

Shrinking Against Sentiment:

Exploiting Latent Asset Demand in Portfolio Selection^{*}

Nathan Lassance Alberto Martín-Utrera

February 6, 2023

Abstract

We examine how sentiment-driven demand, a key component of latent asset demand, can be used to build mean-variance portfolios. We decompose these portfolios into an equally weighted component and an arbitrage component that captures the asset mispricing unexplained by the equally weighted component. Our approach shrinks mean-variance portfolios toward the equally weighted component when investor sentiment is low, i.e., shrinks against sentiment, reducing estimation risk and imposing a tighter bound on the amount of asset mispricing the arbitrage component exploits. The significant economic gains offered by our approach highlight the importance of considering latent demand in building robust investment strategies.

Keywords: Sentiment, mean-variance portfolios, latent demand.

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 under Grant number J.0115.22. The authors thank helpful comments from Mathijs Cosemans, Victor DeMiguel, Olivier Ledoit, André Portela-Santos, Majeed Simaan, Raman Uppal, Rodolphe Vanderveken, Xiaolu Wang, Alex Wassensteiner, Michael Wolf, Guofu Zhou, conference participants at the 3rd Frontiers of Factor Investing, INFORMS Annual Meeting, 16th International Conference on Computational and Financial Econometrics, and seminar participants at Iowa State University, UCLouvain, New Jersey Institute of Technology, and Stevens Institute of Technology.

1 Introduction

Latent asset demand is pervasive in financial markets and affects stock returns. In a recent study, [Koijen and Yogo \(2019\)](#) find that “*stock returns are mostly explained by demand shocks that are unrelated to changes in observed characteristics*”.¹ One important component of latent demand is *sentiment*-driven demand, which [Baker and Wurgler \(2007\)](#) define as the asset demand based on the “*belief about future cash flows and investment risks that is not justified by the facts at hand*”.² The evidence that latent demand affects asset prices begs the question of whether one can utilize it to construct portfolio strategies. Thus, our objective is to link latent asset demand with the performance of mean-variance portfolios and exploit it to construct robust investment strategies.

The link between latent asset demand and the performance of mean-variance portfolios relies on the influence that sentiment (our proxy for latent demand) has on the returns of the equally weighted ($1/N$) portfolio and arbitrage returns. In particular, sentiment predicts $1/N$ returns negatively and arbitrage returns positively.³ Figure 1 provides anecdotal evidence of these effects. For instance, after the substantial increase in investor sentiment during the dot-com bubble, the $1/N$ Sharpe ratio was negative while the an-


¹This is a striking finding because observable characteristics relate to the loadings on common factors, as in [Koijen and Yogo \(2019\)](#), and should explain the cross-section ([Ross, 1976](#)). While surprising, the idea that common sources of return variation are unpriced has received recent empirical support. In a thought-provoking paper, [Lopez-Lira and Roussanov \(2023\)](#) document that measurable common risk factors are not priced, and a long-short portfolio hedging all sources of common variation delivers a Sharpe ratio above one. Our work sheds light on this topic and shows that long-short portfolios uncorrelated to the main drivers of return variation can indeed deliver high Sharpe ratios, but their good performance is concentrated during times of high latent demand.

²[Koijen, Richmond, and Yogo \(2020\)](#) say that latent demand “can arise from investor sentiment or portfolio constraints that are not easily modeled as a function of observed characteristics”.

³[Huang et al. \(2015\)](#) document a negative relation between sentiment and future value-weighted market returns. Unlike them, our analysis focuses on equally weighted ($1/N$) returns, which highly correlate with value-weighted market returns. Indeed, the insights of [Huang et al. \(2015\)](#) hold for $1/N$. We find empirically that, on average, the $1/N$ portfolio delivers an annualized Sharpe ratio of 0.85 in low-sentiment months, and 0.26 in high-sentiment months. The opposite relation holds for arbitrage returns where there is a strong positive relation between investor sentiment and subsequent arbitrage returns, as documented by [Stambaugh, Yu, and Yuan \(2012\)](#).

Figure 1: Sentiment and the two components of mean-variance portfolios

The top panel depicts the Huang, Jiang, Tu, and Zhou (2015) sentiment index from July 1965 to June 2022, which is standardized to have zero mean and unit variance. The bottom panel depicts the annualized Sharpe ratio delivered by the $1/N$ portfolio and the arbitrage portfolio in three different periods: (i) March 2001 to March 2003, (ii) June 2009 to December 2019, and (iii) December 2021 to June 2022. The Sharpe ratios are obtained from the dataset of monthly excess returns on the 25 portfolios of stocks sorted on size and book-to-market.



nualized Sharpe ratio of a long-short arbitrage portfolio constructed from a dataset of stocks sorted on size and book-to-market was around 1.4.⁴ On the contrary, during the post financial crisis period, when sentiment was low, the $1/N$ annualized Sharpe ratio was about 1, whereas that of the arbitrage portfolio was only 0.2. More recently, from December 2021 to June 2022, we see similar evidence to the dot-com bubble, where a sharp increase in investor sentiment is followed by a poor performance of $1/N$ and a strong performance of the arbitrage portfolio.

These contrasting effects have important implications for investment decisions because mean-variance portfolios can be decomposed into the sum of the $1/N$ portfolio and an arbitrage portfolio. Thus, we introduce a novel portfolio framework—*shrinking against*

⁴Section 3 provides details of how we construct the arbitrage portfolio.

sentiment—that tilts mean-variance portfolios toward $1/N$ when sentiment is low and toward the arbitrage portfolio when investor sentiment is high. To the best of our knowledge, our sentiment-driven methodology is the first attempt at incorporating a real-time proxy for latent asset demand in the construction of optimal portfolios.

Our manuscript makes four contributions. First, we decompose mean-variance portfolios as the sum of $1/N$ and an arbitrage portfolio. This mean-variance decomposition motivates our shrinkage methodology that exploits sentiment. Intuitively, mean-variance investors will benefit from an increased exposure toward the $1/N$ (arbitrage) component when sentiment is low (high) and vice versa due to the negative (positive) relation between sentiment demand and $1/N$ (arbitrage) average returns. In Section A of the Internet Appendix, we also provide a theoretical foundation for shrinking against sentiment where we show that in an economy with rational and sentiment investors where asset mispricing (i.e., asset alphas) is a function of sentiment demand, $1/N$ average returns are negatively affected by sentiment, and arbitrage average returns are positively affected by sentiment.

Second, we analytically characterize the mean-variance portfolio exposure to the $1/N$ and arbitrage components. We show that because the mean-variance portfolio correlation with each component is proportional to their in-sample Sharpe ratio, mean-variance portfolios are overly influenced by the arbitrage component, which carries substantial estimation risk. We analytically characterize the out-of-sample performance losses of arbitrage portfolios and document that these are indeed very large and range between 50 and 80% in our empirical tests.

Third, we show that shrinking the covariance matrix of a mean-variance portfolio allows us to increase the correlation between mean-variance and $1/N$ returns, which reduces the impact of estimation risk stemming from the arbitrage component. In addition, we show that shrinking the covariance matrix of stock returns decreases the in-sample Sharpe ratio of the arbitrage portfolio, and this is equivalent to imposing a no-arbitrage

condition in the sense of the Arbitrage Pricing Theory of Ross (1976). Shrinking against sentiment also allows us to shy away from the arbitrage portfolio when sentiment is low, which has two benefits. First, the positive relation between sentiment and subsequent arbitrage returns justifies reducing the importance of the arbitrage portfolio during low-sentiment regimes. Second, we show that the arbitrage portfolio's out-of-sample performance losses due to estimation errors are the largest when its in-sample performance is low, typically during low-sentiment periods.

Our fourth contribution is to illustrate the performance gains of our portfolio methodology across six empirical datasets of characteristic-sorted portfolios. Our analysis shows that the average outperformance in terms of Sharpe ratio of the *shrinking-against-sentiment* (SAS) portfolio relative to the four benchmark mean-variance portfolios we consider is 41%. Moreover, the benchmark mean-variance portfolios require, on average, a turnover that is 461% higher than that of the SAS portfolio. Therefore, the performance gains are even more substantial in the presence of transaction costs. For instance, for transaction costs of 30 basis points, the average outperformance of the SAS portfolio relative to the benchmark mean-variance portfolios is 71%, and 133% relative to the 1/N portfolio. Our empirical analysis illustrates that SAS is a statistically and economically motivated approach to portfolio selection that delivers sizable economic gains.

Finally, we provide further intuition on why exploiting latent asset demand is important for portfolio selection. In particular, we evaluate the out-of-sample Sharpe ratio of the 1/N portfolio, the arbitrage portfolio, and the SAS portfolio following high and low-sentiment periods. We show that, on average, the 1/N portfolio delivers an annualized Sharpe ratio of 0.24 in high-sentiment periods and 0.87 in low-sentiment periods. On the contrary, the arbitrage portfolio delivers a Sharpe ratio of 1.19 in high-sentiment periods and 0.47 in low-sentiment periods. Unlike the 1/N and arbitrage portfolios, the SAS portfolio performs consistently well across both sentiment regimes. Specifically, it deliv-

ers an average Sharpe ratio of 1.46 in high-sentiment periods and 0.87 in low-sentiment periods. Overall, we show that the economic gains of the SAS portfolio relative to the benchmark mean-variance portfolios stem from the substantially stronger performance that SAS delivers during low-sentiment regimes.

2 Literature review

Our work is related to [Kojen and Yogo \(2019\)](#), which shows that latent demand has a dominating effect in explaining the cross-section of asset returns. In a world where asset returns follow a particular one-factor model, our arbitrage portfolio harvests the mispricing associated with the *latent* asset demand of sentiment investors; see Section A of the Internet Appendix. Accordingly, our portfolio approach uses investor sentiment as a proxy for latent asset demand to optimally time the exposure to the component in mean-variance portfolios that harvests sentiment-driven mispricing.

Our work is also closely related to [Raponi, Uppal, and Zaffaroni \(2021\)](#), which also decomposes mean-variance portfolios as the sum of two components: an *alpha* and a *beta* component. Their approach builds on the Arbitrage Pricing Theory (APT) of [Ross \(1976\)](#). We show in Proposition 2 that our mean-variance portfolio decomposition is equivalent to the alpha and beta decomposition of [Raponi, Uppal, and Zaffaroni \(2021\)](#) when asset returns follow a specific one-factor model. Unlike [Raponi, Uppal, and Zaffaroni \(2021\)](#) who develop robust *static* portfolios, our methodology exploits sentiment to build a robust *dynamic* strategy that tilts the performance of mean-variance portfolios toward $1/N$ when sentiment is low, and toward the arbitrage portfolio when sentiment is high.

The portfolio shrinkage approach we consider in this manuscript is a form of regularization that is related to portfolio problems that account for *time-invariant* transaction costs ([Oliveres-Nadal and DeMiguel, 2018](#)). In contrast, our methodology allows us to in-

corporate *time-varying* liquidity considerations because as [Baker and Stein \(2004\)](#) show, sentiment is positively correlated with market liquidity. Therefore, our shrinkage portfolio approach is sensible not only because it harvests a larger premium from $1/N$ when sentiment is low, but also because it allows us to build a low-turnover strategy when liquidity decreases.

Our shrinkage portfolio approach can also be cast as a regression problem with norm constraints ([Britten-Jones, 1999](#); [DeMiguel, Garlappi, Nogales, and Uppal, 2009](#)). Accordingly, our methodology shares elements with the work of [Campbell and Thompson \(2008\)](#), [Pettenuzzo, Timmermann, and Valkanov \(2014\)](#), and [Bakalli, Guerrier, and Scaillet \(2023\)](#), who impose economic constraints in forecasting regressions for the equity premium. Like them, we exploit the theoretical connection between an economic variable and asset prices to impose a particular structure to the solution of the regression problem. In particular, we use sentiment to strategically tilt the performance of mean-variance portfolios toward $1/N$ or the arbitrage component.

Our work is also related to the literature documenting the economic gains in portfolio selection of methods that exploit the predictability of stock returns ([DeMiguel, Nogales, and Uppal, 2014](#); [Gu, Kelly, and Xiu, 2020](#)). Indeed, serial autocorrelation and machine learning techniques can substantially improve performance. However, these techniques come at the expense of higher portfolio turnover. In contrast to this literature, we focus on the sentiment-driven predictability of the two components defining mean-variance portfolios. Our results indicate that our proposed methodology delivers better performance than cutting-edge shrinkage techniques at a substantially lower turnover.

Other papers have used investor sentiment to *explain* the performance of investment strategies. For instance, [Stambaugh, Yu, and Yuan \(2012\)](#) show that many asset pricing anomalies only deliver significant abnormal returns during high-sentiment periods. In contrast to [Stambaugh, Yu, and Yuan \(2012\)](#), our work provides a methodology to utilize

sentiment in the classic mean-variance framework, which accounts for the contrasting effect that sentiment has on $1/N$ and anomaly returns. Also, we show that our proposed methodology delivers sizable economic gains out of sample and net of transaction costs.

Finally, our work also shares some elements with that of [Chu, He, Li, and Tu \(2022\)](#), who show that sentiment has opposing effects on the ability of fundamental and nonfundamental economic variables to predict the equity premium. Unlike [Chu et al. \(2022\)](#), our manuscript shows that sentiment affects differently the performance of the two components that define mean-variance portfolios.

3 A decomposition of mean-variance portfolios

In this section, we decompose the performance of mean-variance portfolios as the sum of two components: a $1/N$ component and an arbitrage component. Our decomposition shows that the performance of the mean-variance portfolio is the sum of a component whose performance is negatively correlated with sentiment (the $1/N$ portfolio) and a component whose performance is positively correlated with sentiment (the arbitrage portfolio). We then show that the mean-variance portfolio is highly exposed to the arbitrage component, which suffers from large out-of-sample performance losses. To alleviate the impact of the estimation errors affecting the arbitrage component and impose a no-arbitrage condition, we shrink the covariance matrix of stock returns.

3.1 $1/N$, arbitrage, and mean-variance portfolios

We begin by defining the portfolios we study in this section. First, we define the $1/N$ portfolio as the equally weighted portfolio of all N stocks in the investment universe,

$$w_{ew} = \iota/N, \tag{1}$$

where ι is an N -dimensional vector of ones. Second, we define the arbitrage portfolio as the zero-cost portfolio maximizing a mean-variance utility:

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

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

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

where $\mu_g = \frac{\iota^\top \Sigma^{-1} \mu}{\iota^\top \Sigma^{-1} \iota}$ is the mean return of the global-minimum-variance portfolio.⁵ Third, the optimal mean-variance portfolio is defined as the unconstrained portfolio maximizing a mean-variance utility, which has the following well-known solution:

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

Throughout this section, we make the following assumption.

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

Assumption 1 allows us to relate the mean-variance portfolio $w^\star$ to the $1/N$ portfolio w_{ew} and the arbitrage portfolio w_a . It is also a mild assumption from an empirical standpoint; we confirm in unreported results that the first principal component (PC) extracted from the datasets considered in this manuscript has a correlation of about 99% with the returns of the $1/N$ portfolio. In the next proposition, we show that under Assumption 1 the mean-variance portfolio $w^\star$ is given by a linear combination of w_{ew} and w_a .⁶

⁵Roll (1980) considers an orthogonal portfolio similar to the arbitrage portfolio in (2).

⁶Section D in the Internet Appendix contains the proofs of all our theoretical results.

Proposition 1 *Let Assumption 1 hold. Then, the optimal mean-variance portfolio w^* in (3) is*

$$w^* = w_a + \frac{1}{\gamma} \frac{\mu_{ew}}{\sigma_{ew}^2} w_{ew}, \quad (4)$$

where $\mu_{ew} = w_{ew}^\top \mu$ and $\sigma_{ew}^2 = w_{ew}^\top \Sigma w_{ew}$ are the mean return and variance of the 1/N portfolio w_{ew} , respectively.

We now show how the mean-variance portfolio decomposition in Proposition 1 relates to the alpha and beta portfolio decomposition of Raponi, Uppal, and Zaffaroni (2021). Let us assume that asset returns follow a one-factor model such that

$$r = \alpha + \beta f + \epsilon, \quad (5)$$

where $r \in \mathbb{R}^N$ is the vector of excess returns, $\alpha \in \mathbb{R}^N$ is the vector of pricing errors (i.e., the alphas), $\beta \in \mathbb{R}^N$ is the vector of factor betas, f is the 1/N portfolio's excess return, and ϵ is a zero-mean innovation. The following proposition shows that our mean-variance portfolio decomposition is equivalent to a decomposition of an alpha and a beta portfolio.

Proposition 2 *Let Assumption 1 hold and assume that excess returns follow the one-factor model in (5). Then, the optimal mean-variance portfolio w^* in (3) is equivalent to the sum of an alpha and a beta portfolio,*

$$w^* = \underbrace{\frac{1}{\gamma} \Sigma^{-1} \alpha}_{\text{Alpha portfolio}} + \underbrace{\frac{1}{\gamma} \Sigma^{-1} \beta \mu_{ew}}_{\text{Beta portfolio}}, \quad (6)$$

where $\alpha = \mu - \beta \mu_{ew}$ and $\beta = \frac{\text{cov}(r, f)}{\text{var}(f)} = \frac{\Sigma w_{ew}}{\sigma_{ew}^2}$. Moreover, the alpha portfolio and the beta portfolio are equivalent to the arbitrage and equally weighted components in (4).

Proposition 2 highlights that when asset returns follow the one-factor model in (5), the arbitrage portfolio captures asset mispricing, i.e., the α 's. In other words, the arbitrage

portfolio captures the cross-sectional variation of average returns unrelated to the main driver of stock return variation. If asset mispricing is driven by the latent demand of sentiment investors, the relative performance of each component must be influenced by the amount of sentiment traders in the economy.⁷ Thus, latent asset demand is an important factor of the cross-section of stock returns, as highlighted by [Kojien and Yogo \(2019\)](#), and it is also an important element driving the performance of optimal portfolios. Motivated by this insight, in Section 4, we use a proxy for latent asset demand—investor sentiment—to strategically change the exposure to the 1/N and arbitrage components of mean-variance portfolios.

3.2 Decomposition of mean-variance performance

We now decompose the mean-variance portfolio's squared Sharpe ratio into a 1/N and an arbitrage component. The squared Sharpe ratio of portfolio w is

$$\text{SR}^2(w) = \frac{\mathbb{E}^2[w^\top r]}{\text{Var}[w^\top r]} = \frac{(w^\top \mu)^2}{w^\top \Sigma w}, \quad (7)$$

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

$$\text{SR}^2(w^*) = \mu^\top \Sigma^{-1} \mu. \quad (8)$$

In the following proposition, we show explicitly that the performance of the mean-variance portfolio in (8) can be decomposed as the sum of a 1/N component and an arbitrage component.

Proposition 3 *Let Assumption 1 hold. Then, the following holds:*

⁷We formally illustrate this point in Section A of the Internet Appendix.

1. The squared Sharpe ratio of the mean-variance portfolio in (3) is the sum of the squared Sharpe ratios of the 1/N portfolio and the arbitrage portfolio in (2):

$$\text{SR}^2(w^*) = \text{SR}^2(w_{ew}) + \text{SR}^2(w_a), \quad (9)$$

where

$$\text{SR}^2(w_{ew}) = \text{SR}_{PC_1}^2 \quad \text{and} \quad \text{SR}^2(w_a) = \sum_{i=2}^N \text{SR}_{PC_i}^2 \quad (10)$$

with $\text{SR}_{PC_i}^2 = \mu_{PC_i}^2 / \sigma_{PC_i}^2 = (v_i^\top \mu)^2 / v_i^\top \Sigma v_i$ being the squared Sharpe ratio of the i th principal component of stock returns, respectively, and v_i the i th eigenvector of Σ .

2. The squared correlation between the return of the 1/N and arbitrage portfolios and that of the mean-variance portfolio is proportional to their squared Sharpe ratios:

$$\text{Corr}^2(w_{ew}^\top r, (w^*)^\top r) = \frac{\text{SR}^2(w_{ew})}{\text{SR}^2(w^*)}, \quad (11)$$

$$\text{Corr}^2(w_a^\top r, (w^*)^\top r) = \frac{\text{SR}^2(w_a)}{\text{SR}^2(w^*)}. \quad (12)$$

Proposition 3 provides two insights. First, the squared Sharpe ratio of the arbitrage portfolio is defined by the squared Sharpe ratio of low-variance principal components that do not explain much of stock return variation. If these low-variance principal components deliver high average returns, such investment strategies correspond with arbitrage opportunities in the sense of Ross (1976).

Second, the squared correlation between the returns of the 1/N (arbitrage) portfolio and those of the mean-variance portfolio is an increasing function of the 1/N (arbitrage) portfolio's squared Sharpe ratio. This result suggests that mean-variance portfolios self-control the exposure to the 1/N and arbitrage components. Accordingly, one could expect that the mean-variance portfolio will be more exposed to 1/N than to the arbitrage component when sentiment is low and 1/N performance is stronger.

However, Figure 2 shows that this is not true in general. The figure depicts the correlation between $1/N$ and arbitrage returns and those of the mean-variance portfolio for the six datasets we consider in our empirical analysis (see Section 4.2). The correlations are evaluated over the whole sample and separately in high-sentiment and low-sentiment regimes using the Huang et al. (2015) index. Two results stand out. First, as discussed above, the correlation with $1/N$ returns is larger in low-sentiment periods and smaller in high-sentiment ones, and vice versa for the correlation with arbitrage returns. Second, the mean-variance portfolio is highly exposed to the arbitrage component. In particular, the correlation with arbitrage portfolio returns is nearly 100% in high-sentiment regimes, and more than 80% in low-sentiment regimes. The reason why the mean-variance portfolio is highly correlated with the arbitrage portfolio is that this portfolio delivers a very large in-sample Sharpe ratio, which defines the correlation with the mean-variance portfolio as highlighted in Proposition 3. The arbitrage portfolio, however, suffers important out-of-sample performance losses due to estimation risk. In Section 3.3, we characterize the arbitrage out-of-sample performance losses, and we show how to mitigate these losses by imposing a no-arbitrage condition in Section 3.4.


3.3 Out-of-sample loss of the arbitrage portfolio

The findings illustrated in Figure 2 show that the mean-variance portfolio is highly exposed to the arbitrage component relative to the $1/N$ component. To understand the extent to which this large exposure can be problematic out of sample, we now derive an analytical expression for the expected out-of-sample Sharpe ratio loss of the in-sample arbitrage portfolio. Specifically, let the sample arbitrage portfolio be

$$\hat{w}_a = \frac{1}{\gamma} \hat{\Sigma}^{-1} (\hat{\mu} - \hat{\mu}_g \iota) \quad (13)$$

Figure 2: Correlation with 1/N and arbitrage portfolio returns

This figure depicts the correlation between the returns of the mean-variance portfolio w^* in (3) and those of the 1/N portfolio w_{ew} and the arbitrage portfolio w_a in (2). The returns are obtained using the whole sample of the six empirical datasets described in Section 4.2. The correlations are computed using Proposition 3 either on the whole sample or separately 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.



where $\hat{\mu}$ and $\hat{\Sigma}$ are the sample mean and covariance matrix of stock returns:

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T r_t \quad \text{and} \quad \hat{\Sigma} = \frac{1}{T} \sum_{t=1}^T (r_t - \hat{\mu})(r_t - \hat{\mu})^\top, \quad (14)$$

and r_t is the N -dimensional vector of excess returns at time t . Moreover, let us define the expected out-of-sample Sharpe ratio of $\hat{w}_a$ as⁸

$$\mathbb{E}[\text{SR}(\hat{w}_a)] = \frac{\mathbb{E}[\hat{w}_a^\top \mu]}{\sqrt{\mathbb{E}[\hat{w}_a^\top \Sigma \hat{w}_a]}}. \quad (15)$$

Using the above expression, we characterize the expected loss for the arbitrage portfolio in the following proposition.

⁸In Section B of the Internet Appendix, we derive a more involved closed-form expression for the *exact* expected out-of-sample Sharpe ratio, $\mathbb{E}[(\hat{w}_a^\top \mu)(\hat{w}_a^\top \Sigma \hat{w}_a)^{-1/2}]$, from which we draw similar insights as those we obtain from (15).

Proposition 4 *Let Assumption 1 hold. In addition, assume that stock returns are normally distributed and that $T > N + 3$. Then, the expected out-of-sample Sharpe ratio loss of the sample arbitrage portfolio $\hat{w}_a$ in (13) is*

$$\frac{\text{SR}(w_a) - \mathbb{E}[\text{SR}(\hat{w}_a)]}{\text{SR}(w_a)} = 1 - \sqrt{\frac{(T - N)(T - N - 3)}{(T - 2)(T - N - 1)} \frac{\text{SR}^2(w_a)}{\text{SR}^2(w_a) + \frac{N-1}{T}}} > 0, \quad (16)$$

where $\mathbb{E}[\text{SR}(\hat{w}_a)]$ is defined in (15). Moreover, this loss decreases with $\text{SR}^2(w_a)$.

Proposition 4 shows that the out-of-sample performance loss of the sample arbitrage portfolio can be substantial when the ratio between the number of assets N and the sample size T is large. Proposition 4 also shows that the out-of-sample loss of the sample arbitrage portfolio is smaller when the arbitrage portfolio's squared Sharpe ratio, $\text{SR}^2(w_A)$, is large, which typically happens following high-sentiment periods. This result suggests that the shrinking-against-sentiment framework proposed in Section 4 is sensible because it maintains a high exposure with the arbitrage portfolio when it performs better *and* when it is subject to less estimation risk.⁹


We confirm this theoretical prediction empirically. In particular, for each dataset in Section 4.2, we estimate the out-of-sample arbitrage portfolio returns using a rolling window of $T = 180$ months.¹⁰ Then, Figure 3 depicts the out-of-sample Sharpe ratio loss defined in (16) for both high and low-sentiment regimes. On average across the six datasets, the loss is 79% following low-sentiment periods but only 49% following high-sentiment periods, which is consistent with the insights we draw from Proposition 4.

⁹Note that in the theoretical analysis presented in this section, we make the assumption that returns are iid for tractability. While this assumption is a practical way of studying the impact of estimation errors on the arbitrage component, the true data generating process is not iid. Indeed, our portfolio framework acknowledges this and uses investor sentiment to account for relevant information about the time-varying dynamics of expected returns. It is likely that with non-iid returns, the impact of estimation risk has a more complex structure. Therefore, our analysis of estimation risk in this section is a conservative approach to the impact of parameter uncertainty on the arbitrage portfolio performance.

¹⁰We use a sample size of $T = 180$ months because the sample covariance matrix is nearly singular for the 100SINV dataset when $T = 120$ months.

Figure 3: Out-of-sample Sharpe ratio loss of the arbitrage portfolio

This figure depicts the out-of-sample Sharpe ratio loss in Equation (16) for the sample arbitrage portfolio $\hat{w}_a$ in (13) (in percentage points). The out-of-sample losses are obtained using the six empirical datasets described in Section 4.2. For each dataset, we estimate the arbitrage portfolio using a fixed estimation window of $T = 180$ months, evaluate its out-of-sample return on the next month, and apply this method recursively to obtain a time series of out-of-sample arbitrage portfolio returns. We also compute the time series of in-sample arbitrage portfolio returns over the same time period. We then split those two time series in high and low-sentiment regimes using the Huang et al. (2015) sentiment index, and evaluate the out-of-sample Sharpe ratio loss for each regime. 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.



3.4 No-arbitrage condition and shrinkage

This section shows how one can impose a no-arbitrage condition using a shrinkage covariance matrix. In addition, we show that the no-arbitrage condition imposed by the shrinkage covariance matrix also allows us to control the mean-variance portfolio's exposure to $1/N$ and the arbitrage portfolio.

First, let us define the arbitrage condition under the one-factor model in Equation (5):

$$\alpha^\top \Sigma^{-1} \alpha < \pi, \quad (17)$$

where $\pi > 0$ bounds the amount of mispricing in the economy. The following proposition links the no-arbitrage condition with the performance of the arbitrage portfolio.

Proposition 5 *Let Assumption 1 hold and assume that excess returns follow the one-factor model in (5). Then, the no-arbitrage condition is equivalent to imposing a bound*

on the maximum squared Sharpe ratio delivered by the arbitrage portfolio. That is,

$$\alpha^\top \Sigma^{-1} \alpha = \text{SR}^2(w_a) < \pi, \quad (18)$$

where $\pi > 0$ bounds the arbitrage portfolio's squared Sharpe ratio and thus the amount of mispricing in the economy.

To impose the no-arbitrage condition highlighted in Proposition 5, we use a shrinkage covariance matrix of stock returns similar to that in Ledoit and Wolf (2004):

$$\Sigma_{sh} = (1 - \delta)\Sigma + \delta\nu I_N, \quad (19)$$

where δ is the shrinkage intensity, $\nu = \frac{1}{N} \text{Trace}(\Sigma)$ is the cross-sectional average of return variances, and I_N is the identity matrix. Using the shrinkage covariance matrix (19) instead of the covariance matrix Σ gives the shrinkage arbitrage portfolio,

$$w_{sa} = \frac{1}{\gamma} \Sigma_{sh}^{-1} \left(\mu - \frac{\iota^\top \Sigma_{sh}^{-1} \mu}{\iota^\top \Sigma_{sh}^{-1} \iota} \iota \right), \quad (20)$$

and the shrinkage mean-variance portfolio,

$$w_s^* = \frac{1}{\gamma} \Sigma_{sh}^{-1} \mu. \quad (21)$$

In the next proposition, we derive in closed form the squared Sharpe ratio of the 1/N portfolio, the shrinkage arbitrage portfolio, and the shrinkage mean-variance portfolio.

Proposition 6 *Let Assumption 1 hold. Then, the squared Sharpe ratios of the 1/N portfolio, the shrinkage arbitrage portfolio in (20), and the shrinkage mean-variance portfolio in (21) are $\text{SR}^2(w_{ew}) = \text{SR}_{PC_1}^2$, $\text{SR}^2(w_{sa}) = S_\delta(2)$, and $\text{SR}^2(w_s^*) = S_\delta(1)$, where*

$$S_\delta(j) = \left(\sum_{i=j}^N \frac{\text{SR}_{PC_i}^2}{1 - \delta + \delta\nu/\sigma_{PC_i}^2} \right)^2 \left(\sum_{i=j}^N \frac{\text{SR}_{PC_i}^2}{(1 - \delta + \delta\nu/\sigma_{PC_i}^2)^2} \right)^{-1}. \quad (22)$$

Proposition 6 extends the result in part 1 of Proposition 3, which we recover when $\delta = 0$, for the case where we apply shrinkage to the covariance matrix. The following corollary builds on the insights in Proposition 6 and connects the shrinkage intensity of the covariance matrix with the no-arbitrage condition.

Corollary 1 *Let Assumption 1 hold. Then, the squared Sharpe ratio of the shrinkage arbitrage portfolio w_{sa} decreases with the shrinkage intensity δ .*

Corollary 1 shows that one can impose a tighter bound on the arbitrage portfolio's maximum squared Sharpe ratio by shrinking the covariance matrix of stock returns, which is equivalent to imposing a no-arbitrage constraint as highlighted in Proposition 5.

Figure 4 depicts the Sharpe ratio of the shrinkage mean-variance portfolio w_s^* in (21), the 1/N portfolio w_{ew} , and the shrinkage arbitrage portfolio w_{sa} in (20) as a function of the shrinkage intensity δ . 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 $\delta = 0$, the mean-variance portfolio attains its largest squared Sharpe ratio, which is close to that of the arbitrage portfolio. However, as δ increases, the squared Sharpe ratio of the mean-variance portfolio converges to that of the 1/N component and shies away from the arbitrage portfolio.

In the next proposition, we formally study the correlation between the return of the shrinkage mean-variance portfolio and that of 1/N and the shrinkage arbitrage portfolio.


Proposition 7 *Let Assumption 1 hold. Then, the squared correlations between the return of the shrinkage mean-variance portfolio in (21) and that of the 1/N portfolio and the shrinkage arbitrage portfolio in (20) are*

$$\text{Corr}^2(w_{ew}^\top r, (w_s^*)^\top r) = \frac{\text{SR}^2(w_{ew})}{\text{SR}^2(w_{ew}) + \sum_{i=2}^N s_i^2(\delta) \text{SR}_{PC_i}^2}, \quad (23)$$

$$\text{Corr}^2(w_{sa}^\top r, (w_s^*)^\top r) = 1 - \text{Corr}^2(w_{ew}^\top r, (w_s^*)^\top r), \quad (24)$$

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^* in (21), the equally weighted portfolio w_{ew} , and the shrinkage arbitrage portfolio w_{sa} in (20) as a function of the shrinkage intensity δ defining the shrinkage covariance matrix in (19). The Sharpe ratios are computed using the results in Proposition 6. We construct the portfolios using the sample moments from the dataset of monthly excess returns on the 25 portfolios of stocks sorted on size and book-to-market from July 1965 to June 2022.



respectively, where $s_i(\delta)$ is a shrinkage factor applied to lower-variance PCs:


$$s_i(\delta) = \frac{\sigma_{PC_i}^2}{\sigma_{PC_i}^2 + \kappa(\delta)(\sigma_{PC_1}^2 - \sigma_{PC_i}^2)} \leq 1 \quad \text{with} \quad \kappa(\delta) = \frac{\delta\nu}{\delta\nu + (1 - \delta)\sigma_{PC_1}^2}. \quad (25)$$

Moreover, the squared correlation with $1/N$, $\text{Corr}^2(w_{ew}^\top r, (w_s^*)^\top r)$, is increasing in the shrinkage intensity δ , and the squared correlation with the shrinkage arbitrage portfolio, $\text{Corr}^2(w_{sa}^\top r, (w_s^*)^\top r)$, is decreasing in δ .

Proposition 7 shows that the squared correlation between the return of any shrinkage mean-variance portfolio and that of the $1/N$ portfolio mechanically increases when the $1/N$ squared Sharpe ratio, $\text{SR}^2(w_{ew})$, increases. However, a key consequence of the results in Proposition 7 is that shrinking the covariance matrix toward the identity allows us to *control* the exposure of mean-variance portfolios to the $1/N$ and arbitrage components. In particular, the larger the shrinkage intensity δ , the larger the correlation between $1/N$ returns and mean-variance returns. This is because the squared Sharpe ratios of lower-

Figure 5: Correlation with 1/N and arbitrage portfolios

This figure depicts the correlation between the return of the shrinkage mean-variance portfolio w_s^* in (21) and the return of the equally weighted portfolio w_{ew} and the shrinkage arbitrage portfolio w_{sa}^* in (20) as a function of the shrinkage intensity δ defining the shrinkage covariance matrix in (19). The correlations are computed using the results in Proposition 7. We construct the portfolios using the sample moments from the dataset of monthly excess returns on the 25 portfolios of stocks sorted on size and book-to-market from July 1965 to June 2022.



variance PCs, $\text{SR}_{PC_i}^2$ for $i = 2, \dots, N$, are toned down by the shrinkage factor $s_i^2(\delta) \leq 1$ in (25) that decreases with the shrinkage intensity δ .

The specific form of the shrinkage factor $s_i(\delta)$ in Proposition 7 is consistent with shrinkage methodologies in the literature. Specifically, $s_i(\delta)$ decreases with the index i , and thus the contributions of low-variance PCs are shrunk more intensively than for high-variance ones. Shrinking low-variance PCs is economically meaningful because it rules out near-arbitrage opportunities that are unlikely to persist out of sample (Kozak, Nagel, and Santosh, 2020). Moreover, Tay, Friedman, and Tibshirani (2021) recently introduced a regularization method for linear regression in which the contributions of PCs are shrunk with an intensity that is equivalent to our $s_i(\delta)$. Tay, Friedman, and Tibshirani (2021) show that their method helps improve predictability for data with a strong factor structure, as it is the case with stock returns.

Figure 5 uses the same data as in Figure 4 to depict the correlation between the return of the shrinkage mean-variance portfolio in (21) and the return of the 1/N portfolio

and the shrinkage arbitrage portfolio in (20) as a function of the shrinkage intensity δ . Consistent with the theoretical findings in Proposition 7, the figure shows that the correlation between $1/N$ and the shrinkage mean-variance portfolio is 32% when $\delta = 0$, but it increases to nearly 100% when $\delta = 1$. On the other hand, the correlation with the shrinkage arbitrage portfolio is 95% when the shrinkage intensity $\delta = 0$, but it decreases to 6.3% when $\delta = 1$.

4 Shrinking against sentiment (SAS)

In this section, we explore the economic gains that a mean-variance investor can obtain by exploiting investor sentiment, our proxy for latent asset demand, in the construction of robust investment strategies. Section 4.1 utilizes 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 to the covariance matrix when sentiment is low (high). Our methodology accounts for the contrasting impact that sentiment has on $1/N$ and arbitrage returns, as well as for the out-of-sample performance losses of the arbitrage component due to estimation risk that we characterize in Section 3.3.

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 the logistic regression (Hastie, Tibshirani, Friedman, and Franklin, 2009). In the simplest setting, the logistic regression estimates the probability of a particular event

happening given some information. Given the level of investor sentiment at time t , our objective is to estimate the probability that the equally weighted portfolio return will be larger than the out-of-sample return of the arbitrage portfolio at time $t + 1$. We use this probability to determine the relevance of $1/N$ and the arbitrage component in the construction of the mean-variance portfolio.

More precisely, we start with a sample of T monthly excess returns. We use the decomposition of the mean-variance portfolio in Proposition 1 —i.e., $w^* = w_a + \frac{\mu_{ew}}{\gamma\sigma_{ew}^2}w_{ew}$ — and estimate the two components, w_a and $\frac{\mu_{ew}}{\gamma\sigma_{ew}^2}w_{ew}$, using the first half of the sample.¹¹ We then evaluate the out-of-sample return of these two portfolios on the second half. Using the out-of-sample returns, we estimate the probability of the event $r_{t+1}^{ew} > r_{t+1}^a$ conditional on lagged sentiment using a logistic function:

$$p_t = \mathbb{P}(r_{t+1}^{ew} > r_{t+1}^a \mid \text{Sentiment}_t) = \frac{e^{\beta_0 + \beta_1 \text{Sentiment}_t}}{1 + e^{\beta_0 + \beta_1 \text{Sentiment}_t}}, \quad (26)$$

where Sentiment_t is the [Huang et al. \(2015\)](#) sentiment index, which we reconstruct using only information up to time t to avoid look-ahead biases. The coefficients β_0 and β_1 in (26) can be easily estimated via maximum likelihood; see, e.g., [Hastie et al. \(2009\)](#).

Our method then exploits the estimate of probability p_t to calibrate the shrinkage intensity δ applied to the covariance matrix (19). We achieve this by using p_t to determine the desired correlation between $1/N$ and the shrinkage mean-variance portfolio. That is, if probability p_t is high (low), then our approach imposes a high (low) correlation between $1/N$ and the shrinkage mean-variance portfolio.


In particular, we use the predicted probability p_t to linearly interpolate the squared correlation between the return of $1/N$ and that of the shrinkage mean-variance portfolio:

$$\text{Corr}^2(p_t) = (1 - p_t) \text{Corr}_0^2 + p_t \text{Corr}_1^2, \quad (27)$$

¹¹Because both components are proportional to $1/\gamma$, we set $\gamma = 1$ without loss of generality.

Figure 6: Shrinkage intensity and sentiment

This figure depicts the shrinking-against-sentiment shrinkage intensity (left axis) obtained from the procedure described in Section 4.1 constructed from the dataset of 25 portfolios of stocks sorted on size and book-to-market, together with the sentiment index proposed by Huang et al. (2015) (right axis).



where Corr_0^2 and Corr_1^2 are the squared correlations in Equation (23) when the shrinkage intensity δ is zero and one, respectively. Then, using (23), we extract the value of the shrinkage intensity δ that gives the desired correlation with $1/N$, $\text{Corr}^2(p_t)$.¹² Given the *negative* relation between $1/N$ returns and sentiment, our method shrinks against sentiment to harvest a larger premium from the $1/N$ component when sentiment is low.¹³

Figure 6 documents the relation between the Huang et al. (2015) sentiment index and the proposed SAS shrinkage intensity δ obtained from the dataset of 25 portfolios of stocks sorted on size and book-to-market. We observe that the shrinkage intensity has a substantial negative correlation with investor sentiment of -40% , as desired.

¹²We interpolate the *squared* correlation because, as shown in Proposition 7, the squared correlations between the return of the shrinkage mean-variance portfolio and that of $1/N$ and the shrinkage arbitrage portfolio add up to one. Therefore, we can show that relation (27) also holds for the squared correlation with the return of the shrinkage arbitrage portfolio.

¹³Although we only shrink the covariance matrix, the criterion we propose to calibrate the shrinkage intensity accounts for the out-of-sample estimation errors in the arbitrage portfolio, which stem from *both* the mean and the covariance matrix.

4.2 Data

We proxy for investor sentiment using the [Huang et al. \(2015\)](#) sentiment index. The index spans the period between July 1965 and June 2022, 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 equity share in new issues, and the dividend premium.¹⁴ The sentiment index is a latent variable extracted through partial least squares. We take these five sentiment variables for the out-of-sample analysis and compute the sentiment index using the method of [Huang et al. \(2015\)](#) iteratively to avoid look-ahead biases.

We use the monthly excess returns of six datasets. The first three datasets are composed of 5×5 double-sorted portfolios: (i) 25 portfolios sorted on size and market beta (25SBETA) from July 1965 to June 2022, (ii) 25 portfolios sorted on size and book-to-market (25SBTM) from July 1965 to June 2022, (iii) 25 portfolios sorted on operating profitability and investment (25OPINV) from July 1965 to June 2022. The next two datasets come from the 23 anomalies considered by [Novy-Marx and Velikov \(2016\)](#) and are downloaded from Robert Novy-Marx's website: (iv) the long and short legs of eight low-turnover anomalies (16ANOM) from July 1965 to December 2013 and (v) the long and short legs of all the 23 anomalies (46ANOM) from July 1973 to December 2013. Our last dataset is composed of 10×10 double-sorted portfolios: (vi) 100 portfolios formed on size and investment (100SINV).

4.3 Out-of-sample performance

In this section, we assess the out-of-sample performance of the SAS portfolio.

¹⁴We download the variables we need to construct the sentiment index from Jeffrey Wurgler's website.

4.3.1 Methodology and portfolio strategies

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 $t = 1, \dots, T - M$.¹⁵ 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 $t = 1, \dots, T - M$ monthly observations.

The analysis focuses on the estimated tangency mean-variance portfolio,

$$w_t^* = \frac{1}{\gamma} \Sigma_t^{-1} \mu_t, \quad (28)$$

where μ_t and Σ_t are the estimated vector of means and covariance matrix of stock returns in excess of the risk-free rate with information up to month t , and γ is set equal to $|\iota^\top \Sigma_t^{-1} \mu_t|$ to control for leverage.¹⁶ We consider seven portfolio strategies that estimate the tangency portfolio w_t^* in different ways. The first five strategies estimate μ_t using the sample mean in (14), while the covariance matrix Σ_t is estimated using different shrinkage methods, which we list below:

1. The proposed covariance matrix in Equation (19) that shrinks the sample covariance matrix toward the identity and the shrinkage intensity δ is estimated from the shrinking-against-sentiment methodology in Section 4.1. Plugging this matrix in (28) delivers the SAS portfolio.

¹⁵In unreported results, we confirm that our results are robust to using rolling windows instead of expanding windows. The latter have the benefit of delivering better estimates of mean returns, which typically require long time series to obtain accurate predictions (Merton, 1980).

¹⁶We only scale the tangency portfolio in (28) when the denominator $|\iota^\top \Sigma_t^{-1} \mu_t|$ is larger than one, which avoids excessive leverage and extreme portfolio weights as noted by Kirby and Ostdiek (2012). In Appendix C.1, we also show that the results are robust to considering alternative values of γ .

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 (28) 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 (28) delivers the MV-MKT portfolio.
4. The nonlinear shrinkage covariance matrix proposed by [Ledoit and Wolf \(2020\)](#). Plugging this matrix in (28) delivers the MV-NL portfolio.
5. The covariance matrix proposed by [Chen and Yuan \(2016\)](#) that fully shrinks the contribution of low-variance PCs. Plugging this matrix in (28) delivers the MV-PCA portfolio.

The last two benchmarks are simple low-turnover strategies:

6. The reward-to-risk (RTR) timing strategy of [Kirby and Ostdiek \(2012\)](#) that uses the sample estimates of μ_t and Σ_t in (14) but setting all the covariances equal to zero and replacing the mean returns μ_t by $\max(0, \mu_t)$.
7. The equally weighted ($1/N$) portfolio w_{ew} , which is a key benchmark because [DeMiguel, Garlappi, and Uppal \(2009\)](#) argue that the benefits of optimal diversification can be offset by estimation errors.

In Section C.2 of the Internet Appendix, we show that our results are robust to constructing mean-variance portfolios with a dynamic covariance matrix that exploits the serial autocorrelation of variances and covariances.

4.3.2 Out-of-sample results

Panel A in Table 1 reports the annualized out-of-sample Sharpe ratios of the SAS portfolio and the six benchmarks across the six datasets. The table shows that SAS portfolios deliver the largest Sharpe ratio in all datasets, and the average outperformance of the SAS portfolio across the six datasets is 37%, 60%, 37%, and 31% relative to the MV-1N, MV-MKT, MV-NL, and MV-PCA mean-variance portfolios, respectively.¹⁷ The improvement is even larger relative to the simple RTR and 1/N policies: 88% and 146%, respectively. The good performance of the SAS portfolio relative to the benchmark mean-variance portfolios does not come at the expense of higher turnover. On the contrary, Panel B in Table 1 shows that the four benchmark mean-variance portfolios require, on average, a turnover 461% higher than that of the SAS portfolio.

The substantially lower turnover that the SAS portfolio requires suggests that its performance is likely to survive the impact of transaction costs. To illustrate this point, we now report the annualized out-of-sample Sharpe ratio of the SAS portfolio and the six benchmark strategies net of proportional transaction costs. We define the net return of portfolio p at time $t + 1$ as

$$r_{t+1}^p = \left(1 - \kappa \sum_{i=1}^N |w_{i,t}^p - w_{i,(t-1)+}^p|\right) \left(1 + (w_t^p)^\top r_{t+1}\right) - 1, \quad (29)$$

where r_{t+1} is the vector of excess returns at time $t + 1$, w_t^p is the portfolio strategy p estimated with return data up to time t , $w_{i,(t-1)+}^p$ is the portfolio weight in stock i at time t before rebalancing, and κ is the level of proportional transaction costs. Similar to Barroso and Saxena (2022), we consider proportional transaction costs of $\kappa = 10, 30$, and 50 basis points. Table 2 shows that the annualized Sharpe ratio of the different portfolio

¹⁷The numbers in parenthesis are bootstrap p -values relative to the 1/N portfolio. To compute the test p -values, we generate 1,000 bootstrap samples using the stationary block bootstrap approach of Politis and Romano (1994) with an average block size of five and use the methodology of Ledoit and Wolf (2008, Remark 3.2.) to produce the resulting p -values.

Table 1: Gross out-of-sample Sharpe ratios and portfolio turnover

This table reports the out-of-sample performance of the SAS portfolio and six benchmark portfolios. The first four benchmarks are mean-variance portfolios constructed with the shrinkage covariance matrices of [Ledoit and Wolf \(2004\)](#) (MV-1N), [Ledoit and Wolf \(2003\)](#) (MV-MKT), [Ledoit and Wolf \(2020\)](#) (MV-NL), and [Chen and Yuan \(2016\)](#) (MV-PCA). The last two benchmarks are the reward-to-risk timing strategy of [Kirby and Ostdiek \(2012\)](#) (RTR), and the equally weighted (1/N) portfolio. Panels A and B report the annualized out-of-sample Sharpe ratio and the average monthly turnover 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 1/N. The six datasets are described in Section 4.2.

Policy	25SBETA	25SBTM	25OPINV	16ANOM	46ANOM	100SINV
Panel A: Sharpe ratios						
SAS	0.84 (0.07)	0.90 (0.05)	0.93 (0.02)	1.08 (0.00)	2.09 (0.00)	1.05 (0.00)
MV-1N	0.53 (0.96)	0.47 (0.72)	0.67 (0.51)	1.00 (0.05)	2.00 (0.00)	0.87 (0.14)
MV-MKT	0.43 (0.64)	0.34 (0.34)	0.63 (0.66)	0.94 (0.07)	1.86 (0.00)	0.84 (0.25)
MV-NL	0.51 (0.87)	0.47 (0.70)	0.65 (0.64)	0.93 (0.06)	2.08 (0.00)	0.99 (0.06)
MV-PCA	0.62 (0.73)	0.73 (0.42)	0.80 (0.14)	0.88 (0.10)	1.26 (0.00)	0.84 (0.15)
RTR	0.64 (0.00)	0.61 (0.00)	0.62 (0.00)	0.58 (0.00)	0.60 (0.00)	0.64 (0.00)
1/N	0.57	0.56	0.55	0.45	0.36	0.58
Panel B: Turnover						
SAS	0.27	0.42	0.35	0.44	1.89	0.46
MV-1N	2.01	2.11	1.40	1.69	12.5	3.64
MV-MKT	2.66	2.99	1.46	2.42	15.1	3.72
MV-NL	2.56	2.51	1.47	2.50	8.72	2.64
MV-PCA	0.55	0.91	0.33	3.25	14.4	0.59
RTR	0.04	0.04	0.04	0.05	0.04	0.04
1/N	0.04	0.04	0.04	0.04	0.04	0.04

strategies deteriorates due to transaction costs. However, the impact of transaction costs on the performance of the SAS portfolio is substantially lower. For instance, for $\kappa = 30$ basis points, the average outperformance in terms of annualized Sharpe ratio of the SAS portfolio relative to the benchmark mean-variance portfolios is 71%. Moreover, even though the RTR and 1/N portfolios only require a very small turnover, the average outperformance in terms of annualized Sharpe ratio of the SAS portfolio relative to the RTR and 1/N portfolios is 78% and 133%, respectively.


Table 2: Net-of-cost out-of-sample Sharpe ratios

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

Policy	25SBETA	25SBTM	25OPINV	16ANOM	46ANOM	100SINV
Panel A: Sharpe ratios net of 10bps						
SAS	0.82 (0.09)	0.88 (0.05)	0.91 (0.02)	1.06 (0.00)	2.03 (0.00)	1.02 (0.00)
MV-1N	0.48 (0.74)	0.43 (0.59)	0.63 (0.62)	0.96 (0.07)	1.95 (0.00)	0.78 (0.36)
MV-MKT	0.38 (0.45)	0.30 (0.25)	0.59 (0.75)	0.89 (0.09)	1.80 (0.00)	0.75 (0.43)
MV-NL	0.46 (0.64)	0.42 (0.58)	0.61 (0.77)	0.88 (0.09)	2.02 (0.00)	0.91 (0.11)
MV-PCA	0.61 (0.84)	0.70 (0.48)	0.78 (0.16)	0.83 (0.15)	1.18 (0.00)	0.82 (0.18)
RTR	0.64 (0.00)	0.61 (0.00)	0.61 (0.00)	0.58 (0.00)	0.59 (0.00)	0.63 (0.00)
1/N	0.57	0.56	0.55	0.45	0.36	0.58
Panel B: Sharpe ratios net of 30bps						
SAS	0.79 (0.13)	0.84 (0.08)	0.87 (0.04)	1.02 (0.00)	1.93 (0.00)	0.97 (0.00)
MV-1N	0.39 (0.45)	0.35 (0.41)	0.55 (0.93)	0.87 (0.10)	1.85 (0.00)	0.60 (0.91)
MV-MKT	0.28 (0.22)	0.21 (0.14)	0.51 (0.89)	0.79 (0.17)	1.66 (0.00)	0.57 (1.00)
MV-NL	0.35 (0.39)	0.33 (0.34)	0.52 (0.96)	0.79 (0.21)	1.91 (0.00)	0.76 (0.35)
MV-PCA	0.57 (0.93)	0.64 (0.63)	0.74 (0.25)	0.72 (0.27)	1.01 (0.05)	0.76 (0.29)
RTR	0.63 (0.00)	0.60 (0.00)	0.61 (0.00)	0.57 (0.00)	0.59 (0.00)	0.63 (0.00)
1/N	0.56	0.55	0.54	0.45	0.35	0.58
Panel C: Sharpe ratios net of 50bps						
SAS	0.75 (0.19)	0.79 (0.10)	0.82 (0.05)	0.98 (0.00)	1.83 (0.00)	0.91 (0.00)
MV-1N	0.29 (0.22)	0.27 (0.20)	0.46 (0.76)	0.78 (0.22)	1.74 (0.00)	0.43 (0.44)
MV-MKT	0.17 (0.07)	0.12 (0.03)	0.43 (0.67)	0.69 (0.25)	1.51 (0.00)	0.38 (0.34)
MV-NL	0.25 (0.15)	0.23 (0.15)	0.44 (0.68)	0.69 (0.35)	1.78 (0.00)	0.60 (0.83)
MV-PCA	0.54 (0.95)	0.58 (0.90)	0.69 (0.36)	0.61 (0.52)	0.84 (0.12)	0.71 (0.43)
RTR	0.63 (0.00)	0.60 (0.00)	0.60 (0.00)	0.57 (0.00)	0.58 (0.00)	0.62 (0.00)
1/N	0.56	0.55	0.53	0.44	0.34	0.57

Figure 7: Impact of sentiment on Sharpe ratio

This figure depicts the out-of-sample annualized gross Sharpe ratio of the equally weighted portfolio, the sample arbitrage portfolio, the SAS portfolio, and the MV-1N portfolio that exploits the shrinkage covariance matrix in [Ledoit and Wolf \(2004\)](#). The Sharpe ratios are obtained following our empirical methodology in Section 4.3.1. The six datasets are described in Section 4.2. The out-of-sample Sharpe ratios are evaluated either on the whole out-of-sample period or separately 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 historical median.



5 Understanding the performance of SAS portfolios

Our proposed portfolio strategy shrinks the covariance matrix of stock returns more following low-sentiment periods and less following high-sentiment periods. To understand why this approach delivers good out-of-sample performance, we depict in Figure 7 the out-of-sample annualized gross Sharpe ratio across sentiment regimes of the 1/N portfolio, the arbitrage portfolio, the proposed SAS portfolio, and the MV-1N portfolio that exploits the shrinkage covariance matrix of [Ledoit and Wolf \(2004\)](#).

We draw four insights from Figure 7. First, over the whole sample, the SAS portfolio performs at least equally well, and in general substantially better, than both the 1/N and arbitrage portfolios. Second, the performance of the 1/N and arbitrage portfolios

varies substantially across sentiment regimes. On average across the six datasets, the 1/N portfolio delivers a Sharpe ratio of 0.24 in high-sentiment periods and 0.87 in low-sentiment months. On the contrary, the sample arbitrage portfolio delivers a Sharpe ratio of 1.19 in high-sentiment periods and 0.47 in low-sentiment months. Third, our SAS methodology performs consistently well across sentiment regimes. In particular, the SAS portfolio delivers, on average, a Sharpe ratio of 1.46 in high-sentiment periods and 0.87 in low-sentiment months. This performance is 34% larger than that of the arbitrage portfolio in high-sentiment periods and more than 100% in low-sentiment months. Therefore, SAS portfolios harvest a substantial premium due to their strategically larger exposure to the 1/N component in low-sentiment periods and larger exposure to the arbitrage component in high-sentiment periods. Fourth, the SAS portfolio consistently outperforms the MV-1N portfolio during high and low-sentiment regimes. On average, the annualized Sharpe ratio of the SAS portfolio is 12% higher than that of the MV-1N portfolio in high-sentiment periods and 111% in low-sentiment months. Therefore, the economic gains of the SAS portfolio relative to benchmark mean-variance portfolios that ignore the effect of sentiment on the performance of its 1/N and arbitrage components come from the substantially stronger performance that SAS delivers during low-sentiment regimes.

In the Internet Appendix, we provide additional intuition for the good performance of SAS portfolios. In particular, we show in Section C.3 that the SAS portfolio typically keeps a higher exposure to 1/N than that of the MV-1N portfolio, which allows the SAS portfolio to harvest a larger premium when sentiment is low. Also, in Section C.4 we confirm the theoretical results of Baker and Stein (2004) and show that sentiment is positively correlated with liquidity measures such as bid-ask spreads. Accordingly, shrinking against sentiment strategically tilts the mean-variance portfolio toward the low-turnover 1/N portfolio when liquidity is low and trading costs are high.

6 Conclusion

We introduce a novel approach to constructing mean-variance portfolios that exploits a proxy for latent asset demand—*investor sentiment*. Our methodology relies on the decomposition of mean-variance portfolios as the sum of $1/N$ and an arbitrage portfolio and employs shrinkage techniques to tilt the optimal portfolio toward $1/N$ and away from the arbitrage component when sentiment is low.

The key benefits of this approach are twofold. First, our shrinkage methodology that tilts mean-variance portfolios to the $1/N$ component is equivalent to imposing a no-arbitrage condition that bounds asset mispricing. We determine the mispricing bound using investor sentiment, which allows us to establish a looser bound when sentiment is high, and therefore, the arbitrage component can harvest higher returns. Second, we demonstrate that shrinking against sentiment alleviates the impact of estimation errors on portfolio performance stemming from the arbitrage component, especially when estimation risk is high. As a consequence of the methodological improvements offered by our shrinkage approach, our portfolios deliver significant economic gains relative to alternative mean-variance benchmark portfolios at a substantially lower turnover.

Overall, our theoretical and empirical analysis illustrates the importance of accounting for latent asset demand in the construction of robust investment strategies.

References

- Bakalli, G., S. Guerrier, and O. Scaillet. 2023. A penalized two-pass regression to predict stock returns with time-varying risk premia. *Journal of Econometrics* (forthcoming).
- Baker, M., and J. Stein. 2004. Market liquidity as a sentiment indicator. *Journal of Financial Markets* 7:271–99.
- Baker, M., and J. Wurgler. 2007. Investor sentiment in the stock market. *Journal of Economic Perspective* 21:129–52.
- Barroso, P., and A. Detzel. 2021. Do limits to arbitrage explain the benefits of volatility-managed portfolios? *Journal of Financial Economics* 140:744–67.
- Barroso, P., and K. Saxena. 2022. Lest we forget: Using out-of-sample errors in portfolio optimization. *The Review of Financial Studies* 35:1222–78.
- Britten-Jones, M. 1999. The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance* 54:655–71.
- Campbell, J., and S. Thompson. 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies* 21:1509–31.
- Chen, J., and M. Yuan. 2016. Efficient portfolio selection in a large market. *Journal of Financial Econometrics* 14:496–524.
- Chu, L., X.-Z. He, K. Li, and J. Tu. 2022. Investor sentiment and paradigm shifts in equity return forecasting. *Management Science* 68:4301–25.
- DeMiguel, V., L. Garlappi, F. Nogales, and R. Uppal. 2009. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science* 55:798–812.
- DeMiguel, V., L. Garlappi, and R. Uppal. 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies* 22:1915–53.
- DeMiguel, V., A. Martin-Utrera, F. Nogales, and R. Uppal. 2020. A transaction-cost perspective on the multitude of firm characteristics. *The Review of Financial Studies* 33:2180–222.
- DeMiguel, V., F. J. Nogales, and R. Uppal. 2014. Stock return serial dependence and out-of-sample portfolio performance. *The Review of Financial Studies* 27:1031–73.
- Gu, S., B. Kelly, and D. Xiu. 2020. Empirical asset pricing via machine learning. *The Review of Financial Studies* 33:2223–73.

- Hastie, T., R. Tibshirani, J. Friedman, and J. Franklin. 2009. *The elements of statistical learning: Data mining, inference and prediction*. Springer.
- Huang, D., F. Jiang, J. Tu, and G. Zhou. 2015. Investor sentiment aligned: A powerful predictor of stock returns. *The Review of Financial Studies* 28:791–837.
- Kirby, C., and B. Ostdiek. 2012. It’s all in the timing: Simple active portfolio strategies that outperform naïve diversification. *Journal of Financial and Quantitative Analysis* 47:437–67.
- Koijen, R., R. Richmond, and M. Yogo. 2020. Which investors matter for equity valuations and expected returns? Working Paper, National Bureau of Economic Research.
- Koijen, R., and M. Yogo. 2019. A demand system approach to asset pricing. *Journal of Political Economy* 127:1475–515.
- Kozak, S., S. Nagel, and S. Santosh. 2020. Shrinking the cross-section. *Journal of Financial Economics* 135:271–92.
- Ledoit, O., and M. Wolf. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10:603–21.
- . 2004. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88:365–411.
- . 2008. Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance* 15:850–9.
- . 2020. Analytical nonlinear shrinkage of large-dimensional covariance matrices. *The Annals of Statistics* 48:3043–65.
- Lopez-Lira, A., and N. Roussanov. 2023. Do common factors really explain the cross-section of stock returns? *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.
- Merton, R. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:323–61.
- Novy-Marx, R., and M. Velikov. 2016. A taxonomy of anomalies and their trading costs. *The Review of Financial Studies* 29:104–47.
- Olivares-Nadal, A., and V. DeMiguel. 2018. A robust perspective on transaction costs in portfolio optimization. *Operations Research* 66:733–9.
- Pettenuzzo, D., A. Timmermann, and R. Valkanov. 2014. Forecasting stock returns under economic constraints. *Journal of Financial Economics* 114:517–53.

- Politis, D., and J. Romano. 1994. The stationary bootstrap. *Journal of the American Statistical Association* 89:1303–13.
- Raponi, V., R. Uppal, and P. Zaffaroni. 2021. Robust portfolio choice. *SSRN working paper*.
- Roll, R. 1980. Orthogonal portfolios. *Journal of Financial and Quantitative Analysis* 15:1005–23.
- Ross, S. A. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13:341–60.
- Stambaugh, R., J. Yu, and Y. Yuan. 2012. The short of it: Investor sentiment and anomalies. *Journal of Financial Economics* 104:288–302.
- Tay, J., J. Friedman, and R. Tibshirani. 2021. Principal component-guided sparse regression. *The Canadian Journal of Statistics* 49:1222–57.

Internet Appendix to

Shrinking Against Sentiment:

Exploiting Latent Asset Demand in Portfolio Selection

This internet appendix contains four sections. Section A presents a theoretical model of the economy with investor sentiment. Section B derives the exact expected Sharpe ratio loss of the arbitrage portfolio. Section C provides additional empirical results. Section D reports the proofs of all theoretical results.

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 analyzes the theoretical relation between sentiment and asset prices in an economy with rational and sentiment-driven investors.

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*

$$d_{t+1} = \bar{d} + u, \tag{A.1}$$

where $\bar{d}_i > 0$ for all $i = 1, \dots, N$ 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 $\mathbb{E}[d_{t+1}] = \bar{d} + \varepsilon$ with $\varepsilon_i \geq 0$ for all $i = 1, \dots, N$.

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. 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 $1/N$ returns and an increase in the returns delivered by arbitrage portfolios that exploit asset overpricing.

A.2 The effect of sentiment on $1/N$ and arbitrage portfolios

We now show how the mass of sentiment investors α , and the level of sentiment demand ε , affect the performance of the $1/N$ and arbitrage portfolios in the economy described in Assumption A.1. First, in the next proposition, we derive the equilibrium Sharpe ratio of the $1/N$ portfolio.

Proposition A.1 *Let Assumption A.1 hold. Then, in equilibrium, the Sharpe ratio of the $1/N$ portfolio w_{ew} is*

$$\text{SR}(w_{ew}) = \gamma\sigma_{ew} - \frac{\alpha}{\sigma_{ew}}w_{ew}^\top\varepsilon, \quad (\text{A.2})$$

where α is the mass of sentiment investors, ε is the sentiment demand vector, and $\sigma_{ew} = \sqrt{w_{ew}^\top \Sigma w_{ew}}$ is the return volatility of the $1/N$ portfolio.

Proposition A.1 shows that the Sharpe ratio of the $1/N$ portfolio is the sum of two terms: 1) the Sharpe ratio when all investors in the market are rational ($\alpha = 0$ or $\varepsilon = 0$) and 2) a term that is increasingly negative as the mass of sentiment investors α and sentiment demand ε increase, consistent with empirical evidence.

Next, we study the effect of sentiment on the performance of the arbitrage portfolio w_a defined in (2). Under Assumption 1, we can characterize theoretically the positive relation between the Sharpe ratio of the arbitrage portfolio w_a and the mass of sentiment investors α and sentiment demand ε . Specifically, the following holds.

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

$$\text{SR}(w_a) = \alpha \sqrt{\sum_{i=2}^N \frac{(v_i^\top \varepsilon)^2}{\sigma_{PC_i}^2}}, \quad (\text{A.3})$$

where α is the mass of sentiment investors, ε is the sentiment demand vector, v_i is the i th eigenvector of the covariance matrix Σ , and $\sigma_{PC_i}^2 = v_i^\top \Sigma v_i$ is the variance of the i th principal component of returns.

Proposition A.2 characterizes the positive theoretical relation between the performance of the mean-variance arbitrage portfolio and the level of investor sentiment characterized by α and ε . In particular, the Sharpe ratio of the arbitrage portfolio is zero

when $\alpha = 0$ and $\varepsilon = 0$. This theoretical positive relation also holds in the data. Proposition A.2 also shows that the arbitrage portfolio captures asset mispricing driven by the latent demand of sentiment investors, ε . Therefore, latent asset demand is not just fundamental to characterize the cross-section of stock returns as pointed out by Kojien and Yogo (2019), but also to understand the sources of portfolio performance.

B Exact expected out-of-sample Sharpe ratio loss

In this section, we derive an analytical expression for the *exact* expected out-of-sample Sharpe ratio loss of the sample arbitrage portfolio $\hat{w}_a$ in (13), which we show agrees closely with our approximation (16). This exact expression builds on the expected out-of-sample Sharpe ratio defined as

$$\mathbb{E}[\text{SR}(\hat{w}_a)] = \mathbb{E}\left[\frac{\hat{w}_a^\top \mu}{\sqrt{\hat{w}_a^\top \Sigma \hat{w}_a}}\right]. \quad (\text{A.4})$$

Proposition A.3 *Consider the same assumptions as in Proposition 4. Then, the exact expected out-of-sample Sharpe ratio loss of the sample arbitrage portfolio $\hat{w}_a$ in (13) is*

$$\begin{aligned} & \frac{\text{SR}(w_a) - \mathbb{E}[\text{SR}(\hat{w}_a)]}{\text{SR}(w_a)} \\ &= 1 - \sqrt{\frac{\text{SR}^2(w_a)T}{2}} \frac{\Gamma(\frac{N}{2})\Gamma(\frac{T}{2})\Gamma(\frac{T-N+3}{2})}{\Gamma(\frac{N+1}{2})\Gamma(\frac{T+1}{2})\Gamma(\frac{T-N+2}{2})} {}_1F_1\left(\frac{1}{2}; \frac{N+1}{2}; -\frac{\text{SR}^2(w_a)T}{2}\right), \end{aligned} \quad (\text{A.5})$$

where ${}_1F_1$ is the confluent hypergeometric function and $\mathbb{E}[\text{SR}(\hat{w}_a)]$ is defined in (A.4). Moreover, this loss decreases with $\text{SR}^2(w_a)$.

Just like the approximation in (16), this loss decreases with the arbitrage portfolio's squared Sharpe ratio, $\text{SR}^2(w_A)$. Moreover, we find that the exact and approximate expressions agree very closely. For example, let $N = 25$, $T = 180$, and $\text{SR}(w_a) = 0.2$. Then, the approximate loss in (16) is 0.5546 while the exact one in (A.5) is 0.5527.

C Additional empirical results

C.1 Other normalization of mean-variance portfolios

In this section, we consider two other normalization of the mean-variance portfolio. First, the maximum-utility normalization,

$$w_t^* = \frac{1}{\gamma} \Sigma_t^{-1} \mu_t, \quad (\text{A.6})$$

where γ is the investor's risk-aversion coefficient that we set equal to three as in [Kan and Zhou \(2007\)](#) and [Tu and Zhou \(2011\)](#). The second normalization is such that the portfolio risk is constant,

$$w_t^* = \frac{\sigma}{\sqrt{\mu_t^\top \Sigma_t^{-1} \mu_t}} \Sigma_t^{-1} \mu_t, \quad (\text{A.7})$$

where we set a target volatility of $\sigma = 0.08$. For the 1/N benchmark, the maximum-utility portfolio is $\frac{\mu_{ew}}{\gamma \sigma_{ew}^2} w_{ew}$ and the constant-risk portfolio $\frac{\sigma}{\sigma_{ew}} w_{ew}$. To avoid leverage, when the sum of portfolio weights is larger than one, we scale down the exposure to the risky assets so that the sum of the weights is capped at one.

Table [A.1](#) reports the net out-of-sample Sharpe ratio using proportional transaction costs of 30 basis points, and the turnover, for the seven portfolio strategies we consider. The conclusions in the main body of the manuscript are robust to using different normalization. Specifically, the average outperformance of the SAS portfolio relative to the four benchmark mean-variance portfolios is 29% and 19% for the maximum-utility and constant-risk normalization, respectively. The SAS portfolio also continues to largely outperform the simple RTR and 1/N portfolios. Moreover, the benchmark mean-variance portfolios require a turnover that is on average 350% and 229% larger than that of the SAS portfolio for the maximum-utility and constant-risk normalization, respectively.

Table A.1: Sharpe ratio and turnover for two other portfolio normalization

This table reports the out-of-sample performance of the SAS portfolio and six benchmark portfolios. The first four benchmarks are mean-variance portfolios constructed with the shrinkage covariance matrices of [Ledoit and Wolf \(2004\)](#) (MV-1N), [Ledoit and Wolf \(2003\)](#) (MV-MKT), [Ledoit and Wolf \(2020\)](#) (MV-NL), and [Chen and Yuan \(2016\)](#) (MV-PCA). The last two benchmarks are the reward-to-risk timing strategy of [Kirby and Ostdiek \(2012\)](#) (RTR), and the equally weighted (1/N) portfolio. Panels A and B report the annualized out-of-sample Sharpe ratio net of proportional transaction costs of 30 basis points and the average monthly turnover of each portfolio for the two normalization described in Section C.1, as well as the p-values (in parenthesis) of the significance test of the difference of each portfolio's Sharpe ratio with that of EW. The six datasets are described in Section 4.2.

Policy	25SBETA	25SBTM	25OPINV	16ANOM	46ANOM	100SINV
Panel A: Maximum utility						
<i>Net Sharpe ratio</i>						
SAS	0.80 (0.05)	0.87 (0.01)	0.89 (0.01)	1.09 (0.00)	2.15 (0.00)	0.97 (0.00)
MV-1N	0.58 (0.64)	0.76 (0.22)	0.68 (0.33)	1.00 (0.03)	2.09 (0.00)	0.62 (0.61)
MV-MKT	0.47 (0.91)	0.58 (0.68)	0.67 (0.33)	0.94 (0.04)	1.93 (0.00)	0.60 (0.64)
MV-NL	0.53 (0.91)	0.69 (0.39)	0.71 (0.26)	0.96 (0.02)	1.86 (0.00)	0.81 (0.14)
MV-PCA	0.69 (0.27)	0.73 (0.19)	0.73 (0.14)	0.90 (0.05)	1.38 (0.00)	0.80 (0.09)
RTR	0.63 (0.00)	0.60 (0.00)	0.61 (0.01)	0.57 (0.00)	0.59 (0.00)	0.63 (0.00)
1/N	0.51	0.49	0.47	0.36	0.28	0.51
<i>Turnover</i>						
SAS	0.24	0.28	0.27	0.29	1.13	0.46
MV-1N	1.10	1.19	1.08	1.05	6.60	3.48
MV-MKT	1.52	1.56	1.10	1.44	7.35	3.43
MV-NL	1.36	1.41	1.02	1.29	3.31	2.37
MV-PCA	0.32	0.54	0.22	1.76	5.84	0.38
RTR	0.04	0.04	0.04	0.05	0.04	0.04
1/N	0.02	0.02	0.02	0.02	0.02	0.03
Panel B: Constant risk						
<i>Net Sharpe ratio</i>						
SAS	0.82 (0.07)	0.87 (0.03)	0.90 (0.02)	1.08 (0.00)	1.96 (0.00)	1.00 (0.00)
MV-1N	0.60 (0.91)	0.90 (0.11)	0.75 (0.27)	1.01 (0.04)	1.94 (0.00)	0.84 (0.16)
MV-MKT	0.51 (0.93)	0.80 (0.28)	0.75 (0.29)	0.97 (0.03)	1.97 (0.00)	0.84 (0.15)
MV-NL	0.52 (0.81)	0.85 (0.17)	0.75 (0.30)	0.95 (0.04)	1.77 (0.00)	0.91 (0.06)
MV-PCA	0.71 (0.40)	0.71 (0.35)	0.75 (0.21)	0.93 (0.07)	1.42 (0.00)	0.80 (0.20)
RTR	0.63 (0.00)	0.60 (0.00)	0.61 (0.00)	0.57 (0.00)	0.59 (0.00)	0.63 (0.00)
1/N	0.56	0.55	0.54	0.45	0.35	0.58
<i>Turnover</i>						
SAS	0.22	0.27	0.28	0.27	0.39	0.40
MV-1N	0.89	0.94	0.83	0.70	1.33	1.94
MV-MKT	1.12	1.14	0.83	0.85	1.44	1.85
MV-NL	1.07	1.07	0.79	0.83	0.97	1.41
MV-PCA	0.30	0.48	0.30	0.99	1.65	0.38
RTR	0.04	0.04	0.04	0.05	0.04	0.04
1/N	0.04	0.04	0.04	0.04	0.04	0.04

C.2 Dynamic covariance matrix

In this section, we assess the performance of mean-variance portfolios constructed using a dynamic covariance matrix. Specifically, given an estimation window of T months, we estimate the covariance matrix using the exponentially weighted moving average model:

$$\hat{\Sigma}_{dyn} = \sum_{t=1}^T (1 - \lambda) \lambda^{t-1} (r_t - \hat{\mu})(r_t - \hat{\mu})^\top, \quad (\text{A.8})$$

where $\lambda = 0.94$ as recommended by RiskMetrics. Then, we shrink $\hat{\Sigma}_{dyn}$ toward the identity matrix using either the SAS shrinkage intensity or that of [Ledoit and Wolf \(2004\)](#) (MV-1N). In [Table A.2](#), we report the Sharpe ratio and turnover of the two portfolios for the three types of mean-variance portfolio normalization we consider in [Sections 4.3.1 and C.1](#). Two results stand out. First, the SAS portfolio continues to outperform the MV-1N portfolio and to deliver a smaller turnover. Second, the performance of the SAS portfolio is comparable to that under the static sample covariance matrix, and thus, our proposed portfolio methodology also delivers good performance when applied to dynamic covariance matrices.

C.3 Time-varying exposure to the 1/N component

Here we study the performance of the SAS portfolio using conditional time-series regressions. In particular, we consider the one-factor model in [\(5\)](#) and regress the out-of-sample returns of SAS portfolios on 1/N portfolio returns, and we allow for the slope coefficient to vary over time with investor sentiment. That is,

$$r_t^{\text{SAS}} = \alpha + \underbrace{(\beta_0 + \beta_1 \text{Sentiment}_{t-1})}_{\beta_t} r_t^{1/N} + \epsilon_t, \quad (\text{A.9})$$

Table A.2: Out-of-sample performance with dynamic covariance matrix

This table reports the out-of-sample performance of the SAS portfolio and the mean-variance portfolio constructed with the shrinkage covariance matrix of [Ledoit and Wolf \(2004\)](#) (MV-1N). The sample covariance matrix is estimated dynamically using the exponentially weighted moving average model in [\(A.8\)](#). Panels A, B, and C report the annualized out-of-sample Sharpe ratios before and after proportional transaction costs of 30 basis points, and the average monthly turnover, for the three normalization described in Sections [4.3.1](#) and [C.1](#), respectively. The six datasets are described in Section [4.2](#).

Policy		25SBETA	25SBTM	25OPINV	16ANOM	46ANOM	100SINV
Panel A: Tangency							
SAS	Gross SR	0.68	0.70	0.85	0.90	2.01	0.76
	Net SR	0.61	0.62	0.76	0.83	1.85	0.62
	Turnover	0.39	0.59	0.47	0.49	1.33	0.93
MV-1N	Gross SR	0.50	0.31	0.31	0.56	1.10	0.40
	Net SR	0.14	-0.07	-0.01	0.28	0.70	-0.06
	Turnover	3.42	5.73	6.98	6.36	28.2	19.0
Panel B: Maximum utility							
SAS	Gross SR	0.76	0.81	0.89	0.93	2.02	0.81
	Net SR	0.69	0.72	0.81	0.87	1.87	0.70
	Turnover	0.32	0.47	0.37	0.32	1.17	0.67
MV-1N	Gross SR	0.67	0.51	0.44	0.71	1.63	0.53
	Net SR	0.30	0.14	0.11	0.43	1.21	0.19
	Turnover	2.60	3.37	3.99	3.42	15.0	9.10
Panel C: Constant risk							
SAS	Gross SR	0.81	0.83	0.89	0.96	2.14	0.95
	Net SR	0.74	0.75	0.81	0.89	2.03	0.84
	Turnover	0.31	0.44	0.37	0.29	0.52	0.57
MV-1N	Gross SR	0.80	0.77	0.56	0.77	1.97	0.83
	Net SR	0.40	0.36	0.22	0.51	1.55	0.39
	Turnover	2.30	2.76	2.85	1.91	4.85	4.23

where Sentiment_t is the [Huang et al. \(2015\)](#) sentiment index at time t and $r_t^{1/N}$ is the $1/N$ portfolio return at time t . The sentiment index is standardized to have zero mean and unit variance, so that β_0 represents the average $1/N$ beta. Regression model [\(A.9\)](#) allows us to capture the time-varying $1/N$ beta of SAS portfolios. Table [A.3](#) 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.

Table A.3: Conditional 1/N betas

This table reports the intercept, slope coefficients, and Newey-West t-statistics (in square brackets) from the conditional time-series regressions in Equation (A.9). We run time-series regressions for the 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.

	25SBETA		25SBTM		25OPINV		16ANOM		46ANOM		100SINV	
	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N	SAS	MV-1N
α	0.01 [4.21]	0.01 [4.33]	0.01 [4.50]	0.02 [2.58]	0.01 [4.73]	0.02 [3.53]	0.02 [5.83]	0.03 [5.61]	0.07 [11.2]	0.50 [11.3]	0.01 [5.45]	0.03 [5.24]
β_0	0.41 [6.61]	0.22 [2.10]	0.78 [13.7]	0.57 [2.68]	0.73 [13.6]	0.49 [3.48]	0.73 [9.23]	0.33 [2.29]	0.18 [1.33]	-2.07 [-2.59]	0.71 [13.6]	0.32 [2.27]
β_1	-0.10 [-1.48]	-0.09 [-0.82]	-0.30 [-3.95]	-0.95 [-2.38]	-0.22 [-3.32]	-0.56 [-3.23]	-0.50 [-4.97]	-0.82 [-4.77]	-0.40 [-3.89]	-1.78 [-3.42]	-0.24 [-5.12]	-0.44 [-3.37]


We observe that in all six datasets, the exposure of SAS portfolios to the 1/N component decreases when sentiment increases (i.e., $\beta_1 < 0$ in (A.9)), and it is statistically significant. Similarly, the conditional betas of MV-1N portfolios also decrease with sentiment, however their statistical significance is lower. In general, as discussed in Section 3.4, the correlation of any mean-variance portfolio with 1/N returns should decrease when 1/N returns are smaller, which is typically the case after periods of high sentiment.

The main insight obtained from Table A.3 is highlighted in Figure A.1, which shows the substantial time-varying nature of 1/N betas, β_t in Equation (A.9), of mean-variance portfolios for the 25SBTM dataset. First, the figure shows that both SAS and MV-1N portfolio betas have a strong negative correlation with investor sentiment. Interestingly, the market beta of mean-variance portfolios can even be negative when investor sentiment is large and therefore when market returns are prone to be low or even negative.

Second, our SAS portfolio keeps a higher exposure to the 1/N component relative to the MV-1N portfolio. Indeed, in Table A.3 the average 1/N beta, β_0 , of the SAS portfolio is larger than that of the MV-1N portfolio across all datasets. This result can be explained because MV-1N portfolios shrink less intensively and become less exposed to 1/N over

Figure A.1: Conditional 1/N betas versus sentiment

This figure depicts the conditional 1/N beta, β_t in Equation (A.9), of the SAS portfolio for the dataset of 25 portfolios of stocks sorted on size and book-to-market described in Section 4.2 (left axis). For comparison purposes, the figure also depicts the market beta of the mean-variance portfolio that exploits the Ledoit and Wolf (2004) shrinkage covariance matrix (MV-1N). The dashed horizontal lines depict the average 1/N beta, β_0 in (A.9). The figure also depicts the Huang et al. (2015) sentiment index (right axis).



time as the size of the expanding window increases and parameter uncertainty in the covariance matrix decreases. In contrast, SAS portfolios shrink not only as a function of parameter uncertainty in the covariance matrix. First, as explained in Section 4.1, SAS portfolios shrink as a function of parameter uncertainty in the vector of means, which affects the out-of-sample performance of the arbitrage portfolio and can be substantial even for large sample sizes. Second, SAS portfolios shrink as a function of sentiment, and find it beneficial to keep a large exposure to the 1/N component when sentiment is low.

The higher average 1/N beta of SAS portfolios is also an important distinction relative to the benchmark mean-variance portfolios because it reduces transaction costs, as evidenced by the substantially lower turnover of our strategy in Table 1. We study the trading-cost benefits of our SAS portfolio approach in more detail in the next section.

Table A.4: 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. The time period spans 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 (A.10). We regress each cross-sectional percentile on the prior-month value of the [Huang et al. \(2015\)](#) sentiment index.

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

C.4 Sentiment and market liquidity

[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 toward 1/N when liquidity is low and vice versa, which reduces transaction costs. To study 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 \(2020, Appendix IA.1\)](#).¹⁸ For each stock i in month t , we define its corresponding bid-ask spread as

$$\hat{s}_{i,t} = \frac{1}{D} \sum_{d=1}^D \hat{s}_{i,d}, \quad \hat{s}_{i,d} = \sqrt{\max\{4(c_{i,d} - \eta_{i,d})(c_{i,d} - \eta_{i,d+1}), 0\}}, \quad (\text{A.10})$$

where D is the number of days in month t , $\hat{s}_{i,d}$ 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 , i.e., the mean of daily high and low log-prices.

¹⁸We download daily price data from CRSP.

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

$$p_{t+1}^l = \alpha + \beta \text{Sentiment}_t + \epsilon_{t+1}, \quad (\text{A.11})$$

where Sentiment_t is the [Huang et al. \(2015\)](#) sentiment index in month t . Table [A.4](#) 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 1/N portfolio when sentiment and liquidity are low. 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 than 1/N, when liquidity is high. Indeed, in unreported results we show that the 1/N portfolio has a dramatically lower turnover than the arbitrage portfolio across all datasets we consider in the out-of-sample analysis. That is, tilting mean-variance portfolios toward the arbitrage component requires a level of transaction costs that is several orders of magnitude larger than those paid by the 1/N portfolio.

D Proofs of all results

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

Proof of Proposition [1](#)

From the definition of the arbitrage portfolio w_a in [\(2\)](#), we have that

$$w_a = w^* - \frac{\mu_g}{\gamma} \Sigma^{-1} l. \quad (\text{A.12})$$

To prove the proposition, we rewrite μ_g and $\Sigma^{-1}\iota$ under Assumption 1. Recall that the eigenvalue decomposition of the covariance matrix Σ is

$$\Sigma = VD V^\top, \quad (\text{A.13})$$

where D is a diagonal matrix whose i th element is the eigenvalue associated to the i th PC, and V is a matrix whose i th column v_i is the eigenvector associated to the i th PC, and $v_1 = \iota/\sqrt{N}$ under Assumption 1. Now, $\mu_g = \frac{\iota^\top \Sigma^{-1} \mu}{\iota^\top \Sigma^{-1} \iota}$, where

$$\iota^\top \Sigma^{-1} \mu = \sqrt{N} v_1^\top V D^{-1} V^\top \mu = \sqrt{N} \frac{\mu_{PC_1}}{\sigma_{PC_1}^2}, \quad (\text{A.14})$$

$$\iota^\top \Sigma^{-1} \iota = N v_1^\top V D^{-1} V^\top v_1 = \frac{N}{\sigma_{PC_1}^2}. \quad (\text{A.15})$$

Therefore, μ_g simplifies to

$$\mu_g = \frac{\iota^\top \Sigma^{-1} \mu}{\iota^\top \Sigma^{-1} \iota} = \frac{\mu_{PC_1}}{\sqrt{N}}. \quad (\text{A.16})$$

Moreover, $\Sigma^{-1}\iota$ simplifies to

$$\Sigma^{-1}\iota = V D^{-1} V^\top \iota = \sqrt{N} V D^{-1} V^\top v_1 = \sqrt{N} \frac{v_1}{\sigma_{PC_1}^2} = N \frac{w_{ew}}{\sigma_{PC_1}^2}. \quad (\text{A.17})$$

Finally, under Assumption 1 it holds that $w_{ew} = v_1/\sqrt{N}$, and thus the mean return and variance of the $1/N$ portfolio are

$$\mu_{ew} = \frac{\mu_{PC_1}}{\sqrt{N}} \quad \text{and} \quad \sigma_{ew}^2 = \frac{\sigma_{PC_1}^2}{N}. \quad (\text{A.18})$$

Equation (A.18) implies that $\mu_g = \mu_{ew}$ and $\Sigma^{-1}\iota = w_{ew}/\sigma_{ew}^2$, and plugging these two equalities in (A.12) yields Equation (4), which completes the proof.

Proof of Proposition 2

The factor f is the return of the $1/N$ portfolio, and thus, the mean of excess returns r is $\mu = \alpha + \beta\mu_{ew}$. Therefore, the mean-variance portfolio $w^* = \frac{1}{\gamma}\Sigma^{-1}\mu$ admits the decomposition in (6). Now, let $\beta = \frac{\Sigma w_{ew}}{\sigma_{ew}^2}$, where $w_{ew} = \iota/N$. Then, it is clear that the beta portfolio $\frac{1}{\gamma}\Sigma^{-1}\beta\mu_{ew} = \frac{1}{\gamma}\frac{\mu_{ew}}{\sigma_{ew}^2}w_{ew}$, which is the $1/N$ component in (4). Moreover, from the proof of Proposition 1, we have that $\mu_g = \mu_{ew}$ under Assumption 1. Therefore, the arbitrage portfolio is $w_a = \frac{1}{\gamma}\Sigma^{-1}(\mu - \mu_{ew}\iota)$, and it is equivalent to the alpha portfolio $\frac{1}{\gamma}\Sigma^{-1}(\mu - \mu_{ew}\beta)$ if $\beta = \iota$ under Assumption 1. The latter holds because, from Equation (A.17), we have that $\Sigma\iota = Nw_{ew}\sigma_{PC_1}^2 = N\sigma_{ew}^2\iota$, and thus $\beta = \frac{\Sigma\iota}{N\sigma_{ew}^2} = \iota$, which completes the proof.

Proof of Proposition 3

Given the eigenvalue decomposition of Σ in (A.13), it is straightforward to show that the mean-variance portfolio's squared Sharpe ratio is

$$\text{SR}^2(w^*) = \mu^\top \Sigma^{-1} \mu = \sum_{i=1}^N \text{SR}_{PC_i}^2, \quad (\text{A.19})$$

where $\text{SR}_{PC_i}^2$ is the squared Sharpe ratio of the i th PC of stock returns.

Part 1. Given the eigenvalue decomposition of the Sharpe ratio of the mean-variance portfolio in (A.19), we need to prove the equalities for the Sharpe ratio of the $1/N$ and arbitrage portfolios in Equation (10). From Equation (A.18), we directly have that the $1/N$ Sharpe ratio is $\text{SR}(w_{ew}) = \mu_{ew}/\sigma_{ew} = \text{SR}_{PC_1}$, as in Equation (10). Turning to the arbitrage portfolio w_a in (2), its mean return and variance are

$$w_a^\top \mu = \frac{1}{\gamma}(\mu - \mu_g \iota)^\top \Sigma^{-1} \mu, \quad (\text{A.20})$$

$$w_a^\top \Sigma w_a = \frac{1}{\gamma^2}(\mu - \mu_g \iota)^\top \Sigma^{-1}(\mu - \mu_g \iota). \quad (\text{A.21})$$

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

$$w_a^\top \mu = \frac{1}{\gamma} \left(\mu^\top \Sigma^{-1} \mu - \frac{(\iota^\top \Sigma^{-1} \mu)^2}{\iota^\top \Sigma^{-1} \iota} \right) \quad \text{and} \quad w_a^\top \Sigma w_a = \frac{1}{\gamma} w_a^\top \mu. \quad (\text{A.22})$$

Therefore, the squared Sharpe ratio of the arbitrage portfolio is

$$\text{SR}^2(w_a) = \frac{(w_a^\top \mu)^2}{w_a^\top \Sigma w_a} = \mu^\top \Sigma^{-1} \mu - \frac{(\iota^\top \Sigma^{-1} \mu)^2}{\iota^\top \Sigma^{-1} \iota}. \quad (\text{A.23})$$

Finally, plugging equations (A.14), (A.15), and (A.19) into (A.23), we obtain Equation (10), which completes the proof.

Part 2. This result is a special case of the result in Proposition 7 when the shrinkage intensity $\delta = 0$.

Proof of Proposition 4

Under the assumption that stock returns are i.i.d. normally distributed and $T > N + 3$, we have from Lemma 1 in Kan, Wang, and Zhou (2021) that the expected out-of-sample mean return and variance of the sample arbitrage portfolio $\hat{w}_a$ are

$$\mathbb{E}[\hat{w}_a^\top \mu] = \frac{1}{\gamma} \frac{T}{T - N - 1} \text{SR}^2(w_a), \quad (\text{A.24})$$

$$\mathbb{E}[\hat{w}_a^\top \Sigma \hat{w}_a] = \frac{1}{\gamma^2} \frac{T^2(T - 2)}{(T - N)(T - N - 1)(T - N - 3)} \left(\text{SR}^2(w_a) + \frac{N - 1}{T} \right), \quad (\text{A.25})$$

where $\text{SR}^2(w_a)$ is given by (A.23) in general, and by (10) under Assumption 1. Therefore, the expected out-of-sample Sharpe ratio of $\hat{w}_a$ is

$$\mathbb{E}[\text{SR}(\hat{w}_a)] = \frac{\mathbb{E}[\hat{w}_a^\top \mu]}{\sqrt{\mathbb{E}[\hat{w}_a^\top \Sigma \hat{w}_a]}} = \sqrt{\frac{(T - N)(T - N - 3)}{(T - 2)(T - N - 1)}} \frac{\text{SR}^2(w_a)}{\sqrt{\text{SR}^2(w_a) + \frac{N - 1}{T}}}, \quad (\text{A.26})$$

which delivers the desired result in Equation (16) and completes the proof.

Proof of Proposition 5

Under Assumption 1, we have from Equation (9) that $\text{SR}^2(w_a) = \text{SR}^2(w^*) - \text{SR}^2(w_{ew}) = \mu^\top \Sigma^{-1} \mu - \mu_{ew}^2 / \sigma_{ew}^2$. Now, let us expand $\alpha^\top \Sigma^{-1} \alpha$ where $\alpha = \mu - \beta \mu_{ew}$:

$$\alpha^\top \Sigma^{-1} \alpha = \mu^\top \Sigma^{-1} \mu + \mu_{ew}^2 \beta^\top \Sigma^{-1} \beta - 2\mu_{ew} \beta^\top \Sigma^{-1} \mu. \quad (\text{A.27})$$

Using $\beta = \Sigma w_{ew} / \sigma_{ew}^2$, it is easy to see that $\mu_{ew}^2 \beta^\top \Sigma^{-1} \beta = \mu_{ew} \beta^\top \Sigma^{-1} \mu = \text{SR}^2(w_{ew})$ and thus $\alpha^\top \Sigma^{-1} \alpha = \text{SR}^2(w_a)$, which completes the proof.

Proof of Proposition 6

The squared Sharpe ratio of the 1/N portfolio is independent of the shrinkage intensity δ and, as shown in (10), is given by $\text{SR}^2(w_{ew}) = \text{SR}_{PC1}^2$.

We turn next to the shrinkage mean-variance portfolio in (21), $w_s^* = \frac{1}{\gamma} \Sigma_{sh}^{-1} \mu$, where $\Sigma_{sh} = V((1 - \delta)D + \delta\nu I_N)V^\top$. Its mean return is

$$(w_s^*)^\top \mu = \frac{1}{\gamma} \mu^\top V((1 - \delta)D + \delta\nu I_N)^{-1} V^\top \mu = \frac{1}{\gamma} \sum_{i=1}^N \frac{\mu_{PCi}^2}{(1 - \delta)\sigma_{PCi}^2 + \delta\nu}. \quad (\text{A.28})$$

Moreover, its variance is

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

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

$$(D_\delta)_{ii} = \frac{\sigma_{PCi}^2}{((1 - \delta)\sigma_{PCi}^2 + \delta\nu)^2}. \quad (\text{A.29})$$

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

$$(w_s^*)^\top \Sigma w_s^* = \frac{1}{\gamma^2} \sum_{i=1}^N \frac{\mu_{PCi}^2 \sigma_{PCi}^2}{((1 - \delta)\sigma_{PCi}^2 + \delta\nu)^2}. \quad (\text{A.30})$$

Finally, using Equations (A.28)–(A.30), the squared Sharpe ratio of the shrinkage mean-variance portfolio is

$$\text{SR}^2(w_s^*) = \frac{((w_s^*)^\top \mu)^2}{(w_s^*)^\top \Sigma w_s^*} = S_\delta(1), \quad (\text{A.31})$$

where $S_\delta(j)$ is defined in (22).

A similar line of reasoning holds for the shrinkage arbitrage portfolio in (20),

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

Under Assumption 1, it holds that $\frac{\iota^\top \Sigma_{sh}^{-1} \mu}{\iota^\top \Sigma_{sh}^{-1} \iota} = \mu_{PC1}/\sqrt{N}$, and thus,

$$w_{sa} = \frac{1}{\gamma} \Sigma_{sh}^{-1} \left(\mu - \frac{\mu_{PC1}}{\sqrt{N}} \iota \right). \quad (\text{A.32})$$

Therefore, the mean return of the shrinkage arbitrage portfolio is

$$w_{sa}^\top \mu = \frac{1}{\gamma} \left(\mu^\top \Sigma_{sh}^{-1} \mu - \mu^\top \Sigma_{sh}^{-1} \iota \frac{\mu_{PC1}}{\sqrt{N}} \right) = \frac{1}{\gamma} \sum_{i=2}^N \frac{\mu_{PCi}^2}{(1-\delta)\sigma_{PCi}^2 + \delta\nu}, \quad (\text{A.33})$$

which holds because $\mu^\top \Sigma_{sh}^{-1} \iota = \sqrt{N} \frac{\mu_{PC1}}{(1-\delta)\sigma_{PCi}^2 + \delta\nu}$ under Assumption 1. Moreover, the variance of the shrinkage arbitrage portfolio is

$$\begin{aligned} w_{sa}^\top \Sigma w_{sa} &= \frac{1}{\gamma^2} \left(\mu - \frac{\mu_{PC1}}{\sqrt{N}} \iota \right)^\top \Sigma_{sh}^{-1} \Sigma \Sigma_{sh}^{-1} \left(\mu - \frac{\mu_{PC1}}{\sqrt{N}} \iota \right) \\ &= (w_s^*)^\top \Sigma w_s^* + \frac{1}{\gamma^2} \left(\frac{\mu_{PC1}^2}{N} \iota^\top V D_\delta V^\top \iota - 2 \frac{\mu_{PC1}}{\sqrt{N}} \iota^\top V D_\delta V^\top \mu \right) \\ &= \frac{1}{\gamma^2} \sum_{i=2}^N \frac{\mu_{PCi}^2 \sigma_{PCi}^2}{((1-\delta)\sigma_{PCi}^2 + \delta\nu)^2}, \end{aligned} \quad (\text{A.34})$$

where the diagonal matrix D_δ is defined in (A.29). The last equality in (A.34) holds because, under Assumption 1,

$$\iota^\top V D_\delta V^\top \iota = N \frac{\sigma_{PC1}^2}{((1-\delta)\sigma_{PCi}^2 + \delta\nu)^2} \quad \text{and} \quad \iota^\top V D_\delta V^\top \mu = \sqrt{N} \frac{\mu_{PC1} \sigma_{PC1}^2}{((1-\delta)\sigma_{PCi}^2 + \delta\nu)^2}.$$

Finally, using Equations (A.33)–(A.34), the squared Sharpe ratio of the shrinkage arbitrage portfolio is

$$\text{SR}^2(w_{sa}) = \frac{(w_{sa}^\top \mu)^2}{w_{sa}^\top \Sigma w_{sa}} = S_\delta(2), \quad (\text{A.35})$$

where $S_\delta(j)$ is defined in (22), which completes the proof.

Proof of Corollary 1

Because the shrinkage arbitrage portfolio w_{sa} in (20) is proportional to $1/\gamma$, any value of γ delivers the same squared Sharpe ratio, and the same mean-variance utility up to a $1/\gamma$ coefficient. That is, the return mean and standard deviation of w_{sa} trace out a linear frontier. Therefore, the squared Sharpe ratio of w_{sa} decreases with δ if its mean-variance utility also decreases with δ . To show that this result holds, denote the squared Sharpe ratio of the i th PC extracted from the shrinkage covariance matrix Σ_{sh} in (19) by

$$\text{SR}_{PC_i}^2(\delta) = \frac{\text{SR}_{PC_i}^2}{1 - \delta + \delta\nu/\sigma_{PC_i}^2}, \quad (\text{A.36})$$

which is positive for all $\delta \in [0, 1]$. Then, given the formulas for the return mean and variance of w_{sa} in (A.33)–(A.34), the mean-variance utility of w_{sa} is

$$w_{sa}^\top \mu - \frac{\gamma}{2} w_{sa}^\top \Sigma w_{sa} = \frac{1}{2\gamma} \sum_{i=2}^N \text{SR}_{PC_i}^2(\delta) \left(2 - \frac{\text{SR}_{PC_i}^2(\delta)}{\text{SR}_{PC_i}^2} \right). \quad (\text{A.37})$$

The derivative of (A.37) with respect to δ is

$$\frac{\partial}{\partial \delta} \left[w_{sa}^\top \mu - \frac{\gamma}{2} w_{sa}^\top \Sigma w_{sa} \right] = -\frac{\delta}{\gamma} \sum_{i=2}^N \text{SR}_{PC_i}^2(\delta) \left(\frac{1 - \nu/\sigma_{PC_i}^2}{1 - \delta + \delta\nu/\sigma_{PC_i}^2} \right)^2, \quad (\text{A.38})$$

and it is negative for all $\delta \in [0, 1]$, which completes the proof.

Proof of Proposition 7

Let us begin with the correlation between the return of the $1/N$ portfolio and that of the shrinkage mean-variance portfolio. This correlation is

$$\text{Corr}(w_{ew}^\top r, (w_s^*)^\top r) = \frac{w_{ew}^\top \Sigma w_s^*}{\sqrt{(w_{ew}^\top \Sigma w_{ew})((w_s^*)^\top \Sigma w_s^*)}}. \quad (\text{A.39})$$

The return variances of w_{ew} and w_s^* are given by Equations (A.18) and (A.30), respectively, and thus we only need to treat the term $w_{ew}^\top \Sigma w_s^*$, which under Assumption 1 simplifies to

$$w_{ew}^\top \Sigma w_s^* = \frac{1}{\gamma N} \iota^\top \Sigma \Sigma_{sh}^{-1} \mu = \frac{1}{\gamma \sqrt{N}} v_1^\top V \tilde{D}_\delta V^\top \mu,$$

where $\tilde{D}_\delta$ is a diagonal matrix with entries

$$(\tilde{D}_\delta)_{ii} = \frac{\sigma_{PC_i}^2}{(1 - \delta)\sigma_{PC_i}^2 + \delta\nu}.$$

Therefore, the quantity $w_{ew}^\top \Sigma w_s^*$ is

$$w_{ew}^\top \Sigma w_s^* = \frac{1}{\gamma \sqrt{N}} \frac{\mu_{PC_1} \sigma_{PC_1}^2}{(1 - \delta)\sigma_{PC_1}^2 + \delta\nu}. \quad (\text{A.40})$$

Plugging (A.18), (A.30), and (A.40) into (A.39), we find that the squared correlation between the return of the $1/N$ portfolio and that of the shrinkage mean-variance portfolio is given by (23), where $s_i(\delta) \leq 1$ because $\sigma_{PC_1}^2 \geq \sigma_{PC_i}^2$ for all $i = 2, \dots, N$.

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

$$\text{Corr}(w_{sa}^\top r, (w_s^*)^\top r) = \frac{w_{sa}^\top \Sigma w_s^*}{\sqrt{(w_{sa}^\top \Sigma w_{sa})((w_s^*)^\top \Sigma w_s^*)}}. \quad (\text{A.41})$$

The return variances of w_{sa} and w_s^* are given by Equations (A.34) and (A.30), respectively, and thus we only need to treat the term $w_{sa}^\top \Sigma w_s^*$, which under Assumption 1 is

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

As shown in the proof of Proposition 6, we have the decomposition $\Sigma_{sh}^{-1} \Sigma \Sigma_{sh}^{-1} = V D_\delta V^\top$, where D_δ is the diagonal matrix in (A.29). Using this result, the quantity $w_{sa}^\top \Sigma w_s^*$ is

$$w_{sa}^\top \Sigma w_s^* = \frac{1}{\gamma^2} \sum_{i=2}^N \frac{\mu_{PC_i}^2 \sigma_{PC_i}^2}{((1-\delta)\sigma_{PC_i}^2 + \delta\nu)^2}. \quad (\text{A.42})$$

Plugging (A.34), (A.30), and (A.42) into (A.41), we find that the correlation between the return of the shrinkage arbitrage portfolio and that of the shrinkage mean-variance portfolio is given by (24).

Finally, we conclude the proof by showing that the squared correlation with $1/N$, $\text{Corr}^2(w_{ew}^\top r, (w_s^*)^\top r)$, is increasing in the shrinkage intensity δ , which implies that the correlation with the shrinkage arbitrage portfolio is decreasing in δ given that the two squared correlations sum up to one in (24). Given Equation (23), to prove this result it is sufficient to show that the quantity $s_i(\delta)$ in (25) is decreasing in δ for all $i = 2, \dots, N$, i.e., that $\kappa(\delta)$ is increasing in δ . The derivative of $\kappa(\delta)$ with respect to δ is

$$\frac{\partial \kappa(\delta)}{\partial \delta} = \frac{\nu \sigma_{PC_1}^2}{(\delta\nu + (1-\delta)\sigma_{PC_1}^2)^2},$$

which is indeed positive and thus completes the proof.

Proof of Proposition A.1

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

$$w_l^* = \frac{1}{\gamma} \Sigma^{-1} \left(\mathbb{E}_l[d_{t+1}] - P_t \right), \quad (\text{A.43})$$

where P_t is the vector of equilibrium prices and $\mathbb{E}_l[d_{t+1}]$ 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

$$\frac{\alpha}{\gamma}\Sigma^{-1}(\bar{d} + \varepsilon - P_t) + \frac{1-\alpha}{\gamma}\Sigma^{-1}(\bar{d} - P_t) = w_{ew}, \quad (\text{A.44})$$

where w_{ew} is the 1/N portfolio. From the clearing condition (A.44), the vector of equilibrium prices P_t is

$$P_t = \bar{d} - \gamma\beta_{ew}\sigma_{ew}^2 + \alpha\varepsilon, \quad (\text{A.45})$$

where $\sigma_{ew}^2 = w_{ew}^\top \Sigma w_{ew}$ and $\beta_{ew} = \frac{\Sigma w_{ew}}{\sigma_{ew}^2}$ is the vector of market betas. Accordingly, the vector of equilibrium mean returns μ is

$$\mu = \mathbb{E}[R_{t+1}] = \mathbb{E}[d_{t+1} - P_t] = \bar{d} - P_t = \gamma\beta_{ew}\sigma_{ew}^2 - \alpha\varepsilon. \quad (\text{A.46})$$

Finally, using expression (A.46), we have that the 1/N mean return is $\mu_{ew} = w_{ew}^\top \mu = \gamma\sigma_{ew}^2 - \alpha\varepsilon^\top w_{ew}$, and dividing by the 1/N volatility σ_{ew} yields the 1/N Sharpe ratio in Equation (A.2), which completes the proof.

Proof of Proposition A.2

The vector of equilibrium mean returns μ in the economy in Assumption A.1 is $\mu = \gamma\beta_{ew}\sigma_{ew}^2 - \alpha\varepsilon$, as shown in Equation (A.46). Further, under Assumption 1, the Sharpe ratio of the arbitrage portfolio admits the decomposition in Equation (10). Plugging μ into this equation yields

$$\text{SR}(w_a) = \sqrt{\sum_{i=2}^N \frac{(v_i^\top (\gamma\beta_{ew}\sigma_{ew}^2 - \alpha\varepsilon))^2}{\sigma_{PCi}^2}}. \quad (\text{A.47})$$

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

Proof of Proposition A.3

Under the assumption that stock returns are i.i.d. normally distributed and $T > N + 3$, Kan, Wang, and Zhou (2021) derive a stochastic representation for the out-of-sample mean return and variance of the fully invested sample mean-variance portfolio, $\hat{w}_g + \hat{w}_a$, where $\hat{w}_g$ is the sample global-minimum-variance portfolio and $\hat{w}_a$ is the sample arbitrage portfolio in (13). Keeping only the part pertaining to $\hat{w}_a$, we find that the out-of-sample mean return and variance of $\hat{w}_a$ have the following stochastic representation:

$$\hat{w}_a^\top \mu = \frac{\text{SR}(w_a)\sqrt{T}}{\gamma} \left(\frac{z_1 + z_2\sqrt{\frac{1-b}{b}}}{u_2} \right), \quad (\text{A.48})$$

$$\hat{w}_a^\top \Sigma \hat{w}_a = \frac{T}{\gamma^2} \left(\frac{z_1^2 + z_2^2 + u_1}{u_2^2 b} \right), \quad (\text{A.49})$$

where $z_1 \sim \mathcal{N}(\text{SR}(w_a)\sqrt{T}, 1)$, $z_2 \sim \mathcal{N}(0, 1)$, $u_1 \sim \chi_{N-3}^2$, $u_2 \sim \chi_{T-N+1}^2$, and $b \sim \text{Beta}(\frac{T-N+2}{2}, \frac{N-1}{2})$. Therefore, the out-of-sample Sharpe ratio of $\hat{w}_a$ has the stochastic representation

$$\text{SR}(\hat{w}_a) = \frac{\hat{w}_a^\top \mu}{\sqrt{\hat{w}_a^\top \Sigma \hat{w}_a}} = \text{SR}(w_a) \frac{z_1\sqrt{b} + z_2\sqrt{1-b}}{\sqrt{z_1^2 + z_2^2 + u_1}}. \quad (\text{A.50})$$

Now, this representation is similar to that for the out-of-sample Sharpe ratio of the sample mean-variance portfolio in Kan, Wang, and Zheng (2022), for which they derive the expectation in their Lemma 1. Therefore, using a similar proof to theirs, we find that

$$\mathbb{E}[\text{SR}(\hat{w}_a)] = \text{SR}^2(w_a) \sqrt{\frac{T}{2}} \frac{\Gamma(\frac{N}{2})\Gamma(\frac{T}{2})\Gamma(\frac{T-N+3}{2})}{\Gamma(\frac{N+1}{2})\Gamma(\frac{T+1}{2})\Gamma(\frac{T-N+2}{2})} {}_1F_1\left(\frac{1}{2}; \frac{N+1}{2}; -\frac{\text{SR}^2(w_a)T}{2}\right), \quad (\text{A.51})$$

which yields Equation (A.5).

To conclude the proof, we need to prove that the loss in Equation (A.5) is a decreasing function of $\text{SR}^2(w_a)$. Proving this is equivalent to showing that the function

$$f(x) = x \times {}_1F_1\left(\frac{1}{2}; a + \frac{1}{2}; -x^2\right), \quad (\text{A.52})$$

where $a \geq 1$, increases with x for $x \geq 0$. It is easy to see that $f(0) = 0$. Taking the derivative of $f(x)$, we obtain

$$\begin{aligned} f'(x) &= {}_1F_1\left(\frac{1}{2}; a + \frac{1}{2}; -x^2\right) - \frac{x^2}{a + \frac{1}{2}} {}_1F_1\left(\frac{3}{2}; a + \frac{3}{2}; -x^2\right) \\ &= e^{-x^2} {}_1F_1\left(a; a + \frac{1}{2}; x^2\right) - \frac{x^2}{a + \frac{1}{2}} e^{-x^2} {}_1F_1\left(a; a + \frac{3}{2}; x^2\right) \\ &= e^{-x^2} \left[{}_1F_1\left(a; a + \frac{1}{2}; x^2\right) - \frac{x^2}{a + \frac{1}{2}} {}_1F_1\left(a; a + \frac{3}{2}; x^2\right) \right], \end{aligned} \quad (\text{A.53})$$

where the second equality follows by applying Kummer's transformation, ${}_1F_1(a; b; z) = e^z {}_1F_1(b - a; b; -z)$. To prove that $f'(x) > 0$ for $x \geq 0$ in (A.53), we need to show that

$${}_1F_1\left(a; a + \frac{1}{2}; y\right) > \frac{y}{a + \frac{1}{2}} {}_1F_1\left(a; a + \frac{3}{2}; y\right) \quad (\text{A.54})$$

for $y \geq 0$ and $a \geq 1$. Using the definition ${}_1F_1(a; b; y) = \sum_{k=0}^{\infty} \frac{(a)_k y^k}{(b)_k k!}$, where $(x)_k = x(x+1) \cdots (x+k-1)$ is the Pochhammer symbol, Inequality (A.54) indeed holds:

$$\begin{aligned} &{}_1F_1\left(a; a + \frac{1}{2}; y\right) - \frac{y}{a + \frac{1}{2}} {}_1F_1\left(a; a + \frac{3}{2}; y\right) \\ &= 1 + \sum_{k=0}^{\infty} \frac{(a)_{k+1} y^{k+1}}{(a + \frac{1}{2})_{k+1} (k+1)!} - \sum_{k=0}^{\infty} \frac{(a)_k y^{k+1}}{(a + \frac{1}{2})_{k+1} k!} \\ &> \sum_{k=0}^{\infty} \frac{(a)_{k+1} y^{k+1}}{(a + \frac{1}{2})_{k+1} (k+1)!} - \sum_{k=0}^{\infty} \frac{(a)_k y^{k+1}}{(a + \frac{1}{2})_{k+1} k!} \\ &= \sum_{k=0}^{\infty} \frac{(a)_k y^{k+1}}{(a + \frac{1}{2})_{k+1} k!} \left(\frac{a+k}{k+1} - 1 \right) \geq 0 \end{aligned} \quad (\text{A.55})$$

for $a \geq 1$, which completes the proof.

References in Internet Appendix

- Abdi, F., and A. Rinaldo. 2017. A simple estimation of bid-ask spreads from daily close, high, and low prices. *The Review of Financial Studies* 30:4437–80.
- Baker, M., and J. Stein. 2004. Market liquidity as a sentiment indicator. *Journal of Financial Markets* 7:271–99.
- Hong, H., and D. Sraer. 2016. Speculative betas. *The Journal of Finance* 71:2095–144.
- Huang, D., F. Jiang, J. Tu, and G. Zhou. 2015. Investor sentiment aligned: A powerful predictor of stock returns. *The Review of Financial Studies* 28:791–837.
- Kan, R., X. Wang, and X. Zheng. 2022. In-sample and out-of-sample Sharpe ratios of multi-factor asset pricing models. *SSRN working paper*.
- Kan, R., X. Wang, and G. Zhou. 2021. Optimal portfolio choice with estimation risk: No risk-free asset case. *Management Science* 68:2047–68.
- Kan, R., and G. Zhou. 2007. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42:621–56.
- Koijen, R., and M. Yogo. 2019. A demand system approach to asset pricing. *Journal of Political Economy* 127:1475–515.
- Ledoit, O., and M. Wolf. 2004. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88:365–411.
- Miller, E. 1977. Risk, uncertainty, and divergence of opinion. *The Journal of Finance* 32:1151–68.
- Tu, J., and G. Zhou. 2011. Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies. *Journal of Financial Economics* 99:204–15.