

NBER WORKING PAPER SERIES

THE STATISTICAL LIMIT OF ARBITRAGE

Rui Da
Stefan Nagel
Dacheng Xiu

Working Paper 33070
<http://www.nber.org/papers/w33070>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2024

We are grateful for comments from Adem Atmaz, Lars Hansen, Ye Luo, Andreas Neuhierl, Markus Pelger, Seth Pruitt, Yao Zeng, and seminar and conference participants at Columbia University, Indiana University, Louisiana State University, Princeton University, ITAM Business School, University of Cincinnati, University of Chicago, Yale School of Management, EPFL, ETH Zurich, Stockholm Business School, University of Gothenburg, University of Liverpool, University of Oxford, City University of Hong Kong, Gregory Chow Seminar Series in Econometrics and Statistics, HKUST, KAIST, Peking University, Shanghai University of Finance and Economics, Southern University of Science and Technology, Tsinghua University, University of Melbourne, Sao Paulo School of Economics, American Finance Association Annual Meetings, NBER Summer Institute, Jacobs Levy Center Frontiers in Quantitative Finance Conference, Stanford Institute for Theoretical Economics, SFS Cavalcade North America, Annual SoFiE Conference, Global AI Finance Research Conference in Singapore, China International Conference in Finance, and Wabash River Finance Conference. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w33070>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Rui Da, Stefan Nagel, and Dacheng Xiu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Statistical Limit of Arbitrage
Rui Da, Stefan Nagel, and Dacheng Xiu
NBER Working Paper No. 33070
October 2024
JEL No. C55, C58, G11, G14

ABSTRACT

We investigate the economic consequences of statistical learning for arbitrage pricing in a high-dimensional setting. Arbitrageurs learn about alphas from historical data. When alphas are weak and rare, estimation errors hinder arbitrageurs—even those employing optimal machine learning techniques—from fully exploiting all true pricing errors. This statistical limit to arbitrage widens the equilibrium bounds of alphas beyond what traditional arbitrage pricing theory predicts, leading to a significant divergence between the feasible Sharpe ratio achievable by arbitrageurs and the unattainable theoretical maximum under perfect knowledge of alphas.

Rui Da
Kelley School of Business
Indiana University
HH 6100
1309 E. 10th St.
Bloomington, IN 47405
ruida@iu.edu

Dacheng Xiu
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
dacheng.xiu@chicagobooth.edu

Stefan Nagel
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
stefan.nagel@chicagobooth.edu

A supplemental appendix is available at <http://www.nber.org/data-appendix/w33070>

1 Introduction

Full-information rational expectations models assume that economic agents have perfect knowledge of the functional form and parameters of the data-generating process (DGP), thereby bypassing the statistical challenges real-world agents encounter when estimating the law of motion. While this assumption might be reasonable in a low-dimensional environment with ample data for learning, many economic settings are more realistically high-dimensional, with an unknown DGP and limited data. This raises important questions: How should agents optimally learn in such environments, and what are the implications for equilibrium outcomes? In this paper, we explore these questions in the context of asset market arbitrage.

The absence of near-arbitrage opportunities is a fundamental principle in most asset pricing theories, including the Arbitrage Pricing Theory (APT). These theories implicitly assume that if investment opportunities with exceptionally high reward-to-risk ratios were to exist, they would attract arbitrageurs. Armed with knowledge of the functional form and the parameters governing the DGP of returns, these arbitrageurs would exploit and, in turn, eliminate such opportunities. In reality, however, sophisticated investors seeking near-arbitrage opportunities lack perfect knowledge of the true functional form and parameters. Instead, they rely on statistical analysis to infer the existence of such opportunities from historical data, introducing statistical uncertainty. In some cases, such as derivatives pricing, this uncertainty may be small enough that it does not significantly hinder arbitrage activity. However, in noisier, high-dimensional settings like the cross-section of stock returns, statistical uncertainty can be substantial, imposing a meaningful limit on arbitrage.

To analyze the effects of arbitrageur learning, we consider a setting in which returns follow a statistical linear factor model. Near-arbitrage opportunities take the form of factor-neutral trading strategies with high Sharpe ratios. Exploiting these opportunities requires arbitrageurs to have knowledge of factor model alphas, which they must learn from historical return data. We derive the optimal Sharpe ratio achievable by any feasible arbitrage trading strategy, which is strictly dominated by the infeasible optimal Sharpe ratio that arbitrageurs could achieve if they were endowed with perfect knowledge of alphas. This, in turn, provides a new no-feasible-near-arbitrage bound on the Sharpe ratio that accounts for the statistical limit to arbitrage.

The difficulty of the learning problem depends on the DGP of alpha signals. While our theory does not generally rely on specific cross-sectional distributions of alpha signals, we use simple special cases to illustrate how the optimal Sharpe ratio varies with the strength and sparsity of alphas. When alphas are strong and not too rare relative to the dimensionality of the cross-section and the sample size, arbitrageurs can, in the limit, perfectly learn the distri-

bution of alpha. However, when alphas are weaker and more rare, inferring their distribution becomes more challenging, creating a gap between the optimal feasible Sharpe ratio and the infeasible Sharpe ratio that requires perfect knowledge of alphas. For instance, the infeasible Sharpe ratio may explode asymptotically, while the feasible Sharpe ratio remains bounded.

The existence of this statistical limit to arbitrage implies that the range within which mispricing can survive in equilibrium is broader compared to a scenario where arbitrageurs have full knowledge of the DGP and its parameters. Mispricing can survive because it is obscured by statistical uncertainty. Empirically, it is the feasible Sharpe ratio—not the unattainable, infeasible one—that reveals the minimum reward-to-risk compensation arbitrageurs require.

We further demonstrate how arbitrageurs can construct a feasible trading strategy that achieves the theoretically optimal feasible Sharpe ratio, uniformly over DGPs of alphas, regardless of the strength and sparsity of alphas. This means that the feasible Sharpe ratio bound is in fact sharp. A uniformly valid trading strategy is desirable because, in practice, arbitrageurs do not know which DGP accurately represents the observed data. The optimal strategy involves estimating the empirical distribution of alpha signals and assigning portfolio weights based on the relative magnitudes and associated uncertainty of the alpha estimates. Assets with high alpha t -statistics receive portfolio weights proportional to their signal strength. Although weaker alphas are more difficult to exploit, ignoring them would lead to a suboptimal strategy. Instead, the optimal approach constructs portfolio weights for weak signals by locally smoothing alpha signals cross-sectionally.

To empirically contrast feasible and infeasible Sharpe ratios, we propose an estimator of the infeasible Sharpe ratio that a hypothetical arbitrageur endowed with perfect knowledge of DGP parameters would perceive. While this Sharpe ratio can be estimated consistently, it cannot be achieved by any feasible portfolio with weights constructed using historical data. The infeasible Sharpe ratio often serves as the foundation for tests of the APT, see, e.g., [Gibbons et al. \(1989\)](#), [Gagliardini et al. \(2016\)](#), [Fan et al. \(2015\)](#), and [Pesaran and Yamagata \(2017\)](#). While these tests are powerful and can lead to the discovery of alpha signals, they are not relevant for arbitrageurs limited to feasible trading strategies. Our effort in constructing the optimal feasible arbitrage portfolio and evaluating its economic performance directly responds to Shanken’s call to empirically assess the APT by “characterizing the investment opportunities that are available as a consequence of the observed expected return deviation” and “examining the extent to which we can approximate an arbitrage with existing assets” ([Shanken, 1992](#)). Here we do so taking into account the statistical limits to arbitrage.

Under certain conditions, the optimal feasible arbitrage strategy can be approximated using alternative approaches such as multiple testing, shrinkage, and selection applied to alphas estimated from cross-sectional regressions. These alternative approaches are used widely by

empirical asset pricing researchers and practitioners, making it valuable to understand how closely they approximate the optimal feasible strategy. A multiple-testing procedure as in [Benjamini and Hochberg \(1995\)](#) (BH) seeks to prevent false alpha discoveries by applying a p -value threshold that controls the false discovery rate (FDR), setting alphas that do not pass this threshold to zero. The BH procedure achieves optimal performance only when a small number of true alpha signals are strong; otherwise, it tends to be overly conservative, falling short of optimality. Applying ridge-style shrinkage to portfolio weights yields a portfolio that is equivalent to one constructed with alpha estimates from plain cross-sectional regressions without shrinkage. This approach can achieve optimality when almost all true alphas are either uniformly strong or uniformly weak. Lasso-based selection aims to balance the strengths of the two previous methods, yielding a portfolio that closely approaches the theoretically optimal feasible Sharpe ratio, provided an optimal tuning parameter is chosen.

Finally, we demonstrate the empirical implications of the statistical limits of arbitrage by examining monthly equity returns in US stock market from 1965 to 2020. Our empirical analysis has two parts. In the first part, we examine individual stock returns. We construct a multi-factor model that uses observed stock characteristics as risk exposures. The characteristics include 16 commonly used return predictors and 11 industry indicators. These characteristics and industry dummies capture similar equity factors in the MSCI Barra model widely-used among practitioners. In the second part, we analyze returns on 1,273 characteristics-sorted portfolios and 49 industry portfolios, for which we construct latent factor models via singular value decomposition, following [Giglio and Xiu \(2021\)](#). Data on characteristics and portfolio returns in our empirical analysis is from [Chen and Zimmermann \(2020\)](#).¹

For individual stocks, we estimate alphas as the averaged residuals in cross-sectional regressions of stock returns on characteristics in 60-month rolling windows. The cross-sectional R^2 s are relatively low, averaging around 8% over our sample period. This suggests that learning alphas from realized returns is difficult, as the alphas are obscured by large amounts of noise. Accordingly, only 7.58% and 1.12% of individual stocks' alpha estimates have t -statistics greater than 2.0 and 3.0, respectively, in absolute values. Even without accounting for multiple testing, these estimates indicate that non-zero alphas are both rare and weak.

One may be concerned that these results may overstate the difficulty of learning alphas because they use only realized returns as the only signal of alpha. Our analysis of portfolio returns addresses these concerns. The large number of characteristics used to construct these portfolios can be interpreted as signals of alphas, and the portfolios as managed investment strategies that exploit the information in these signals. Although a latent factor model accounts for a larger portion (around 35%) of the cross-sectional variation in realized portfolio

¹See openassetpricing.com (Stock-level Signal Datasets August 2023 Release).

returns than the factor model for individual stock returns, we find that portfolio alphas still exhibit notable rarity and weakness. This is particularly evident once we adjust for publication years, ensuring that portfolios are considered only after their sorting characteristics were published in the previous years.

We then find, across both individual stocks and portfolios and using various methods, that the optimal feasible arbitrage portfolios yield moderately low annualized Sharpe ratios below 0.7. In contrast, the infeasible Sharpe ratios are considerably higher, averaging more than 4.8 and reaching as high as 16 in some sample periods for individual stocks, and ranging from 5 to 20 for portfolios. These infeasible Sharpe ratios represent what arbitrageurs could achieve with perfect knowledge of the DGP parameters. The large gap between feasible and infeasible Sharpe ratios highlights the empirical significance of the statistical limits to arbitrage. When these statistical limits to arbitrage are taken into account, the empirical facts are in line with the implication of the APT that near-arbitrage opportunities should be absent.

Our paper builds on a large literature on the arbitrage pricing theory (APT) developed by [Ross \(1976\)](#) and later refined by [Huberman \(1982\)](#), [Chamberlain and Rothschild \(1983\)](#), and [Ingersoll \(1984\)](#). As in these foundational studies, we employ asymptotic arguments that avoid imposing strong structural assumptions about the economy. The results of this analysis should be interpreted as an asymptotic approximation for a more realistic setting with a finite number of assets. The key point is that weak economic restrictions rule out Sharpe ratios far above the Sharpe ratios of diversified factor portfolios. Different from earlier work on the APT, we show in this paper that the statistical limit to arbitrage dramatically widens this Sharpe ratio bound compared with an economy in which arbitrageurs are endowed with perfect knowledge of DGP parameters. In this regard, our paper is also related to the literature on the limits of arbitrage reviewed in [Gromb and Vayanos \(2010\)](#). Complementary to this literature, the arbitrage limit in our setting stems from statistical uncertainty, rather than from risk, costs, frictions, and other constraints faced by arbitrageurs.

Our paper provides a solution to a long standing problem in optimal portfolio choice under parameter uncertainty. Historically, the plug-in mean-variance portfolio, which is based on sample means and covariance matrices, has been criticized for underperforming. Assuming normally distributed returns, [Kan and Zhou \(2007\)](#) study the expected performance of this portfolio and observe that its Sharpe ratio is below its theoretical value. Further studies by [Tu and Zhou \(2010\)](#) and [Kan et al. \(2022\)](#) expand their analysis by exploring alternative portfolio rules that aim to minimize utility loss under estimation uncertainty. However, these studies only identify optimal solutions within a predetermined class of strategies, which does not reflect the flexibility investors have in practice to choose any perceived superior strategy. In contrast, our study identifies an upper bound for all feasible Sharpe ratios without imposing

restrictions on the choice of strategies. Our analytical framework is primarily concerned with arbitrage portfolios, but our results are broadly applicable to optimal portfolios in general, thus providing a solution to this classic problem.

A closely related paper to ours is [Kim et al. \(2020\)](#), which proposes a characteristics-based factor model to construct feasible arbitrage portfolios. Their asymptotic theory does not preclude arbitrage opportunities with a theoretically infinite Sharpe ratio, which implies a rather strong signal-to-noise ratio in their alpha signals. Our setting is considerably different in that the premise of our framework rules out infinite feasible Sharpe ratios, which enforces weak and rare signals. In our setting, alphas cannot possibly be recovered with certainty even when the sample size is large.²

Our paper also contributes to the evolving literature on applications of statistical and machine learning in asset pricing, and in particular on the topic of testing the APT, e.g., [Gibbons et al. \(1989\)](#), [Gagliardini et al. \(2016\)](#), and [Fan et al. \(2015\)](#), as well as on testing for alphas, e.g., [Barras et al. \(2010\)](#), [Harvey and Liu \(2020\)](#), and [Giglio et al. \(2021\)](#). The first literature focuses on testing a null that all alphas are equal to zero. This is certainly an interesting null hypothesis. However, as we emphasize in this paper, an economically sensible interpretation of the APT should allow for statistical limits to arbitrage. This means that the APT does allow for non-zero alphas as long as they do not induce an explosive feasible Sharpe ratio. The second literature focuses on detecting strong alphas, applying multiple testing methods, such as the BH method, or extensions thereof, to control the FDR. In contrast, we allow for rare and weak alpha signals such that any procedure aiming to control the FDR is too conservative with too few or no discoveries.³ Our objective here is not on model testing or signal detection. Rather, we strive for the optimal economic performance of arbitrage portfolios. We show that even if signals were so weak that they are undetectable by multiple testing methods, they may lead to a portfolio with a considerable Sharpe ratio.

In macroeconomics and finance, rational expectation models in which decision makers are not confronted with statistical uncertainty about the structure of the economy have attracted criticism ([Hansen, 2007](#)). Bayesian learning is one way to expose model decision makers to statistical uncertainty. [Pastor and Veronesi \(2009\)](#) survey the literature on learning in financial markets. In many settings, e.g., [Collin-Dufresne et al. \(2016\)](#), learning can be sufficiently slow such that its effects persist in empirically realistic sample sizes, even though

²On the empirical side, [Guijarro-Ordóñez et al. \(2022\)](#) propose a deep learning approach to statistical arbitrage that achieves a sizable out-of-sample Sharpe ratio. The profits of their trading strategy stem from generalized return reversals at daily to weekly frequencies, potentially due to liquidity provision and other microstructure channels. Our empirical analysis is not targeted towards characterizing the reward-to-risk ratios for high frequency traders, nor for traders that turnover a large portion of their portfolios daily.

³[Donoho and Jin \(2004\)](#) adopt the so-called higher criticism approach, dating back to [Tukey \(1976\)](#), to detect rare and weak signals in a stylized multiple testing problem.

convergence to rational expectations takes place in the long-run. An exception is [Martin and Nagel \(2021\)](#) where learning effects persist because investors face a high-dimensional inference problem about the process generating firm cash flows. Similarly, arbitrageurs in our model attempt to make inference on a high-dimensional parameter vector with a potentially insufficient sample size, but they learn about returns, not firms' underlying cash flows. We examine different sequences of DGPs and in most scenarios, our learning system does not converge to a rational expectations limit.⁴

Our paper is also related to [Chen et al. \(2021b\)](#) and [Chen et al. \(2021a\)](#) in that they also account for the distinction between beliefs of economic agents and the DGP revealed by empirical evidence. They model belief distortions as a change of measure in moment conditions, use statistical measures of divergence relative to rational expectations to bound the set of subjective probabilities, and seek robust inference with respect to this form of misspecification. In the spirit of [Hansen \(2014\)](#), we calculate an optimal feasible Sharpe ratio for arbitrageurs inside the economic model, which differs from the infeasible one as it appears to an econometrician outside the model. In our setting, the deviation from rational expectations stems naturally from the statistical obstacles faced by economic agents. A more subtle but important point is that the outcomes attainable by economic agents embracing machine learning methods in a high-dimensional environment differ from those that rational expectation agents endowed with knowledge of the economy's structure could achieve asymptotically.

The literature on portfolio choice through Bayesian and economic decision-theoretical perspectives is extensive, as outlined in the survey by [Avramov and Zhou \(2010\)](#). Early works like those by [Jorion \(1986\)](#) and [Frost and Savarino \(1986\)](#) highlight the advantages of using Bayes-Stein shrinkage solutions. Further contributions in this topic incorporate informative and economically motivated priors, as in [Black and Litterman \(1992\)](#), [Pastor \(2000\)](#), and [Pastor and Stambaugh \(2000\)](#). In our context, the optimal portfolio weights are proportional to the posterior mean of alpha. This resembles the classical normal mean problem in empirical Bayes, dating back to [Robbins \(1956\)](#), where the unknown parameters (alphas) are regarded as random draws from some common distribution, and only a noisy version (realized ex-factor returns) is observed. Our nonparametric approach thereby shares the same spirit of nonparametric empirical Bayes, see, e.g., [Johns \(1957\)](#), [Zhang \(1997\)](#), and [Brown and Greenshtein \(2009\)](#). Yet unlike the classical empirical Bayes inference, our analysis allows for weak and rare alphas as motivated from economic restrictions, and also considers Sharpe ratios, not only the posterior mean of alphas.

⁴Our analysis is related to a large literature in econometrics and statistics that discuss uniform validity of asymptotic approximations, see, e.g., [Staiger and Stock \(1997\)](#), [Imbens and Manski \(2004\)](#), [Leeb and Pötscher \(2005\)](#), [Andrews et al. \(2020\)](#).

Our paper proceeds as follows. Section 2 develops our main result on the statistical limit to arbitrage. Specifically, Section 2.1 sets up the model, Section 2.2 motivates and then defines the feasibility constraint facing arbitrageurs, Sections 2.3 - 2.4 specify arbitrageurs' decision problem, derive the optimal strategy, illustrate the Bayes correction for alpha, and demonstrate the gap between feasible and infeasible Sharpe ratios, Section 2.5 constructs a feasible trading strategy that achieves the bound, Section 2.6 proposes an estimator of the infeasible Sharpe ratio, and finally Section 2.7 analyzes alternative strategies. Section 3 provides simulation evidence, followed by an empirical analysis in Section 4. Section 5 concludes. The appendix provides additional empirical and theoretical results, as well as all the technical details.

2 Main Theoretical Results

We start by revisiting the arbitrage pricing theory framework developed by Ross (1976). This theory is primarily based on a reduced-form statistical model for asset returns, which, despite its stylized nature, offers significant theoretical insights and remains relevant for guiding empirical investment decisions.

2.1 Factor Model Setup

The factor economy has N assets in the investment universe. The $N \times 1$ vector of excess returns r_t follows a reduced-form linear factor model:

$$r_t = \alpha + \beta\gamma + \beta v_t + u_t, \quad (1)$$

where β is an $N \times K$ matrix of factor exposures (with the first column being a vector of 1s), α is an $N \times 1$ vector of pricing errors, v_t is a $K \times 1$ vector of zero-mean factor innovations with covariance matrix Σ_v , γ is a $K \times 1$ vector of risk premia (first entry corresponding to the column of 1s is the zero beta rate), and u_t is a vector of idiosyncratic returns with a diagonal covariance matrix Σ_u and zero mean conditional on all the variables before time- t .⁵

We assume at any given time t , arbitrageurs examine a sample of size T , derived from Equation (1), spanning from $t - T + 1$ to t . Throughout we will consider asymptotic limits as N and T increase while K and t are fixed.⁶ To facilitate our asymptotic analysis along

⁵While approximate factor models become more prevalent following Chamberlain and Rothschild (1983), allowing for off-diagonal entries in the covariance matrix Σ_u would introduce additional statistical obstacles due to the estimation of large covariance matrix for inference on alpha and for building optimal portfolios. For simplicity, we illustrate the economic consequences of statistical limits to arbitrage using a strict factor model. We discuss violations of these model assumptions later.

⁶ This framework presents a slight deviation from the conventional scenario in which arbitrageurs observe

the cross-sectional dimension, N , we regard high dimensional objects such as α , β , and Σ_u as random variables drawn from some cross-sectional distributions, whereas γ and Σ_v are regarded as deterministic parameters. This distinction is made because γ and Σ_v serve as nuisance parameters in our analysis and their dimensions remain fixed. We assume that α has mean zero, and is cross-sectionally independent of β . These conditions are essential for identification of γ in a model that allows for pricing errors. We formalize these conditions below.

To allow for a more general DGP, we may consider a conditional version of (1):

$$r_t = \alpha_{t-1} + \beta_{t-1}\gamma_{t-1} + \beta_{t-1}v_t + u_t, \quad (2)$$

where β_t is a vector of time-varying factor loadings and γ_t is a vector of time-varying risk premia.⁷ Despite the conditional model's greater generality, it does not provide additional economic insights in our analysis of the theoretical limits of arbitrage, compared with the unconditional model. While our primary theoretical findings, such as those provided in Theorems 1 and 2, remain applicable when β is substituted with β_t , we concentrate, for clarity, on the more stylized model (1) to illustrate our theory.

There are at least three variations of the factor model, depending on what econometricians assume to be observable. The most common setup in the academic finance literature imposes that factors are observable as, e.g., in Fama and French (1993).⁸ The second setting, which has gained more popularity recently since its debut in Connor and Korajczyk (1986), assumes that factors are latent. The third setting, arguably most prevalent among practitioners, is the MSCI Barra model originally proposed by Rosenberg (1974), where factor exposures are equal to stock characteristics and hence observable.⁹ One notable advantage of this model is that it sidesteps the statistically and computationally demanding task of estimating a large number of potentially time-varying stock-level factor exposures. In our empirical work, we prefer to use the MSCI Barra model for its practicality in analyzing individual stock returns,

a sample spanning $t = 1, \dots, T$ and make their investment decision at time $T + 1$. Essentially, our approach highlights the continuous process of making investment decisions, rather than conceiving these decisions as occurring at a single, distant future point, T . That said, this conceptual difference does not result in any tangible difference in our theoretical and empirical results.

⁷This model is overly parametrized that parameters are not identifiable without additional restrictions. Some examples of parsimonious conditional factor models include Connor et al. (2012), Gagliardini et al. (2016), and Kelly et al. (2019).

⁸This is different from saying factor innovations, v_t , are observable. The setting of observable factors typically involves another equation that $f_t = \mu + v_t$, where μ are the population means of the observed factors f_t , which are not necessarily identical to the factor risk premia, γ . Since μ is an unknown parameter, v_t is still not observable.

⁹See Kozak and Nagel (2023) on the conditions on factor construction and return covariance matrix that must hold for characteristics to equal factor loadings.

adhering to (2) but with a constant alpha.¹⁰

The constant alpha assumption in our theoretical analysis may seem restrictive. In practice, arbitrageurs may make use of trading signals that capture time-variation in individual stock's alphas. Our framework can in fact easily accommodate this if we assume that (2) holds with $\alpha_t = c_t \alpha$ and $\beta_t = c_t \beta$ for certain characteristics c_t that represent arbitrageurs' signals, then projecting returns onto lagged characteristics results in:

$$(c_{t-1}^\top c_{t-1})^{-1} c_{t-1}^\top r_t = \alpha + \beta(\gamma_{t-1} + v_t) + (c_{t-1}^\top c_{t-1})^{-1} c_{t-1}^\top u_t. \quad (3)$$

This transformation converts a conditional model for individual stocks with time-varying alphas into an unconditional latent factor model for managed portfolios, $(c_{t-1}^\top c_{t-1})^{-1} c_{t-1}^\top r_t$, with constant alphas as in (1). Thus, via these managed portfolios, our theoretical framework also applies to a setting in which arbitrageurs make use of signals about time-varying individual stock alphas. Empirically, we therefore also examine returns on portfolios formed on a large number of stock characteristics, in addition to the individual stock analysis.

2.2 Feasible Near-Arbitrage Opportunities

Building upon the insight of Ross (1976), Huberman (1982) and Ingersoll (1984) established the concept of near-arbitrage, which can be formalized in a more general setting as below:

Definition 1. A portfolio strategy w at time t is said to generate a near-arbitrage under a sequence of data-generating processes, such as (1), defined in a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_s\}_{s \leq t}, P)$, where $\mathcal{F}_t$ contains information generated by $\alpha, \beta, \{r_s, v_s, u_s\}_{s \leq t}$, if it satisfies $w \in \mathcal{F}_t$, and along some diverging subsequence,¹¹ with probability approaching one,

$$\text{Var}(w^\top r_{t+1} | \mathcal{F}_t) \rightarrow 0, \quad E(w^\top r_{t+1} | \mathcal{F}_t) \geq \delta > 0.$$

Intuitively, no near-arbitrage means there exists no sequence of portfolios that earn positive expected returns with vanishing risks. Ingersoll (1984) established that a sufficient and necessary condition for the absence of near-arbitrage is that¹²

$$S^* := \sqrt{\alpha^\top \Sigma_u^{-1} \alpha} \lesssim_P 1. \quad (4)$$

¹⁰Empirically, we use a moving window method for alpha estimation, which effectively allows alphas to vary slowly over time.

¹¹We adopt the same subsequence definition as that used in Ingersoll (1984). The subsequence typically depends on the count of investment opportunities, i.e., N , though we do not need make this explicit in this definition. For simplicity of notation and without ambiguity, we omit the dependence of w on N and t .

¹²We use the notation $a \lesssim b$ to denote $a \leq Cb$ for some constant $C > 0$, $a \lesssim_P b$ to denote $a = O_P(b)$, $a \approx b$ if $a \lesssim b$ and $b \lesssim a$, and use $a \approx_P b$ accordingly. Throughout, c_N represent some sequence vanishing with N .

Here, S^* is the theoretically optimal Sharpe ratio arbitrageurs can achieve in this economy using a portfolio strategy that has zero exposure to factor risks, namely, a “statistical arbitrage” strategy in the jargon of practitioners. This result suggests that moderate mispricing in the form of nonzero alphas is permitted in an economy without near-arbitrage opportunities, but there cannot be too many alphas that are too large, to the extent that S^* explodes.¹³

To achieve this optimal Sharpe ratio, arbitrageurs should hold a factor-neutral portfolio with weights given by $w^* \propto \Sigma_u^{-1}\alpha$, according to Ingersoll (1984).¹⁴ His analysis assumes that arbitrageurs know the true (population) parameters: α and Σ_u . In reality, however, the true parameters are unknown to arbitrageurs as they only have a finite sample of data to learn about these parameters. The consequences of such parameter uncertainty can be minor in some contexts if the true parameters are revealed asymptotically and model predictions converge to those of a rational expectations model in which agents are endowed with perfect knowledge of parameters. Fundamentally, this requires that the learning problem in the limiting experiment becomes increasingly simpler as the sample size increases.

In the current context, the difficulty of the learning problem also hinges on the number of investment opportunities, N . As N increases, for a given sample size T , it becomes increasingly difficult for arbitrageurs to determine which assets truly have nonzero alphas. If the learning problem remains difficult as N and T increase, the learning effect persists, which could lead to limiting implications that differ from the rational expectation case. In our setting, it turns out that the rational expectation limit S^* is only relevant for rather restrictive scenarios. In more realistic settings, e.g., N is much larger than T , the optimal Sharpe ratio arbitrageurs can achieve without factor exposures is far smaller than S^* because of their inability to make error-free inference. Therefore, the condition (4) could be excessively restrictive in such scenarios.

To illustrate this intuition, we consider a simple and specific example.

Example 1. *Suppose the cross-section of alphas is drawn from the following distribution:*

¹³Assuming α_i is *i.i.d.* and $\lambda_{\max}(\Sigma_u) \lesssim_P 1$, by equation (4), we have $\alpha^\top \alpha \lesssim_P \alpha^\top \lambda_{\min}(\Sigma_u^{-1}) \alpha \lesssim_P \alpha^\top \Sigma_u^{-1} \alpha \lesssim_P 1$, so that $E(\alpha_i^2) = o(1)$ by the law of large numbers.

¹⁴In Ingersoll (1984), α is defined to be the cross-sectional projection of the expected returns onto β in the population model, such that $\alpha^\top \Sigma_u^{-1} \beta = 0$. As a result, his arbitrage portfolio weights are factor neutral, i.e., $w^* \beta = 0$ by construction. In contrast, our paper sets forth a predetermined DGP as specified in (1), where α is modeled as a random variable with the property $E(\alpha^\top \beta) = 0$. Consequently, the factor-neutral optimal strategy derived in Theorem 1 and illustrated by (10) does not exactly align with the formula presented by Ingersoll (1984). Under our DGP assumptions, Ingersoll’s portfolio weights are not strictly factor-neutral. Nonetheless, we can show that our optimal portfolio weights remain factor-neutral, and achieve the same Sharpe ratio S^* asymptotically as N increases.

$$\alpha_i \stackrel{i.i.d.}{\sim} \begin{cases} \mu & \text{with prob. } \rho/2 \\ -\mu & \text{with prob. } \rho/2 \\ 0 & \text{with prob. } 1 - \rho \end{cases}, \quad 1 \leq i \leq N, \quad (5)$$

where $\mu \geq 0$ and $0 \leq \rho \leq 1$, and they potentially vary with N and T . In addition, we also assume $\beta = 0$, $\Sigma_u = \sigma^2 \mathbb{I}_N$, for some $\sigma > 0$. Consequently, equation (1) becomes $r_t = \alpha + u_t$.

In this example, μ dictates the strength of alphas, ρ describes how rare alphas are, whereas σ is a nuisance parameter. By modeling parameters μ and ρ as functions of the sample size and dimension of the investment set, we introduce greater flexibility in depicting the challenges arbitrageurs encounter in finite sample situations.¹⁵ To emphasize the role of signal strength and count, we impose in this example that all assets share the same alpha distribution and the same idiosyncratic variance.

Now suppose, more specifically, that the magnitude of (μ, ρ) satisfies

$$\mu \approx T^{-1/2} \quad \text{and} \quad \rho \approx N^{-1/2}. \quad (6)$$

This condition (6) implies that the signal strength μ vanishes as the sample size increases ($T \rightarrow \infty$) and the signal percentage count ρ decays as the investment universe expands ($N \rightarrow \infty$). This setup is used to approximate a situation in which only a small portion of assets have a nonzero and small alpha. On the other hand, σ is assumed to be fixed, since in reality idiosyncratic risks never vanish, whereas alphas can be small due to competition among arbitrageurs. This model resides in an uncommon territory in the existing literature of asset pricing: weak and rare alphas. In fact, the classical no-near-arbitrage condition (4) imposes, implicitly, weakness or rareness on alphas; otherwise, if alphas were strong and dense, $\alpha^\top \alpha$ would explode rather rapidly. Even in the current setting, in light of the fact that $E(\alpha^\top \alpha) = \rho \mu^2 N$, we still have $\alpha^\top \alpha \xrightarrow{P} \infty$ as long as $N^{1/2}/T \rightarrow \infty$. In other words, a near-arbitrage opportunity arises according to (4), with a strategy $w = \sigma^{-2} \alpha$.

However, statistical uncertainty prevents arbitrageurs from having this “free lunch.” In general, it is only possible to recover any element of alpha up to some estimation error of magnitude $T^{-1/2}$.¹⁶ Since, by design, the true alpha is of the same order of magnitude as its level of statistical uncertainty, i.e., $\mu \approx T^{-1/2}$, it is impossible for arbitrageurs to determine

¹⁵Adopting a drifting sequence for parameters is a common trick in econometrics to provide more accurate finite sample approximations. As [Bekker \(1994\)](#) put, “in evaluating the results, it is important to keep in mind that the parameter sequence is designed to make the asymptotic distribution fit the finite sample distribution better. It is completely irrelevant whether or not further sampling will lead to samples conforming to this sequence or not.”

¹⁶[Giglio et al. \(2021\)](#) develop the asymptotic normality result for alpha estimates via a Fama-MacBeth procedure in various scenarios, in which factors are (partially) observable or latent whereas β is unknown. The CLTs in these scenarios share the same form: for any $1 \leq i \leq N$,

precisely which assets among all have nonzero alpha.

For illustration purpose, suppose that arbitrageurs adopt the strategy $\hat{w} = \sigma^{-2}\hat{\alpha}$,¹⁷ replacing α in w with $\hat{\alpha} = \bar{r} = \alpha + \bar{u}$.¹⁸ Out of sample, this portfolio's conditional expected return and conditional variance can be written as:

$$\begin{aligned} \mathbb{E}(\sigma^{-2}(\alpha + \bar{u})^\top(\alpha + u_{t+1})|\mathcal{F}_t) &= \sigma^{-2}(\alpha^\top\alpha + \bar{u}^\top\alpha), \\ \text{Var}(\sigma^{-2}(\alpha + \bar{u})^\top(\alpha + u_{t+1})|\mathcal{F}_t) &= \sigma^{-2}(\alpha^\top\alpha + 2\alpha^\top\bar{u} + \bar{u}^\top\bar{u}), \end{aligned}$$

where u_{t+1} denotes a future de-meaned return at $t + 1$, that shares the same distribution as $\{u_s\}_{s \leq t}$, but is independent of $\bar{u}$ which belongs to the information set up to t , $\mathcal{F}_t$. The resulting squared conditional Sharpe ratio is given by:

$$S^2 = \frac{\sigma^{-4}(\alpha^\top\alpha + \bar{u}^\top\alpha)^2}{\sigma^{-2}(\alpha^\top\alpha + 2\alpha^\top\bar{u} + \bar{u}^\top\bar{u})} \lesssim_P T^{-1} \rightarrow 0, \quad (8)$$

where we use the fact that $\bar{u}^\top\bar{u} \asymp_P N/T$ when u_t is i.i.d.. In other words, this portfolio achieves a Sharpe ratio equal to zero asymptotically.

Is there a superior trading strategy capable of maintaining a non-vanishing Sharpe ratio? The straightforward answer is no. Our discussion below will elucidate that, within this context, namely, (6) holds true, the highest achievable Sharpe ratio for all *feasible* trading strategies employed by arbitrageurs, represented as S^{OPT} , vanishes asymptotically as $N, T \rightarrow \infty$. Conversely, the *infeasible* optimal Sharpe ratio, denoted S^* , diverges under the condition that $N^{1/2}/T \rightarrow \infty$. There is a vast disparity, as shown by this example, between S^{OPT} and S^* .

The difference between feasible and infeasible strategies is primarily driven by the information set accessible to arbitrageurs when they implement their trading strategies. In this example, the infeasible strategy assumes that arbitrageurs have access to a comprehensive information set, $\mathcal{F}_t$, that encompasses the knowledge of α . This knowledge proves to be extremely valuable when α is both rare and weak, which creates significant gap between this strategy and others that lack access to such information.

To clarify, we formalize the definitions as follows:

$$\sqrt{T}(\hat{\alpha}_i - \alpha_i) \xrightarrow{d} \mathcal{N}(0, \sigma_i^2(1 + \gamma^\top(\Sigma_v)^{-1}\gamma)), \quad (7)$$

where σ_i^2 is the i th entry of Σ_u . In the case that β is observable (but factors are not), we can show that the CLT has a similar form except that the scalar $(1 + \gamma^\top(\Sigma_v)^{-1}\gamma)$ disappears.

¹⁷The knowledge of σ is ultimately inconsequential for our purposes, as we will demonstrate subsequently for a more general setting. Despite σ being known, this strategy fails to yield any positive Sharpe ratio.

¹⁸For any time series of random vector a_t , we use $\bar{a}$ to denote its sample average. As we will point out later in the paper, this strategy $\hat{w}$, which we will denote by $\hat{w}^{\text{CSR}}$, fails to achieve the optimal Sharpe ratio in all scenarios. We will discuss the optimal strategy in Section 2.5.

Definition 2. *A strategy is deemed feasible when it relies solely on observable data, in stark contrast to an infeasible strategy that presupposes complete knowledge of the DGP, encapsulated by the full information set, $\mathcal{F}_t$.*

Recognizing the distinction between feasible and infeasible strategies, we proceed to investigate the maximum Sharpe ratio attainable by feasible strategies. This exploration necessitates defining the decision-making problem faced by arbitrageurs, which we turn to next.

2.3 Arbitrageurs' Decision Problem and Feasible Sharpe Ratio Bound

In the broader context of the DGP, given by equation (1), the information set $\mathcal{F}_t$ not only includes α and $\{r_s\}_{s \leq t}$ but also other variables such as Σ_u , β , and $\{v_s\}_{s \leq t}$, alongside unknown parameters like Σ_v and γ . For a strategy to be considered feasible, it must refrain from leveraging these unknown variables and parameters. This constraint introduces complexity to the arbitrageurs' decision-making process by intertwining optimization with statistical techniques, resulting in solutions that are not uniquely defined.

Recognizing that the disparity in Sharpe ratios between infeasible and feasible strategies primarily stems from the challenges of learning about α , we seek to navigate this complexity by considering a broader set of trading strategies. These strategies deliberately exclude reliance on information about α , although they may not all be feasible. Given this constraint, we derive an upper bound for the Sharpe ratio that applies to all such strategies, and by extension, to all feasible strategies, since they similarly avoid using information on α . However, this upper bound may not be achievable by a feasible strategy, as the derivation does not preclude the use of potentially infeasible information on other variables, such as Σ_u . In Section 2.5, we demonstrate that it is indeed possible to attain this upper bound with a well-crafted feasible trading strategy. Consequently, even without perfect knowledge about other variables and parameters, arbitrageurs can effectively use statistical inference to estimate these unknowns, with the statistical error of such estimates diminishing asymptotically. This highlights that the main hurdle for arbitrageurs is not in learning about factors, volatilities, or risk premia, but in acquiring knowledge about α .

We are now ready to characterize the decision problem arbitrageurs confront. Operating under a certain information set, $\mathcal{G}$, a subset of $\mathcal{F}_t$, arbitrageurs at time t are tasked with solving a mean-variance optimization problem expressed as:

$$\tilde{w} = \arg \max_{w \in \mathcal{G}: w^\top \beta = 0} U(w), \quad \text{where} \quad U(w) = \mathbb{E}(w^\top r_{t+1} | \mathcal{G}) - \frac{\kappa}{2} \text{Var}(w^\top r_{t+1} | \mathcal{G}),$$

and κ measures the degree of risk aversion. Our objective is to identify the optimal strategy subject to the measurability constraint: $w \in \mathcal{G}$. Given the necessity for statistical arbitrage

to be factor-neutral, the strategy w must be orthogonal to β , with β being predetermined, i.e., $\beta \in \mathcal{G}$.

To streamline the portfolio formation process, we exclude transaction cost considerations, which yields a static, single-period optimization problem. Thus, we omit subscript t whenever possible. This approach aligns with the APT framework introduced by Ross (1976), where the arbitrageurs' chosen strategy aims to eliminate exposure to systematic factors, concentrating instead on capitalizing on alphas and balancing idiosyncratic risks. Arbitrageurs' utility is modeled as a Bayesian posterior expected loss, a concept discussed in Berger (1985). The next theorem presents a solution with minimal conditions on $\mathcal{G}$.

Theorem 1. *Suppose that r_t follows (1) and that β is $\mathcal{G}$ -measurable. It holds that*

$$\tilde{w} = \frac{1}{\kappa} \tilde{\Sigma}_u^{-1/2} \mathbb{M}_{\tilde{\Sigma}_u^{-1/2} \beta} \tilde{\Sigma}_u^{-1/2} \tilde{\alpha}, \quad (9)$$

where $\tilde{\alpha} := E(\alpha|\mathcal{G})$ and $\tilde{\Sigma}_u := E(\Sigma_u|\mathcal{G}) + \text{Var}(\alpha|\mathcal{G})$.¹⁹

This result associates the optimal strategy with the posterior summary statistics of α and Σ_u . Essentially, the choice of the optimal $\mathcal{G}$ -measurable strategy is determined by arbitrageurs' best parameter estimates based on their information set, rather than the true values of the unknown parameters. In the special scenario where $\mathcal{G} = \mathcal{F}_t$, meaning arbitrageurs possess complete information up to time t , we have $\tilde{\alpha} = \alpha$, and $\tilde{\Sigma}_u = \Sigma_u$. Under these conditions, the optimal infeasible strategy, denoted by w^* , is given by

$$w^* = \frac{1}{\kappa} \Sigma_u^{-1/2} \mathbb{M}_{\Sigma_u^{-1/2} \beta} \Sigma_u^{-1/2} \alpha. \quad (10)$$

Moving forward, our attention is directed towards a specific choice of $\mathcal{G}$ as generated by $\{(r_s, \beta, v_s, \Sigma_u) : t - T + 1 \leq s \leq t\}$.²⁰ Accordingly, $\tilde{w}$ is assumed associated with this particular choice by default hereafter.

Next, we examine the investment performance of arbitrageurs' optimal strategy, $\tilde{w}$. It is important to note that in evaluating a strategy, the conditional Sharpe ratio is calculated based on the comprehensive information set, $\mathcal{F}_t$, rather than the arbitrageurs' specific information set, $\mathcal{G}$. For any given strategy w , its conditional Sharpe ratio is defined as:

$$S(w) := E(w^\top r_{t+1} | \mathcal{F}_t) / \text{Var}(w^\top r_{t+1} | \mathcal{F}_t)^{1/2}.$$

¹⁹ $\mathbb{M}_A = \mathbb{I} - A(A^\top A)^{-1}A^\top$ for any matrix A .

²⁰The nuisance parameters γ and Σ_v are deterministic constants and, as such, are encompassed within any information set. Although they are part of $\mathcal{G}$, their presence does not influence the optimal strategy of arbitrageurs.

To describe the asymptotic behavior of a strategy, we need to impose assumptions on the return generating process:

Assumption 1. *For each $N \geq 1$, the following conditions hold:*

- (a) (α_i, u_i) is i.i.d. across i , and satisfies $E(\alpha_i | (\Sigma_u)_{i,i}) = 0$ and $E\|\alpha\|_{\text{MAX}}^2 = o(1)$. Moreover, it holds that $1 \lesssim_P \lambda_{\min}(\Sigma_u) \leq \lambda_{\max}(\Sigma_u) \lesssim_P 1$.²¹
- (b) The pricing errors α , factors v_t , factor loadings β , and idiosyncratic errors u_t are, conditionally on Σ_u , mutually independent.

Condition (a) suggests that the alphas in our model are inherently weak; as the number of assets, N , increases, their magnitudes diminish.²² Moreover, this condition ensures that volatilities remain within upper and lower bounds. Condition (a) therefore implies that learning about alpha is a more arduous task than learning about volatilities. Condition (b) is imperative for the model's identification. For instance, assuming independence between α and β is key to identify the risk premia, γ , see, e.g., Giglio et al. (2022).

Given this assumption, the subsequent theorem illustrates that the optimal strategy $\tilde{w}$ yields a Sharpe ratio that is nearly at its maximum.

Theorem 2. *Suppose that r_t follows (1) and that Assumption 1 holds. Let $\mathcal{G}$ represent the information set generated by $\{(r_s, \beta, v_s, \Sigma_u) : t - T + 1 \leq s \leq t\}$. For any strategy w that is both $\mathcal{G}$ -measurable and factor-neutral, it follows that as $N \rightarrow \infty$,*

$$S(w) \leq S(\mathcal{G}) + o_P(1 + S(\mathcal{G})), \quad (11)$$

where $S(\mathcal{G})^2 = \tilde{\alpha}^\top \Sigma_u^{-1} \tilde{\alpha}$, and the equality in (11) holds if $w = \tilde{w}$. Moreover, the optimal strategy for arbitrageurs, $\tilde{w}$, has a simple approximation:

$$\left\| \tilde{w} - \frac{1}{\kappa} \mathbb{M}_\beta \Sigma_u^{-1} \tilde{\alpha} \right\| = o_P(1 + S(\mathcal{G})). \quad (12)$$

Accordingly, their optimal utility satisfies $\sqrt{2\kappa U(\tilde{w})} = S(\mathcal{G}) + o_P(1 + S(\mathcal{G}))$.

Theorem 2 derives an upper bound, denoted as $S(\mathcal{G})$, on the Sharpe ratio that any $\mathcal{G}$ -measurable strategy can attain. This upper bound also connects with arbitrageurs' optimal

²¹For a matrix A , we use $\|A\|$ and $\|A\|_{\text{MAX}} = \max_{i,j} |a_{ij}|$ to denote the operator norm (or ℓ_2 norm) and the ℓ_∞ norm of A on the vector space. We use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote the minimum and maximum eigenvalues of A .

²²By Assumption 1(a), $\text{Var}(\alpha_i) = E(\alpha_i^2) = o(1)$. Referencing our earlier discussion (footnote 13), a diminishing variance in α is essential for precluding near-arbitrage opportunities within Ross' APT framework.

utility. By the definition of $S(\mathcal{G})$, this upper bound satisfies

$$E(S(\mathcal{G})^2) \leq E(\alpha^\top \Sigma_u^{-1} \alpha), \quad (13)$$

where the equality holds only when $\tilde{\alpha} = \alpha$ almost surely. The right-hand side corresponds to the infeasible scenario in which arbitrageurs have perfect knowledge of α , which echoes (4).

Furthermore, the result demonstrates that the optimal strategy, as articulated in Theorem 1, can reach this upper bound with a negligible approximation error. In scenarios where $S(\mathcal{G})$ is finite, that is $O_P(1)$, the error term $o_P(1 + S(\mathcal{G}))$ reduces to $o_P(1)$ and vanishes in the limit. The theorem additionally addresses situations where $S(\mathcal{G})$ may diverge. In this case, the approximation error is reduced to $o_P(S(\mathcal{G}))$, remaining inconsequentially small relative to $S(\mathcal{G})$ itself.

Also, Equation (12) offers a more straightforward and intuitive formula for approximating the optimal feasible strategy, $\tilde{w}$. Additionally, as a byproduct of Theorem 2, the optimal infeasible strategy in Equation (10) can also be approximated in a simpler form:

$$\left\| w^* - \frac{1}{\kappa} \mathbb{M}_\beta \Sigma_u^{-1} \alpha \right\| = o_P(1 + S^*), \quad (14)$$

which represents a special case of Equation (12) where $\mathcal{G}$ is set as $\mathcal{F}_t$. At its core, it is $\tilde{\alpha}$, the posterior estimate of the pricing errors that dictates the optimal Sharpe ratio achievable by arbitrageurs, rather than the true α . Intuitively, part of the construction in (12), $\Sigma_u^{-1} \tilde{\alpha}$, is the optimal allocation to the ex-factor returns, $\alpha + u_t = r_t - \beta(\gamma + v_t)$, based on a simple mean-variance analysis, except for the use of $\tilde{\alpha}$, because arbitrageurs do not observe true alphas in the DGP. Pre-multiplying $\Sigma_u^{-1} \tilde{\alpha}$ by $\mathbb{M}_\beta$ simply eliminates the portfolio's factor exposures, because $(\mathbb{M}_\beta \Sigma_u^{-1} \tilde{\alpha})^\top \beta = 0$.

Theorems 1 and 2 are concerned with factor-neutral strategies, as our main results focus on the limits of arbitrage. Proposition B1 in our appendix broadens this scope to encompass any strategy, revealing that any $\mathcal{G}$ -measurable strategy, w , adheres to a similar upper limit:

$$S(w) \leq (S(\mathcal{G})^2 + \gamma^\top \Sigma_v^{-1} \gamma)^{1/2} + o_P(1 + S(\mathcal{G})), \quad (15)$$

where $\gamma^\top \Sigma_v^{-1} \gamma$ is the squared optimal Sharpe ratio earned from the factor portfolios.

This result addresses a longstanding challenge in optimal portfolio choice under parameter uncertainty. Kan and Zhou (2007) examine the expected performance of the plug-in mean-variance portfolio and find that its Sharpe ratio is smaller than the infeasible Sharpe ratio, S^* . To construct a solution, they consider a limited class of trading strategies that does not accommodate the broad range of strategies investors might consider in real-world scenarios.

In contrast, by imposing only mild restrictions on the DGP of returns, equation (15) identifies an upper bound for all feasible Sharpe ratios without limiting the choice of strategies.

The result in Theorem 2 appears to require that arbitrageurs rely on the information set $\mathcal{G}$, which embodies perfect knowledge of factors, v_t , and their exposures, β , in addition to past asset returns, r_t . Moreover, arbitrageurs appear to have perfect knowledge of the covariance matrix of idiosyncratic errors, Σ_u . In fact, the upper bound in (15) still holds if arbitrageurs are endowed with less information, because for any information sets $\mathcal{G}'$ and $\mathcal{G}$ such that $\mathcal{G}' \subseteq \mathcal{G}$, we have $E(S(\mathcal{G}')^2) \leq E(S(\mathcal{G})^2)$. Since all feasible strategies are measurable with respect to $\mathcal{G}$, $S(\mathcal{G})$ establishes an upper bound for the Sharpe ratio applicable to this set of strategies, although it may not be sharp. In light of this and Definitions 1 and 2, we immediately obtain a sufficient condition of the absence of near-arbitrage with feasible strategies:

Corollary 1. *Suppose the same assumptions as in Theorem 2 hold. For any given return-generating process satisfying (1), there exists no feasible factor-neutral strategy that leads to a near-arbitrage, if*

$$S(\mathcal{G}) \lesssim_P 1, \quad \text{as } N \rightarrow \infty. \quad (16)$$

2.4 Bayes Correction for Selection Bias

While the previous discussion underscored the role of the posterior mean of α , denoted as $\tilde{\alpha} = E(\alpha|\mathcal{G})$, on the optimal $\mathcal{G}$ -measurable portfolio, this expression remains implicit and not directly actionable. Our focus now shifts to reducing the information set $\mathcal{G}$ to its sufficient statistics for α . This step is crucial for understanding why $\tilde{\alpha}$ leads to enhanced portfolio performance. To facilitate this analysis, we introduce further assumptions:

Assumption 2. *For each $N \geq 1$, the following conditions hold:*

- (a) $u_{i,t} = \sigma_i \varepsilon_{i,t}$, where $\varepsilon_{i,t}$ follows a standard normal distribution, and is i.i.d. across (i, t) and independent of Σ_u .
- (b) $s_i := \alpha_i / \sigma_i$ is independent of σ_i .

Based on the DGP described in Equation (1) and given the information set $\mathcal{G}$, Assumptions 1 and 2(a) together ensure that the key summary statistics for α_i are the volatility σ_i and the sample average of the ex-factor returns, expressed as $\check{\alpha}_i := \bar{r}_i - \beta_i(\gamma + \bar{v}) = \alpha_i + \bar{u}_i$. In other words, $E(\alpha_i|\mathcal{G}) = E(\alpha_i|\check{\alpha}_i, \sigma_i)$. Consequently, this assumption simplifies the conditional information set $\mathcal{G}$ in the posterior distribution of α to merely a two-dimensional vector comprising these conditioning variables. To evaluate this conditional expectation, it becomes necessary to assume a specific form of dependence between α_i and σ_i .

In light of this, we introduce Assumption 2(b), which allows us to further express $E(\alpha_i|\check{\alpha}_i, \sigma_i) = \sigma_i E(s_i|\check{\alpha}_i, \sigma_i) = \sigma_i E(s_i|\check{\alpha}_i/\sigma_i)$.²³ This leads to $\tilde{s}_i := E(s_i|\mathcal{G}) = E(s_i|\check{s}_i)$, where $\check{s}_i := \check{\alpha}_i/\sigma_i$. Consequently, in terms of the scaled version of α , namely, s , the conditioning information set is now a single scalar variable, which simplifies the estimation problem later. Further, in light of Theorem 2, $\tilde{\alpha}$, $\tilde{w}$, and $S(\mathcal{G})^2$ can all be represented in relation to $\tilde{s}$.²⁴

$$\tilde{\alpha} = \Sigma_u^{-1/2}\tilde{s}, \quad \tilde{w} \propto \mathbb{M}_\beta \Sigma_u^{-1/2}\tilde{s}, \quad S(\mathcal{G})^2 = \tilde{s}^\top \tilde{s}. \quad (17)$$

Finally, note that

$$\check{s}_i = s_i + \bar{\varepsilon}_i \sim \mathcal{N}(s_i, T^{-1}), \quad \text{conditional on } s_i.$$

This formulation casts the original posterior inference problem into the framework of the classical Gaussian sequence model for recovering a high-dimensional mean vector, s , from noisy observations $\check{s}$, which has been extensively studied in the statistics literature (see, e.g., Robbins (1956), Efron (2011), and Efron (2019)). Although the assumption that ε_i follows a Gaussian distribution may seem restrictive, this framework is sufficiently versatile to accommodate a wide range of distributions for s_i .

Arbitrageurs face the challenge of identifying the true underlying signal, s , from an observed noisy signal $\check{s}$. This task is complicated by what Efron (2011) described as selection bias or “the winner’s curse”: A high observed signal, $\check{s}_i$, could reflect a high s_i , or it could be the result of “luck”, with an unusually large noise realization $\bar{\varepsilon}_i$. Consequently, arbitrageurs must carefully adjust their investment strategies to mitigate this potential bias. The correction involves relying on $\tilde{s}$, which accounts for this bias. To see this, we present an explicit formula for $\tilde{s}$:

Theorem 3. *Suppose that r_t follows (1) and Assumptions 1 and 2 hold. We introduce a function $\psi(a) = E(s_i|\check{s}_i = a)$, with which we have $\tilde{\alpha}_i = \sigma_i \tilde{s}_i$, where $\tilde{s}_i = \psi(\check{s}_i)$. Moreover, it holds that*

$$\psi(a) = a + \frac{1}{T} \frac{d}{da} \log p(a), \quad (18)$$

where $p(a) = E(\phi_{1/T}(a - s_i))$ is the probability distribution function of $\check{s}_i$.²⁵

The preceding discussion directly leads to the first assertion of the theorem. Equation (18) is rooted in Tweedie’s formula (Robbins, 1956), which establishes a connection between the posterior mean of s , given $\check{s} = a$, denoted $\psi(a)$, and the posterior probability distribution of

²³This equality relies on the result that conditional on $\hat{\alpha}_i/\sigma_i$, α_i/σ_i is independent of σ_i . We impose this condition primarily for clarity of exposition and simplicity of Algorithm 1 below.

²⁴Given that risk aversion does not impact the out-of-sample Sharpe ratio, we will use $\propto$ to substitute for κ^{-1} in our subsequent discussions on the portfolio strategy.

²⁵ $\phi_x(a)$ is the distribution function for $\mathcal{N}(0, x)$.

š. The formula’s second component, $T^{-1}d \log p(a)/da$, plays a crucial role in adjusting for the selection bias in the observed signal, $\check{s}_i$, which introduces several intriguing properties. An example of these properties, as shown by Andrews et al. (1972), is that $\psi(a)$ is a nondecreasing function of a . This property ensures that the relative magnitude of various signals is preserved post-correction. Moreover, under mild assumptions about the prior distribution of s_i , such as being symmetric and unimodal, the posterior mean $\tilde{s}$ induces a shrinkage effect towards the prior mean, which, in our context, is zero.

The posterior shrinkage on α , or equivalently on s , consequently induces a “shrinkage” effect on the optimal Sharpe ratio achievable by arbitrageurs, as demonstrated by the inequality (13). Based on the results of Theorem 3, and under a technical condition concerning the tail behavior of s_i , we can obtain a more explicit formula for $S(\mathcal{G})$:

Corollary 2. *Under the same assumptions outlined in Theorem 3, and with the additional condition that $E(s_i^2 \mathbb{1}_{\{|s_i| \geq c_N\}}) = o(N^{-1})$, we arrive at the conclusion that:*

$$S(\mathcal{G}) = S^{\text{OPT}} + o_P(1), \quad \text{with} \quad S^{\text{OPT}} = \left(N \int \psi(a)^2 p(a) da \right)^{1/2}.$$

This result enables us to compute S^{OPT} in various examples. For instance, by utilizing this result, we can compare S^{OPT} with S^* in Example 1. This comparison illuminates the feasibility of attaining non-vanishing Sharpe ratios and how different DGP conditions influence the attainability.

Corollary 3. *Suppose that the same assumptions as in Corollary 2 hold. In addition, we assume alpha follows (5) as in Example 1. Then we have $S^* = \sigma^{-1} \mu(\rho N)^{1/2} + o_P(1)$. Further, assuming that $\sigma^{-1} \mu(\rho N)^{1/2}$ does not vanish, then it holds that $S^{\text{OPT}} \leq (1 - \epsilon) \sigma^{-1} \mu(\rho N)^{1/2}$ for some $\epsilon > 0$, if and only if*

$$T^{1/2} \mu / \sigma - \sqrt{-2 \log \rho} \lesssim 1. \tag{19}$$

Corollary 3 suggests that when $T^{1/2} \mu / \sigma$ is large enough that the constraint (19) is violated, $S^* \approx_P S^{\text{OPT}}$ —that is, in the limit, learning does not play any role—arbitrageurs in this scenario achieve the same Sharpe ratio as under perfect knowledge of alphas. Furthermore, the rareness parameter ρ does not make much difference if $T^{1/2} \mu / \sigma$ gets sufficiently large. That said, if ρ approaches to zero so fast that $\sqrt{-2 \log \rho}$ dominates $T^{1/2} \mu / \sigma$ —that is, alpha is extremely rare and sufficiently weak—the learning problem becomes rather challenging and the Sharpe ratio attainable by arbitrageurs, S^{OPT} , is dominated, in the limit, by the infeasible Sharpe ratio S^* .

To give a concrete example of Corollary 3, consider an alternative DGP assumption instead of (6):²⁶

$$\mu \asymp N^{-\eta} \quad \text{and} \quad \rho > 0 \quad \text{is fixed.} \quad (20)$$

In this scenario, $(S^*)^2 \asymp_P N^{1-2\eta}$, which explodes unless $\eta > 1/2$. If we further assume that $N/T \rightarrow \psi > 0$, then the left-hand-side of condition (19) is of order $N^{1/2-\eta} \vee 1$, so that (19) holds if and only if $\eta \geq 1/2$. Therefore, $\eta < 1/2$ is not consistent with absence of (feasible) near arbitrage because S^* explodes and, since $S^{\text{OPT}} = S^*$ by Corollary 3, S^{OPT} explodes, too. If $\eta > 1/2$, S^* (and hence S^{OPT}) vanishes, which does not seem like an economically plausible case. If we think that asset demand distortions are sufficiently big so that arbitrageur activity is required to prevent substantial mispricing, and that arbitrageurs require some compensation for holding arbitrage positions, then a setting where true alphas disappear asymptotically is not plausible. This suggests that under the DGP (20), the only economically plausible case with absence of near-arbitrage, but not absence of mispricing, is $\eta = 1/2$. That is, η can be thought as determined in equilibrium. Mispricing should be big enough that arbitrageurs are active, earning a non-vanishing Sharpe ratio, and at the same time small enough that near-arbitrage opportunities do not exist.

We now illustrate the behavior of S^{OPT} numerically and verify the theoretical predictions of Corollary 3 using the DGP specified in Example 1. Figure 1 reports the Sharpe ratio, S^{OPT} , of optimal feasible arbitrage portfolios for a range of μ/σ and ρ values in the case of $N = 1,000$ and $T = 20$ years. Recall that according to model (5), a fraction ρ of assets have alphas with a Sharpe ratio μ/σ . That is, ρ characterizes the rareness of the alpha signal, whereas μ/σ captures its strength. We intentionally choose a wide range of μ/σ (with annualized Sharpe ratios from 0.11 to 10.95) and ρ (from 0.12% to 50%) to shed light on the dependence of Sharpe ratios on signal weakness and rareness, although some of the resulting portfolio Sharpe ratios (the top left corner of Figure 1) are unrealistically high. Note that when $\mu/\sigma \times \sqrt{12}$ hits 0.44, its corresponding t-statistic based on a 20-year sample exceeds 1.96, the typical t-hurdle for a standard student-t test.

The pattern of Sharpe ratios agrees with our intuition and theoretical predictions. For any fixed ρ , as the alpha signal weakens (i.e., μ/σ decreases), the optimal Sharpe ratio drops. The same is true if we decrease the signal count (i.e., ρ vanishes), for any fixed value of μ/σ . The arbitrageur's learning problem is easiest when signal is strong and count is large (top left corner), and most challenging towards the right bottom corner, where the optimal Sharpe ratios drop to near zero.

²⁶It is easy to show that the setup (20) satisfies all assumptions of Corollary 2 for all fixed $\eta > 0$.




Figure 1: Optimal Sharpe Ratios (S^{OPT}) of Feasible Arbitrage Portfolios

Note: The figure reports optimal Sharpe ratios of feasible arbitrage portfolios in model (5), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

The reported Sharpe ratios in Figure 1 are only a fraction of the corresponding (infeasible) Sharpe ratios, $S^* = \sqrt{\alpha^\top (\Sigma_u)^{-1} \alpha} = \mu/\sigma \sqrt{\rho N}$, as shown by Figure 2. The pattern we see from Figure 2 agrees with theoretical predictions of Corollary 3. When the annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$ is larger than 2.74, regardless of the values of ρ , the signal-to-noise ratio of the learning problem is sufficiently strong that the statistical limit to arbitrage does not matter much, and hence S^{OPT}/S^* is close to 1. Nonetheless, this regime is irrelevant in practice, since it is mostly associated with unrealistically high Sharpe ratios (see Figure 1). In contrast, as μ/σ diminishes, the gap between S^* and S^{OPT} widens. In almost all empirically relevant scenarios, S^* is much larger than the feasible Sharpe ratio than arbitrageurs can actually earn with an optimal strategy.

2.5 Constructing the Optimal Arbitrage Portfolio

In our previous discussion, Theorem 2 establishes that the optimal Sharpe ratio attainable by any feasible strategy is capped by $S(\mathcal{G})$, which, by Corollary 2, is approximately equal to S^{OPT} under additional conditions. Following this, Corollary 3 illustrates the dependence of S^{OPT} on the unknown parameters in the DGP of alpha, i.e., (5). In addition, (17) and Theorem 3 characterize the strategy $\tilde{w}$ that achieves the optimal Sharpe ratio attainable by any feasible strategy.

However, this characterization assumes access to the information set $\mathcal{G}$, which includes




Figure 2: Ratios between S^{OPT} and S^*

Note: The figure reports the ratios of optimal Sharpe ratios between feasible and infeasible arbitrage portfolios. The simulation setting is based on model (5), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

not only past returns, r_t , but also potentially unknown variables, such as β , v_t , and Σ_u . Furthermore, it relies on knowledge of γ and the probability distribution function of $\check{s}$, i.e., $p(a)$. As a result, while the Sharpe ratio bound remains a valid upper ceiling—since having less information generally leads to lower achievable Sharpe ratios—the strategy $\tilde{w}$ is infeasible. Its reliance on Σ_u and $\tilde{s}$, which in turn depend on ex-factor returns, $\bar{r}_i - \beta_i(\gamma + \bar{v})$, and the probability distribution of $\check{s}$ via (18), makes its implementation impractical. Therefore, arbitrageurs encounter additional challenges in this context. They must conduct statistical inference concerning unknown parameters and variables to shape their investment strategies. Moreover, the devised optimal strategy must be flexible and adapt to the nuances of different DGPs.

Despite these challenges, arbitrageurs can indeed construct a uniformly optimal strategy that attains S^{OPT} across a broad spectrum of DGPs, even without access to all information in $\mathcal{G}$ and perfect knowledge of γ and the probability distribution of $\check{s}$. In fact, this strategy relies only on observable data, making it entirely feasible. Specifically, we formulate this strategy within a framework where factors remain latent, but factor exposures are observable—a scenario that aligns with our empirical analysis.

We describe this portfolio strategy as “all weather,” signifying its adaptability to all considered DGP scenarios. Moreover, the fact that this strategy achieves S^{OPT} implies that

the feasible Sharpe ratio upper bound we derive is sharp. We outline the following algorithm to construct this strategy:

Algorithm 1 (Optimal Arbitrage Portfolio via Empirical Bayes).

Inputs: r_t , $t \in \mathcal{T} = \{t - T + 1, \dots, t\}$ and β .

S1. Construct cross-sectional regression estimates of alpha and idiosyncratic volatilities, for each $i = 1, 2, \dots, N$:

$$\hat{\alpha} = T^{-1} \sum_{s \in \mathcal{T}} \mathbb{M}_{\beta} r_s, \quad \hat{\sigma}_i^2 = T^{-1} \sum_{s \in \mathcal{T}} ((\mathbb{M}_{\beta} r_s)_i - \hat{\alpha}_i)^2, \quad \text{and} \quad \hat{s}_i := \hat{\alpha}_i / \hat{\sigma}_i.$$

S2. Construct a nonparametric estimate of the marginal density of $\hat{s}$ using Gaussian kernel function $\phi_{1/T}(x)$ and bandwidth $k_N \asymp (\log N)^{-1}$:

$$\hat{p}(a) = \frac{1}{N k_N} \sum_i \phi_{1/T}\left(\frac{\hat{s}_i - a}{k_N}\right).$$

S3. Estimate $\tilde{s}$ by Tweedie's formula (18) and isotonic regression:

$$\check{\psi}(a) = a + \frac{1 + k_N^2}{T} \frac{d}{da} \log \hat{p}(a). \\ \hat{\psi} = \arg \min_{x \in \mathbb{R}^N} \|x - \check{\psi}\|^2, \text{ s.t. } x_i \leq x_j \text{ if } \hat{s}_i \leq \hat{s}_j, \text{ for } 1 \leq i, j \leq N, \text{ where } \check{\psi}_i := \check{\psi}(\hat{s}_i).$$

S4. Construct the arbitrage portfolio weights as $\hat{w}^{\text{OPT}} \propto \mathbb{M}_{\beta} \hat{\Sigma}_u^{-1/2} \hat{\psi}$.

Outputs: $\hat{w}^{\text{OPT}}$.

Step S1 of Algorithm 1 provides feasible estimates of $\hat{\alpha}$ and $\hat{\Sigma}_u = \text{Diag}(\hat{\sigma}_i^2)$, which in turn leads to the sufficient statistic, $\hat{s}$. The motivation behind Steps S2 and S3 stems from Tweedie's formula (18); here, we employ a nonparametric empirical Bayes method for estimating the posterior mean function using kernel density estimation, as suggested by Brown and Greenshtein (2009). The incorporation of the factor $(1 + k_N^2)$ serves to adjust for finite sample biases introduced by estimation errors in $\hat{p}(a)$. To enhance the nonparametric estimator's performance in finite samples, enforcing monotonicity on $\psi(\cdot)$ proves beneficial. For this purpose, isotonic regression is utilized (see Robertson et al. (1988)), yielding a monotonic piece-wise linear approximation of $\psi(\cdot)$. Step S4 constructs the optimal portfolio weights, $\hat{w}^{\text{OPT}}$, following (17).

An essential step towards achieving optimality involves aggregating information from assets with comparable $\hat{s}_i$, as done in Step S2. This strategy outperforms the alternatives, some

of which directly use estimated $\widehat{s}$ as if these estimates are not susceptible to estimation errors even when they are rather weak, or simply ignore the contribution of all weaker signals. Like any machine learning method, the proposed approach requires a tuning parameter k_N , which can be selected in a validation sample.

The following theorem demonstrates the optimality of $\widehat{w}^{\text{OPT}}$:

Theorem 4. *Let $\mathbb{P}$ denote the collection of all data-generating processes under which the same assumptions outlined in Corollary 2 hold. Moreover, we assume $\|\beta\|_{\text{MAX}} \lesssim_{\mathbb{P}} 1$, $\lambda_{\min}(\beta^\top \beta) \gtrsim_{\mathbb{P}} N$, and that the distribution of s_i is symmetric. In addition, suppose that $N^d \lesssim T \lesssim N^{d'}$ for fixed constants $d > 1/2$ and $d' < 1$. We denote the Sharpe ratio achieved by $\widehat{w}^{\text{OPT}}$ as $\widehat{S}^{\text{OPT}} := \mathbb{E}(r_{t+1}^\top \widehat{w}^{\text{OPT}} | \mathcal{F}_t) / \text{Var}(r_{t+1}^\top \widehat{w}^{\text{OPT}} | \mathcal{F}_t)^{1/2}$. Then it holds that $\widehat{w}^{\text{OPT}}$ achieves, asymptotically, the upper bound S^{OPT} , uniformly over all sequences of data-generating processes. That is, for any $\epsilon > 0$,*

$$\lim_{N, T \rightarrow \infty} \sup_{\mathbb{P} \in \mathbb{P}} \mathbb{P}(|\widehat{S}^{\text{OPT}} - S^{\text{OPT}}| \geq \epsilon S^{\text{OPT}} + \epsilon) = 0.$$

Theorem 4 imposes a mild assumption on $\lambda_{\min}(\beta^\top \beta)$, which requires that all factors are pervasive.²⁷ This condition is frequently utilized in the factor model literature and is particularly relevant here, given our assumption that the factors within our model are latent. A significant difference between our context and the traditional empirical Bayes literature (e.g., Brown and Greenshtein (2009)) lies in our lack of direct access to the signals, s_i 's. Instead, we estimate these signals via cross-sectional regressions. This step results in our estimators, $\widehat{s}_i$'s, being inevitably polluted by estimation errors. The additional conditions imposed by Theorem 4 ensure that such estimation errors become asymptotically negligible.

Theorem 4 concludes that in the context of a linear factor model, arbitrageurs can construct this strategy, without any knowledge besides past returns and risk exposures (beta), to achieve the maximal Sharpe ratio over all feasible trading strategies that have zero exposure to factor risks. This Sharpe ratio precisely characterizes the limit of feasible arbitrages in economic terms. A key reason for this result is that idiosyncratic variances represented by Σ_u remain bounded both from below and above as N and T increase, unlike alphas. This difference is empirically supported; alphas are expected to be small and rare, a result of the competitive nature of arbitrage activities, in contrast to the consistent presence of idiosyncratic risks. Consequently, uncertainty about Σ_u becomes negligible in the context of statistical arbitrage limitations. Furthermore, given the low-dimensional nature of v_t and γ and arbitrageurs' knowledge of β , the problem of learning about v_t and γ is negligible as well.

²⁷See, e.g., Assumption I.1 of Giglio and Xiu (2021). While our theoretical results may extend to certain weak factor settings, this is not our emphasis here.

The term $\epsilon S^{\text{OPT}} + \epsilon$ accommodates both small and large values of S^{OPT} . If $S^{\text{OPT}} \lesssim 1$, then ϵ dominates and the estimation error inside the probability is characterized by the absolute difference between $\widehat{S}^{\text{OPT}}$ and S^{OPT} . Otherwise, if $S^{\text{OPT}} \rightarrow \infty$, the estimation error is described in percentage terms. This is necessary because we simultaneously consider a large class of models where the Sharpe ratios may become extremely high or approach zero.

With Theorem 4, we establish the necessity for the no near-arbitrage condition given by (16).

Corollary 4. *Suppose the same assumptions as in Theorem 4 hold. The portfolio weights $\widehat{w}^{\text{OPT}}$ yield a near-arbitrage strategy under any sequences of data-generating processes for which condition (16) does not hold.*

The ability of arbitrageurs to construct an optimal strategy that realizes S^{OPT} carries important economic implications. In an economy with statistical limits to arbitrage, the equilibrium compensation that arbitrageurs demand for executing such a strategy equals S^{OPT} . If it were otherwise, arbitrageurs would continue trading until all profit opportunities were exhausted. Therefore, $\widehat{S}^{\text{OPT}}$ can be interpreted as an empirical estimate of this equilibrium compensation, which we seek to pin down empirically.

2.6 Estimating Optimal Infeasible Sharpe Ratio

We are also interested in estimating the optimal infeasible Sharpe ratio, S^* . This Sharpe ratio can be estimated by an econometrician ex-post and in-sample, but it cannot be achieved by any feasible portfolio. The existing literature on testing the APT often constructs test statistics in the spirit of Gibbons et al. (1989), which are effectively based on S^* (see, e.g., Pesaran and Yamagata (2017) and Fan et al. (2015)). While such tests are powerful and may lead to detection of alphas, their relevance for arbitrageurs might be limited. The challenge for arbitrageurs lies in translating statistical evidence of alpha discoveries into a feasible portfolio strategy that realizes profits. Our proposed S^{OPT} tackles this issue, providing an economically more meaningful evaluation of the APT.

To provide a benchmark for $\widehat{S}^{\text{OPT}}$, we propose an estimator for S^* inspired by its sample analog:

$$\widetilde{S}^* = \left(\bar{r}^\top \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \mathbb{M}_\beta \bar{r} \right)^{1/2}. \quad (21)$$

Unfortunately, this estimator has a non-vanishing asymptotic bias for certain DGPs we consider, as we will show later. To fix this issue, we propose a new estimator that is uniformly consistent:

$$\widehat{S}^* = \left(\bar{r}^\top \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \mathbb{M}_\beta \bar{r} - N/T \right)^{1/2}. \quad (22)$$

The next proposition summarizes the asymptotic properties of both estimators.

Proposition 1. *Suppose that the same assumptions outlined in Corollary 2 hold. Moreover, we assume $\|\beta\|_{\text{MAX}} \lesssim_P 1$ and $\lambda_{\min}(\beta^\top \beta) \gtrsim_P N$. We also impose $T \lesssim N$ and $T^{-1}N^{1/2} \log N \leq c_N$. Then we have*

$$\begin{aligned} \left| \hat{S}^* - S^* \right| / (1 + S^*) &= o_P \left(T^{-1/2} N^{1/4} \sqrt{\log N} \right), \\ \left| \tilde{S}^* - \left((S^*)^2 + NT^{-1} \right)^{1/2} \right| / (1 + S^*) &= o_P \left(T^{-1} N^{1/2} \log N \right). \end{aligned}$$

Similar to Theorem 4, the estimation error is relative when S^* dominates 1.0 asymptotically, and in absolute terms if S^* is dominated by 1.0.²⁸ This accommodates a large class of models, some of which have an exploding or a shrinking S^* . While it is possible to estimate S^* , it is not possible to build a portfolio that realizes it, unless the signal-to-noise ratio is sufficiently large such that $S^* = S^{\text{OPT}}$. Empirically, the difference between $\hat{S}^*$ and $\hat{S}^{\text{OPT}}$ thereby informs us about the economic consequences of learning.

2.7 Alternative Strategies for Arbitrage Portfolios

Algorithm 1 introduces a nuanced approach that allows arbitrageurs to achieve feasible optimality. In practice, many empirical asset pricing researchers and practitioners often rely on simpler strategies. Given the widespread adoption of these alternative approaches, it is useful to assess how closely they approximate the optimal feasible strategy. In this section, we examine their strengths and weaknesses in different DGP scenarios.

2.7.1 Cross-Sectional Regression

One of the most common strategies involves forming portfolios based directly on the cross-sectional regression estimates of α , obtained in Step S1 of Algorithm 1. This method is referred to as CSR, with its portfolio represented by $\hat{w}^{\text{CSR}}$:

$$\hat{w}^{\text{CSR}} \propto \mathbb{M}_\beta \hat{\Sigma}_u^{-1} \hat{\alpha}. \quad (23)$$

This is effectively a sample analogue of the approximated infeasible strategy given by (14). A specific instance of this strategy was discussed following Example 1.

We now exploit Example 1 to illustrate the pros and cons of the CSR strategy. Figure 3 illustrates the relationship between S^{CSR} , the theoretical Sharpe ratio achieved by $\hat{w}^{\text{CSR}}$, and S^{OPT} across a spectrum of DGPs. According to Proposition B2 in the online appendix, S^{CSR}

²⁸Obviously, the threshold 1.0 can be replaced by any fixed constant.

tends to be dominated by S^{OPT} in scenarios when alpha signals are sparse ($\sqrt{\rho T}\mu/\sigma$ is not excessively large) and simultaneously strong ($\sqrt{T}\mu/\sigma$ is not exceedingly small). This specific regime of dominance is clearly marked with black numerals within the heatmap in Figure 3. As the ratio $\mu/\sigma \times \sqrt{12}$ edges closer to 1.0—either moving towards this vertical threshold from the right hand side or descending from the upper left corner—the gap between S^{CSR} and S^{OPT} enlarges increasingly.

The CSR approach takes all signal estimates directly, without differentiating the significant ones from the insignificant ones. Consequently, even fake signals (pure noise) are assigned non-zero weights. This hurts the portfolio’s performance. On the other hand, the CSR strategy can achieve optimality when the strong signals are abundant (so that portfolio weights allocated to noise are inconsequential) or when all signals are weak (so that they do not differ too much from fake ones). The latter case is interesting, as it also suggests that simply ignoring weaker signals is not optimal. That said, Figure 1 shows that the DGPs for which the cross-sectional regression approach is strongly dominated by our optimal strategy are associated with realistic Sharpe ratios.




Figure 3: Ratios between S^{CSR} and S^{OPT}

Note: The figure reports the ratios between the Sharpe ratios of the OLS based portfolio and the feasible optimal arbitrage portfolio. The theoretical Sharpe ratio achieved by CSR is denoted as S^{CSR} , given explicitly by Proposition B2. The simulation setting is based on model (5), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

The CSR approach is a simple benchmark as it does not rely on any advanced statistical techniques to detect signals or distinguish their strength. The strategy we discuss next is

more advanced, in that it controls false discoveries among selected strong signals using the B-H procedure proposed by [Benjamini and Hochberg \(1995\)](#).

2.7.2 False Discovery Rate Control

To address the aforementioned selection bias in identifying profitable alpha signals, an alternative methodology conceptualizes the search for alpha as a multiple testing problem. In this case, with N assets that are each potentially associated with a nonzero α_i , we can establish for each asset i a null hypothesis $\mathbb{H}_0^i : \alpha_i = 0$. Rejection of this null hypothesis leads to the discovery of a non-zero alpha. Rather than focusing on the significance level of individual tests, a more appropriate strategy involves controlling the FDR, an approach recommended by [Barras et al. \(2010\)](#), [Bajgrowicz and Scaillet \(2012\)](#), and [Harvey et al. \(2016\)](#) in various asset pricing contexts. [Giglio et al. \(2021\)](#) have proved the validity of the B-H procedure for FDR control in a general factor model setting for alpha detection. Below we describe the necessary steps to prepare alpha estimates for constructing an arbitrage portfolio.

Begin with a series of p-values, p_i , where each is the result of a t-test on the cross-sectional regression estimate of α_i , $\sqrt{T}\hat{s}_i$, for $i = 1, 2, \dots, N$. These p-values assess the significance of each α_i 's deviation from zero. Arrange these p-values in ascending order, from the smallest to the largest, resulting in a sorted sequence $p_{(1)} \leq \dots \leq p_{(N)}$. Identify a critical index, $\hat{k}$, defined as the maximum i such that $p_{(i)} \leq \tau i/N$ where τ is a predetermined significance level, commonly set at 5%.

The threshold $\hat{k}$ is selected such that, on average, at least a fraction $(1 - \tau)$ of the alpha estimates identified as significant (i.e., those with p-values smaller than $p_{(\hat{k})}$) are truly non-zero. This B-H criterion proves effective in guarding against false discoveries irrespective of the overall proportion of non-zero alphas present in the DGP. The selected alpha estimates are then used as inputs for constructing an arbitrage portfolio, as illustrated by the following equation:

$$\hat{w}^{\text{BH}}(\tau) \propto \mathbb{M}_\beta \hat{\Sigma}_u^{-1} \hat{\alpha}^{\text{BH}}(\tau), \text{ where } \hat{\alpha}_i^{\text{BH}}(\tau) = \hat{\alpha}_i \mathbb{1}_{\{p_i \leq p_{(\hat{k})}\}}. \quad (24)$$

This strategy introduces a hard-thresholding mechanism to the alpha estimates, effectively nullifying the impact of alphas deemed insignificant.

Controlling the FDR on top of the CSR estimates is intuitively appealing, but doing so incurs a potential loss of power that may hurt investment performance. Indeed, [Proposition B3](#) in the appendix shows, in the context of [Example 1](#), arbitrageurs who adopt the B-H based-trading strategy do not achieve the optimal portfolio for a large class of DGP sequences.

As shown by [Proposition B3](#) and illustrated numerically by [Figure 4](#), the discrepancy between the optimal Sharpe ratio and that achieved through the B-H method is largely

determined by signal strength. The B-H procedure nears optimality when the signal strength, quantified by $\sqrt{T}\mu/\sigma$, is substantial—exceeding the threshold of $\sqrt{-2\log\rho}$. The instances where the B-H strategy reaches optimality are depicted by the white values on Figure 4, with the boundary of this optimal region approximated by the line where $\mu/\sigma\sqrt{12} = 2.19$. This demonstrates that the B-H method excels in identifying strong signals, leading to near-optimal portfolios when signals are strong. Conversely, in the presence of weak signals, the B-H procedure, which amounts to hard-thresholding, tends to underperform compared to CSR.

This point is further elaborated in Figure 1, which underscores that even if individual alphas are weak, their aggregated effect on a portfolio’s Sharpe ratio can be non-trivial. The B-H approach takes a notably conservative stance towards signal selection, especially in contexts where signals are weak. This cautious approach ensures the reliability of selected alphas by focusing on those that are truly significant. However, this method might not fully capitalize on the potential cumulative impact of weaker signals. In contrast, our optimal arbitrage portfolio leverages the full spectrum of alpha estimates, including false positives, extending beyond the significant selections made through the B-H procedure.

At its core, this delineates a subtle yet critical divergence between two objectives: alpha testing and portfolio construction. Alpha testing prioritizes the identification of statistically significant alphas while controlling FDR, whereas portfolio construction concentrates on using all available information to optimize performance. These objectives do not always align.

The CSR and the B-H approaches represent two typical strategies in practice. The former trades all signals without distinguishing their strength, whereas the latter only trades the stronger signals. Neither approach always achieves optimality.

2.7.3 Shrinkage Approaches

In the approximated infeasible strategy given by (14), the portfolio weights with respect to ex-factor returns ($\mathbb{M}_\beta r_{t+1} \approx \alpha + u_t$) are $\Sigma_u^{-1}\alpha$. The CSR approach replaces the weights $\Sigma_u^{-1}\alpha$ with the sample analogue $\widehat{\Sigma}_u^{-1}\widehat{\alpha}$, and the B-H approach additionally imposes hard-thresholding regularization on those weights.²⁹ Multiplying the ex-factor return weights by $\mathbb{M}_\beta$ yields the portfolio weights in terms of raw returns (r_{t+1}). Besides the hard-thresholding of B-H, we can consider shrinkage-type regularization:

$$\max_w \left\{ w^\top \widehat{\alpha} - \frac{1}{2} w^\top \widehat{\Sigma}_u w - p_\lambda(\widehat{\Sigma}_u^{1/2} w) \right\},$$

²⁹Regularizing those weights amounts to incorporating priors onto the alpha estimates.




Figure 4: Ratios between S^{BH} and S^{OPT}

Note: The figure reports the ratios between the Sharpe ratios of the multiple testing based portfolio (via B-H procedure) and the feasible optimal arbitrage portfolio. The theoretical Sharpe ratio achieved by B-H is denoted as S^{BH} , given explicitly by Proposition B3. The simulation setting is based on model (5), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

where $p_\lambda(x) = \lambda\|x\|_1$ or $\lambda\|x\|_2^2$, for some $\lambda > 0$. Since $\widehat{\Sigma}_u$ is diagonal, this optimization problem has a closed-form solution of weights with respect to ex-factor returns: $\check{w}_{q,i}(\lambda) = \widehat{\sigma}_i^{-1}\psi_q(\widehat{s}_i, \lambda)$, for $i = 1, 2, \dots, N$, where $q = 1$ corresponds to the Lasso penalty and $q = 2$ the ridge, and $\psi_q(s, \lambda)$ is

$$\psi_1(s, \lambda) = \text{sgn}(s)(|s| - \lambda)_+, \quad \psi_2(s, \lambda) = (1 + 2\lambda)^{-1}s.$$

This leads to the optimal portfolio weight on r_t :³⁰

$$\widehat{w}_q(\lambda) \propto \mathbb{M}_\beta \check{w}_q(\lambda), \quad q = 1, 2.$$

Depending on the magnitude of λ , the Lasso approach replaces all smaller signals (i.e., $\widehat{s}_i$) by zero and shrinks the larger ones by λ in absolute terms. In other words, the Lasso approach is the soft-thresholding alternative to the B-H method. In contrast, the ridge penalty shrinks all signals proportionally. Since proportional scaling of portfolio weights does not affect the Sharpe ratio, this means that ridge is equivalent to CSR! This “embedded” shrinkage effect

³⁰An alternative strategy is to impose sparsity directly on the portfolio weights with respect to raw returns. While this approach might be appealing from the transaction cost point of view, it does not associate with an explicit prior on alpha, hence is more difficult to interpret.

of CSR explains why it performs well in the case of small signals.

Proposition B4, along with Proposition B2 in the online appendix, offers explicit formulae for the optimal feasible Sharpe ratios in the shrinkage case. The Sharpe ratio of ridge is not affected by the tuning parameter, but Lasso’s performance is contingent on its tuning parameter. However, even with the optimal choice of tuning parameter, it cannot achieve the optimal feasible Sharpe ratio in all DGPs.

Figure 5 compares $S_{\lambda}^{\text{Lasso}}$ with S^{OPT} , where $S_{\lambda}^{\text{Lasso}}$ denotes the theoretical Sharpe ratio achieved by Lasso for a given tuning parameter choice λ . In practice, the tuning parameter choice would require a cross-validation procedure. Here we adopt the theoretically optimal tuning parameter that maximizes $S_{\lambda}^{\text{Lasso}}$ in Figure 5. This strategy, unlike the B-H approach where the tuning parameter τ is selected based on a clear statistical criterion, is not fully feasible in practice, since the optimal tuning depends on the DGP, which is not fully known. Hence, Figure 5 should be understood as showcasing the optimal performance of a class of trading strategies indexed by the tuning parameter of Lasso, and only reports the best performance within this class. This is a fair comparison to OPT, which represents the optimal performance of all strategies, beyond the Lasso class, as the implementation of our algorithm to achieve S^{OPT} would also require choice of a tuning parameter. Although Proposition B4 suggests that Lasso is not uniformly optimal, it performs remarkably well, achieving the optimal Sharpe ratio in almost all regimes. Intuitively, when signals are very strong, Lasso behaves like a hard-thresholding selector, as shrinkage has minimal effect. Conversely, when signals are weak, ridge (and hence CSR) approach optimality, and Lasso can replicate their performance by setting the tuning parameter close to zero.

3 Simulation Evidence

This section illustrates the applicability of our theory through simulations and assesses the finite sample performance of our proposed portfolio strategies.

For simplicity and clarity, we simulate a one-factor model of returns as specified by (1). We set the factor risk premium at 5% per annum with an annualized volatility of 25%. The distribution of beta across assets is modeled as a normal distribution with a mean and variance of one. Given our focus on arbitrage portfolios, the parameters of the factor component (including the number of factors) are inconsequential, because factors are eliminated by $\mathbb{M}_{\beta}$ in the initial steps of constructing these trading strategies. In addition, we use the model in (5) from Example 1 for the cross-sectional distribution of alpha. We maintain a constant idiosyncratic volatility, denoted as σ for all assets, because the ratio α/σ determines signal strength, and so there is no need to vary both α and σ in the cross section. While our Monte




Figure 5: Ratios between $S_{\lambda}^{\text{Lasso}}$ and S^{OPT}

Note: The figure reports the ratios between the Sharpe ratios of the Lasso based portfolio and the feasible optimal arbitrage portfolio. The theoretical Sharpe ratio achieved by Lasso is treated as the maximum of $S_{\lambda}^{\text{Lasso}}$ over all λ , given explicitly by Proposition B4. The simulation setting is based on model (5), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$. The tuning parameter λ is selected to maximize $S_{\lambda}^{\text{Lasso}}$.

Carlo experiment is stylized, it effectively demonstrates the impact of estimation errors in factors and volatilities and the finite sample performance of our theoretical predictions.

We compare the finite sample performance of our portfolio estimators across various DGPs. We examine a broad spectrum of signal strength (μ/σ) and sparsity (ρ) values to cover a range of potential empirical scenarios. For each set of parameters $(\mu/\sigma, \rho)$, we simulate the corresponding DGP, construct portfolio weights, $\hat{w}^A$, where A represents OPT, CSR, BH, or Lasso, and compute the corresponding theoretical Sharpe ratio: $\hat{S}^A = \hat{w}^{A\top} \mu / \sqrt{\hat{w}^{A\top} \Sigma_u^{-1} \hat{w}^A}$.

Our algorithm requires a tuning parameter. We choose it with a validation procedure that divides the in-sample data into two parts: the first 80% is used for training, and the remaining 20% is reserved for validation. The optimal parameter maximizes the sample Sharpe ratio in this validation subset. Once identified, this tuning parameter is then used with the entire in-sample (training and validation) data to estimate portfolio weights. We use the same validation approach to determine the tuning parameter for the Lasso method.³¹

In light of Theorem 4, the following performance measure is sensible:

³¹The specific grid we use for tuning our method consists of the set $\{0.25, 0.5, \dots, 4\} \times (\log N)^{-1}$. For the Lasso method, the parameter λ is selected from the set $\{2^{-10}, 2^{-9}, \dots, 2^{-1}\}$. For our method, theoretical guidance from Theorem 4 helps — we know a tuning parameter $\simeq (\log N)^{-1}$ would be optimal. For Lasso, however, the optimal rate depends on the DGP, and we lack simple theoretical results to guide us.

$$\text{Gap}^A(\mu/\sigma, \rho) = \widehat{\mathbb{E}} \left(|\widehat{S}^A - S^{\text{OPT}}| / (1 + S^{\text{OPT}}) \right),$$

where the explicit dependence of $\widehat{S}^A$ and S^{OPT} on μ/σ and ρ is not specified, and $\widehat{\mathbb{E}}(\cdot)$ denotes the sample average computed over Monte Carlo simulations. This formula measures error as relative percentages of S^{OPT} when S^{OPT} is significantly large (i.e., significantly greater than 1). Conversely, when S^{OPT} is small (i.e., $o_P(1)$), the error is gauged in absolute terms. This approach ensures meaningful error measurement if S^{OPT} diverges or diminishes depending on the parameters involved.

To evaluate the asymptotic behavior of various estimators, we also report a measure of the estimation error:

$$\text{RMSE}^A(\mu/\sigma, \rho) = \left(\left(\widehat{\mathbb{E}}(\widehat{S}^A - S^A) / (1 + S^{\text{OPT}}) \right)^2 + \widehat{\text{Var}} \left(\widehat{S}^A / (1 + S^{\text{OPT}}) \right) \right)^{1/2},$$

where $\widehat{\text{Var}}(\cdot)$ denotes the sample variance across Monte Carlo repetitions.

Table 1 presents the maximum performance gap and estimation error across all values of μ/σ and ρ . The results consistently show that OPT has the smallest gap to optimal performance in comparison to CSR, BH, or Lasso. Lasso ranks second, while CSR and BH exhibit significantly poorer performance. Notably, as T increases from 2 years to 20 years, the maximum gap of OPT decreases from 0.381 to 0.152 for $N = 1,000$, and from 0.276 to 0.127 for $N = 3,000$. In contrast, as N increases, the performance of the latter three methods tends to deteriorate. This trend can be attributed to an increasing theoretical performance gap, quantified as $|S^A - S^{\text{OPT}}| / (1 + S^{\text{OPT}})$, that plays a dominating role. Indeed, the bottom panel of the table illustrate that the estimation error asymptotically diminishes with increasing N .

Finally, Figure 6 reports the sample average of estimation error $|\widehat{S}^* - S^*| / (1 + S^*)$ over Monte Carlo simulations for each value of μ/σ and ρ . The results confirm the consistency result in Proposition 1. The relative error remains small when S^* is large or moderate (significantly greater than 1). Conversely, near the bottom right corner of Figure 6, where S^* is essentially zero as shown by Figures 1 and 2, the estimation proves more challenging. In this scenario, the error, becoming absolute ($S^* \ll 1$), is notably larger due to a significant upward bias that is difficult to mitigate in finite sample.

4 Empirical Analysis of US Equities

To demonstrate the empirical relevance of the statistical limit of arbitrage, we analyze US monthly equity returns. Our analysis is divided into two parts: the first focuses on individual equity returns, and the second utilizes portfolios as test assets. Employing portfolios allows us

	$N = 1,000$				$N = 3,000$			
	$T = 2$	$T = 5$	$T = 10$	$T = 20$	$T = 2$	$T = 5$	$T = 10$	$T = 20$
	$\sup_{\mu/\sigma, \rho} \text{Gap}^A(\mu/\sigma, \rho)$							
OPT	0.381	0.243	0.171	0.152	0.276	0.228	0.143	0.127
CSR	0.622	0.586	0.532	0.477	0.462	0.604	0.610	0.562
BH	0.739	0.759	0.743	0.704	0.781	0.822	0.814	0.791
Lasso	0.396	0.259	0.176	0.154	0.298	0.265	0.159	0.172
	$\sup_{\mu/\sigma, \rho} \text{RMSE}^A(\mu/\sigma, \rho)$							
OPT	0.521	0.445	0.441	0.446	0.333	0.346	0.273	0.236
CSR	128	0.165	0.211	0.264	0.093	0.075	0.086	0.105
BH	0.492	0.436	0.440	0.442	0.359	0.347	0.263	0.248
Lasso	0.521	0.435	0.439	0.444	0.336	0.356	0.302	0.241

Table 1: Simulation Results

Note: The top panel displays the maximum Sharpe Ratio gap, defined as $\sup_{\mu/\sigma, \rho} \text{Gap}^A(\mu/\sigma, \rho)$, across all values of μ/σ and ρ shown in Figure 1. Similarly, the lower panel details the maximum root-mean-squared error, defined as $\sup_{\mu/\sigma, \rho} \text{RMSE}^A(\mu/\sigma, \rho)$. Here, A represents the methods OPT, CSR, BH, or Lasso for various combinations of N and T (measured in years). The OPT and Lasso methods employ a validation procedure to determine the optimal tuning parameters. The BH method ensures control of the false discovery rate at a 5% level.




Figure 6: Comparison between $\hat{S}^*$ and S^*

Note: The figure reports the sample average of $|\hat{S}^* - S^*|/(1 + S^*)$ over Monte Carlo repetitions. The simulation setting is based on model (5), in which a $100 \times \rho\%$ of assets have α s that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$. In this experiment, $N = 1,000$ and $T = 20$ years.

to explore scenarios where alphas are linear functions of characteristics, aligning our approach with the common empirical practice of predicting alphas based on such characteristics. We

begin our discussion with a description of the datasets.

4.1 US Equity Data

Our monthly equity sample covers the period from January 1965 to December 2020. For individual equities, we utilize a multi-factor model with observable factor loadings, closely resembling the widely used MSCI Barra model. Specifically, we use 16 characteristics and 11 GICS sectors, drawing on both empirical insights from existing asset pricing literature and industry practice. The selected characteristics include market beta, size, operating profits/book equity, book equity/market equity, asset growth, momentum, short-term reversal, industry momentum, illiquidity, leverage, return seasonality, sales growth, accruals, dividend yield, tangibility, and idiosyncratic risk, which are downloaded directly from the website openassetpricing.com (Stock-level Signal Datasets August 2023 Release). Details on the construction of these characteristics can be found in [Chen and Zimmermann \(2020\)](#).

We download monthly return data for individual equities from CRSP and apply several preprocessing steps. First, in the case of delistings, we use the delisting return as the final return in the delisting month. Next, we merge the returns data with the aforementioned characteristics database using permnos, which yields an average of 6,540 unique permnos per month. We then apply standard filters (require share codes 10 and 11, and exchange codes 1, 2, and 3) to refine the dataset. This process selectively removes returns for certain months from stocks that fail to meet these criteria during those periods. After applying these filters, the average number of stocks remaining per month totals 4,756.

We address missing characteristics in our dataset with an approach that avoids forward-looking bias. If GICS codes are missing, we use the most recent records available before the data is missing, but we do not copy records backwards in time. Observations that do not have a GICS code after this procedure are excluded, which predominantly affects records before 1990. With GICS codes in place, we implement a two-step procedure to address other missing characteristics. We fill each missing characteristic with the sector-wise median for that characteristic each month. If a characteristic’s values are missing for an entire sector in a given month, we use the cross-sectional median from all stocks for that month. After these preprocessing steps, the average number of stocks per month is reduced to 4,095.

Characteristics-sorted portfolio returns are directly downloaded from the website openassetpricing.com under the “Portfolio Return Datasets” section. The dataset comprises 1,322 portfolios, spanning the years from 1965 to 2020. The portfolios are long-only and sorted based on 212 characteristics, including 49 industry indicators. The number of portfolios per characteristic varies depending on the specifications in the original papers that introduced each characteristic. Additionally, some portfolios may contain missing values due to a lack of

observations in certain categories. We use value-weighted portfolios to avoid high exposure to illiquid small stocks.

4.2 Analysis of Individual Equity Returns

We consider a specific version of the model as defined in (2), setting α_t as constant and β_t as observable. Each month, we regress next month's returns on the 27 predictors, using all stocks present in the current month's cross-section, and including an intercept. We normalize the 16 characteristics within each cross-section to mitigate the impact of extreme outliers. This normalization transforms the characteristics to follow a normal distribution.³²

Figure 7 plots the time series of the cross-sectional regression R^2 s over time. The R^2 has declined since the beginning of the sample until the 1990s. This coincides with an increase in the number of stocks in US equity markets. The R^2 s are moderately low, with an average of 8.07%, which suggests that a substantial portion of cross-sectional variation of individual equity returns is idiosyncratic. Therefore, learning alphas from residuals of the factor model is a difficult statistical task.




Figure 7: Time-series of the Cross-sectional R^2 s

Note: The figure provides a time series of the cross-sectional R^2 s for individual equity returns, derived from regressing next month's returns against 27 firm-level characteristics in cross-sectional regressions.

³²For each characteristic $c_{i,t}$, the normalization applies the functional form $\Phi^{-1}(\text{rank}(c_{i,t}))$, where $\Phi^{-1}(\cdot)$ is the inverse of the cumulative distribution function of the standard normal distribution. Compared to uniform normalization, this approach leads to slightly improved cross-sectional R^2 values, thereby providing a more accurate representation of realized returns. It is important to note, however, that the specific form of normalization does not significantly impact subsequent results on arbitrage Sharpe ratios.

4.2.1 Rare and Weak Alphas

We now study the statistical properties of individual equity alphas using the full sample data. For each stock, we collect its regression residuals and take their average as an estimate for its alpha. We require at least 60 observations. This ensures a sufficiently large sample size for inference on alpha, although the empirical distribution of alphas' t-statistics turns out to be insensitive to this requirement. Figure 8 provides histograms of the t-statistics and Sharpe ratios for alphas of all 12,734 stocks in our sample that meet this criterion. Because these stocks have different sample sizes, the histograms of the Sharpe ratios are not simply the scaled version of the histogram of the t-statistics.

Only 7.58% of the t-statistics exceed 2.0 in magnitude, and more than 1.12% exceed 3.0. This suggests that truly significant alphas are extremely rare. Moreover, the largest Sharpe ratio among all individual stocks' alphas is approximately 1.86. Notably, only 0.71% of the alphas have a Sharpe ratio greater than 1.0. These summary statistics suggest that rare and weak alpha is perhaps the most relevant scenario in practice.




Figure 8: Histograms of the t-Statistics and Sharpe Ratios of Estimated Alphas

Note: The figure provides the histograms of the Student t-statistics (left) and Sharpe ratios (right) of estimated alphas for all tickers in our sample with at least 60 months of data. The total number of tickers available is 12,734.

Alphas are meaningless without reference to a specific factor model. While our analysis includes only 27 firm-level characteristics, constructing a factor model with additional characteristics would transform “alpha” into risk premia. Put it differently, extracting additional “factors” from returns would result in even rarer and weaker alphas.

4.2.2 Modest Feasible vs. Large Infeasible Sharpe Ratios

We now compare arbitrage portfolios based on various strategies discussed in Section 2.7. At the end of each month, we construct optimal portfolio weights using these strategies. These weights are estimated monthly with a 60-month rolling window, and the portfolios are rebalanced accordingly.³³ Both Lasso and OPT methods require a tuning parameter, which is selected annually. For this purpose, the final year of the rolling window is set aside as the validation sample to optimize the tuning parameter selection.

All these strategies yield modest Sharpe ratios. OPT reaches the top of the chart, yielding 0.82, followed by CSR at 0.70. The BH and Lasso approaches obtain Sharpe ratios of 0.65 and 0.62, respectively. To compare cumulative returns, we normalize all strategies to have the same (ex-post) volatility. The resulting time-series of normalized cumulative returns are shown in Figure 14. Notably, most returns are generated during the late 1980s to early 2000s. After 2005, the cumulative returns of all strategies tend to plateau, indicating a decrease in profitability of statistical arbitrage.

A detailed examination of these strategies offers further insights. The BH strategy is highly conservative; over 51 years of out-of-sample trading (from January 1970 to December 2020), there are 284 months with zero holdings. The maximum number of stocks held in any month is 32, with an average of only 6 stocks during months when holdings are non-zero. In contrast, the CSR and OPT strategies hold all stocks that meet the sample criteria, averaging 2,833 stocks per month. The number of stocks held by the Lasso strategy is notably volatile, ranging from none to almost all stocks in a given month, with an average of 722 stocks per month. This volatility reflects the underlying weakness in their alphas, whose estimates are prone to fluctuate with different samples.

In contrast to the simulation study where performance of strategies are evaluated over a large class of DGPs, results presented in the empirical analysis here all originate from a single DGP, i.e. the real market. Given the strong performance of CSR and the weaker performance of BH, it is likely that the alphas in reality are rare and weak to the extent that CSR is rather close to the optimal strategy. In this case, Lasso could only perform equally well as CSR if we were able to select the optimal tuning, i.e., zero, (no selection or shrinkage), which cross-validation cannot reliably achieve in all rolling windows. If we lived in a market where its DGP makes CSR suboptimal in many rolling windows, Lasso would have a chance to outperform, given its ability to achieve optimality within a larger class of DGPs.

We also estimate the perceived or infeasible Sharpe ratios using (21) within the same 60-month rolling window. Our estimate, $\tilde{S}^*$, averages 4.81 (with negative estimates truncated

³³Empirical results based on 36- and 84-month windows are reported in Table A2 in the appendix for robustness checks.




Figure 9: Normalized Cumulative Returns of Arbitrage Portfolios

Note: This figure compares the cumulative returns of OPT (black solid), CSR (red dotted), BH (green dashed), and Lasso (blue dot-dashed) strategies. We normalize all returns by their realized volatilities calculated by the square root of the sum of the squared returns over the entire sample, only for comparison purpose.

at 0), but it can occasionally exceed 16.0. These estimates are much higher than the feasible Sharpe ratios ($0.6 \sim 0.8$) we obtain for any of these strategies. This supports our theoretical prediction that the impact of statistical learning alone can eliminate a substantial portion of investment opportunities.

4.3 Analysis of Portfolio Returns

When portfolios are chosen as test assets, we opt for latent factor models due to the absence of natural choices of their factors or factor exposures. This approach aligns with the return generating process presented in (1) but with unknown β and v_t . Our factor model estimation step follows Algorithm 4 outlined in Giglio et al. (2021). We estimate this model using a 60-month rolling window, consistent with the main empirical analysis above.³⁴ Using the estimated factor loadings from each rolling window, we then conduct a cross-sectional regression of the average returns over this window against these loadings, augmented with a column of ones, to determine the risk premia of the latent factors, the zero-beta rate, and the residuals, which serve as estimates for α 's. It is important to note that these loadings

³⁴Results based on alternative 3-year and 7-year windows, which show similar outcomes, are reported in Tables A1 and A2 in the appendix for reference.

are identifiable only up to a rotation; hence, we report only those results that are invariant to such rotations, for instance, R^2 values and summary statistics for α 's.

4.3.1 Portfolio Alphas

These latent factor models demonstrate substantial in-sample cross-sectional explanatory power for expected returns,³⁵ accounting for an average of approximately 48% of the cross-section variation in average returns across all rolling windows, with 10 factors.³⁶ This results in notably smaller Sharpe ratios for the alphas compared to the portfolios' own Sharpe ratios. Figure 10 contrasts the histograms of their respective Sharpe ratios. Notably, the distribution of the alphas' Sharpe ratios in darker gray exhibits thinner tails compared to those of the portfolios' Sharpe ratios in lighter gray.

We then apply the same algorithms to the portfolios' alphas that we previously used for individual stocks. The normalized cumulative returns are displayed in Figure 11. All strategies earn substantially higher Sharpe ratios than those achieved with individual stocks. Notably, the OPT strategy once again leads, yielding a Sharpe ratio of 1.71, closely followed by CSR at 1.67. Meanwhile, the BH and Lasso strategies register Sharpe ratios of 1.49 and 1.18, respectively.

Next, we visualize the average portfolio weights over time using a heatmap in Figure 12. The portfolios in this figure are sorted by the average weight they receive with the OPT method, with the highest positive weights on top and the biggest negative weights at the bottom. The figure shows that both OPT and CSR exhibit similar allocation patterns, with non-zero weights assigned to many portfolios. However, OPT, which applies shrinkage to portfolio alphas, tends to assign smaller weights than CSR. Due to the absence of shrinkage, CSR has consistently higher weights than all other methods. As previously mentioned, CSR achieves the same Sharpe ratio as ridge, because uniform shrinkage of all assets' alphas does not affect the Sharpe ratio. BH selects a narrower range of portfolios, most of which have zero or near-zero weights, indicating a highly selective strategy. Lasso, although less conservative than BH, also shows lighter colors on the heatmap, reflecting a significant level of selection and shrinkage in the portfolio weights.

Among all characteristics, those assigned higher positive weights by these methods include portfolios sorted by dividend seasonality identified by [Hartzmark and Solomon \(2013\)](#), price as discussed in [Blume and Husic \(1973\)](#), intangible return from [Daniel and Titman \(2006\)](#), gross profits to total assets analyzed by [Novy-Marx \(2013\)](#), and industry return of big firms

³⁵This is different from the cross-section R^2 with observable characteristics reported for individual equities, which represents the (in-sample) explanatory power for realized returns.

³⁶Results for models with varying numbers of factors are included in Tables A1 and A2 in the appendix for comparison.




Figure 10: Sharpe Ratios of Portfolio Returns and Alphas

Note: This figure compares the histograms of the Sharpe ratios of portfolios' returns and their alphas.




Figure 11: Normalized Cumulative Returns of Arbitrage Portfolios

Note: This figure compares the cumulative returns of OPT (black solid), CSR (red dotted), BH (green dashed), and Lasso (blue dot-dashed) strategies. We normalize all returns by their realized volatilities calculated by the square root of the sum of the squared returns over the entire sample, only for comparison purpose. The base assets for each strategy are all portfolios sorted by characteristics.

studied in [Hou \(2007\)](#). Conversely, the characteristics with the most negative weights in-




Figure 12: Average Weights Assigned to Portfolios by Different Methods

Note: This figure displays a heatmap that shows the average portfolio weights over time for portfolios sorted by characteristics, formed using BH, CSR, Lasso, and OPT methods.

clude portfolios sorted by spinoffs documented by [Cusatis et al. \(1993\)](#), past trading volume from [Brennan et al. \(1998\)](#), initial public offerings from [Ritter \(1991\)](#), size from [Dharan and Ikenberry \(1995\)](#), and exchange switch as in [Banz \(1981\)](#).

4.3.2 Accounting for Publication Effects

Many of the characteristics used in the construction of the portfolio return data set were introduced to the academic literature later in the sample period. For this reason, the above analysis may overstate achievable Sharpe ratios. If arbitrageurs were not aware of the predictive power of these characteristics in the early parts of the sample, these Sharpe ratios were not truly achievable. Moreover, academics' discovery of the predictive power of these characteristics could potentially reflect ex-post selection bias, rendering the in-sample predictive power spurious. To address this issue, we revise our portfolio analysis to only include characteristics from the year following their publication. Except for the 49 industry portfolios (assumed) available since 1965, few portfolios were introduced in earlier years. As illustrated in [Figure 13](#), a significant number of these characteristics were introduced after 2000.

With this adjusted approach, the Sharpe ratios for the OPT and CSR are now substantially lower at 0.61 and 0.66, respectively. These magnitudes align more closely with the Sharpe ratios we obtained from individual stocks. The Lasso and BH strategies now achieve only 0.45 and 0.35, respectively. Consistent with [McLean and Pontiff \(2016\)](#), this suggests

that some of the initial discoveries of signals may have been “lucky,” performing well in-sample but then failing to replicate this success out-of-sample, or that competition among arbitrageurs shrinks the alphas once the discoveries become public.




Figure 13: Number of Portfolios Before and After Adjusting for Publication Years

Note: This figure plots two curves: one showing the total monthly count of characteristics-sorted portfolios and the other displaying the count adjusted for only those portfolios whose characteristics have been previously published.

Figure 15 shows infeasible Sharpe ratios estimated both before and after adjusting for publication effects. Before this adjustment, the infeasible Sharpe ratio is consistently above 20, but after this adjustment, including characteristics only after their publication year, it is much lower in the early part of the sample, starting around 5 in the early 1970s. As more signals are discovered over time, the infeasible Sharpe ratio increases, eventually reaching levels similar to those estimated without adjustment.

Despite these adjustments, a significant gap still exists between the theoretical infeasible Sharpe ratios that range from 5 to 20 and the achievable Sharpe ratios, which are smaller than 0.7. This highlights the quantitative relevance of statistical limits of arbitrage in real-world scenarios. The large reduction in Sharpe ratios suggests that statistical limits may play an even more critical role than publication effects for characterizing the investment opportunities that are actually feasible for arbitrageurs.




Figure 14: Normalized Cumulative Returns of Arbitrage Portfolios

Note: This figure compares the cumulative returns of OPT (black solid), CSR (red dotted), BH (green dashed), and Lasso (blue dot-dashed) strategies. We normalize all returns by their realized volatilities calculated by the square root of the sum of the squared returns over the entire sample, only for comparison purpose. The base assets only include portfolios sorted by characteristics that were previously published.




Figure 15: Time Series Plots of Infeasible Sharpe Ratios

Note: This figure compares the Infeasible Sharpe ratios on the basis of portfolios before and after adjusting for publication years.

5 Conclusion

Taking stock, our paper provides a new theoretical framework for understanding the implications of statistical learning in asset pricing. Rather than endowing arbitrageurs with perfect knowledge of DGP parameters, we expose them to the challenge of learning about alphas. Arbitrageurs in reality arguably face a high-dimensional world, with many assets and many alpha signals, which can make statistical learning difficult. These difficulties give rise to statistical limits of arbitrage, which manifest as a gap between the infeasible Sharpe ratio that arbitrageurs could earn with perfect knowledge of alphas and the feasible Sharpe ratio that arbitrageurs can attain with statistical learning.

While high-dimensional, the linear factor model environment that arbitrageurs face in our analysis is still quite simple. The gap between feasible and infeasible Sharpe ratios would further increase if arbitrageurs faced additional statistical challenges, e.g., model misspecification, omitted factors, weak factors, or a large non-sparse idiosyncratic covariance matrix.

Our approach to characterizing the feasible Sharpe ratio is in important ways different from other analyses of statistical learning in asset pricing. Many papers applying machine learning methods in the construction of trading strategies have documented impressive Sharpe ratios. Such strategies often rely on ad-hoc model design (e.g., a neural network with a specific architecture) and tuning parameters selection. In this regard, the empirical analysis can at best provide a “lower bound” on the performance of machine learning strategies. Our paper provides a theoretical framework to understand the “upper bound” on the performance of any statistical learning strategy in the specific context of arbitrage pricing theory. In an equilibrium model, this upper bound in turn would be tied to the compensation that arbitrageurs require in equilibrium, giving it an important economic interpretation.

In addition to the optimal strategy that attains the feasible Sharpe ratio, we also examine, theoretically and empirically, the performance of other portfolio construction approaches that rely multiple testing correction, variable selection, and shrinkage. While these approaches are well motivated statistically, they do not perform as well as the optimal strategy in economic terms, as measured by the Sharpe ratio. This finding highlights that good statistical properties, such as, for example, with regards to Type I and Type II errors or the false discovery rate, do not necessarily translate into good economic performance. For instance, a statistical procedure that guards against false discoveries may be overly conservative for investment purposes. This divergence between statistical and economic objectives suggests some caution when applying statistical tools imported from other areas—as in the current machine learning literature in asset pricing—without proper consideration of the economic objectives.

References

- Andrews, D. W. K., X. Cheng, and P. Guggenberger (2020). Generic results for establishing the asymptotic size of confidence sets and tests. *Journal of Econometrics* 218(2), 496–531.
- Andrews, R. L., J. C. Arnold, and R. G. Krutchkoff (1972). Shrinkage of the posterior mean in the normal case. *Biometrika* 59(3), 693–695.
- Avramov, D. and G. Zhou (2010). Bayesian portfolio analysis. *Annual Review of Financial Economics* 2(1), 25–47.
- Bajgrowicz, P. and O. Scaillet (2012, December). Technical trading revisited False discoveries, persistence tests, and transaction costs. *Journal of Financial Economics* 106(3), 473–491.
- Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of financial economics* 9(1), 3–18.
- Barras, L., O. Scaillet, and R. Wermers (2010). False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance* 65(1), 179–216.
- Bekker, P. A. (1994). Alternative approximations to the distributions of instrumental variable estimators. *Econometrica* 62(3), 657–681.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57(1), 289–300.
- Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer.
- Black, F. and R. Litterman (1992). Global portfolio optimization. *Financial Analysts Journal* 48(5), 28–43.
- Blume, M. E. and F. Husic (1973). Price, beta, and exchange listing. *The Journal of Finance* 28(2), 283–299.
- Brennan, M. J., T. Chordia, and A. Subrahmanyam (1998). Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of financial Economics* 49(3), 345–373.
- Brown, L. D. and E. Greenshtein (2009). Nonparametric empirical Bayes and compound decision approaches to estimation of a high-dimensional vector of normal means. *The Annals of Statistics* 37(4), 1685 – 1704.

- Chamberlain, G. and M. Rothschild (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica* 51, 1281–1304.
- Chen, A. Y. and T. Zimmermann (2020). Open source cross-sectional asset pricing. *Available at SSRN*.
- Chen, X., L. P. Hansen, and P. G. Hansen (2021a). Robust identification of investor beliefs. *Proceedings of the National Academy of Sciences* 117(52), 33130–33140.
- Chen, X., L. P. Hansen, and P. G. Hansen (2021b). Robust inference for moment condition models without rational expectations. *Journal of Econometrics*, forthcoming.
- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review* 106(3), 664–698.
- Connor, G., M. Hagmann, and O. Linton (2012). Efficient semiparametric estimation of the fama-french model and extensions. *Econometrica* 80(2), 713–754.
- Connor, G. and R. A. Korajczyk (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics* 15(3), 373–394.
- Cusatis, P. J., J. A. Miles, and J. R. Woolridge (1993). Restructuring through spinoffs: The stock market evidence. *Journal of financial economics* 33(3), 293–311.
- Daniel, K. and S. Titman (2006). Market reactions to tangible and intangible information. *The Journal of Finance* 61(4), 1605–1643.
- Dharan, B. G. and D. L. Ikenberry (1995). The long-run negative drift of post-listing stock returns. *The Journal of Finance* 50(5), 1547–1574.
- Donoho, D. and J. Jin (2004). Higher criticism for detecting sparse heterogeneous mixtures. *Annals of Statistics* 32(3), 962–994.
- Efron, B. (2011). Tweedie’s formula and selection bias. *Journal of the American Statistical Association* 106(496), 1602–1614.
- Efron, B. (2019). Bayes, Oracle Bayes and Empirical Bayes. *Statistical Science* 34(2), 177 – 201.
- Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33(1), 3–56.

- Fan, J., Y. Liao, and J. Yao (2015). Power enhancement in high-dimensional cross-sectional tests. *Econometrica* 83(4), 1497–1541.
- Frost, P. A. and J. E. Savarino (1986). An empirical bayes approach to efficient portfolio selection. *The Journal of Financial and Quantitative Analysis* 21(3), 293–305.
- Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity datasets. *Econometrica* 84(3), 985–1046.
- Gibbons, M. R., S. A. Ross, and J. Shanken (1989). A test of the efficiency of a given portfolio. *Econometrica* 29, 1121–1152.
- Giglio, S., B. Kelly, and D. Xiu (2022). Factor models, machine learning, and asset pricing. *Annual Review of Financial Economics* 14, 337–368.
- Giglio, S., Y. Liao, and D. Xiu (2021). Thousands of alpha tests. *Review of Financial Studies* 34(7), 3456–3496.
- Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy* 129(7), 1947–1990.
- Gromb, D. and D. Vayanos (2010). Limits of arbitrage. *Annual Review of Financial Economics* 2, 251–275.
- Guijarro-Ordóñez, J., M. Pelger, and G. Zanolini (2022). Deep learning statistical arbitrage. Technical report, Stanford University.
- Hansen, L. P. (2007). Beliefs, doubts, and learning: Valuing macroeconomic risk. *American Economic Review* 97(2), 1–30.
- Hansen, L. P. (2014). Nobel lecture: Uncertainty outside and inside economic models. *Journal of Political Economy* 122(51), 945–987.
- Hartzmark, S. M. and D. H. Solomon (2013). The dividend month premium. *Journal of Financial Economics* 109(3), 640–660.
- Harvey, C. R. and Y. Liu (2020). False (and missed) discoveries in financial economics. *Journal of Finance*, forthcoming.
- Harvey, C. R., Y. Liu, and H. Zhu (2016). ... and the cross-section of expected returns. *Review of Financial Studies* 29(1), 5–68.

- Hou, K. (2007). Industry information diffusion and the lead-lag effect in stock returns. *The review of financial studies* 20(4), 1113–1138.
- Huberman, G. (1982). A simple approach to arbitrage pricing theory. *Journal of Economic Thoery* 28(1), 183–191.
- Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica* 72(6), 1845–1857.
- Ingersoll, J. E. (1984). Some results in the theory of arbitrage pricing. *Journal of Finance* 39(4), 1021–1039.
- Johns, M. V. (1957). Non-parametric empirical bayes procedures. *Annals of Mathematical Statistics* 28, 649–669.
- Jorion, P. (1986). Bayes-stein estimation for portfolio analysis. *The Journal of Financial and Quantitative Analysis* 21(3), 279–292.
- Kan, R., X. Wang, and G. Zhou (2022). Optimal portfolio choice with estimation risk: No risk-free asset case. *Management Science*, *forthcoming*.
- Kan, R. and G. Zhou (2007). Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42(3), 621–656.
- Kelly, B., S. Pruitt, and Y. Su (2019). Some characteristics are risk exposures, and the rest are irrelevant. *Journal of Financial Economics*, *forthcoming*.
- Kim, S., R. Korajczyk, and A. Neuhierl (2020). Arbitrage portfolios. *Review of Financial Studies*, *forthcoming*.
- Kozak, S. and S. Nagel (2023). When do cross-sectional asset pricing factors span the stochastic discount factor? Technical report, National Bureau of Economic Research.
- Leeb, H. and B. M. Pötscher (2005). Model selection and inference: Facts and fiction. *Econometric Theory* 21(1), 21–59.
- Martin, I. and S. Nagel (2021). Market efficiency in the age of big data. *Journal of Financial Economics*, *forthcoming*.
- McLean, R. D. and J. Pontiff (2016). Does academic research destroy stock return predictability? *The Journal of Finance* 71(1), 5–32.

- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of financial economics* 108(1), 1–28.
- Pastor, L. (2000). Portfolio selection and asset pricing models. *The Journal of Finance* 55(1), 179–223.
- Pastor, L. and R. F. Stambaugh (2000). Comparing asset pricing models: An investment perspective. *Journal of Financial Economics* 56, 335–381.
- Pastor, L. and P. Veronesi (2009). Learning in financial markets. *Annual Review of Financial Economics* 1(1), 361–381.
- Pesaran, H. and T. Yamagata (2017). Testing for alpha in linear factor pricing models with a large number of securities. Technical report.
- Ritter, J. R. (1991). The long-run performance of initial public offerings. *The journal of finance* 46(1), 3–27.
- Robbins, H. (1956). An empirical bayes approach to statistics. *Berkeley Symposium on Mathematical Statistics and Probability* 3, 157–163.
- Robertson, T., R. Dykstra, and F. Wright (1988). Order restricted statistical inference. *New York: Wiley*.
- Rosenberg, B. (1974). Extra-market components of covariance in security returns. *Journal of Financial and Quantitative Analysis* 9(2), 263–274.
- Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13(3), 341–360.
- Shanken, J. (1992). The current state of the arbitrage pricing theory. *Journal of Finance* 47(4), 1569–1574.
- Staiger, D. and J. H. Stock (1997). Instrumental variables regression with weak instruments. *Econometrica* 65(3), 557–586.
- Tu, J. and G. Zhou (2010). Incorporating economic objectives into bayesian priors: Portfolio choice under parameter uncertainty. *The Journal of Financial and Quantitative Analysis* 45(4), 959–986.
- Tukey, J. W. (1976). The higher criticism. course notes, statistics 411. Technical report, Princeton University.

Zhang, C.-H. (1997). Empirical bayes and compound estimation of normal means. *Statistica Sinica* 7(1), 181–193.