Shrinking Against Sentiment:

Exploiting Behavioral Biases in Portfolio Optimization*

Nathan Lassance Alberto Martín-Utrera

May 14, 2022

Abstract

High sentiment predicts lower market returns, higher arbitrage returns, and lower transaction costs. We propose a shrinkage methodology that exploits this empirical evidence to construct mean-variance portfolios. Exploiting the eigenvalue decomposition of the covariance matrix of stock returns, we show that mean-variance portfolio performance is the sum of two components: a market and an arbitrage component. Shrinking the sample covariance matrix toward the identity in the construction of mean-variance portfolios gives more relevance to the market component as the shrinkage intensity increases. We time the exposure to each component by shrinking more (less) when sentiment is low (high), which provides sizable economic gains even net of transaction costs.

Keywords: Shrinkage, sentiment, parameter uncertainty.

JEL Classification: G11, G12.

*Lassance, LFIN/LIDAM, UCLouvain, e-mail: nathan.lassance@uclouvain.be; Martín-Utrera, Iowa State University, corresponding author, e-mail: amutrera@iastate.edu. N. Lassance is supported by the Fonds de la Recherche Scientifique (FNRS) under Grant number J.0115.22. The authors thank helpful comments from Victor DeMiguel, Olivier Ledoit, André Portela-Santos, Majeed Simaan, Raman Uppal, Michael Wolf, Guofu Zhou, and seminar participants at INFORMS Annual Meeting, New Jersey Institute of Technology, and Stevens Institute of Technology.

1 Introduction

The seminal work of Markowitz (1952) provides an intuitive mathematical framework to construct well-diversified portfolios and is the cornerstone for the classic capital asset pricing model of Sharpe (1964) and Lintner (1965). These highly influential papers rely on the assumptions that investors know the true distributional properties of asset returns, are rational, and have homogeneous beliefs. While these assumptions provide a tractable framework to construct equilibrium models, they represent an important challenge for the practical implementation of mean-variance portfolios. In this paper, we address these limitations and propose a framework to account for parameter uncertainty and behavioral biases in the construction of optimal portfolios.

There is an extensive literature on portfolio selection documenting the large negative impact that parameter uncertainty has on the performance of optimal mean-variance portfolios (Michaud, 1989; DeMiguel, Garlappi, and Uppal, 2009b). Shrinking the sample covariance matrix of stock returns is a popular approach to overcome this challenge, and one breakthrough in this discipline is the work of Ledoit and Wolf (2004). Their proposed shrinkage estimator is a linear combination of the sample covariance matrix and the identity matrix, and has been very successful at improving the out-of-sample performance of diversified portfolios (DeMiguel, Garlappi, Nogales, and Uppal, 2009a).

A key element of a shrinkage estimator is the intensity with which the sample estimator is shrunk toward the chosen target (DeMiguel, Martin-Utrera, and Nogales, 2013). The calibration of this shrinkage intensity typically relies on statistical rather than economic arguments; for example, Ledoit and Wolf (2004) calibrate this intensity to minimize the mean squared error of the covariance matrix. In this paper, we instead investigate the economic gains of a mean-variance portfolio that exploits a shrinkage covariance matrix whose calibration is based on an economically motivated criterion.

The shrinkage method we propose hinges on the eigenvalue decomposition of the covariance matrix of stock returns used in the construction of mean-variance portfolios. Using this decomposition, we show that the performance of mean-variance portfolios can be characterized as the sum of two components: the performance of the market portfolio and the performance of an arbitrage portfolio. In our theory, the performance of the market portfolio is defined by the first principal component of stock returns, and the performance of the arbitrage portfolio is defined by the remaining lower-variance principal components.

We show theoretically that shrinking the covariance matrix toward the identity gives more relevance to the market component, and this relevance increases with the degree of shrinkage. We then calibrate the shrinkage intensity, and thus the exposure to the market and arbitrage portfolios, from the level of *investor sentiment* in the economy, which Baker and Wurgler (2007) define as the demand for risky assets not justified by the evidence at hand. Specifically, we build on the empirical and theoretical evidence that while market performance is negatively correlated with investor sentiment, the performance of arbitrage portfolios correlates positively with sentiment (Stambaugh, Yu, and Yuan, 2012; Huang, Jiang, Tu, and Zhou, 2015). Accordingly, the method we propose assigns a higher shrinkage intensity when investor sentiment is low and a more moderate shrinkage intensity when sentiment is high, that is, we *shrink against sentiment*.

The proposed methodology does not only allow us to time the return premia of the market and arbitrage portfolios, but it also allows us to exploit the relation between market liquidity an investor sentiment highlighted by Baker and Stein (2004). In particular, Baker and Stein (2004) show that high-sentiment periods are associated with higher-liquidity periods and vice versa. Therefore, our shrinkage approach that tilts the performance of mean-variance portfolios toward the market during times of low sentiment is sensible because it harvests a larger premium from the market and because it allows us to tilt our portfolio toward a low-turnover strategy when liquidity decreases.

To set the stage for our novel methodology, we first examine the empirical relationship between investor sentiment and portfolio returns. As a proxy for investor sentiment, we use the sentiment index proposed by Huang et al. (2015).[1] We document that investor

[1]The Huang et al. (2015) sentiment index builds on the work of Baker and Wurgler (2006). In particular, they extract a latent variable using the same sentiment proxies for investor sentiment. However,

sentiment correlates negatively with market returns and positively with arbitrage returns. The top panel in Figure 1 shows that from July 1965 to December 2018, the annualized Sharpe ratio of the market portfolio in low-sentiment regimes is about one, whereas that in high-sentiment months is about -0.05.

To evaluate the relationship between sentiment and arbitrage returns, we consider prominent arbitrage factor strategies: value (HML), profitability (RMW), investment (CMA), and the optimal long-short mean-variance portfolio that combines all the necessary portfolios for the construction of the HML, RMW, and CMA factors. The top panel in Figure 1 shows that in high-sentiment months, the arbitrage portfolios deliver an annualized Sharpe ratio of 0.60, 0.63, 0.95, and 1.43, respectively. In low-sentiment months, the arbitrage portfolios deliver a substantially lower annualized Sharpe ratio of 0.11, 0.12, -0.04, and 0.60, respectively.[2] In addition, the bottom panel in Figure 1 shows that the CAPM alphas of these arbitrage portfolios are dramatically lower during low-sentiment months. Accordingly, the empirical evidence suggests that the economic gains relative to the market from exploiting arbitrage strategies are small during low-sentiment periods. We argue that this insight has important implications for portfolio construction.

Our paper makes three contributions to the portfolio selection literature. First, we use the eigenvalue decomposition of the covariance matrix to demonstrate that the squared Sharpe ratio of the mean-variance portfolio can be decomposed as the sum of the squared Sharpe ratio of the market and the squared Sharpe ratio of an arbitrage portfolio. In addition, we show that the contribution of the market (arbitrage) portfolio to the squared Sharpe ratio of the mean-variance portfolio is equal to the squared correlation between the returns of the market (arbitrage) portfolio and those of the mean-variance portfolio. We exploit this result to show that one can shrink the covariance matrix toward

unlike Baker and Wurgler (2006), Huang et al. (2015) do not extract the latent sentiment index as the first principal component, and instead they use partial least squares. The advantage of this method over principal components is that it is designed to extract the unobservable sentiment component that better explains stock returns. In unreported results, we confirm that our results are robust to using the Baker and Wurgler (2006) sentiment index.

[2] In Section A of the Internet Appendix, we present a theoretical model of the economy with investor sentiment that captures those empirical findings in Figure 1.

Figure 1: Sentiment and portfolio returns

This figure illustrates the relationship between investor sentiment and portfolio returns. The top panel depicts the annualized Sharpe ratio of five portfolios: 1) the market portfolio, 2) the value factor (HML), 3) the profitability factor (RMW), 4) the investment factor (CMA), and 5) the optimal mean-variance portfolio that combines all the portfolios used in the construction of the HML, RMW and CMA factors, subject to the constraint that the weights of the portfolio add up to zero, and considering a risk-aversion coefficient of five. The bottom panel depicts the monthly CAPM alphas of the non-market factors. The figure depicts the performance of these portfolios in high- and low-sentiment regimes. Like Barroso and Detzel (2021), we define high-sentiment regimes as those years for which the sentiment index at the end of the prior year is above its median value for the entire sample. The monthly factor returns and sentiment index span the period July 1965 through December 2018.

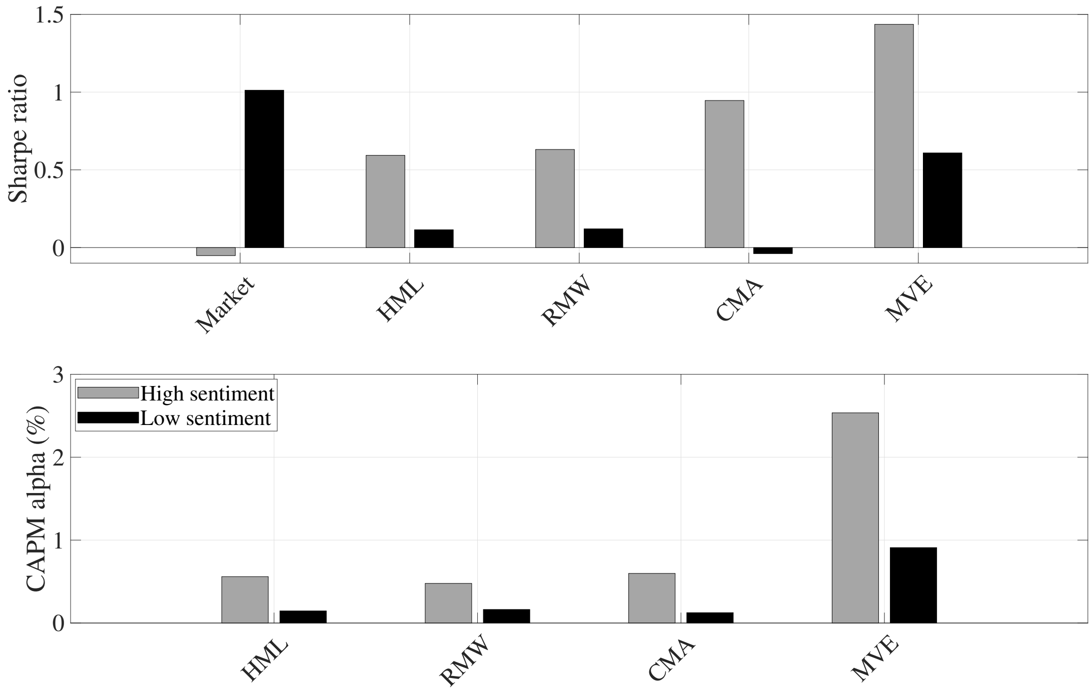

the identity matrix, as in Ledoit and Wolf (2004), and increase the correlation between the mean-variance portfolio and the market portfolio. Accordingly, employing a shrinkage covariance matrix allows us to strike a balance between the underlying market and arbitrage components that determine the performance of the mean-variance portfolio.

Our second contribution is to propose a novel shrinkage technique for the covariance matrix of stock returns that relies on *economic* instead of *statistical* arguments. Our method relies on the estimate of the probability that the market will deliver a positive excess return next period, which we determine from a logistic regression that utilizes investor sentiment as our forecasting variable. We then use this probability to linearly interpolate the desired correlation between the market and the mean-variance portfolio that exploits a shrinkage covariance matrix. Under this approach, if investor sentiment predicts that the probability of having a positive market excess return is high, our proposed methodology imposes a high correlation between the market and the mean-variance portfolio. To achieve the desired correlation, we use our analytical expression for the correlation between the market and the shrinkage mean-variance portfolio to obtain the shrinkage intensity that provides that correlation.

Our third contribution is to validate the performance of our proposed shrinkage methodology for portfolio selection across six empirical datasets. Our analysis shows that the median outperformance in terms of annualized Sharpe ratio of the *shrinking-against-sentiment* (SAS) portfolio relative to the benchmark mean-variance portfolios is about 15%. The excellent performance of the SAS portfolios does not come at the expense of higher turnover. Indeed, we show that the median turnover increase required by the benchmark mean-variance portfolios relative to the SAS portfolio is about 450%. This results in a more considerable outperformance of the SAS portfolio relative to the benchmark mean-variance portfolios in the presence of transaction costs. For instance, for transaction costs of 50 basis points, the median outperformance in terms of annualized Sharpe ratio of the SAS portfolio relative to the benchmark mean-variance portfolios is about 37%, and 66% relative to the equally weighted portfolio. Accordingly, shrinking against sentiment is an economically motivated approach to portfolio selection that brings sizable performance gains both in the absence and in the presence of trading costs.

We then provide further intuition on why shrinking against sentiment is a good investment approach. First, we run time-series regressions of SAS portfolio returns on the market conditional on lagged sentiment. These regressions show that the SAS portfolio has a larger exposure to the market than competing mean-variance portfolios, which allows our SAS portfolio to harvest a larger market premium. However, the SAS portfolio strategically reduces its market beta when sentiment increases and thus when market returns are low. Second, we run time-series predictive regressions of stock-level bid-ask spreads on lagged sentiment and find that higher sentiment predicts lower bid-ask spreads, consistent with Baker and Stein (2004). Therefore, shrinking against sentiment allows us to strategically tilt our mean-variance portfolio toward the market portfolio when the market premium is high and liquidity is low, which allows us to harvest higher return at substantially lower turnover and trading costs than standard mean-variance portfolios.

The methodology proposed in this paper shares elements with the literature on conditional mean-variance portfolios. Ferson and Siegel (2001) show how to construct the unconditionally optimal portfolio weights of mean-variance investors exploiting conditioning information. They show that the optimal portfolio depends on conditional mean and covariances, which are a function of state variables. Brandt and Santa-Clara (2006) show how to explicitly incorporate conditioning information in the construction of an optimal mean-variance portfolio. Their proposed method assumes that the optimal vector of portfolio weights is a function of a set of common state variables such as the dividend yield or short-term interest rates. Similar to these papers, our work exploits conditioning information (i.e., investor sentiment) in the construction of mean-variance portfolios. A distinctive feature of our work is that we exploit conditioning information *only* to estimate the shrinkage covariance matrix used in the construction of mean-variance portfolios.

Our work is closely related to the literature on shrinkage estimators in portfolio optimization.[3] Ledoit and Wolf (2003) propose a class of linear shrinkage estimators for the covariance matrix that combines the sample covariance matrix with the CAPM-implied covariance matrix of stock returns. Ledoit and Wolf (2004) propose another linear shrinkage estimator that combines the sample covariance matrix with a multiple of the identity matrix. This estimator improves the out-of-sample performance of mean-variance portfolios (DeMiguel et al., 2009a). Ledoit and Wolf (2017, 2020) introduce

[3]Other methodologies than shrinkage exist to alleviate the impact of parameter uncertainty, such as Bayesian approaches with diffuse priors (Klein and Bawa, 1976; Brown, 1978), priors based on asset pricing models (MacKinlay and Pastor, 2000; Pastor, 2000; Pastor and Stambaugh, 2000), priors based on economic objectives (Tu and Zhou, 2010), portfolio combinations (Tu and Zhou, 2011; Kan and Zhou, 2007; Kan, Wang, and Zhou, 2021), robust optimization methods (Goldfarb and Iyengar, 2003; Garlappi, Uppal, and Wang, 2007), mean-variance timing rules (Kirby and Ostdiek, 2012), and methods based on imposing constraints (Best and Grauer, 1992; Jagannathan and Ma, 2003; DeMiguel et al., 2009a).

a class of nonlinear shrinkage methods for the covariance matrix of stock returns that dominates linear shrinkage.[4] A key differentiating feature between our work and the aforementioned papers is that while they use statistical arguments to define the optimal degree of shrinkage, we introduce an economically motivated criterion based on sentiment to shrink the covariance matrix of stock returns.

Our work is also related to the literature on investor sentiment and asset prices; see, for instance, Yu and Yuan (2011), Stambaugh et al. (2012), Huang et al. (2015), and Shen, Yu, and Zhao (2017). These papers characterize the empirical and theoretical relationship between investor sentiment and asset prices. In particular, they show that the level of investor sentiment predicts market (arbitrage) returns negatively (positively). We complement and substantiate the findings in these papers and show in Section 2 that sentiment has a long-lasting effect on returns across several arbitrage portfolios, subsamples, and international markets. Moreover, unlike these papers, our focus is on portfolio optimization, and we introduce a simple approach to accommodate investor sentiment in the construction of optimal mean-variance portfolios.

Other papers use the eigenvalue decomposition of the covariance matrix to characterize portfolio performance. Pedersen, Babu, and Levine (2021) map the optimal mean-variance portfolio in the space of eigenvectors to provide a battery of solutions to mitigate the impact of estimation error. Zhao, Chakrabarti, and Muthuraman (2019) decompose the global-minimum-variance portfolio as a linear combination of a signal-only and a noise-only portfolio that are constructed from high-variance and low-variance principal components, respectively. Lassance, DeMiguel, and Vrins (2022) characterize portfolio weights as a linear combination of factor exposures defined by the principal and independent components of stock returns. In this paper, we use the eigenvalue decomposition of the covariance matrix to decompose the performance of the shrinkage mean-variance portfolio into that of the market and arbitrage portfolios, which are driven by the first and remaining principal components, respectively. Our objective with this decomposi-

[4] See Ledoit and Wolf (2022) for an excellent review of the literature on shrinkage estimators.

tion is to blend machine learning methods with economic theory by exploiting investor sentiment in the construction of mean-variance portfolios.[5]

Raponi, Uppal, and Zaffaroni (2020) and Kelly, Malamud, and Pedersen (2021) also characterize the performance of optimal portfolios as the sum of two components: an "alpha" component and a "beta" component. While the alpha and beta components are defined differently in the two papers, they both originate from linear trading strategies that invest in a number of predictive characteristics. In contrast, the decomposition of the performance of shrinkage mean-variance portfolios that we present in this paper relies on the eigenvalue decomposition of the covariance matrix of stock returns. Raponi et al. (2020) improve portfolio performance by correcting the misspecification of the alpha and beta components, while Kelly et al. (2021) do so by constraining the norm of a matrix that determines portfolio weights from exposures to characteristics. Unlike these two papers, we propose a shrinkage method that exploits sentiment to establish an optimal exposure to the two components that define the performance of mean-variance portfolios.

2 Sentiment and stock returns

To set the stage for our main contribution, here we reproduce and substantiate the results in the existing literature that document a strong relation between stock returns and investor sentiment. Baker and Wurgler (2007) define investor sentiment as *"the belief about future cash flows and investment risks that is not justified by the facts at hand."* We measure investor sentiment with the Huang et al. (2015) sentiment index and study how it affects the returns of several portfolios.[6] We consider five portfolios, which are the Fama-

[5]Our work is also related to the recent strand of the literature that focuses on the application of machine learning methods for asset pricing; see, for instance, Feng, Giglio, and Xiu (2020), Freyberger, Neuhierl, and Weber (2020), Gu, Kelly, and Xiu (2020), Kozak, Nagel, and Santosh (2020), and Bryzgalova, Pelger, and Zhu (2020). The shrinkage method presented in this paper can be linked to the machine learning literature. In particular, shrinkage estimators are a form of regularization (DeMiguel et al., 2009a), which is a widely used technique in machine learning (Schlkopf, Smola, and Bach, 2018). Unlike the aforementioned papers, whose objective is to characterize the cross-section, we combine shrinkage methods with economic theory to design profitable mean-variance portfolios.

[6]Huang et al. (2015) extract a latent sentiment variable from five sentiment proxies used in Baker and Wurgler (2006) using partial least squares (PLS). The five proxies are the closed-end fund discount, the number of IPOs, the first-day returns on IPOs, the share of equity issues, and the dividend premium.

French market portfolio and four prominent long-short portfolios from the empirical asset pricing literature: 1) the value factor (HML), 2) the profitability factor (RMW), 3) the investment factor (CMA), and 4) the optimal mean-variance portfolio[7] that combines all the characteristic portfolios used in the construction of the HML, RMW and CMA factors, subject to the constraint that the weights of the portfolio add up to zero.[8] We provide evidence on the relation between sentiment and stock returns in the U.S. in Section 2.1, and show that the results are robust to considering different subsamples and international returns in Section 2.2. In Section A of the Internet Appendix, we also provide theoretical support for the relation between investor sentiment and stock returns with a model of the economy in which investors have heterogeneous beliefs similar to that considered by Hong and Sraer (2016).

2.1 U.S. evidence

Figure 1 depicts the annualized Sharpe ratio of the five portfolios in high and low-sentiment months. We observe that while the performance of the market portfolio is the strongest in low-sentiment regimes, the performance of the four arbitrage portfolios is the strongest in high-sentiment regimes. These results illustrate the negative (positive) relation between sentiment and market (arbitrage) returns.

In addition, Table 1 reports the intercept and slope coefficients of long-run predictive regressions of the cumulative returns of a particular portfolio strategy on lagged values of sentiment. The results in this table show that while the slope coefficient of the sentiment variable is negative and statistically significant for the market, it is positive and statistically significant for the considered arbitrage portfolio returns. The slope coefficients are significant at different horizons. In particular, sentiment can have statistically significant

[7]Throughout the analysis, we consider an investor with a risk-aversion coefficient of $\gamma = 5$.

[8]The characteristic portfolios that produce the HML, RMW and CMA factors are the small-value, big-value, small-growth, big-growth, small-robust, big-robust, small-weak, big-weak, small-conservative, big-conservative, small-aggressive, and big-aggressive portfolios, all of which can be downloaded from Kenneth French's website.

effects on the returns of portfolio strategies for up to two years. In summary, Table 1 documents that sentiment can have long-lasting effects on stock returns.

These results are consistent with the literature. Huang et al. (2015) document that investor sentiment is negatively correlated with future market returns, and Stambaugh et al. (2012) find that the returns of arbitrage portfolios are stronger after episodes of high investor sentiment. In addition to the results offered by Huang et al. (2015) and Stambaugh et al. (2012), we document that investor sentiment has long-lasting effects on both market returns and arbitrage returns.

2.2 Subsamples and international evidence

We now provide additional evidence on the relation between investor sentiment and stock returns. First, we study the robustness of the results in the previous section to different subsamples. In particular, we split the sample into two periods of equal length: 1) July 1965 - March 1992, and 2) April 1992 - December 2018. Figure 2 depicts the annualized Sharpe ratio in these two different periods of the five portfolios considered in the analysis.

Figure 2 shows that the results are consistent with those in Figure 1. The market portfolio delivers an annualized Sharpe ratio of about 0.38 and 1.18 in low-sentiment regimes for the first and second subperiods, respectively. On the contrary, the market Sharpe ratios in high-sentiment regimes for the same subperiods are about 0.15 and -0.05, respectively. The impact of investor sentiment on arbitrage portfolio returns (i.e., HML, RMW, CMA, and MVE) is similar in the subsample analysis to that in the entire sample. In particular, arbitrage portfolios consistently deliver better performance in high-sentiment regimes. For instance, for the first subperiod, the MVE portfolio delivers a Sharpe ratio in high-sentiment regimes that is 17% higher than that in low-sentiment regimes. In the second subperiod, the difference is even larger and the MVE portfolio delivers a Sharpe ratio in high-sentiment regimes that is over four times larger than that in low-sentiment regimes. These results suggest that the negative (positive) correlation between market (arbitrage) returns and sentiment is more prominent in recent periods.

Figure 2: Sentiment and returns: subsample analysis

This figure depicts the annualized Sharpe ratio in two different periods of five portfolios: 1) the market portfolio, 2) the value factor (HML), 3) the profitability factor (RMW), 4) the investment factor (CMA), and 5) the optimal mean-variance portfolio that combines all the portfolios used in the construction of the HML, RMW and CMA factors, subject to the constraint that the weights of the portfolio add up to zero, and considering a risk-aversion coefficient of five. The figure depicts the performance of these portfolios in high- and low-sentiment regimes. Like Barroso and Detzel (2021), we define high-sentiment regimes as those years for which the sentiment index at end of the prior year is above its median value for the entire sample.

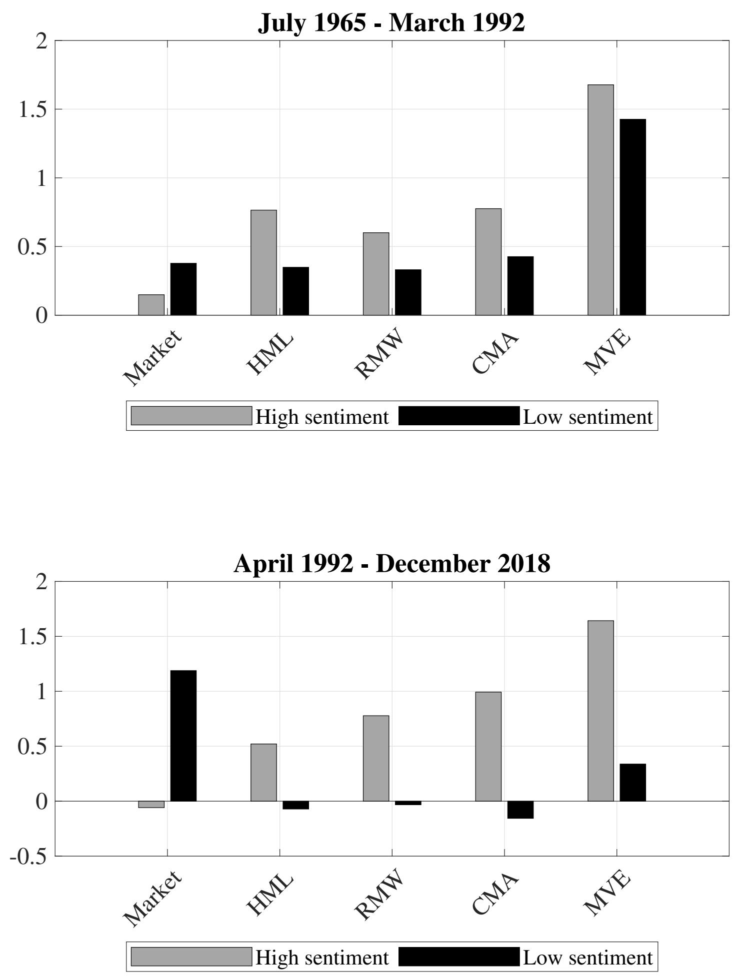

Figure 3: Sentiment and returns: international evidence

This figure depicts the annualized Sharpe ratio in international equity markets of five portfolios: 1) the market portfolio, 2) the value factor (HML), 3) the profitability factor (RMW), 4) the investment factor (CMA), and 5) the optimal mean-variance portfolio that combines all the portfolios used in the construction of the HML, RMW and CMA factors, subject to the constraint that the weights of the portfolio add up to zero, and considering a risk-aversion coefficient of five. The figure depicts the performance of these portfolios in high- and low-sentiment regimes. Like Barroso and Detzel (2021), we define high-sentiment regimes as those years for which the sentiment index at end of the prior year is above its median value for the entire sample. The market and arbitrage portfolio returns are obtained from six different international equity markets: 1) Asian countries excluding Japan, 2) Japan, 3) North America, 4) developed countries excluding the US, 5) Europe, and 6) emerging markets. The monthly market and factor returns are downloaded from Kenneth French's website and span the 28-year period from July 1990 to December 2018.

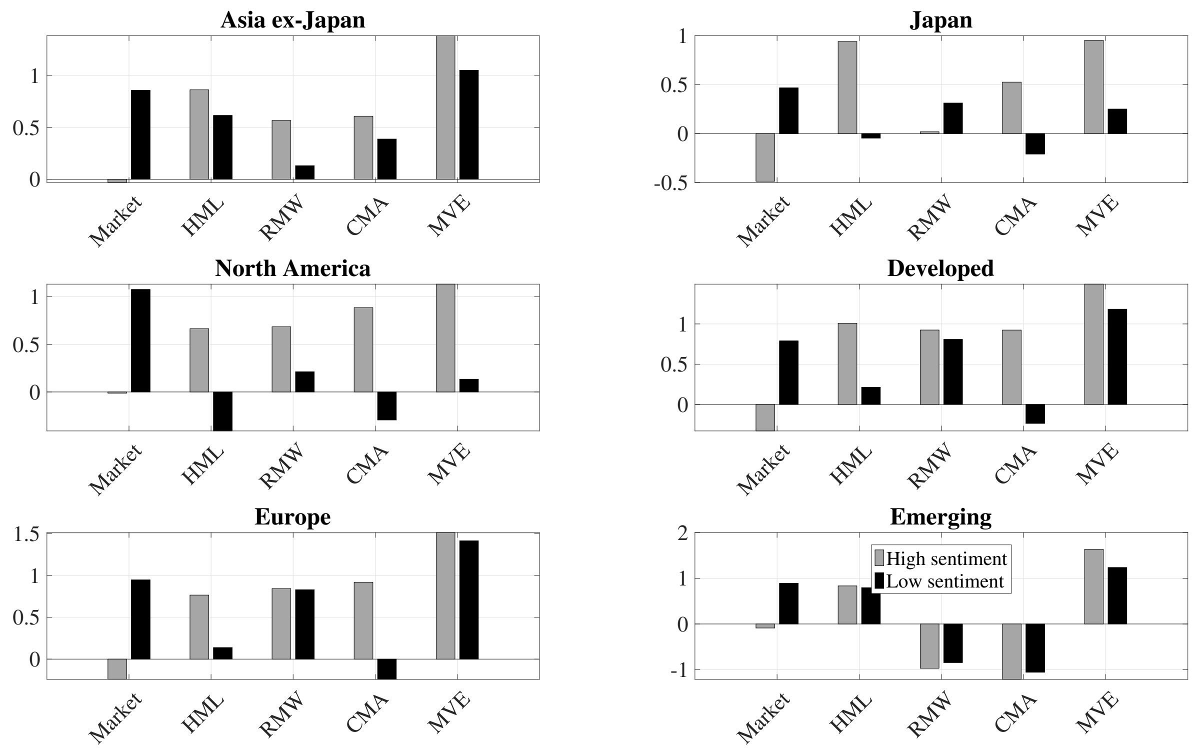

One concern is that the results presented in the previous section are sample-specific, and hence the relation between sentiment and stock returns is spurious (Fama, 1998). To tackle this concern, Baker, Wurgler, and Yuan (2012) show that sentiment demand correlates with stock returns also in other countries. We provide supporting evidence for the findings of Baker et al. (2012) in international equity markets in Figure 3. This figure shows that sentiment, measured with the Huang et al. (2015) sentiment index,[9] correlates with market and arbitrage returns in six different international equity markets. In particular, market returns are remarkably higher during low-sentiment regimes across the six different international equity markets. The international evidence for the relation between sentiment and arbitrage returns is equally strong as that in the US. We observe that in all six international markets, the performance of the MVE arbitrage portfolio that combines HML, RMW, and CMA is stronger during high-sentiment regimes.

Overall, there is a strong evidence, robust to different subsamples and across international equity markets, that investor sentiment correlates negatively with future market returns and positively with future arbitrage returns. In this paper, we exploit this finding to construct shrinkage mean-variance portfolios with improved performance.

3 A decomposition of mean-variance portfolios

In this section, we show that the performance of the optimal mean-variance portfolio can be decomposed as the sum of the performance of the market and an arbitrage portfolio. Given the empirical evidence presented in Section 2, this decomposition allows us to split the performance of the mean-variance portfolio into a component that is negatively correlated with sentiment (the market) and a component that is positively correlated with sentiment (the arbitrage portfolio). We also show how one can control the exposure of the mean-variance portfolio to the market and the arbitrage portfolios by shrinking the covariance matrix of stock returns.

[9] While the Huang et al. (2015) sentiment index is constructed with US sentiment proxies, international sentiment variables are highly correlated, and therefore the U.S. sentiment variable can be used as a proxy for a global sentiment variable. In particular, Baker et al. (2012) document that the U.S. and a global sentiment index have a correlation of nearly 90%.

3.1 Market, arbitrage, and mean-variance portfolios

We begin by laying out the notation we use in the next sections. First, we define the market portfolio as the equally weighted portfolio of all stocks in the investment universe,

$$w_M = \iota/N, \tag{1}$$

where ι is an N-dimensional vector of ones. Second, we define the arbitrage portfolio as the long-short portfolio maximizing mean-variance utility:[10]

$$w_A = \arg\max_{w} \; w^\top \mu - \frac{\gamma}{2} w^\top \Sigma w \quad \text{s.t.} \quad w^\top \iota = 0,$$

where μ and Σ are the vector of means and covariance matrix of stock returns in excess of the risk-free rate, respectively, and γ is the investor's risk-aversion coefficient. The solution to the arbitrage portfolio is

$$w_A = \frac{1}{\gamma}\Sigma^{-1}(\mu - \mu_g \iota), \tag{2}$$

where $\mu_g = \frac{\iota^\top \Sigma^{-1}\mu}{\iota^\top \Sigma^{-1}\iota}$ is the average return of the global-minimum-variance portfolio. Third, the optimal mean-variance portfolio is defined as the unconstrained portfolio maximizing mean-variance utility, which has the well-known solution

$$w^\star = \frac{1}{\gamma}\Sigma^{-1}\mu. \tag{3}$$

3.2 The effect of sentiment on mean-variance performance

We now decompose the mean-variance portfolio squared Sharpe ratio into a market and an arbitrage component. The general definition of portfolio's w squared Sharpe ratio is

$$\text{SR}^2(w) = \frac{\text{E}(w^\top R)^2}{\text{Var}(w^\top R)} = \frac{(w^\top \mu)^2}{w^\top \Sigma w}, \tag{4}$$

[10] This is the MVE arbitrage portfolio in Section 2 that optimally combines the portfolios used in the construction of the HML, RMW and CMA factors.

where R is the vector of stock returns in excess of the risk-free rate. It is well known that the squared Sharpe ratio of the mean-variance portfolio (3) is

$$\mathrm{SR}^2(w^\star) = \mu^\top \Sigma^{-1} \mu. \tag{5}$$

We make the following simplifying assumption to decompose the mean-variance portfolio squared Sharpe ratio in (5) as the sum of two components.

Assumption 1 *The first eigenvector of the covariance matrix of stock returns, Σ, is proportional to the equally weighted market portfolio, $v_1 = \iota/\sqrt{N}$.*

Assumption 1 is mild from an empirical standpoint; we confirm in unreported results that the first principal component extracted from the empirical datasets considered in this paper has a correlation of about 99% with the returns of the equally weighted market portfolio. Moreover, under Assumption 1 it is straightforward to show that the arbitrage portfolio (2) is orthogonal to the market portfolio, i.e., $w_A^\top w_M = 0$. Under Assumption 1 the market portfolio is the component that explains most of the time-series variability of stock returns, and therefore the dynamics of arbitrage portfolio returns should be explained by lower-variance principal components because the arbitrage portfolio is orthogonal to the market. Indeed, in the following proposition, we explicitly show that the performance of the mean-variance portfolio in (5) can be decomposed as the sum of a market and an arbitrage component.

Proposition 1 *Let Assumption 1 hold. Then, the following holds:*

1. *The squared Sharpe ratio of the mean-variance portfolio in (3) can be decomposed as the sum of the squared Sharpe ratios of the equally weighted market portfolio and the arbitrage portfolio in (2):*

$$\mathrm{SR}^2(w^\star) = \mathrm{SR}^2(w_M) + \mathrm{SR}^2(w_A), \tag{6}$$

where

$$\mathrm{SR}^2(w_M) = \frac{\mu_{PC_1}^2}{\sigma_{PC_1}^2}, \tag{7}$$

$$\mathrm{SR}^2(w_A) = \sum_{i=2}^{N} \frac{\mu_{PC_i}^2}{\sigma_{PC_i}^2}, \tag{8}$$

with $\mu_{PC_i} = v_i^\top \mu$ *and* $\sigma_{PC_i}^2 = v_i^\top \Sigma v_i$ *being the average return and variance of the* i*th principal component of stock returns, respectively, and* v_i *the* i*th eigenvector of* Σ.

2. *The contributions of the market and arbitrage portfolios to the squared Sharpe ratio of the mean-variance portfolio are equal to the squared correlations between their returns:*

$$\frac{\mathrm{SR}^2(w_M)}{\mathrm{SR}^2(w^\star)} = \mathrm{Corr}^2\big(w_M^\top R, (w^\star)^\top R\big), \tag{9}$$

$$\frac{\mathrm{SR}^2(w_A)}{\mathrm{SR}^2(w^\star)} = \mathrm{Corr}^2\big(w_A^\top R, (w^\star)^\top R\big) = 1 - \mathrm{Corr}^2\big(w_M^\top R, (w^\star)^\top R\big). \tag{10}$$

Proposition 1 shows that the mean-variance portfolio performance is the sum of the market *and* the arbitrage portfolio performance. Moreover, Proposition 1 shows that the market (arbitrage portfolio) contribution to the squared Sharpe ratio of the mean-variance portfolio is equal to the squared correlation between the returns of the market (arbitrage portfolio) and those of the mean-variance portfolio. This result is intriguing because the level of investor sentiment in the economy has two opposing effects in mean-variance portfolios. On the one hand, it reduces the Sharpe ratio of the market component, and on the other hand, it increases the Sharpe ratio of the arbitrage component. Therefore, it is of interest to develop methods to tilt the performance of mean-variance portfolios toward the market (arbitrage) component when investor sentiment is low (high). In the next section, we show that exploiting a particular shrinkage covariance matrix in the construction of mean-variance portfolios allows us to strike a balance between the market and the arbitrage components.

3.3 Shrinking the covariance matrix

Shrinking the covariance matrix is a powerful method to improve the out-of-sample performance of mean-variance portfolios in the presence of parameter uncertainty. In this section, we explain how one can balance the relative importance of the market and the arbitrage components in the performance of a mean-variance portfolio by shrinking the covariance matrix of stock returns. First, let us define the shrinkage covariance matrix of stock returns as in Ledoit and Wolf (2004):

$$\Sigma_{sh} = (1-a)\Sigma + a\nu I_N, \tag{11}$$

where a is the shrinkage intensity, ν is a positive scaling parameter, and I_N is the identity matrix of order N. The eigenvalue decomposition of the shrinkage covariance matrix is

$$\Sigma_{sh} = V((1-a)D + a\nu I_N)V^\top, \tag{12}$$

where D and V are the eigenvalue and eigenvector matrices that define the covariance matrix $\Sigma = V\,D\,V^\top$. This identity implies that the eigenvectors of matrices Σ and Σ_{sh} are identical, however the eigenvalues of matrix Σ_{sh} are defined by a convex combination of the eigenvalue matrices D and νI_N. Using the shrinkage covariance matrix (12) instead of the covariance matrix Σ gives the shrinkage arbitrage portfolio,

$$w_{SA} = \frac{1}{\gamma}\Sigma_{sh}^{-1}\left(\mu - \frac{\iota^\top\Sigma_{sh}^{-1}\mu}{\iota^\top\Sigma_{sh}^{-1}\iota}\iota\right). \tag{13}$$

and the shrinkage mean-variance portfolio,

$$w_S^\star = \frac{1}{\gamma}\Sigma_{sh}^{-1}\mu, \tag{14}$$

In the next proposition, we derive analytical expressions for the Sharpe ratios of the market portfolio, the shrinkage arbitrage portfolio, and the shrinkage mean-variance portfolio.

Proposition 2 *Let Assumption 1 hold. Then, the squared Sharpe ratios of the market portfolio, the shrinkage arbitrage portfolio in* (13)*, and the shrinkage mean-variance portfolio in* (14) *are*

$$\mathrm{SR}^2(w_M) = S_a(1,1), \tag{15}$$

$$\mathrm{SR}^2(w_{SA}) = S_a(2,N), \tag{16}$$

$$\mathrm{SR}^2(w_S^\star) = S_a(1,N), \tag{17}$$

Figure 4: Impact of shrinkage intensity on squared Sharpe ratio

This figure depicts the monthly squared Sharpe ratio of the shrinkage mean-variance portfolio $w_S^\star$ in (14), the equally weighted market portfolio w_M, and the shrinkage arbitrage portfolio w_{SA} in (13) as a function of the shrinkage intensity a defining the shrinkage covariance matrix in (11). The Sharpe ratios are computed using the results in Proposition 2. We construct the portfolios using the sample moments from the dataset of monthly returns on the 25 portfolios of stocks sorted on size and book-to-market downloaded from Kenneth French's website spanning July 1965 through December 2018. The value of ν for the shrinkage covariance matrix is equal to the cross-sectional average of the variance of the 25 portfolio returns.

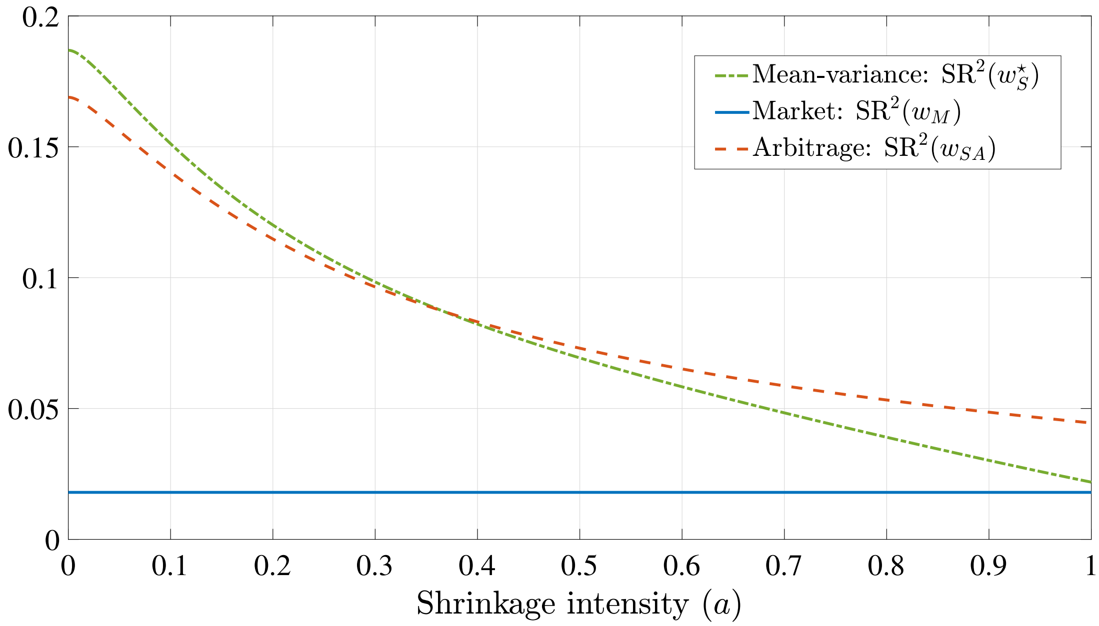

where

$$S_a(j,n) = \left(\sum_{i=j}^{n} \frac{\mu_{PC_i}^2}{(1-a)\sigma_{PC_i}^2 + a\nu}\right)^2 \left(\sum_{i=j}^{n} \frac{\mu_{PC_i}^2 \sigma_{PC_i}^2}{((1-a)\sigma_{PC_i}^2 + a\nu)^2}\right)^{-1}. \qquad (18)$$

When the shrinkage intensity a is zero, we recover the decomposition of the mean-variance portfolio performance in Proposition 1. For a shrinkage intensity $a > 0$, the squared Sharpe ratios of the arbitrage and mean-variance portfolios change with a. Figure 4 depicts the Sharpe ratio of the shrinkage mean-variance portfolio $w_S^\star$ in (14), the equally weighted market portfolio w_M, and the shrinkage arbitrage portfolio w_{SA} in (13) as a function of the shrinkage intensity a. We calibrate μ and Σ with the sample moments from the dataset of 25 portfolios of stocks sorted on size and book-to-market. The figure shows that when the shrinkage intensity a is zero, the arbitrage portfolio has the largest contribution to the squared Sharpe ratio of the mean-variance portfolio. However, as the shrinkage intensity increases, the squared Sharpe ratio of the mean-variance portfolio gets closer to that of the market portfolio, and they are nearly equal when $a = 1$. A fundamental insight from this figure is that shrinking the covariance matrix mutes the impact of low-variance principal components and, in the limit when $a = 1$, the shrinkage mean-variance portfolio performance is nearly identical to that of the market.

We now study more formally the impact of the shrinkage intensity a on the contribution of the market and shrinkage arbitrage portfolios to the performance of the shrinkage mean-variance portfolio. Specifically, in the next proposition we derive closed-form expressions for the correlation between the returns of the shrinkage mean-variance portfolio and the returns of the market and shrinkage arbitrage portfolios. Similar to Proposition 1, we show that the two squared correlations sum up to one for any shrinkage intensity a, and thus the market and arbitrage components fully determine the performance of the shrinkage mean-variance portfolio.

Proposition 3 *Let Assumption 1 hold. Then, the squared correlations between the return of the shrinkage mean-variance portfolio and that of the market portfolio and the shrinkage arbitrage portfolio are*

$$\mathrm{Corr}^2\big(w_M^\top R, (w_S^\star)^\top R\big) = \frac{\mu_{PC_1}^2 \sigma_{PC_1}^2}{\mu_{PC_1}^2 \sigma_{PC_1}^2 + \sum_{i=2}^N \mu_{PC_i}^2 \sigma_{PC_i}^2 \left(\frac{(1-a)\sigma_{PC_1}^2 + a\nu}{(1-a)\sigma_{PC_i}^2 + a\nu}\right)^2}, \tag{19}$$

$$\mathrm{Corr}^2\big(w_{SA}^\top R, (w_S^\star)^\top R\big) = 1 - \mathrm{Corr}^2\big(w_M^\top R, (w_S^\star)^\top R\big), \tag{20}$$

respectively. Moreover, the squared correlation with the market, $\mathrm{Corr}^2\big(w_M^\top R, (w_S^\star)^\top R\big)$*, is increasing in the shrinkage intensity a, and the squared correlation with the shrinkage arbitrage portfolio,* $\mathrm{Corr}^2\big(w_{SA}^\top R, (w_S^\star)^\top R\big)$*, is decreasing in a.*

Proposition 3 shows that shrinking the covariance matrix toward the identity allows us to control the exposure of mean-variance portfolios to the market and arbitrage port-

Figure 5: Correlation with market and arbitrage portfolios

This figure depicts the correlation between the returns of the shrinkage mean-variance portfolio $w_S^\star$ in (14) and the returns of the equally weighted market portfolio w_M and the shrinkage arbitrage portfolio w_{SA} in (13) as a function of the shrinkage intensity a defining the shrinkage covariance matrix in (11). The correlations are computed using the results in Proposition 3. We construct the portfolios using the sample moments from the dataset of monthly returns on the 25 portfolios of stocks sorted on size and book-to-market downloaded from Kenneth French's website spanning July 1965 through December 2018. The value of ν for the shrinkage covariance matrix is equal to the cross-sectional average of the variance of the 25 portfolio returns.

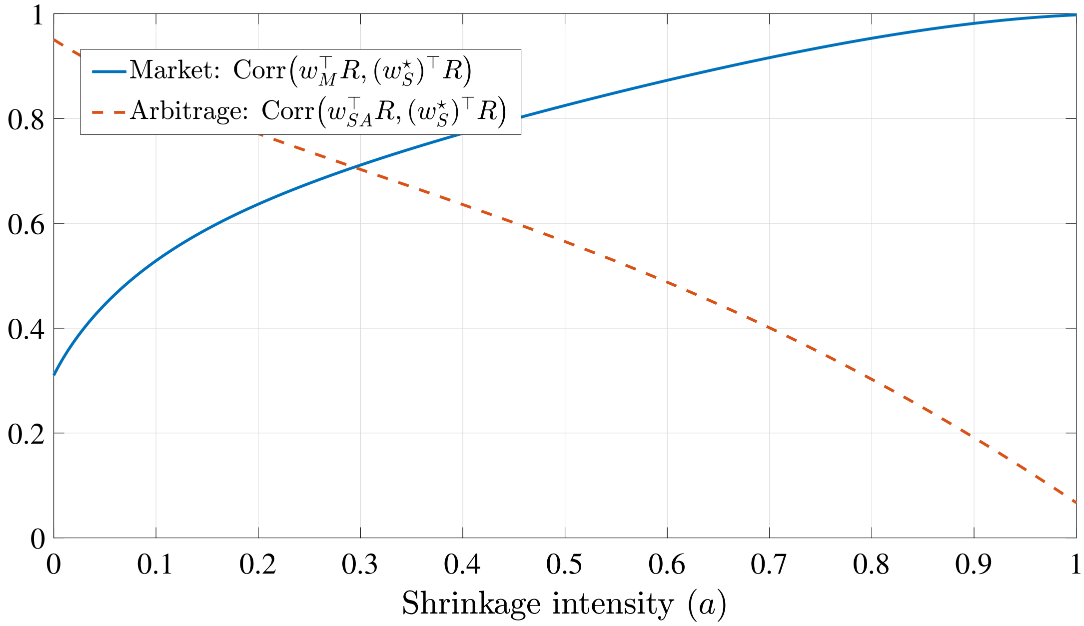

folios. In particular, the larger the shrinkage intensity a, the larger (lower) the correlation between market (arbitrage) returns and shrinkage mean-variance portfolio returns.

Figure 5 depicts the correlation between the returns of the shrinkage mean-variance portfolio in (14) and the returns of the equally weighted market portfolio and the shrinkage arbitrage portfolio as a function of the shrinkage intensity a using the results in Proposition 3. We calibrate μ and Σ with the sample moments from the dataset of 25 portfolios of stocks sorted on size and book-to-market. The figure shows that the correlation between the market and the shrinkage mean-variance portfolio is 31% when the shrinkage intensity a is zero, but it increases to almost 100% when a is equal to one. On the other hand, the correlation with the shrinkage arbitrage portfolio is 95% when the shrinkage intensity a is zero, but it decreases to 6.7% when a is equal to one.

The results presented in this section demonstrate the role of shrinkage on mean-variance portfolio performance. In the next section, we show how to account for investor sentiment in the construction of shrinkage mean-variance portfolios.

4 Shrinking against sentiment (SAS)

The empirical analysis presented in Section 2 documents a strong negative (positive) relationship between investor sentiment and market (arbitrage) portfolio returns. In this section, we illustrate the economic gains that a mean-variance investor can obtain by exploiting the empirical connection between investor sentiment and the performance of the market and arbitrage portfolios. In Section 4.1, we exploit the theory developed in Section 3 to incorporate investor sentiment in the construction of shrinkage covariance matrices for portfolio selection. Section 4.2 describes the data used in the empirical analysis and Section 4.3 documents the performance gains of our proposed methodology.

4.1 Shrinkage criterion

Our proposed mean-variance portfolio *shrinks against sentiment* (SAS) by applying a higher (lower) shrinkage intensity when sentiment is low (high). The criterion we propose in this section exploits classification methods to map sentiment onto a probability space. A popular classification method in the statistical learning literature is logistic regression (Hastie, Tibshirani, Friedman, and Franklin, 2009). In the simplest setting, a logistic regression estimates the probability of a particular event happening given some information (e.g., $P(\text{tomorrow rains} \mid \text{today is sunny})$). In particular, given the level of investor sentiment at time t, our objective is to estimate the probability that the market delivers a positive excess return over the next period (i.e., $r_{t+1}^{\text{MKT}} > 0$). The probability of this event happening determines the relevance of the market component in the overall performance of the mean-variance portfolio. We define the probability of the event

$r_{t+1}^{\mathrm{MKT}} > \ 0$ as

$$p = P\big(Y_{t+1} = 1 \bigm| \text{Sentiment}_t\big), \text{ where } Y_{t+1} = \begin{cases} 1 & \text{if } r_{t+1}^{\mathrm{MKT}} > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{21}$$

The above probability is conditioned on lagged sentiment, which is used in the following logistic function:

$$p = P\big(Y_{t+1} = 1 \bigm| \text{Sentiment}_t\big) = \frac{e^{\beta_0 + \beta_1 \text{Sentiment}_t}}{1 + e^{\beta_0 + \beta_1 \text{Sentiment}_t}}. \tag{22}$$

The coefficients β_0 and β_1 in (22) can be easily estimated via maximum likelihood; see, for instance, James, Witten, Hastie, and Tibshirani (2013).

Our method then exploits the estimate of probability p to back out the value of the shrinkage intensity a applied to the covariance matrix. We achieve this by using p to determine the desired correlation between the market and the shrinkage mean-variance portfolio. If the level of investor sentiment predicts that the probability p of having a positive market excess return is high (low), our approach imposes a high (low) correlation between the market and the shrinkage mean-variance portfolio. This is, our proposed method shrinks against sentiment.

More precisely, we use the predicted probability p that the market will deliver a positive return using the logistic function (22) to linearly interpolate the correlation between the return of the market and that of the shrinkage mean-variance portfolio:

$$\text{Corr}(p) = (1 - p)\,\text{Corr}_0 + p\,\text{Corr}_1, \tag{23}$$

where Corr_0 and Corr_1 are the correlations between the market and the shrinkage mean-variance portfolio in Proposition 3 when the shrinkage intensity a is zero and one, respectively. Then, using our theoretical result in Proposition 3, we can back out the value of the shrinkage intensity a from $\text{Corr}(p)$. To the best of our knowledge, our paper is the first to propose a shrinkage criterion for the covariance matrix of stock returns based on economic arguments instead of statistical arguments, which is the standard approach adopted in the construction of shrinkage matrices.

Figure 6: Shrinkage intensity and sentiment demand

This figure depicts the shrinking-against-sentiment shrinkage intensity obtained from the procedure described in Section 4.1 constructed from a dataset of 25 portfolios of stocks sorted on size and book-to-market that we download from Kenneth French's website, together with the investor sentiment variable proposed by Huang et al. (2015). The gray bars correspond to the NBER recession periods.

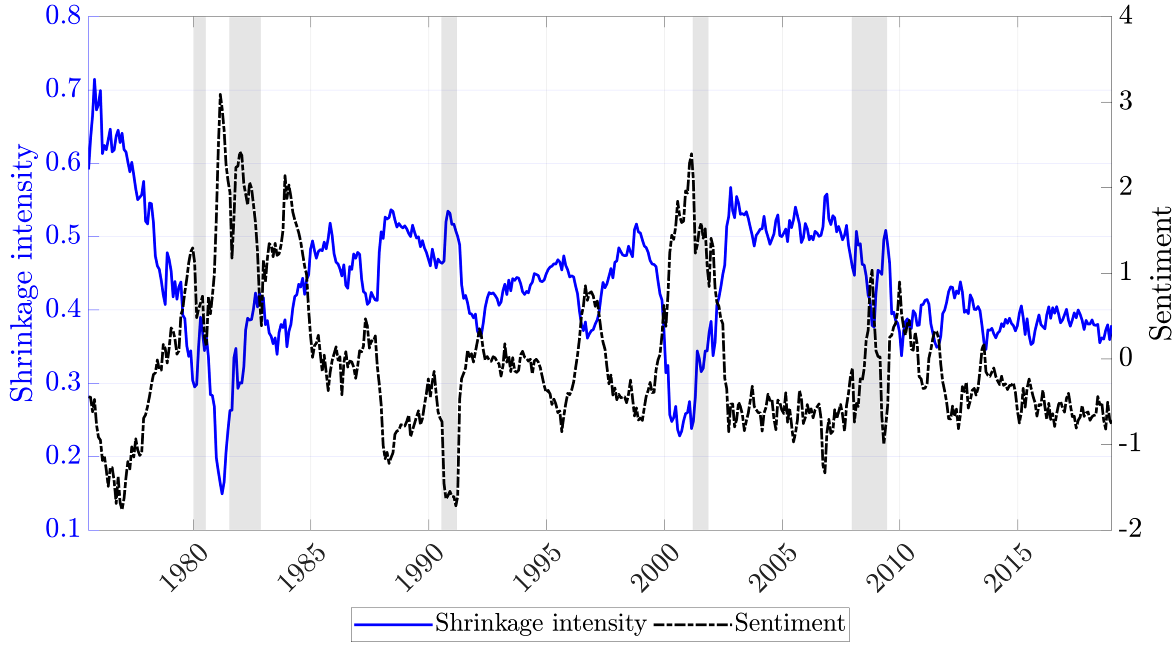

Figure 6 documents the relationship between the Huang et al. (2015) sentiment index and the proposed SAS shrinkage intensity a obtained from the dataset of 25 portfolios of stocks sorted on size and book-to-market. We observe that the shrinkage intensity has a substantially negative correlation with investor sentiment of -74%.

4.2 Data

We proxy for investor sentiment using the Huang et al. (2015) sentiment index, which we download from Guofu Zhou's website.[11] The sentiment index spans the period July 1965−December 2018, and it is a composite of five individual sentiment variables: the closed-end fund discount, the number of IPOs, the first-day returns on IPOs, the share

[11]We thank Guofu Zhou for making his data publicly available in his website.

of equity issues, and the dividend premium. The sentiment index is a latent variable extracted through partial least squares.

We use the monthly excess returns of six datasets. The first four datasets are downloaded from Kenneth French's website: (i) 10 momentum portfolios (*10MOM*) from July 1965 to December 2018, (ii) 25 portfolios sorted on size and book-to-market (*25SBTM*) from July 1965 to December 2018, (iii) 25 portfolios sorted on operating profitability and investment (*25OPINV*) from July 1965 to December 2018, (iv) 49 industry portfolios (*49IND*) from July 1969 to December 2018. The last two datasets come from the 23 anomalies considered by Novy-Marx and Velikov (2016) and are downloaded from Robert Novy-Marx's website: (v) the long and short legs of eight low-turnover anomalies (*16ANOM*) from July 1965 to December 2013 and (vi) the long and short legs of all the 23 anomalies (*46ANOM*) from July 1973 to December 2013.[12]

4.3 Out-of-sample performance

In this section, we assess the out-of-sample performance of SAS portfolios and several benchmark portfolios.

4.3.1 Methodology

For each dataset containing a total of T monthly observations, we construct portfolios every month with an expanding window of $M + t - 1$ observations, where $M = 120$ and $t = 1, \ldots, T - M$.[13] Then, for each estimated portfolio, we evaluate its performance using the next-month return observation at time $M + t$. We continue this process for all $t = 1, \ldots, T - M$ monthly observations. The analysis in this section focuses on the

[12] We thank Kenneth French and Robert Novy-Marx for making their data publicly available.

[13] In unreported results, we confirm that the findings in this section are robust to considering a rolling window of $M = 120$ observations instead of the expanding window approach we adopt here. Expanding windows have the benefit of delivering better estimates of mean returns, which typically require long time series to obtain accurate predictions (Merton, 1980).

tangency mean-variance portfolio, which is defined as[14]

$$w^\star = \frac{\Sigma^{-1}\mu}{|\iota^\top \Sigma^{-1}\mu|}, \tag{24}$$

where μ and Σ are the vector of means and covariance matrix of stock returns in excess of the risk-free rate, respectively, and ι is an N-dimensional vector of ones. A key element in this formulation is the covariance matrix Σ. Our analysis focuses on different portfolio strategies that estimate the covariance matrix in different ways. Below, we provide a list of the five different methods we consider for the estimation of the covariance matrix:

1. The shrinkage covariance matrix in Equation (11) that shrinks the sample covariance matrix toward the identity, where the shrinkage intensity a is estimated from the shrinking-against-sentiment methodology in Section 4.1. Plugging this matrix in (24) delivers the SAS portfolio.

2. The shrinkage covariance matrix proposed by Ledoit and Wolf (2004). This matrix shrinks the sample covariance matrix toward the identity to minimize its mean-squared error. Plugging this matrix in (24) delivers the MV-1N portfolio.

3. The shrinkage covariance matrix proposed by Ledoit and Wolf (2003). This matrix shrinks the sample covariance matrix toward the CAPM-implied covariance matrix. Plugging this matrix in (24) delivers the MV-MKT portfolio.

4. The nonlinear shrinkage covariance matrix proposed by Ledoit and Wolf (2020). The previous two shrinkage matrices proposed by Ledoit and Wolf are linear combinations of the sample covariance matrix and a target matrix. In Ledoit and Wolf (2020), the authors propose a nonlinear shrinkage covariance matrix calibrated to optimize the out-of-sample mean-variance portfolio performance; see also Ledoit and Wolf (2017). Plugging this matrix in (24) delivers the MV-NL portfolio.

[14] We only scale the tangency portfolio in (24) when the denominator $|\iota^\top \Sigma^{-1}\mu|$ is larger than one. This procedure helps avoid excessive leverage and extreme portfolio weights (Kirby and Ostdiek, 2012)

5. The covariance matrix proposed by Chen and Yuan (2016) that relies on principal component analysis. The authors consider the tangency portfolio in (24) and estimate the inverse covariance matrix Σ^{-1} by $V_K D_K^{-1} V_K^\top$, where V_K and D_K are the eigenvectors and eigenvalues matrices when keeping only the first $K \leq N$ principal components. The number K is determined via the popular method of Bai and Ng (2002), which is consistent in high dimension.[15] Plugging this matrix in (24) delivers the MV-PCA portfolio.

In addition to the mean-variance portfolios that exploit different estimates of the covariance matrix, we study the performance of two simple low-turnover strategies. First, the reward-to-risk (RTR) timing strategy of Kirby and Ostdiek (2012, Equation (13)), which corresponds to the tangency portfolio in (24) with all covariances set equal to zero and mean returns μ_i replaced by $\max(0, \mu_i)$. Second, the equally weighted portfolio, which is the market portfolio in our theory of Section 3 and corresponds to the tangency portfolio in (24) when all means, variances, and covariances are the same across assets. DeMiguel et al. (2009b) argue that the benefits of optimal diversification can be offset by estimation errors. Therefore, the naïve portfolio that assigns an equal weight to each asset emerges as a key benchmark in the assessment of portfolio strategies.

4.3.2 Out-of-sample results

Table 2 reports the annualized out-of-sample Sharpe ratios of the SAS portfolio and the six benchmarks across the six datasets described in Section 4.2. The table shows that SAS portfolios deliver a good performance. In particular, the median outperformance of the SAS portfolio across the six datasets is about 28%, 37%, 31%, and 8% relative to the MV-1N, MV-MKT, MV-NL, and MV-PCA mean-variance portfolios, respectively. The improvement is even larger relative to the two simple RTR and 1/N policies, for which the median outperformance is about 44% and 78%, respectively. The good performance of the SAS portfolios does not come at the expense of higher turnover. On the contrary,

[15] As in Bai and Ng (2002) and Chen and Yuan (2016), we impose a maximum value of $K = 8$.

the turnover of the SAS portfolio is the lowest among the competing mean-variance portfolios.[16] More precisely, the median turnover increase required by the benchmark mean-variance portfolios relative to the SAS portfolio is about 450%.[17]

The fact that SAS portfolios do not require much trading to harvest high risk-adjusted returns implies that their performance is likely to survive the impact of transaction costs. To address this issue, we compute the Sharpe ratio of the SAS portfolio and the six benchmarks net of proportional transaction costs. We define the net return of portfolio strategy p estimated with historical return data up to time t as

$$r_{t+1}^{p} = \Bigg(1 - \kappa \underbrace{\sum_{i=1}^{N} \left| w_{i,t}^{p} - w_{i,(t-1)^+}^{p} \right|}_{\text{Portfolio turnover}} \Bigg) \Big(1 + \left(w_t^p\right)^\top r_{t+1}\Big) - 1, \tag{25}$$

where r_{t+1} is the N-dimensional vector of excess returns at time $t+1$, w_t^p is the portfolio strategy p estimated with historical return data up to time t, $w_{i,(t-1)^+}$ is the portfolio weight in stock i at time t before rebalancing, and κ is the level of proportional transaction costs. Table 3 reports the annualized Sharpe ratio net of proportional transaction costs of the different portfolio strategies. Similar to Barroso and Saxena (2022), we consider proportional transaction costs of $\kappa = 10$, 30, and 50 basis points. The analysis in this table strengthens the relative performance of the SAS portfolio. We see that as the level of proportional transaction costs increases, the annualized Sharpe ratio of the different portfolios strategies deteriorates. However, the impact of transaction costs on the performance of SAS portfolios is much softer. For instance, for $\kappa = 50$ basis points as in DeMiguel et al. (2009b), the median outperformance in terms of annualized Sharpe ratio of the SAS portfolio relative to the benchmark mean-variance portfolios is about 36%. Moreover, even though the RTR and 1/N portfolios only require a very small turnover,

[16] In unreported results, we find that the lower turnover of the SAS portfolio comes from a more intensive shrinkage of the sample covariance matrix of returns relative to the benchmarks.

[17] The increase in turnover is particularly substantial for the 46ANOM dataset. However, even without the 46ANOM dataset the median turnover increase required by the benchmark mean-variance portfolios relative to the SAS portfolio is about 435%.

the median outperformance in terms of annualized Sharpe ratio of the SAS portfolio relative to the RTR and 1/N portfolios is about 34% and 66%, respectively.

5 Understanding the performance of SAS portfolios

This section provides intuition about the good out-of-sample performance of SAS portfolios. First, we run conditional time-series regressions of SAS portfolio returns on the market to explore the sources of the good performance of our proposed SAS portfolios.

5.1 Time-varying exposure to the market

Here we study the performance of SAS portfolios using conditional time-series regressions. In particular, we regress the out-of-sample returns of SAS portfolios on market returns, and we allow for the slope coefficient to vary over time with investor sentiment. That is,

$$r_t^{\text{SAS}} = \alpha + \underbrace{(\beta_0 + \beta_1 \,\text{Sentiment}_{t-1})}_{\beta_t} r_t^{MKT} + \epsilon_t, \tag{26}$$

where Sentiment_t is the Huang et al. (2015) sentiment index. Regression model (26) allows us to capture the time-varying market beta of SAS portfolios. Table 4 reports the intercept, slope coefficients, and t-statistics from the conditional time-series regression. For comparison purposes, we report the results for both the SAS portfolio and the mean-variance portfolio that exploits the Ledoit and Wolf (2004) covariance matrix, which appears with the acronym MV-1N.

We observe that in five out of the six datasets, the systematic risk of SAS portfolios decreases when sentiment increases (i.e., $\beta_1 < 0$), and it is statistically or nearly statistically significant in four of these datasets. Similarly, the conditional market betas of MV-1N portfolios also decrease with sentiment in the same datasets, however their statistical significance is systematically lower. In general, the correlation between any mean-variance portfolio with the market should decrease when market returns are lower,

Figure 7: Conditional market betas

This figure depicts the conditional CAPM market beta, β_t in Equation (26), of the SAS portfolio for the dataset of 25 portfolios of stocks sorted on size and book-to-market described in Section 4.2. For comparison purposes, the figure also depicts the market beta of the mean-variance portfolio exploiting the Ledoit and Wolf (2004) shrinkage covariance matrix, which appears with the acronym MV-1N. The figure also depicts the Huang et al. (2015) sentiment index (right axis).

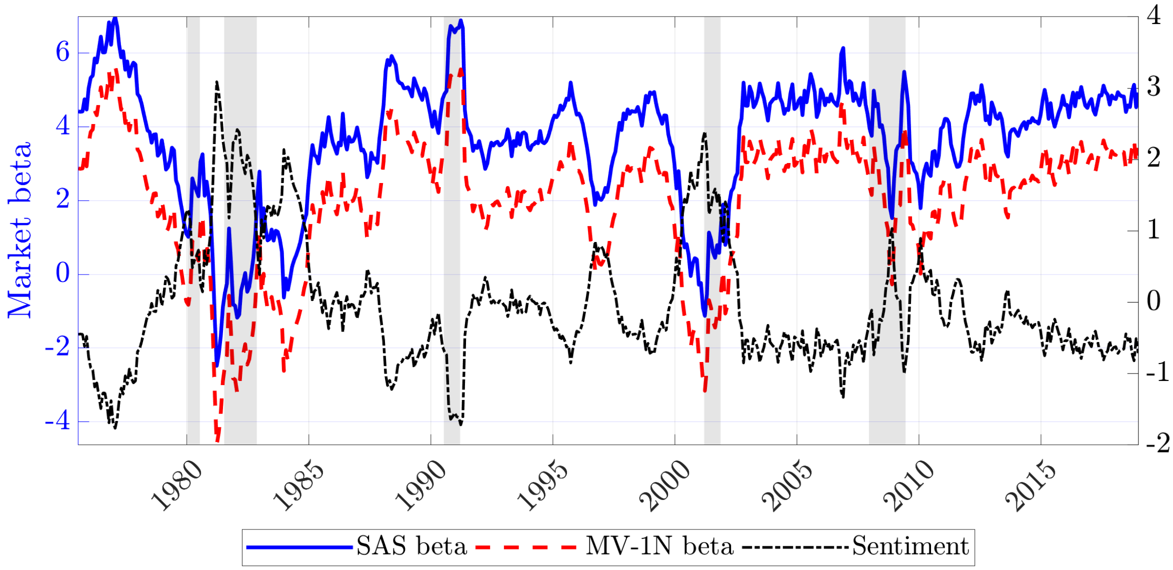

which typically happens when sentiment increases.[18] This implies that mean-variance portfolios reduce the exposure to the market as investor sentiment increases.

The main insight obtained from Table 4 is highlighted in Figure 7, which shows the substantial time-varying nature of market betas, β_t in Equation (26), of mean-variance portfolios for the 25SBTM dataset. First, the figure shows that both mean-variance portfolio betas have a strong negative correlation with investor sentiment. In particular, as sentiment increases, market beta decreases. Interestingly, the market beta of mean-variance portfolios can be negative when investor sentiment is large and therefore when market returns are prone to be low or even negative. However, our SAS portfolio typically keeps a higher exposure to the market than that of the MV-1N portfolio, as evidenced

[18]More formally, Equation (19) shows that the correlation between the return of any shrinkage mean-variance portfolio and the market portfolio decreases with the mean return of the first principal component, μ_{PC_1}, which is proportional to the mean return of the market portfolio under Assumption 1.

by the systematically larger β_0's in Table 4. The higher average market beta allows SAS portfolios to harvest a larger premium from the market than that of the MV-1N portfolio, which is particularly useful when sentiment is low and market returns are high.

The higher average market beta of SAS portfolios is also an important distinction relative to the benchmark mean-variance portfolios because it reduces transaction costs. Indeed, as noted in Proposition 3, one can achieve less exposure to the market by reducing the shrinkage intensity of the covariance matrix, however this gives a higher relevance to the arbitrage component, which requires a higher turnover and transaction costs than the market portfolio. We study the trading-cost benefits of our SAS portfolio approach in more detail in the next section.

5.2 Additional reason to shrink against sentiment

Baker and Stein (2004) argue that investor sentiment is positively related to market liquidity. We now confirm that sentiment predicts liquidity and therefore our shrinkage approach allows us to tilt our portfolio towards the market when liquidity is low and vice versa, which reduces transaction costs. To address this issue, we estimate stock-level bid-ask spreads using the *two-day corrected* method proposed in Abdi and Ranaldo (2017) and also used by DeMiguel, Martin-Utrera, Nogales, and Uppal (2020, Appendix IA.1).[19] For each stock i in month t, we define its corresponding bid-ask spread as

$$\widehat{s}_{i,t} = \frac{1}{D}\sum_{d=1}^{D}\widehat{s}_{i,d}, \quad \widehat{s}_{i,d} = \sqrt{\max\{4(c_{i,d}-\eta_{i,d})(c_{i,d}-\eta_{i,d+1}),0\}}, \tag{27}$$

where D is the number of days in month t, $\widehat{s}_{i,d}$ is the *two-day* bid-ask spread estimate, $c_{i,d}$ is the closing log-price on day d, and $\eta_{i,d}$ is the mid-range log-price on day d; that is, the mean of daily high and low log-prices.

Now, let us define p^l_{t+1} as the lth cross-sectional percentile at time $t+1$ of bid-ask spreads. Then, we run the following regression model:

$$p^l_{t+1} = \alpha + \beta\,\text{Sentiment}_t + \epsilon_{t+1}, \tag{28}$$

[19]We download daily price data from CRSP.

Figure 8: Turnover of market and arbitrage portfolios

This figure depicts the average monthly turnover of the equally weighted market portfolio and the arbitrage portfolio in Equation (2) using a risk-aversion coefficient of five. The table depicts the turnover of these portfolios across the six datasets described in Section 4.2 and following the rebalancing methodology in Section 4.3.1.

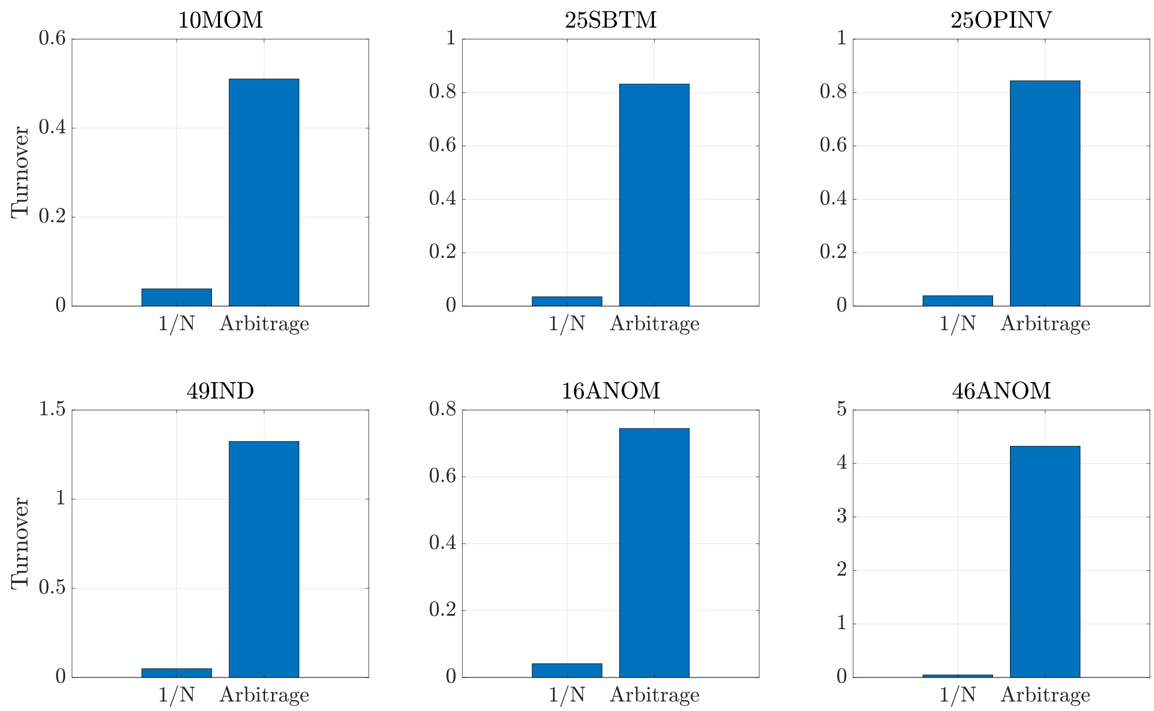

where Sentiment$_t$ is the sentiment index in month t. Table 5 shows that the slope coefficient, β, is negative and statistically significant across all the cross-sectional percentiles we consider. This implies that sentiment can predict next-period liquidity, and therefore, that shrinking against sentiment allows us to strategically tilt our mean-variance portfolio toward the market portfolio, which has a low turnover, when sentiment is low, and thus, when liquidity decreases. Similarly, by shrinking against sentiment we can increase the exposure of our mean-variance portfolio to the arbitrage component, which requires a much higher turnover, when liquidity is high. Indeed, Figure 8 shows that the equally weighted market portfolio has a dramatically lower turnover than the arbitrage portfolio.

6 Conclusion

The real-time implementation of Markowitz mean-variance portfolio theory is challenging because optimal portfolios require the estimation of the covariance matrix of stock returns as well as the vector of mean returns, and these estimates carry substantial statistical errors that can dramatically worsen the performance of mean-variance portfolios. Shrinkage estimators of the covariance matrix are an effective and popular method to mitigate the impact of statistical errors on mean-variance portfolio performance.

In this paper, we adopt this shrinkage approach and show that the performance of the optimal mean-variance portfolio can be decomposed into that of two components: the market portfolio and an arbitrage portfolio. The relevance of these two components in the overall performance of mean-variance portfolios is linked to the shrinkage intensity of the covariance matrix. More specifically, we show theoretically that by shrinking the covariance matrix of returns toward the identity matrix as in Ledoit and Wolf (2004), one can increase the relevance of the market component in the overall performance of the mean-variance portfolio. On the other hand, by shrinking less, one can increase the relevance of the arbitrage component in the overall performance.

Motivated by this decomposition of mean-variance portfolio performance, we introduce a novel methodology to calibrate the shrinkage intensity in the covariance matrix of stock returns, *shrinking against sentiment*, which shrinks more when investor sentiment is low and vice versa. We show that investor sentiment predicts lower market returns, higher arbitrage returns, and higher liquidity. Accordingly, our empirical results provide evidence that shrinking against sentiment delivers important economic gains relative to several benchmark portfolios both in the absence and in the presence of transaction costs.

Tables

Table 1: Predicting returns with sentiment

This table reports the intercept, slope coefficient, and Newey-West t-statistics (in square brackets) of long-run predictive regressions of the cumulative returns of a particular portfolio strategy on sentiment. In particular, we report the slope coefficients of the model:

$$r^f_{t+1\to t+k} = \alpha + \beta\,\text{Sentiment}_t + \epsilon_t,$$

where $r^f_{t+1\to t+k}$ are the cumulative returns of portfolio f during the period $(t+1, t+k)$, and Sentiment_t is the Huang et al. (2015) sentiment index at time t. We consider five portfolio strategies: 1) the Fama-French market portfolio, 2) the value factor (HML), 3) the profitability factor (RMW), 4) the investment factor (CMA), and 5) the optimal mean-variance portfolio that combines all the portfolios used in the construction of the HML, RMW and CMA factors, subject to the constraint that the weights of the portfolio add up to zero (MVE). The time period spans January 1965 to December 2018.

	1 month		1 quarter		1 year		2 years	
	Intercept	Slope	Intercept	Slope	Intercept	Slope	Intercept	Slope
Market	0.005	−0.007	0.015	−0.019	0.063	−0.052	0.126	−0.049
	[2.825]	[−3.708]	[3.442]	[−4.391]	[3.715]	[−2.647]	[3.870]	[−1.599]
HML	0.003	0.004	0.010	0.011	0.038	0.039	0.079	0.046
	[2.634]	[2.702]	[2.986]	[2.792]	[3.006]	[3.038]	[3.442]	[2.254]
RMW	0.003	0.003	0.008	0.008	0.032	0.025	0.066	0.026
	[2.884]	[2.260]	[3.277]	[2.636]	[3.167]	[1.731]	[3.416]	[1.204]
CMA	0.003	0.003	0.009	0.010	0.035	0.037	0.072	0.058
	[3.488]	[3.050]	[3.985]	[3.186]	[4.182]	[4.911]	[4.264]	[3.870]
MVE	0.016	0.009	0.049	0.027	0.196	0.114	0.397	0.175
	[7.342]	[3.221]	[8.195]	[3.706]	[7.289]	[3.333]	[6.546]	[2.967]

Table 2: Out-of-sample performance

This table reports the out-of-sample performance of the SAS portfolio and six benchmark portfolios: the mean-variance portfolio constructed with the shrinkage covariance matrix proposed by Ledoit and Wolf (2004) (MV-1N), the shrinkage covariance matrix proposed by Ledoit and Wolf (2003) (MV-MKT), the shrinkage covariance matrix proposed by Ledoit and Wolf (2020) (MV-NL), the principal-component-analysis covariance matrix proposed by Chen and Yuan (2016) (MV-PCA), the reward-to-risk timing strategy of Kirby and Ostdiek (2012) (RTR), and the equally weighted portfolio (1/N). Panel A reports the annualized out-of-sample Sharpe ratios of each portfolio, as well as the p-values (in parenthesis) of the significance test of the difference of each portfolio's Sharpe ratio with that of the equally weighted portfolio. Panel B reports the average monthly turnover of each portfolio. The six datasets are described in Section 4.2.

Policy	10MOM	25SBTM	25OPINV	49IND	16ANOM	46ANOM
Panel A: Sharpe ratios						
SAS	0.88	0.78	0.86	0.41	1.05	2.18
	(0.01)	(0.10)	(0.01)	(0.79)	(0.00)	(0.00)
MV-1N	0.89	0.47	0.59	0.19	1.00	2.00
	(0.02)	(0.66)	(0.39)	(0.94)	(0.01)	(0.00)
MV-MKT	0.89	0.34	0.55	0.20	0.94	1.86
	(0.02)	(0.83)	(0.45)	(0.94)	(0.02)	(0.00)
MV-NL	0.94	0.47	0.58	0.21	0.93	2.08
	(0.01)	(0.66)	(0.42)	(0.94)	(0.02)	(0.00)
MV-PCA	0.95	0.85	0.81	0.07	0.96	1.52
	(0.01)	(0.06)	(0.04)	(1.00)	(0.01)	(0.00)
RTR	0.61	0.62	0.60	0.56	0.58	0.60
	(0.00)	(0.00)	(0.00)	(0.08)	(0.00)	(0.00)
1/N	0.45	0.56	0.53	0.53	0.45	0.36
Panel B: Turnover						
SAS	0.27	0.48	0.40	0.33	0.45	0.53
MV-1N	1.73	2.11	1.66	2.51	1.68	12.56
MV-MKT	2.21	2.95	1.74	2.16	2.42	15.15
MV-NL	1.99	2.54	1.62	2.14	2.50	8.72
MV-PCA	2.31	0.90	0.33	0.32	0.84	1.63
RTR	0.05	0.04	0.04	0.05	0.05	0.04
1/N	0.04	0.04	0.04	0.05	0.04	0.04

Table 3: Out-of-sample performance net of transaction costs

This table reports the out-of-sample performance of the SAS portfolio and six benchmark portfolios: the mean-variance portfolio constructed with the shrinkage covariance matrix proposed by Ledoit and Wolf (2004) (MV-1N), the shrinkage covariance matrix proposed by Ledoit and Wolf (2003) (MV-MKT), the shrinkage covariance matrix proposed by Ledoit and Wolf (2020) (MV-NL), the principal-component-analysis covariance matrix proposed by Chen and Yuan (2016) (MV-PCA), the reward-to-risk timing strategy of Kirby and Ostdiek (2012) (RTR), and the equally weighted portfolio (1/N). Panel A, B, and C report the annualized out-of-sample Sharpe ratios of each portfolio net of proportional transaction costs of 10, 30, and 50 basis points, respectively, as well as the p-values (in parenthesis) of the significance test of the difference of each portfolio's Sharpe ratio with that of the equally weighted portfolio. The six datasets are described in Section 4.2.

Policy	10MOM	25SBTM	25OPINV	49IND	16ANOM	46ANOM
Panel A: Sharpe ratios net of 10bps						
SAS	0.87	0.76	0.84	0.39	1.03	2.16
	(0.01)	(0.15)	(0.03)	(0.79)	(0.00)	(0.00)
MV-1N	0.87	0.44	0.55	0.14	0.96	1.97
	(0.03)	(0.66)	(0.45)	(0.94)	(0.03)	(0.00)
MV-MKT	0.86	0.31	0.51	0.16	0.90	1.82
	(0.04)	(0.84)	(0.51)	(0.94)	(0.05)	(0.00)
MV-NL	0.91	0.42	0.54	0.17	0.89	2.03
	(0.02)	(0.69)	(0.47)	(0.93)	(0.05)	(0.00)
MV-PCA	0.92	0.82	0.79	0.05	0.93	1.50
	(0.05)	(0.22)	(0.12)	(0.02)	(0.03)	(0.00)
RTR	0.60	0.61	0.60	0.56	0.58	0.59
	(0.00)	(0.00)	(0.00)	(0.25)	(0.00)	(0.00)
1/N	0.45	0.56	0.53	0.53	0.45	0.36
Panel B: Sharpe ratios net of 30bps						
SAS	0.84	0.72	0.80	0.35	0.99	2.12
	(0.01)	(0.19)	(0.05)	(0.85)	(0.00)	(0.00)
MV-1N	0.81	0.36	0.46	0.05	0.87	1.89
	(0.05)	(0.76)	(0.60)	(0.97)	(0.05)	(0.00)
MV-MKT	0.80	0.22	0.43	0.07	0.80	1.70
	(0.06)	(0.91)	(0.66)	(0.97)	(0.09)	(0.00)
MV-NL	0.86	0.34	0.46	0.09	0.79	1.93
	(0.04)	(0.80)	(0.60)	(0.96)	(0.10)	(0.00)
MV-PCA	0.86	0.76	0.74	0.02	0.87	1.45
	(0.09)	(0.28)	(0.18)	(0.01)	(0.07)	(0.00)
RTR	0.60	0.61	0.59	0.55	0.57	0.59
	(0.00)	(0.00)	(0.00)	(0.25)	(0.00)	(0.00)
1/N	0.45	0.55	0.52	0.52	0.45	0.35
Panel C: Sharpe ratios net of 50bps						
SAS	0.82	0.68	0.76	0.30	0.95	2.08
	(0.01)	(0.25)	(0.08)	(0.90)	(0.00)	(0.00)
MV-1N	0.76	0.29	0.38	−0.05	0.79	1.78
	(0.08)	(0.84)	(0.74)	(0.99)	(0.10)	(0.00)
MV-MKT	0.74	0.14	0.35	−0.03	0.70	1.53
	(0.10)	(0.95)	(0.77)	(0.98)	(0.17)	(0.00)
MV-NL	0.80	0.24	0.39	−0.00	0.69	1.82
	(0.06)	(0.89)	(0.71)	(0.98)	(0.18)	(0.00)
MV-PCA	0.80	0.71	0.70	−0.02	0.81	1.40
	(0.14)	(0.44)	(0.33)	(0.00)	(0.12)	(0.00)
RTR	0.59	0.60	0.59	0.54	0.57	0.58
	(0.00)	(0.00)	(0.00)	(0.25)	(0.00)	(0.00)
1/N	0.44	0.55	0.52	0.51	0.44	0.34

Table 4: Time-series regressions

This table reports the intercept, slope coefficients, and Newey-West t-statistics (in square brackets) from the conditional CAPM time-series regressions in Equation (26). We run time-series regressions for the monthly out-of-sample returns of the SAS portfolio and the mean-variance portfolio that exploits the Ledoit and Wolf (2004) shrinkage covariance matrix (MV-1N) across the six datasets described in Section 4.2.

	10MOM		25SBTM		25OPINV		49IND		16ANOM		46ANOM	
	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N
α (%)	0.05	0.10	0.04	0.11	0.04	0.09	0.01	0.04	0.04	0.11	0.29	1.07
	[4.03]	[5.07]	[3.82]	[5.71]	[3.96]	[3.90]	[1.20]	[1.69]	[5.01]	[5.74]	[9.99]	[13.46]
β_0	3.74	1.52	3.55	1.92	3.53	2.01	4.84	2.80	2.82	1.65	4.07	−6.45
	[10.58]	[2.68]	[11.21]	[3.80]	[10.94]	[3.00]	[12.85]	[3.69]	[11.15]	[3.05]	[4.50]	[−2.86]
β_1	−0.35	−0.42	−1.95	−2.12	−1.90	−2.58	0.21	1.96	−1.49	−2.52	−2.62	−3.99
	[−0.82]	[−0.71]	[−5.51]	[−4.10]	[−3.57]	[−2.84]	[0.57]	[2.48]	[−4.94]	[−3.95]	[−1.72]	[−1.13]

Table 5: Liquidity and sentiment

This table reports the intercept, slope coefficient, and Newey-West t-statistics (in square brackets) of the time series regressions of cross-sectional measures of liquidity on lagged values of sentiment for the period January 1965 to December 2018. In each column, we report the results for the regressions of the cross-sectional 95, 75, 50, 25 and 5th percentiles of the estimate of bid-ask spreads in (27). We regress each cross-sectional percentile on the prior-month value of the sentiment index proposed by Huang et al. (2015).

	P95	P75	P50	P25	P5
Intercept	404	171	97.9	55.1	23.2
	[56.4]	[63.1]	[62.8]	[53.7]	[41.3]
Slope	−25.3	−6.23	−3.23	−3.00	−2.20
	[−3.43]	[−2.21]	[−2.00]	[−2.84]	[−4.20]

References

Abdi, Farshid, and Angelo Ranaldo, 2017, A simple estimation of bid-ask spreads from daily close, high, and low prices, *The Review of Financial Studies* 30, 4437–4480.

Bai, Jushan, and Serena Ng, 2002, Determining the number of factors in approximate factor models, *Econometrica* 70, 496–524.

Baker, Malcolm, and Jeremy C Stein, 2004, Market liquidity as a sentiment indicator, *Journal of Financial Markets* 7, 271–299.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *The Journal of Finance* 61, 1645–1680.

Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor sentiment in the stock market, *Journal of Economic Perspective* 21, 129–152.

Baker, Malcolm, Jeffrey Wurgler, and Yu Yuan, 2012, Global, local, and contagious investor sentiment, *Journal of Financial Economics* 104, 272–287.

Barroso, Pedro, and Andrew L. Detzel, 2021, Do limits to arbitrage explain the benefits of volatility-managed portfolios?, *Journal of Financial Economics* 140, 744–767.

Barroso, Pedro, and Konark Saxena, 2022, Lest we forget: using out-of-sample errors in portfolio optimization, *The Review of Financial Studies* 35, 1222–1278.

Best, Michael J., and Robert R. Grauer, 1992, Positively weighted minimum-variance portfolios and the structure of asset expected returns, *Journal of Financial and Quantitative Analysis* 27, 513–537.

Brandt, Michael W., and Pedro Santa-Clara, 2006, Dynamic portfolio selection by augmenting the asset space, *The Journal of Finance* 61, 2187–2217.

Brown, Stephen J., 1978, The portfolio choice problem: comparison of certainty equivalence and optimal bayes portfolios, *Communications in Statistics - Simulation and Computation* 7, 321–334.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2020, Forest through the trees: Building cross-sections of stock returns, *Available at SSRN 3493458*.

Chen, Jiaqin, and Ming Yuan, 2016, Efficient portfolio selection in a large market, *Journal of Financial Econometrics* 14, 496–524.

DeMiguel, Victor, Lorenzo Garlappi, Francisco J. Nogales, and Raman Uppal, 2009a, A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms, *Management Science* 55, 798–812.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2009b, Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?, *The Review of Financial Studies* 22, 1915–1953.

DeMiguel, Victor, Alberto Martin-Utrera, and Francisco J. Nogales, 2013, Size matters: Optimal calibration of shrinkage estimators for portfolio selection, *Journal of Banking and Finance* 37, 3018–3034.

DeMiguel, Victor, Alberto Martin-Utrera, Francisco J Nogales, and Raman Uppal, 2020, A transaction-cost perspective on the multitude of firm characteristics, *The Review of Financial Studies* 33, 2180–2222.

Fama, Eugene F, 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283–306.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, 2020, Taming the factor zoo: A test of new factors, *The Journal of Finance* 75, 1327–1370.

Ferson, Wayne E, and Andrew F Siegel, 2001, The efficient use of conditioning information in portfolios, *The Journal of Finance* 56, 967–982.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, 2020, Dissecting characteristics nonparametrically, *The Review of Financial Studies* 33, 2326–2377.

Garlappi, Lorenzo, Raman Uppal, and Tan Wang, 2007, Portfolio selection with parameter and model uncertainty: A multi-prior approach, *The Review of Financial Studies* 20, 41–81.

Goldfarb, D., and G. Iyengar, 2003, Robust portfolio selection problems, *Mathematics of Operations Research* 28, 1–38.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.

Hastie, Trevor, Robert Tibshirani, Jerome Friedman, and James Franklin, 2009, *The Elements of Statistical Learning: Data Mining, Inference and Prediction (2nd edition)* (Springer).

Hong, Harrison, and David A Sraer, 2016, Speculative betas, *The Journal of Finance* 71, 2095–2144.

Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou, 2015, Investor sentiment aligned: A powerful predictor of stock returns, *The Review of Financial Studies* 28, 791–837.

Jagannathan, Ravi, and Tongshu Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *The Journal of Finance* 58, 1651–1684.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani, 2013, *An Introduction to Statistical Learning* (Springer).

Kan, Raymond, Xiaolu Wang, and Guofu Zhou, 2021, Optimal portfolio choice with estimation risk: No risk-free asset case, *Management Science* 68, 2047–2068.

Kan, Raymond, and Guofu Zhou, 2007, Optimal portfolio choice with parameter uncertainty, *Journal of Financial and Quantitative Analysis* 42, 621–656.

Kelly, Brian, Semyon Malamud, and Lasse Heje Pedersen, 2021, Principal portfolios, *The Journal of Finance* (forthcoming).

Kirby, Chris, and Barbara Ostdiek, 2012, It's all in the timing: Simple active portfolio strategies that outperform naïve diversification, *Journal of Financial and Quantitative Analysis* 47, 437–467.

Klein, Roger W., and Vijay S. Bawa, 1976, The effect of estimation risk on optimal portfolio choice, *Journal of Financial Economics* 3, 215 – 231.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, Shrinking the cross-section, *Journal of Financial Economics* 135, 271–292.

Lassance, Nathan, Victor DeMiguel, and Fréderic Vrins, 2022, Optimal portfolio diversification via independent component analysis, *Operations Research* 70, 55–72.

Ledoit, Olivier, and Michael Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, *Journal of Empirical Finance* 10, 603–621.

Ledoit, Olivier, and Michael Wolf, 2004, A well-conditioned estimator for large-dimensional covariance matrices, *Journal of Multivariate Analysis* 88, 365–411.

Ledoit, Olivier, and Michael Wolf, 2017, Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks, *The Review of Financial Studies* 30, 4349–4388.

Ledoit, Olivier, and Michael Wolf, 2020, Analytical nonlinear shrinkage of large-dimensional covariance matrices, *The Annals of Statistics* 48, 3043–3065.

Ledoit, Olivier, and Michael Wolf, 2022, The power of (non-) linear shrinking: a review and guide to covariance matrix estimation, *Journal of Financial Econometrics* 20, 187–218.

Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *The Review of Economics and Statistics* 47, 13–37.

MacKinlay, A. Craig, and Lubos Pastor, 2000, Asset pricing models: Implications for expected returns and portfolio selection, *The Review of Financial Studies* 13, 883–916.

Markowitz, Harry, 1952, Portfolio selection, *The Journal of Finance* 7, 77–91.

Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Michaud, Richard O., 1989, The markowitz optimization enigma: Is 'optimized' optimal?, *Financial Analysts Journal* 45, 31–42.

Novy-Marx, Robert, and Mihail Velikov, 2016, A taxonomy of anomalies and their trading costs, *The Review of Financial Studies* 29, 104–147.

Pastor, Lubos, 2000, Portfolio selection and asset pricing models, *The Journal of Finance* 55, 179–223.

Pastor, Lubos, and Robert F. Stambaugh, 2000, Comparing asset pricing models: An investment perspective, *Journal of Financial Economics* 56, 335–381.

Pedersen, Lasse Heje, Abhilash Babu, and Ari Levine, 2021, Enhanced portfolio optimization, *Financial Analysts Journal* 77, 124–151.

Raponi, Valentina, Raman Uppal, and Paolo Zaffaroni, 2020, Robust portfolio choice, *Imperial College Working Paper.*

Schlkopf, Bernhard, Alexander Smola, and Francis Bach, 2018, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (MIT Press).

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium condition of risk, *The Journal of Finance* 19, 425–442.

Shen, Junyan, Jianfeng Yu, and Shen Zhao, 2017, Investor sentiment and economic forces, *Journal of Monetary Economics* 86, 1–21.

Stambaugh, Robert F, Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288–302.

Tu, Jun, and Guofu Zhou, 2010, Incorporating economic objectives into bayesian priors: Portfolio choice under parameter uncertainty, *Journal of Financial and Quantitative Analysis* 45, 959–986.

Tu, Jun, and Guofu Zhou, 2011, Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies, *Journal of Financial Economics* 99, 204 – 215.

Yu, Jianfeng, and Yu Yuan, 2011, Investor sentiment and the mean–variance relation, *Journal of Financial Economics* 100, 367–381.

Zhao, Long, Deepayan Chakrabarti, and Kumar Muthuraman, 2019, Portfolio construction by mitigating error amplification: The bounded-noise portfolio, *Operations Research* 67, 965–983.

Internet Appendix to

Shrinking Against Sentiment:

Exploiting Behavioral Biases in Portfolio Optimization

This internet appendix to the main paper contains two sections. Section A presents a theoretical model of the economy with investor sentiment. Section B details the proofs of all results in the paper and this appendix.

A A model of the economy with investor sentiment

In this section, we present a simple model similar to that considered by Hong and Sraer (2016) that builds intuition for the empirical results exposed in Section 2 by analyzing the theoretical relation between sentiment and asset prices in an economy where investors have heterogeneous beliefs.

A.1 The economy

First, we present the assumptions that define the economy.

Assumption A.1 *Let us assume an economy where:*

1. *Assets live for two periods, t and $t+1$.*
2. *There are N risky assets that pay random dividends at time $t+1$. Dividends d_{t+1} follow the process*
$$d_{t+1} = \overline{d} + u, \tag{A.1}$$
where $\overline{d}_i > 0$ for all $i = 1, \ldots, N$ and u follows a zero-mean multivariate distribution with covariance matrix Σ.
3. *All investors have mean-variance preferences with risk-aversion coefficient $\gamma > 0$.*
4. *A mass $\alpha \in [0, 1]$ of investors are optimistic about future cashflows. In particular, optimistic investors believe expected future dividends are $E[d_{t+1}] = \overline{d} + \delta$ with $\delta_i \geq 0$ for all $i = 1, \ldots, N$.*

Assumption A.1 lays out the main features of an economy plagued by mean-variance investors where a fraction α have *optimistic* expectations about future cashflows, while the remaining fraction $1 - \alpha$ are sophisticated investors with unbiased beliefs.

We assume there are only optimistic investors in the economy for tractability. However, this can be interpreted as the net effect of the excessive demand of risky assets from *optimistic* investors and the low demand of risky assets from *pessimistic* investors. The combination of investor disagreement and short-sale constraints leads to a situation where only the optimistic demand prevails. Indeed, Hong and Sraer (2016) note that in the presence of short-sale constraints, stocks subject to a higher level of disagreement are only held by optimistic investors in equilibrium because, due to short-sale constraints, the pessimistic investors are sidelined.[20] This leads to a situation where equilibrium prices are too high relative to the situation where all investors have homogeneous beliefs, and there are no short-sale constraints (Miller, 1977). As explained below, this effect is the reason why an increase of sentiment demand leads to lower market returns and an increase in the returns delivered by arbitrage portfolios that exploit asset overpricing.

A.2 The effect of sentiment on market and arbitrage portfolios

The fraction α of optimistic investors can be understood as the level of sentiment demand. We now show how the level of sentiment demand α affects the performance of the market and arbitrage portfolios in the economy described in Assumption A.1. First, in the next proposition, we derive the Sharpe ratio of the market portfolio that arises in equilibrium of the economy.

Proposition A.1 *Let Assumption A.1 hold. Then, in equilibrium, the Sharpe ratio of the market portfolio* w_M *is*

$$\mathrm{SR}(w_M) = \gamma\sigma_M - \frac{\alpha}{\sigma_M} w_M^\top \delta, \tag{A.2}$$

where α *is the mass of optimistic investors,* δ *is the sentiment demand vector, and* $\sigma_M = \sqrt{w_M^\top \Sigma w_M}$ *is the market volatility.*

Proposition A.1 shows that the equilibrium Sharpe ratio of the market portfolio is the sum of two terms: 1) the equilibrium Sharpe ratio when all investors in the market are

[20]It is straightforward to account for short-sale constraints at the expense of a more complex analysis. For simplicity, we omit this restriction in the analysis presented in this section.

rational ($\alpha = 0$) and 2) a term that gets increasingly negative as the level of sentiment demand α increases. In particular, Proposition A.1 shows that as the mass of sentiment investors α in the economy increases, the market Sharpe ratio decreases, just like we observe empirically in Section 2.

Next, we study the effect of sentiment demand on the performance of the arbitrage portfolio w_A defined in (2). Under Assumption 1, we can characterize theoretically the positive relationship between the level of sentiment demand α and the Sharpe ratio of the arbitrage portfolio w_A. Specifically, in the assumed economy, the following holds.

Proposition A.2 *Let Assumptions 1 and A.1 hold. Then, in equilibrium, the Sharpe ratio of the arbitrage portfolio is*

$$\mathrm{SR}(w_A) = \alpha \sqrt{\sum_{i=2}^{N} \frac{\left(v_i^\top \delta\right)^2}{\sigma_{PC_i}^2}}, \tag{A.3}$$

where α is the mass of optimistic investors, v_i is the ith eigenvector of the covariance matrix Σ, and $\sigma_{PC_i}^2 = v_i^\top \Sigma v_i$ is the variance of the ith principal component of returns.

Proposition A.2 characterizes the positive theoretical relationship between the level of sentiment demand α and the performance of the mean-variance arbitrage portfolio. In particular, the Sharpe ratio of the arbitrage portfolio is linear in α and attains its lowest value, $\mathrm{SR}(w_A) = 0$, when $\alpha = 0$. This theoretical prediction holds also in the data. Specifically, in Section 2 we show that the mean-variance arbitrage portfolio that combines all the portfolios used in the construction of the HML, RMW, and CMA factors is positively correlated with sentiment and that changes in sentiment can have long-lasting effects on the performance of the mean-variance arbitrage portfolio.

B Proofs of all results

This section contains the proofs of all the propositions in the main body of the paper and this internet appendix

Proof of Proposition 1

Using the eigenvalue decomposition of the covariance matrix, we can easily characterize the squared Sharpe ratio of the mean-variance portfolio as the sum of the squared Sharpe ratio of each of the principal components of stock returns. In particular, we can write the covariance matrix Σ as

$$\Sigma = VDV^\top, \tag{A.4}$$

where D is a diagonal matrix whose ith element is the eigenvalue associated to the ith principal component, and V is a matrix whose ith column v_i is the eigenvector associated to the ith principal component. Given this decomposition, it is straightforward to show that the mean-variance portfolio's squared Sharpe ratio is[21]

$$\mathrm{SR}^2(w^\star) = \sum_{i=1}^{N} \frac{\mu_{PC_i}^2}{\sigma_{PC_i}^2}, \tag{A.5}$$

where $\mu_{PC_i} = v_i^\top \mu$ and $\sigma_{PC_i}^2 = v_i^\top \Sigma v_i$ are the average return and variance of the ith principal component of stock returns, respectively.

Part 1. Given the eigenvalue decomposition of the Sharpe ratio of the mean-variance portfolio in (A.5), we need to prove the equalities for the Sharpe ratio of the market and arbitrage portfolios in Equations (7) and (8). Let us begin with the equally weighted market portfolio, $w_M = \iota/N$. Under Assumption 1, the first eigenvector is $v_1 = \iota/\sqrt{N}$. Consequently, the market average return and variance are

$$\mu_M = w_M^\top \mu = \frac{\mu^\top v_1}{\sqrt{N}} = \frac{\mu_{PC_1}}{\sqrt{N}}, \tag{A.6}$$

[21] This result is similar to the SDF-variance decomposition introduced by Kozak, Nagel, and Santosh (2018, Equation (4)).

$$\sigma_M^2 = w_M^\top \Sigma w_M = \frac{v_1^\top V\, D\, V^\top v_1}{N} = \frac{\sigma_{PC_1}^2}{N}, \tag{A.7}$$

and the market Sharpe ratio is $\mathrm{SR}(w_M) = \mu_M/\sigma_M = \mu_{PC_1}/\sigma_{PC_1}$, as in Equation (8).

Second, we derive the eigenvalue decomposition of the Sharpe ratio of the arbitrage portfolio, w_A in (2). The average return and variance of the arbitrage portfolio are

$$w_A^\top \mu = \frac{1}{\gamma}(\mu - \mu_g \iota)^\top \Sigma^{-1} \mu, \tag{A.8}$$

$$w_A^\top \Sigma w_A = \frac{1}{\gamma^2}(\mu - \mu_g \iota)^\top \Sigma^{-1} (\mu - \mu_g \iota). \tag{A.9}$$

Now, using the definition of $\mu_g = \frac{\iota^\top \Sigma^{-1} \mu}{\iota^\top \Sigma^{-1} \iota}$, and plugging it into the above expressions, we have

$$w_A^\top \mu = \frac{1}{\gamma}\left(\mu^\top \Sigma^{-1} \mu - \frac{\left(\iota^\top \Sigma^{-1} \mu\right)^2}{\iota^\top \Sigma^{-1} \iota}\right) \quad \text{and} \quad w_A^\top \Sigma w_A = \frac{1}{\gamma} w_A^\top \mu. \tag{A.10}$$

Therefore, the squared Sharpe ratio of the arbitrage portfolio is

$$\mathrm{SR}^2(w_A) = \frac{\left(w_A^\top \mu\right)^2}{w_A^\top \Sigma w_A} = \mu^\top \Sigma^{-1} \mu - \frac{\left(\iota^\top \Sigma^{-1} \mu\right)^2}{\iota^\top \Sigma^{-1} \iota}. \tag{A.11}$$

Let us now recall the eigenvalue decomposition of the covariance matrix as $\Sigma = V\, D\, V^\top$, where V is the matrix of the eigenvectors associated to the principal components of stock returns, and D is a diagonal matrix that contains the eigenvalues associated to each of the principal component of stock returns, sorted in decreasing order. We can use the eigenvalue decomposition of the covariance matrix to further decompose the Sharpe ratio of the arbitrage portfolio. In particular,

$$\mu^\top \Sigma^{-1} \mu = \mu^\top V\, D^{-1}\, V^\top \mu = \sum_{i=1}^{N} \frac{\mu_{PC_i}^2}{\sigma_{PC_i}^2}, \tag{A.12}$$

$$\iota^\top \Sigma^{-1} \mu = \sqrt{N} v_1^\top V\, D^{-1}\, V^\top \mu = \sqrt{N} \frac{\mu_{PC_1}}{\sigma_{PC_1}^2}, \tag{A.13}$$

$$\iota^\top \Sigma^{-1} \iota = N v_1^\top V\, D^{-1}\, V^\top v_1 = \frac{N}{\sigma_{PC_1}^2}, \tag{A.14}$$

where $\mu_{PC_i} = v_i^\top \mu$ and $\sigma^2_{PC_i} = v_i^\top \Sigma v_i$ are the average return and variance of the ith principal component. Plugging expressions (A.12)-(A.14) into (A.11), we obtain Equation (8), which completes the proof.

Part 2. This result is a special case of the result in Proposition 3 when the shrinkage intensity $a = 0$.

Proof of Proposition 2

The squared Sharpe ratio of the equally weighted market portfolio is independent of the shrinkage intensity a and, as shown in (7), is given by $\mathrm{SR}^2(w_M) = \mu^2_{PC_1}/\sigma^2_{PC_1}$, which corresponds to $S_a(1,1)$ in (18) for any value of a.

We turn next to the shrinkage mean-variance portfolio in (14), $w_S^\star = \frac{1}{\gamma}\Sigma_{sh}^{-1}\mu$, where $\Sigma_{sh} = V((1-a)D + a\nu I_N)V^\top$. Its average return is

$$(w_S^\star)^\top \mu = \frac{1}{\gamma}\mu^\top V((1-a)D + a\nu I_N)^{-1}V^\top \mu = \frac{1}{\gamma}\sum_{i=1}^{N} \frac{\mu^2_{PC_i}}{(1-a)\sigma^2_{PC_i} + a\nu}. \tag{A.15}$$

Moreover, its variance is

$$(w_S^\star)^\top \Sigma w_S^\star = \frac{1}{\gamma^2}\mu^\top \Sigma_{sh}^{-1}\Sigma\Sigma_{sh}^{-1}\mu.$$

Given the eigenvalue decompositions of Σ and Σ_{sh}, we have that $\Sigma_{sh}^{-1}\Sigma\Sigma_{sh}^{-1} = VD_aV^\top$, where D_a is a diagonal matrix with entries

$$(D_a)_{ii} = \frac{\sigma^2_{PC_i}}{((1-a)\sigma^2_{PC_i} + a\nu)^2}. \tag{A.16}$$

Therefore, the variance of the shrinkage mean-variance portfolio becomes

$$(w_S^\star)^\top \Sigma w_S^\star = \frac{1}{\gamma^2}\sum_{i=1}^{N} \frac{\mu^2_{PC_i}\sigma^2_{PC_i}}{((1-a)\sigma^2_{PC_i} + a\nu)^2}. \tag{A.17}$$

Finally, using Equations (A.15)–(A.17), the squared Sharpe ratio of the shrinkage mean-variance portfolio is

$$\mathrm{SR}^2(w_S^\star) = \frac{((w_S^\star)^\top \mu)^2}{(w_S^\star)^\top \Sigma w_S^\star} = S_a(1,N), \tag{A.18}$$

where S_a is defined in (18).

A similar line of reasoning holds for the shrinkage arbitrage portfolio in (13),

$$w_{SA} = \frac{1}{\gamma}\Sigma_{sh}^{-1}\left(\mu - \frac{\iota^\top \Sigma_{sh}^{-1}\mu}{\iota^\top \Sigma_{sh}^{-1}\iota}\iota\right).$$

Under Assumption 1, it holds that $\frac{\iota^\top \Sigma_{sh}^{-1}\mu}{\iota^\top \Sigma_{sh}^{-1}\iota} = \mu_{PC_1}/\sqrt{N}$, and thus,

$$w_{SA} = \frac{1}{\gamma}\Sigma_{sh}^{-1}\left(\mu - \frac{\mu_{PC_1}}{\sqrt{N}}\iota\right). \tag{A.19}$$

Therefore, the average return of the shrinkage arbitrage portfolio is

$$w_{SA}^\top \mu = \frac{1}{\gamma}\left(\mu^\top \Sigma_{sh}^{-1}\mu - \mu^\top \Sigma_{sh}^{-1}\iota \frac{\mu_{PC_1}}{\sqrt{N}}\right) = \frac{1}{\gamma}\sum_{i=2}^{N} \frac{\mu_{PC_i}^2}{(1-a)\sigma_{PC_i} + a\nu}, \tag{A.20}$$

which holds because $\mu^\top \Sigma_{sh}^{-1}\iota = \sqrt{N}\frac{\mu_{PC_1}}{(1-a)\sigma_{PC_i}+a\nu}$ under Assumption 1. Moreover, the variance of the shrinkage arbitrage portfolio is

$$\begin{aligned} w_{SA}^\top \Sigma w_{SA}^\top &= \frac{1}{\gamma^2}\left(\mu - \frac{\mu_{PC_1}}{\sqrt{N}}\iota\right)^\top \Sigma_{sh}^{-1}\Sigma\Sigma_{sh}^{-1}\left(\mu - \frac{\mu_{PC_1}}{\sqrt{N}}\iota\right) \\ &= (w_S^\star)^\top \Sigma w_S^\star + \frac{1}{\gamma^2}\left(\frac{\mu_{PC_1}^2}{N}\iota^\top V D_a V^\top \iota - 2\frac{\mu_{PC_1}}{\sqrt{N}}\iota^\top V D_a V^\top \mu\right) \\ &= \frac{1}{\gamma^2}\sum_{i=2}^{N} \frac{\mu_{PC_i}^2 \sigma_{PC_i}^2}{((1-a)\sigma_{PC_i}^2 + a\nu)^2}, \end{aligned} \tag{A.21}$$

where the diagonal matrix D_a is defined in (A.16). The last equality in (A.21) holds because, under Assumption 1,

$$\iota^\top V D_a V^\top \iota = N\frac{\sigma_{PC_1}^2}{((1-a)\sigma_{PC_i}^2 + a\nu)^2} \quad \text{and} \quad \iota^\top V D_a V^\top \mu = \sqrt{N}\frac{\mu_{PC_1}\sigma_{PC_1}^2}{((1-a)\sigma_{PC_i}^2 + a\nu)^2}.$$

Finally, using Equations (A.20)–(A.21), the squared Sharpe ratio of the shrinkage arbitrage portfolio is

$$\mathrm{SR}^2(w_{SA}) = \frac{(w_{SA}^\top \mu)^2}{w_{SA}^\top \Sigma w_{SA}} = S_a(2, N), \tag{A.22}$$

where S_a is defined in (18), which completes the proof.

Proof of Proposition 3

Let us begin with the correlation between the returns of the market portfolio and those of the shrinkage mean-variance portfolio. This correlation is

$$\mathrm{Corr}(w_M^\top R, (w_S^\star)^\top R) = \frac{w_M^\top \Sigma w_S^\star}{\sqrt{(w_M^\top \Sigma w_M)((w_S^\star)^\top \Sigma w_S^\star)}}. \tag{A.23}$$

The variance of w_M and $w_S^\star$ are given by Equations (A.7) and (A.17), respectively, and thus we only need to treat the term $w_M^\top \Sigma w_S^\star$, which under Assumption 1 simplifies to

$$w_M^\top \Sigma w_S^\star = \frac{1}{\gamma N} \iota^\top \Sigma \Sigma_{sh}^{-1} \mu = \frac{1}{\gamma \sqrt{N}} v_1^\top V \widetilde{D}_a V^\top \mu,$$

where $\widetilde{D}_a$ is a diagonal matrix with entries

$$\left(\widetilde{D}_a\right)_{ii} = \frac{\sigma_{PC_i}^2}{(1-a)\sigma_{PC_i}^2 + a\nu}.$$

Therefore, the quantity $w_M^\top \Sigma w_S^\star$ is

$$w_M^\top \Sigma w_S^\star = \frac{1}{\gamma \sqrt{N}} \frac{\mu_{PC_1} \sigma_{PC_1}^2}{(1-a)\sigma_{PC_1}^2 + a\nu}. \tag{A.24}$$

Plugging (A.7), (A.17), and (A.24) into (A.23), we find that the squared correlation between the returns of the market portfolio and those of the shrinkage mean-variance portfolio is given by (19).

We turn next to the correlation between the returns of the shrinkage arbitrage portfolio and those of the shrinkage mean-variance portfolio. This correlation is

$$\mathrm{Corr}(w_{SA}^\top R, (w_S^\star)^\top R) = \frac{w_{SA}^\top \Sigma w_S^\star}{\sqrt{(w_{SA}^\top \Sigma w_{SA})((w_S^\star)^\top \Sigma w_S^\star)}}. \tag{A.25}$$

The variance of w_{SA} and $w_S^\star$ are given by Equations (A.21) and (A.17), respectively, and thus we only need to treat the term $w_{SA}^\top \Sigma w_S^\star$, which under Assumption 1 simplifies to

$$w_{SA}^\top \Sigma w_S^\star = \frac{1}{\gamma^2} \left(\mu - \frac{\mu_{PC_1}}{\sqrt{N}} \iota \right)^\top \Sigma_{sh}^{-1} \Sigma \Sigma_{sh}^{-1} \mu.$$

As shown in the proof of Proposition 2, we have decomposition $\Sigma_{sh}^{-1}\Sigma\Sigma_{sh}^{-1} = VD_aV^\top$, where D_a is a diagonal matrix defined in (A.16). Using this result, the quantity $w_{SA}^\top \Sigma w_S^\star$ is

$$w_{SA}^\top \Sigma w_S^\star = \frac{1}{\gamma^2} \sum_{i=2}^{N} \frac{\mu_{PC_i}^2 \sigma_{PC_i}^2}{((1-a)\sigma_{PC_i}^2 + a\nu)^2}. \tag{A.26}$$

Plugging (A.21), (A.17), and (A.26) into (A.25), we find that the correlation between the returns of the shrinkage arbitrage portfolio and those of the shrinkage mean-variance portfolio is given by (20).

Finally, we conclude the proof by showing that the squared correlation with the market $\text{Corr}^2\big(w_M^\top R, (w_S^\star)^\top R\big)$, is increasing in the shrinkage intensity a, which implies that the correlation with the shrinkage arbitrage portfolio is decreasing in a given that the two squared correlations sum up to one in (20). Given Equation (19), to prove this result it is sufficient to show that the quantity

$$\frac{(1-a)\sigma_{PC_1}^2 + a\nu}{(1-a)\sigma_{PC_i}^2 + a\nu}$$

is decreasing in a for all $i = 2, \ldots, N$. The derivative of this quantity with respect to a is

$$\frac{\partial}{\partial a} \frac{(1-a)\sigma_{PC_1}^2 + a\nu}{(1-a)\sigma_{PC_i}^2 + a\nu} = \frac{(\nu - \sigma_{PC_1}^2)((1-a)\sigma_{PC_i}^2 + a\nu) - (\nu - \sigma_{PC_i}^2)((1-a)\sigma_{PC_1}^2 + a\nu)}{(((1-a)\sigma_{PC_i}^2 + a\nu)^2}$$

and is negative under the condition

$$\frac{\partial}{\partial a} \frac{(1-a)\sigma_{PC_1}^2 + a\nu}{(1-a)\sigma_{PC_i}^2 + a\nu} \leq 0 \Longleftrightarrow \sigma_{PC_i}^2 \leq \sigma_{PC_1}^2,$$

which holds indeed for all $i = 2, \ldots, N$ because the eigenvalues are sorted in decreasing order, thus completing the proof.

Proof of Proposition A.1

The optimality conditions of the sophisticated and the optimistic investors give the following optimal mean-variance portfolios:

$$w_l^\star = \frac{1}{\gamma} \Sigma^{-1} \Big(E_l[d_{t+1}] - P_t \Big), \tag{A.27}$$

where P_t is the vector of equilibrium prices and $E_l[d_{t+1}]$ for $l = \{s, o\}$ is the vector of expected cashflows for the sophisticated (s) and optimistic (o) investors. Under Assumption A.1, the market clearing condition is

$$\frac{\alpha}{\gamma}\Sigma^{-1}\left(\overline{d} + \delta - P_t\right) + \frac{1-\alpha}{\gamma}\Sigma^{-1}\left(\overline{d} - P_t\right) = w_M, \tag{A.28}$$

where w_M is the market portfolio. From the clearing condition (A.28), the vector of equilibrium prices P_t is

$$P_t = \overline{d} - \gamma\beta_M\sigma_M^2 + \alpha\delta, \tag{A.29}$$

where $\sigma_M^2 = w_M^\top \Sigma w_M$ and $\beta_M = \frac{\Sigma w_M}{\sigma_M^2}$ is the vector of market betas. Accordingly, the vector of equilibrium average returns μ is

$$\mu = E[R_{t+1}] = E[d_{t+1} - P_t] = \overline{d} - P_t = \gamma\beta_M\sigma_M^2 - \alpha\delta. \tag{A.30}$$

The first term in Equation (A.30) corresponds to the equilibrium average returns when all investors in the market are rational (i.e., when $\alpha = 0$), and the second element, $\alpha\delta$, accounts for the effect that heterogeneous beliefs have on average returns. Finally, using expression (A.30), we have that the market average return is $\mu_M = w_M^\top \mu = \gamma\sigma_M^2 - \alpha\delta^\top w_M$, and dividing by the market volatility σ_M yields the market Sharpe ratio in Equation (A.2), which completes the proof.

Proof of Proposition A.2

The vector of equilibrium average returns μ in the economy in Assumption A.1 is $\mu = \gamma\beta_M\sigma_M^2 - \alpha\delta$, as shown in Equation (A.30). Further, under Assumption 1, the Sharpe ratio of the arbitrage portfolio admits the decomposition in Equation (8). Plugging μ into this equation yields

$$\text{SR}(w_A) = \sqrt{\sum_{i=2}^{N} \frac{\left(v_i^\top(\gamma\beta_M\sigma_M^2 - \alpha\delta)\right)^2}{\sigma_{PC_i}^2}}, \tag{A.31}$$

where v_i is the ith eigenvector of the covariance matrix Σ. The result follows by noticing that $v_i^\top \beta_M = 0$ for all $i > 1$ from the assumption that the first eigenvector is $v_1 = \iota/\sqrt{N}$, which completes the proof.

References in Internet Appendix

Hong, Harrison, and David A Sraer, 2016, Speculative betas, *The Journal of Finance* 71, 2095–2144.

Miller, Edward M, 1977, Risk, uncertainty, and divergence of opinion, *The Journal of Finance* 32, 1151–1168.