estimates across regimes historically are also representative of that in the future. There is large empirical evidence documenting that asset returns behave consistently through time in various economic regimes.[9] On this basis, our setup—although simplistic—would demonstrate the benefits of our methodology.

EXAMPLE: A COMPARISON WITH BL

We compare the performance of our methodology and the BL model using Monte Carlo simulations. Our goal is to examine how forecast errors affect the performance of the respective methodologies. Specifically, we assume investors' views are correct on average but are subject to forecast errors in each simulation. To apply our proposed

[9] For example, equity returns are more correlated and lower in bear markets than that in bull markets (e.g., Longin and Solnik 2001; Ang and Bekaert 2002; Ang and Chen 2002; Patton 2004), periods of high volatility are associated with low excess equity returns and high excess long-term bond returns (e.g., Turner, Startz, and Nelson 1989; Hamilton and Susmel 1994), and real rates tend to drop during recessions (Ang, Bekaert, and Wei 2008). The literature on the topic is vast and beyond the scope of this article.

methodology outlined in the previous section, we use four economic regimes based on inflation and economic growth.

For each regime, we calibrate the annualized covariances and expected returns to empirical data on an equity index, long-term Treasuries, and an index of commodities.[10] For this exercise, we set the true expected returns and covariances equal to the average of the historical conditional values in the rising growth scenario (i.e., average of rising-growth/rising-inflation, and rising-growth/falling-inflation) or equivalently that the probabilities of those two regimes are 50% each.[11]

The forecast errors are defined as follows. For BL, the expected returns' forecasts are drawn from a very narrow multivariate normal distribution with means equal to the true expected returns and covariance equal to the true covariance (i.e., annualized historical covariances) times 1/50, representing very precise forecasts. For our framework, we draw from a uniform distribution centered around the true values of probabilities and with a range of 20% (i.e., +/– 10% around the true value), representing nonprecise forecasts. The parameters chosen are admittedly arbitrary but they possess an important feature. The error we set in the BL model is very small while we allow for a considerable larger error in our framework. As we discuss in the following, we also consider various combinations of such parameters and show the implications for the relative performance of the two methods.

In order to evaluate the performance of the two approaches, we use the ratio between the SR of the strategy with views and without views. When investors have no views, the average ratios are equal to one for both our approach and BL. For either approach, an average ratio greater than one means the portfolio performs better with views on average.

We present our results in Exhibit 2. Panel A plots the distribution of SR ratios with views over SR without views. The blue and orange lines show the results for our model and the BL model, respectively. Both our model and the BL model have averages of SR ratio greater than one. This is expected because the views are on average correct for both methodologies. However, our framework generates a distribution that has a higher mean and is considerably more positively skewed. Panel A clearly shows that our approach generates more robust results.

To further show that our approach is more robust, in Exhibit 3 we plot the distribution of weights allocated to equities, fixed income (10-year government bonds), and commodities from the 1,000 simulations discussed earlier. Panels A and B show the distribution of weights when the BL and our approaches are used, respectively.

Exhibit 3 shows clearly that the dispersion of the weights is considerably larger for BL. For instance, the weights for equity are widely distributed between 0% and 60%, while the weight distribution for our model is smaller between 8% and 30%. For fixed income and commodities, we observe qualitatively similar patterns. Our results show the relative stability of our approach. This is a noteworthy result especially because the views (on expected return) used for BL are much more precise than the views (on probabilities) used for our approach.

The results presented in Panel A of Exhibit 2 and of Exhibit 3 are a function of our chosen parameters. To create an equivalence between the error in the BL model and the error in our framework, we conduct the following analysis. We fix the error range in our framework (e.g., 20% in the example described earlier in Panel A), and we seek the error in the BL model that generates the same average SR from 1,000

[10] We describe the data in detail in the section "Empirical Examples".

[11] Making an assumption on the true expected returns and covariances is required for the exercise, but we point out that the results are independent from the assumption made. In other words, the results presented here are not affected by the assumption made about the true expected returns and covariances.

EXHIBIT 2
Distribution of Out-of-Sample SRs under Error

Panel A: Distribution of SRs

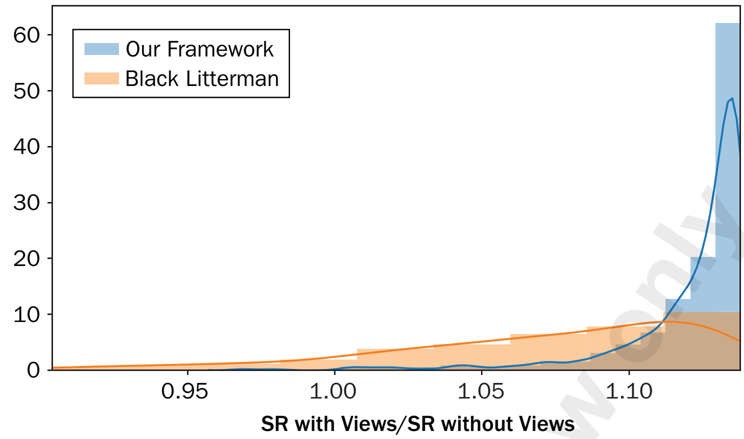

Panel B: Mapping between Errors

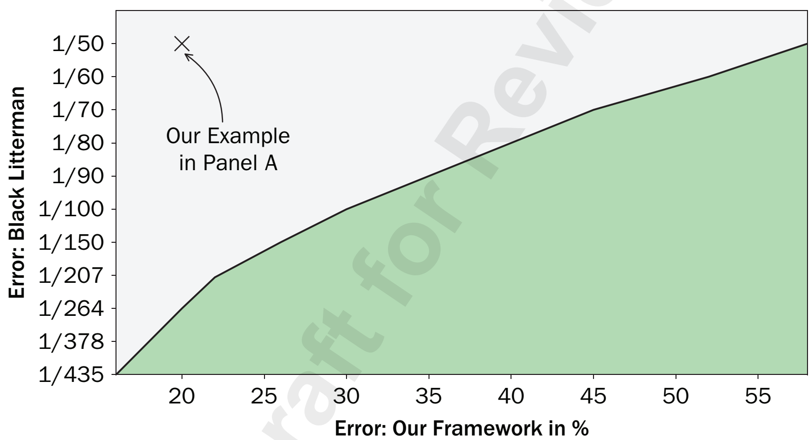

NOTES: Panel A plots the distribution of SRs implied by both our framework and BL. We use 1,000 simulations in this exercise. We assume that investors have views that are correct on average but subject to forecast error. For BL, the forecast error entails drawing expected returns from a multivariate normal distribution centered around the true values and covariance equal to the true covariance times 1/50. For our framework, we draw from a uniform distribution centered around the true values of probabilities and with a range of 20% (+/– 10%). For ease of readability, we report the distribution of the ratio between the SR using the views and the SR without views. We calibrate the covariance matrices and expected returns to empirical data on equities, fixed income, and commodities. Panel B shows the mapping between the error in our framework that leads to the same average SR using BL. For example, a +/– 10% range in our framework is equivalent to having a precision in BL approximately equal to 1/264 of the true covariance matrix (i.e., 264 times more precise than the actual distribution of returns).

simulations using the same methodology described earlier. We repeat this exercise for various levels of the error in our framework. Panel B shows that, when we use an error range of 20% for the error in our framework, investors using the BL model should have the error covariances of the expected return views equal to 1/264 of the annualized historical covariances of expected returns. When the error covariances in the BL model is 1/50 (as is the case in Panel A) of the historical covariances, it corresponds to an error range of approximately 55% in our framework.

In summary, our results clearly show that our methodology—which uses views on probabilities—can likely generate more robust results than BL because our approach is less sensitive to errors in views.

EXHIBIT 3

Distributions in Portfolio Weights between Methodologies

Panel A: BL Approach

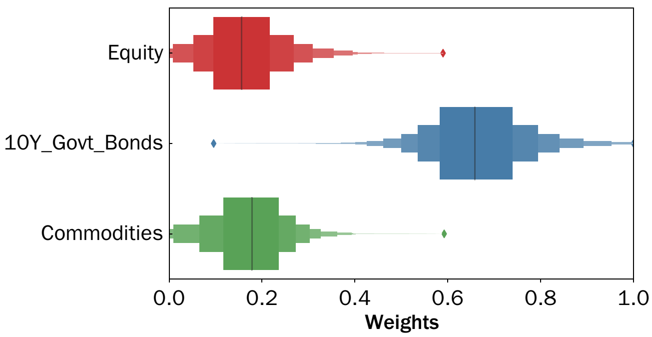

Panel B: This Article's Framework

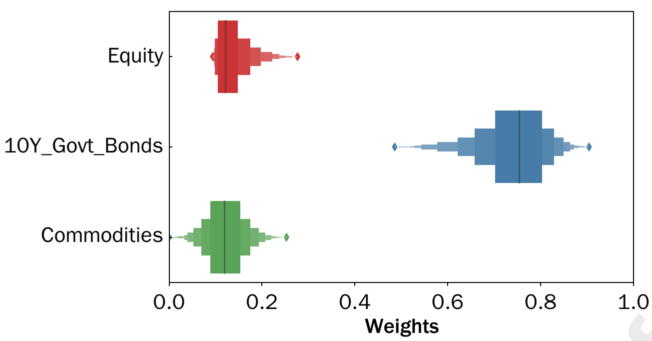

NOTES: This exhibit compares the distribution in weights between our framework and BL's from 1,000 simulations. The distributions of weights are generated using the same simulations as in Exhibit 2. Investors' views are drawn from a distribution centered around the true value (i.e., their views are on average true but subject to error). Expected returns and covariances are from Exhibit 4. Panel A shows the distribution of weights when the BL approach is used. Panel B shows the distribution of weights when our framework is used. The colored boxplot (red for equity, blue for fixed income, and green for commodities) shows the distribution of weights. Specifically, for each asset class, the median is shown by a vertical line segment, and the innermost box covers the interquartile range (50% of the probability mass from the 25th to the 75th percentile); the two incrementally smaller boxes cover additional 25% of the probability mass (i.e., together with the innermost box, they cover 75% of the probability mass); the next narrower box covers an additional 12.5% of the distribution, and so on until 0.1% of the probability mass is left. All remaining observations are considered outliers and are drawn as individual points.

EMPIRICAL EXAMPLES

Data Description and Setup

We provide an empirical example to demonstrate our framework using returns of the S&P 500 (equity), the 10-year US government bonds (fixed income), and the Bloomberg Commodity Index (commodities).[12] We calculate excess returns using the three-month T-bill rate.

For our empirical example, we define economic regimes in terms of inflation and economic growth surprises. We define inflation as the year-over-year (YoY) change in the Consumer Price Index and economic growth rate as the YoY change in the US real GDP. For both of these two variables, we use vintage data when available and realized values otherwise. The vintage data are sourced from the Philadelphia Fed. For estimates of inflation expectation, we use the Survey of Professional Forecasters published by the Philadelphia Fed. For estimates of economic growth expectation, we use a simple five-year rolling average of YoY real GDP growth rate.[13] We define economic surprises as the realized value minus its expectation. For brevity, in this article we also refer positive and negative surprises as rising and falling, respectively (e.g., a positive surprise in inflation can be referred to as rising inflation).

We report the summary statistics for average excess returns, volatilities, and correlations across the four regimes in Exhibit 4. Results are generally intuitive. For example, one would expect equities to perform the best in the rising growth and falling inflation environment. As for (nominal) government bonds, they perform best in falling inflation environments and worst in rising inflation environments. Finally, commodities are commonly thought to perform well when the economy is growing and inflation is high. Our results are consistent with these intuitions.

It is important that assets have a considerably different behavior in the various regimes for our framework to be effective. The appendix provides the results of a formal test for the difference in means and volatilities of the assets across regimes. Our analysis shows that returns for equities, fixed income, and commodities are (in most cases) statistically different across regimes, whereas only the volatilities of equities and fixed income are statistically different from each other in the various regimes.

[12] Our approach can be extended to international countries. For brevity, we choose to focus on US data in this article, and we leave the analysis for international countries for future research.

[13] We set the expected growth rate at time t equal to the average economic growth rate for the past five years. We also evaluated the Survey of Professional Forecasters for real GDP but found that they would not discriminate returns as well as using the five-year moving average. Because this is only an example, we elected to use the five-year moving average for real GDP.

EXHIBIT 4

Conditional Excess Returns, Volatilities, and Correlations

			Rising Growth and Rising Inflation			
	Excess Ret	**Volatility**	**Correlation Matrix**	**Equity**	**10Y Govt Bonds**	**Commodities**
Equity	13.23%	10.48%	Equity	1	–0.084	0.215
10Y Govt Bonds	–0.39%	6.64%	10Y Govt Bonds	–0.084	1	–0.247
Commodities	11.07%	12.21%	Commodities	0.215	–0.247	1
			Rising Growth and Falling Inflation			
	Excess Ret	**Volatility**	**Correlation Matrix**	**Equity**	**10Y Govt Bonds**	**Commodities**
Equity	15.03%	11.82%	Equity	1	0.314	0.277
10Y Govt Bonds	5.12%	9.20%	10Y Govt Bonds	0.314	1	–0.124
Commodities	–2.34%	12.49%	Commodities	0.277	–0.124	1
			Falling Growth and Rising Inflation			
	Excess Ret	**Volatility**	**Correlation Matrix**	**Equity**	**10Y Govt Bonds**	**Commodities**
Equity	–0.22%	15.69%	Equity	1	–0.560	–0.137
10Y Govt Bonds	2.69%	5.93%	10Y Govt Bonds	–0.560	1	0.199
Commodities	6.74%	12.84%	Commodities	–0.137	0.199	1
			Falling Growth and Falling Inflation			
	Excess Ret	**Volatility**	**Correlation Matrix**	**Equity**	**10Y Govt Bonds**	**Commodities**
Equity	1.43%	18.56%	Equity	1	0.244	0.481
10Y Govt Bonds	8.82%	7.35%	10Y Govt Bonds	0.244	1	0.215
Commodities	–8.37%	15.22%	Commodities	0.481	0.215	1

NOTES: This exhibit shows the conditional excess returns, volatilities, and correlations for the assets used in this study. Equity indicates the S&P 500 Index, 10Y Govt Bonds indicates the constant maturity index of 10-year US Treasuries, and commodities indicates the Bloomberg Commodity index. The four conditional regimes are: rising growth and rising inflation, rising growth and falling inflation, falling growth and rising inflation, and falling growth and falling inflation. For details on the estimation of conditional regimes, see "Data Description and Setup".

Demonstration of Our Framework with Prescient Views

The goal of our empirical example is to demonstrate the historical benefits of tilting a portfolio using our approach when investors have prescient views on the future probabilities of different economic regimes. More importantly, our approach is without the use of forward-looking expected returns and covariances. Of course, if investors had reliable forecasts of expected returns, they could use BL directly instead. However, when investors do not have confident views of expected returns but have confident views in the likelihood of various economic regimes, our approach would allow them to include their views in portfolio decisions.

Although it is necessary that investors need ways to come up with those prescient views, such discussion is not the focus of this article. Our aim is to provide a framework for tilting portfolios given views on the likelihood of economic regimes, similar to how BL provided a framework for tilting portfolios given views on assets' expected returns.

Our demonstration procedure is as follows. We rebalance the portfolio yearly at the end of January for each year *t*. On every rebalancing date, we compute the monthly economic surprises from 1947 up to the end of December of the previous year (*t* minus one year) and use them to define historical economic regimes. We use the historical frequencies of those regimes to calibrate their prior probabilities (see Equation 1). We then calculate the conditional expected returns and

EXHIBIT 5

Demonstration for Our Framework: Total Returns

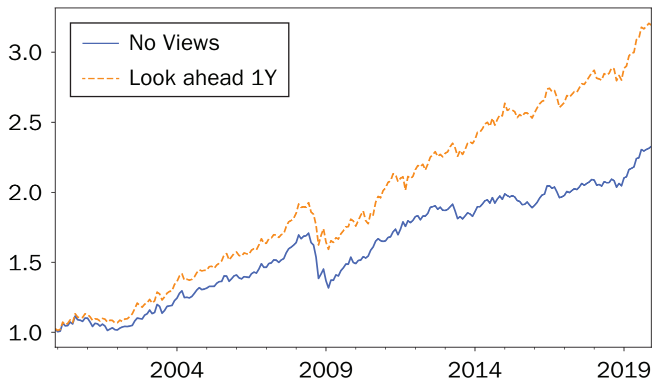

NOTES: This exhibit plots the total returns for the demonstration of our framework using mean–variance. The assets used for this exercise are the S&P 500 Index (equity), the 10-year Treasury returns (fixed income), and the Bloomberg Commodity index (commodities). We assume that investors have prescient views on the probability of future economic regimes (growth and inflation) one year ahead. The orange dashed line and the blue solid line show the performance of a portfolio with and without views, respectively. We test our strategy from 1999 to 2019, and weights are rebalanced yearly.

covariances for various assets under those economic regimes. When investors have no views, prior probabilities are used to calculate the optimal portfolio weights denoted as w_{prior}. When investors have prescient views, we assume their views are equal to the probabilities of economic regimes in the following 12 months and use them to derive the posterior probabilities of economic regimes.[14] Based on those posterior probabilities, investors compute portfolio weights w_{view}, which are tilted with respect to w_{prior}.

We start our demonstration in 1999, which is the first full year for which we have vintage data for inflation expectations.[15] For simplicity, we use the sample-based estimates of both expected returns and covariances. We use the standard mean–variance with short-sale constraints (MVO). Given the estimated probability-weighted expected returns $\hat{\mu}$ and covariances $\hat{\Sigma}$, the weights according to the MVO strategy are calculated by maximizing the SR subject to nonnegativity constraint[16]

$$w_{MVO}: \left\{ \max_{w} \frac{w'\hat{\mu}}{\sqrt{w'\hat{\Sigma}w}} \text{ s.t. } \sum_{i=1}^{N} w_i = 1 \text{ and } w_i \geq 0\ \forall i \right\} \quad (9)$$

where N is the number of assets.

Exhibit 5 displays the cumulative total returns for our demonstration using the strategies described earlier. We test our strategy from 1999 to 2019, and weights are rebalanced yearly.[17] Exhibit 5 shows that, using MVO, our proposed methodology with prescient views leads to a portfolio strategy (shown in orange dashed line) that outperforms a portfolio built with no views (shown in blue solid line).

Over the demonstration period, the portfolio with and without views has an annualized monthly excess return of 4.5% and 2.8%, respectively. The realized volatilities (annualized) of the portfolios with and without views are 6.6% and 6.3%, respectively. The outperformance achieved by our methodology with prescient views is 1.7% per year, and the SRs with and without views are 0.68 and 0.44, respectively. The portfolio with prescient views also consistently outperforms the portfolio without views in five-year subperiods: 1999–2004, 2004–2009, 2009–2014, and 2014–2019. The outperformance (measured as excess return of the portfolio with views versus the portfolio without views) is 2.0%, 1.6%, 1.8%, and 1.0%, respectively.

To further analyze the sources of outperformance, we examine the portfolio tilts as a result of having views on probabilities of economic regimes, we compare the weights with and without views in Exhibit 6. The shaded areas in this exhibit depict the four economic regimes. The exhibit shows that, given prescient views, our methodology intuitively underweight equities and overweight bonds through the Dot-com

[14] For this exercise, we assume a value of $\gamma = 0.5$.

[15] For the calibration of the prior probabilities until 1999, we used realized data.

[16] For ease of notation, we omit the subscript t for the vector of expected returns, covariances, and assets' weights.

[17] We use monthly data starting in 1947 with our demonstration starting in 1999 so that we have a long history to estimate the historical economic regimes and their assets' statistics. Because strategic asset allocation tilts are meant to be infrequent, we elect to rebalance the portfolio annually.

EXHIBIT 6

Change in Weights over Time

NOTES: This exhibit plots the allocation to S&P 500 (equity), 10-year Treasuries (fixed income) and the Bloomberg Commodity index (commodities) for the demonstration described in Exhibit 5. Panels A, B, and C show the weights allocated to equity, fixed income, and commodities, respectively. Weights are rebalanced yearly, and the demonstration period is from 1999 to 2019. The shaded areas show the four economic regimes. The abbreviations R, F, and Inf stand for rising, falling, and inflation, respectively.

bubble of the early 2000s and the Global Financial Crisis of 2007–2008, both of which are falling growth environments. Our methodology overweights equities after 2010 during the post–Global Financial Crisis recovery (e.g., rising growth). There is also an overweight on commodities from 2004 to 2008 during which inflation was elevated (e.g., rising inflation) and commodities performed well.

The demonstration suggests that our simple choice for the definitions of economic regimes leads to intuitive portfolio tilts given correct macroeconomic outlook (e.g., if investors' views are correct and that the likelihood of falling economic growth is increasing, our methodology will likely tilt their portfolio away from equities and into bonds).

CONCLUSION

BL (1992) is arguably the most known approach to incorporate investors' views on asset returns in portfolio construction decisions, and many have improved upon BL's seminal idea to allow for views on assets' volatilities, correlations, or factor returns and covariances. A common characteristic of all these methodologies is that they require views to be expressed directly on expected returns or covariances, but this poses a serious limitation: Investors need to provide point estimates of either expected returns, which are notoriously difficult to forecast, or covariances.

In this article, we bypass this problem by providing a novel methodology that requires investors to instead express their views in terms of probabilities of economic regimes, which—for some investors—could be more intuitive to determine. Our approach could be more readily applicable for investors because asset managers often discuss the future economic outlook by analyzing the likelihood of various macroeconomic scenarios. Using our methodology, asset managers can input their macroeconomic views directly into our framework without the need to translate them to expected returns and covariances as required by existing methodologies. Last, our methodology endogenously translates views on economic regimes to the portfolio weights using investors' preferred portfolio construction method.

Although our methodology is flexible and can be applied to any definition of economic regimes, we focus on two macroeconomic variables—inflation and real GDP growth—in our empirical demonstration. Using four economic regimes based on inflation and economic growth surprises, we demonstrate empirically that our methodology can improve portfolio performance when investors have prescient views on the likelihood of economic regimes. Our simulation results show that our approach is robust against forecast errors, which is important for practical applications.

APPENDIX

TESTING RETURNS AND VOLATILITIES ACROSS DIFFERENT REGIMES

This appendix provides a formal test for the difference in means and volatilities of the assets across regimes. In Panel A of Exhibit A1, we test whether the mean expected returns of the three asset classes—equity, 10-year bonds, and commodities—are different across regimes using the Tukey Honest Significant Difference (HSD) test at the 5% significance level.[18] Panel A shows that the null hypothesis is rejected most of the time for our asset classes, suggesting that their means are likely different across regimes.

To test whether the volatilities of equity, 10-year bonds, and commodities vary across economic regimes, we use Levene's test. The null hypothesis is that the volatilities of a given asset class across economic regimes are jointly equal: $\sigma_{X,HGHI} = \sigma_{X,HGLI} = \sigma_{X,LGHI} = \sigma_{X,LGLI}$, where $\sigma_{X,r}$ is the volatility of asset class X in economic regime r. HGHI denotes positive surprise (H) to growth (G) and positive surprise (H) to inflation (I), while HGLI denotes, positive surprise (H) to growth (G) and negative surprise (L) to inflation (I), vice versa. Panel B of Exhibit A1 shows that the test rejects the null hypothesis for both equity and 10-year bonds because both asset classes exhibit

[18] We elect to use the HSD test rather than pairwise *t*-tests because the HSD test was developed to adjust to the significance level for individual tests when simultaneous statistical inference for several tests is being performed. It is well-known that, when testing multiple hypotheses, the chance of observing a rare event increases, which also increases the Type I error (i.e., incorrectly rejecting a null hypothesis). Our null hypothesis is that all means for a given asset class are the same across economic regimes.

EXHIBIT A1

Statistics Tests for Differences in Conditional Excess Returns and Volatilities

Panel A: Test for Differences in Average Returns

		Reject $\mu_{Regime1} = \mu_{Regime2}$		
Regime 1	**Regime 2**	**Equity**	**10Y Govt Bonds**	**Commodities**
HGHI	HGLI	True	True	True
HGHI	LGHI	True	False	True
HGHI	LGLI	True	True	True
HGLI	LGHI	True	True	False
HGLI	LGLI	True	True	True
LGHI	LGLI	False	True	True

Panel B: Test for Differences in Volatilities

Regime 1	Statistic	P-Value
Equity	13.90	0.000
10Y Govt Bonds	6.59	0.000
Commodities	0.65	0.582

NOTES: This exhibit presents the statistical tests for differences in conditional means and volatilities. In Panel A, we use the Tukey HSD test for whether the assets have equal average returns across regimes. The HSD test applies simultaneously to the set of all pairwise differences in conditional means. We report whether the test rejects (true) or not (false) the null hypothesis that $\mu_{Regime1} = \mu_{Regime2}$. Panel B shows the test statistic and P-values for the Levene (1960) test of whether the assets have equal volatilities across regimes.

very low P-values. For commodities, the test cannot reject the null hypothesis of volatilities being equal across regimes.

Overall, because our definitions for economic regimes are able to discriminate asset returns and risks reasonably well, they are appropriate for the empirical demonstration of our framework, which is described in the section "Demonstration of Our Framework with Prescient Views".

ACKNOWLEDGMENTS

The authors thank the International Centre for Pension Management (ICPM) and the Kroner Center for Financial Research (KCFR) for their financial support and comments on this article. The authors also thank Michael Wissell and Robert Goobie for their support on innovation and research and the Global Peer Financing Association (GPFA) for its interest.

REFERENCES

Ang, A., and G. Bekaert. 2002. "International Asset Allocation with Regime Shifts." *The Review of Financial Studies* 15 (4): 1137–1187.

Ang, A., G. Bekaert, and M. Wei. 2008. "The Term Structure of Real Rates and Expected Inflation." *The Journal of Finance* 63 (2): 797–849.

Ang, A., and J. Chen. 2002. "Asymmetric Correlations of Equity Portfolios." *Journal of Financial Economics* 63 (3): 443–494.

Bansal, R., and I. Shaliastovich. 2013. "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets." *The Review of Financial Studies* 26 (1): 1–33.

Bansal, R., and A. Yaron. 2004. “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” *The Journal of Finance* 59 (4): 1481–1509.

Best, M. J., and R. R. Grauer. 1991. “On the Sensitivity of Mean–Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results.” *The Review of Financial Studies* 4 (2): 315–342.

Black, F., and R. Litterman. 1992. “Global Portfolio Optimization.” *Financial Analysts Journal* 48 (5): 28–43.

Boons, M., F. Duarte, F. De Roon, and M. Szymanowska. 2020. “Time-Varying Inflation Risk and Stock Returns.” *Journal of Financial Economics* 136 (2): 444–470.

Britten-Jones, M. 1999. “The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights.” *The Journal of Finance* 54 (2): 655–671.

Campbell, J. Y., and L. M. Viceira. 1999. “Consumption and Portfolio Decisions When Expected Returns Are Time Varying.” *The Quarterly Journal of Economics* 114 (2): 433–495.

——. 2001. “Who Should Buy Long-Term Bonds?” *American Economic Review* 91 (1): 99–127.

Campbell, J. Y., and T. Vuolteenaho. 2004. “Inflation Illusion and Stock Prices.” *American Economic Review* 94 (2): 19–23.

DeMiguel, V., L. Garlappi, and R. Uppal. 2009. “Optimal versus Naive Diversification: How Inefficient Is the 1/*n* Portfolio Strategy?” *The Review of Financial Studies* 22 (5): 1915–1953.

Figelman, I. 2017. “Black–Litterman with a Factor Structure Applied to Multi-Asset Portfolios.” *The Journal of Portfolio Management* 44 (2): 136–155.

Hamilton, J. D., and R. Susmel. 1994. “Autoregressive Conditional Heteroskedasticity and Changes in Regime.” *Journal of Econometrics* 64 (1–2): 307–333.

Jones, R. C., T. Lim, and P. J. Zangari. 2007. “The Black–Litterman Model for Structured Equity Portfolios.” *The Journal of Portfolio Management* 33 (2): 24–33.

Kolm, P. N., and G. Ritter. 2020. “Factor Investing with Black–Litterman–Bayes: Incorporating Factor Views and Priors in Portfolio Construction.” *The Journal of Portfolio Management* 47 (2): 113–126.

Kung, H. 2015. “Macroeconomic Linkages between Monetary Policy and the Term Structure of Interest Rates.” *Journal of Financial Economics* 115 (1): 42–57.

Levene, H. 1960. “Contributions to Probability and Statistics.” In *Essays in Honor of Harold Hotelling*, pp. 278–292. Redwood City, CA: Stanford University Press.

Longin, F., and B. Solnik. 2001. “Extreme Correlation of International Equity Markets.” *The Journal of Finance* 56 (2): 649–676.

Merton, R. C. 1980. “On Estimating the Expected Return on the Market: An Exploratory Investigation.” *Journal of Financial Economics* 8 (4): 323–361.

Nystrup, P., B. W. Hansen, H. O. Larsen, H. Madsen, and E. Lindström. 2017. “Dynamic Allocation or Diversification: A Regime-Based Approach to Multiple Assets.” *The Journal of Portfolio Management* 44 (2): 62–73.

Nystrup, P., B. W. Hansen, H. Madsen, and E. Lindström. 2015. “Regime-Based versus Static Asset Allocation: Letting the Data Speak.” *The Journal of Portfolio Management* 42 (1): 103–109.

Nystrup, P., H. Madsen, and E. Lindström. 2018. “Dynamic Portfolio Optimization across Hidden Market Regimes.” *Quantitative Finance* 18 (1): 83–95.

Patton, A. J. 2004. “On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation.” *Journal of Financial Econometrics* 2 (1): 130–168.

Qian, E. 2011. "Risk Parity and Diversification." *The Journal of Investing* 20 (1): 119–127.

Schmieder, C., and M. Kollár. *EIB Working Papers 2019/11—Macro-Based Asset Allocation: An Empirical Analysis*, vol. 2019/11. European Investment Bank (EIB), Luxembourg. 2020.

Sheikh, A. Z., and J. Sun. 2012. "Regime Change: Implications of Macroeconomic Shifts on Asset Class and Portfolio Performance." *The Journal of Investing* 21 (3): 36–54.

Turner, C. M., R. Startz, and C. R. Nelson. 1989. "A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market." *Journal of Financial Economics* 25 (1): 3–22.

Zheng, K., W. Xu, and X. Zhang. 2021. "Multivariate Regime Switching Model Estimation and Asset Allocation." *Computational Economics* 1–32. https://doi.org/10.1007/s10614-021-10203-9.