The Statistical Limit of Arbitrage*

Rui Da[†]
University of Chicago

Stefan Nagel[‡]
University of Chicago, NBER, and CEPR

Dacheng Xiu[§]
University of Chicago

This version: July 12, 2022

Abstract

When alphas are weak and rare, and arbitrageurs have to learn about alphas from historical data, there is a gap between Sharpe ratio that is feasible for them to achieve and the infeasible Sharpe ratio that could be obtained with perfect knowledge of the true return distribution. This statistical limit to arbitrage widens the bounds within which true alphas can survive in equilibrium relative to the arbitrage pricing theory (APT) in which arbitrageurs are endowed with perfect knowledge of the return distribution. We derive the optimal Sharpe ratio achievable by any feasible arbitrage strategy, and illustrate in a simple model how this Sharpe ratio varies with the strength and sparsity of alpha signals, which characterize the difficulty of arbitrageurs' learning problem. Furthermore, we design an "all-weather" arbitrage strategy that achieves this optimal Sharpe ratio regardless of the conditions of alpha signals. We also show how arbitrageurs can adopt multiple-testing, LASSO, and Ridge methods to achieve optimality under distinct conditions of alpha signals, respectively. Our empirical analysis of equity returns shows that all strategies we consider achieve a moderately low Sharpe ratio out of sample, in spite of a considerably higher infeasible Sharpe ratio, consistent with absence of feasible near-arbitrage opportunities and relevance of statistical limits to arbitrage.

Keywords: Learning about Alphas, Rational Expectation, Portfolio Choice, Rare and Weak Signal, False Negatives, testing APT, Machine Learning

*This version is preliminary and incomplete. Please do not circulate. We benefited tremendously from discussions with seminar and conference participants at Princeton University, Oxford University, EPFL, University of Liverpool, Sao Paulo School of Economics, Peking University, Tsinghua University, Shanghai University of Finance and Economics, University of Gothenburg, and Stockholm Business School.

[†]Address: 5807 S Woodlawn Avenue, Chicago, IL 60637, USA. E-mail address: rui.da@chicagobooth.edu.

[‡]Address: 5807 S Woodlawn Avenue, Chicago, IL 60637, USA. E-mail address: stefan.nagel@chicagobooth.edu.

[§]Address: 5807 S Woodlawn Avenue, Chicago, IL 60637, USA. E-mail address: dacheng.xiu@chicagobooth.edu.

1 Introduction

It is a fundamental underlying principle of most asset pricing theories, including the Arbitrage Pricing Theory (APT), that investment opportunities with extremely high ratios of reward to risk do not exist in financial markets. Implicitly, these theories rest on the premise that such near-arbitrage opportunities would attract arbitrageurs who exploit and thereby eliminate these opportunities. An important assumption in these theories is that the statistical distribution of returns is known to arbitrageurs. Therefore, near-arbitrage opportunities in the true distribution of returns are ruled out.

In practice, however, sophisticated investors searching for near-arbitrage opportunities do not know this distribution. Instead, they commonly conduct statistical analyses to learn about the existence of such opportunities from historical returns data. As a consequence, they face statistical uncertainty. In some settings, such as in some derivatives pricing applications, for instance, the statistical uncertainty may be sufficiently small that it is not a significant impediment to arbitrageur activity. But in noisy, high-dimensional settings such as the cross-section of stock returns, statistical uncertainty can be substantial and it can can constitute a statistical limit to arbitrage.

To analyze the effects of arbitrageur learning, we consider a setting in which returns follow a statistical linear factor model. Near-arbitrage opportunities are characterized by high Sharpe ratios. To exploit such opportunities, arbitrageurs need knowledge of factor model alphas, but they must learn about these from historical realizations of alpha signals. We derive the optimal Sharpe ratio achievable by any feasible arbitrage trading strategies, 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-near-feasible-arbitrage bound on the Sharpe ratio that accounts for the statistical limit to arbitrage.

The difficulty of the learning problem hinges on the data generating process (DGP) of alpha signals. While our theory generally does not rely on specific cross-sectional distributions of alpha signals, we use simple special cases to demonstrate 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 learn the distribution of alpha perfectly in the limit. But when alpha is weaker and more rare, its inference becomes more challenging and a gap arises 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 in the asymptotically, while the feasible Sharpe ratio stays bounded.

The existence of this statistical limit to arbitrage implies a widening of the bounds in which mispricing can survive in equilibrium compared with a situation in which arbitrageurs know the DGP and its parameters. Some mispricing may survive it is clouded by too much statistical uncertainty. Empirically therefore, the feasible, not the infeasible, Sharpe ratio tells us about the minimum reward-to-risk compensation that 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 reality arbitrageurs do not know which DGP is a correct description of the observed data. The optimal strategy estimates the empirical distribution of alpha signals and assigns weights based on the relative magnitudes and associated uncertainty of the alpha estimates. Assets with with high alpha t-statistics get portfolio weights proportional to their signal strength. Weaker alphas are more difficult to exploit, yet simply ignoring them would lead to a suboptimal trading strategy. The optimal strategy constructs portfolio weights for weak signals by locally smoothing alpha signals cross-sectionally.

To empirically contrast feasible and infeasible Sharpe ratios, we also 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 realized by any feasible portfolio with weights constructed using historical data. The infeasible Sharpe ratio often serves as the building block for tests of APT, see, e.g., **?**, **?**, **?**, and **?**. While such tests are powerful and may lead to discoveries of alpha signals, they are not relevant for arbitrageurs who are confined to feasible trading strategies. Our effort in constructing the optimal feasible arbitrage portfolio and evaluating its economic performance directly responds to Shanken's call (**?**): "... practical content is given to the notion of 'approximate arbitrage,' by characterizing the investment opportunities that are available as a consequence of the observed expected return deviation ... Far more will be learned, I believe, by examining the extent to which we can approximate an arbitrage with existing assets."

Next, we examine whether alternative strategies that exploit multiple testing, shrinkage, and selection techniques to build arbitrage portfolios can attain the optimal feasible Sharpe ratio. With alphas estimated from cross-sectional regressions, one strategy adopts a multiple-testing (BH) procedure as in **?** on the individual p-values of t-statistics for alpha, in order to guard against potential false discoveries among significant alphas, before building the optimal portfolio weights using selected alphas. Other strategies use either LASSO or Ridge penalties to regularize portfolio weights based on alpha estimates. Such strategies amount to imposing a prior distribution on the alphas. We illustrate with a simple example that these strategies can achieve optimal Sharpe ratio under distinct alpha assumptions. In particular, BH procedure achieves optimal performance only when few true alpha signals are substantially strong. Its failure to achieve optimality is precisely due to its conservativeness nature against the less potent alphas. In contrast, the ridge-based portfolio is equivalent to that constructed by alpha estimates from plain cross-sectional regressions. This approach can achieve optimality when almost all true alphas are either uniformly strong or uniformly weak. The LASSO approach attempts to strike a balance between the aforementioned two methods, with a small gap to achieving the theoretically optimal Sharpe ratio.

Finally, we demonstrate the empirical implications of the statistical limits of arbitrage by examining 56 years of monthly individual equity returns in US stock market from 1965 to 2020. The average number of stocks over this period exceeds 4000. We construct residuals via cross-sectional

regressions from a multi-factor model that directly uses observed characteristics as risk exposures. These 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 (**?**), as well as 11 Global industry Classification Standard (GICS) sectors. These characteristics and industry dummies capture similar equity factors in the MSCI Barra model widely-used among practitioners.

A few interesting findings emerge. First, the cross-sectional R^2s are rather low, with a time-series average 8.25% over our sample period from Jan 1965 to Dec 2020. These R^2s are in similar magnitudes compared to existing estimates in the literature, e.g., 7.8% average R^2s from May 1964 to Dec 2009 reported in **?** using 15 factors that largely overlap with ours, but lower than 12-14% over 1987 - 2016 reported in **?** based on latent factor models. This indicates that there exists a considerable amount of idiosyncratic noise in the cross-section of individual equities, which makes learning about alphas an arduous statistical task.

Second, we obtain the t-statistics corresponding to alpha estimates of all individual stocks based on their full record in our sample. Among 12,415 test statistics in total, only 6.35% (0.63%) of the t-statistics are greater than 2.0 (3.0) in absolute values. Only 0.505% of these t-statistics translate to Sharpe ratios greater than 1.0 in magnitude. Even without controlling for multiple testing, these estimates suggest that non-zero alphas are rather rare and weak.

Third, we find that the optimal feasible arbitrage portfolio with different methods achieve a moderately low annualized Sharpe ratio, about 0.5, whereas the infeasible Sharpe ratios over time are considerably higher—beyond 2.5—on average, and can reach as high as 7.5 for some sample periods. The estimated infeasible Sharpe ratio is an estimate of what arbitrageurs could attain if they had perfect knowledge of DGP parameters, but it is not attainable by constructing a feasible arbitrage portfolio. The large gap between feasible and infeasible Sharpe ratios suggests the empirical relevance of the statistical limit of arbitrage. Moreover, the fact that the feasible Sharpe ratio is small suggests the empirical success of APT.

Fourth, among all feasible strategies, BH and our optimal strategy achieve the best performance, around a Sharpe ratio of 0.5, followed by CSR (0.450) and LASSO (0.384), though the differences are not substantial.However, the BH approach is overly conservative that it eliminates almost all weak signals and trade less than 10 stocks each month, with zero trading activities for over half of the entire sample. CSR and out strategy exploit weak signals. CSR trades all stocks, but receive a slightly lower Sharpe ratio, potentially due to misallocation of portfolio weights to fake signals. Our optimal strategy trade almost all stocks, but with weights adaptive to the signal strength. LASSO is most disappointing, but this is due to the uncertainty in optimal tuning parameters. The resulting number of traded stocks per month varies considerably from none to all.

Our paper builds on a large literature on the arbitrage pricing theory (APT) developed by **?** and later refined by **?**, **?**, and **?**. As in these papers, we rely on asymptotic arguments that do not rely on assumptions about investor preferences, but these results should be seen as an asymptotic

approximation for a more realistic setting with a finite number of assets in which weak preference restrictions rule out Sharpe ratios far above the Sharpe ratios of diversified factor portfolios. The statistical limits to arbitrage that we highlight in this paper relax 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 another large strand of literature on the limit of arbitrage, see **?** for a comprehensive review. Complementary to the existing literature, the arbitrage limit in our setting stems from statistical uncertainty, instead of being induced from risk, costs, frictions, and other constraints rational expectation investors are facing.

? argue that the absence of near-arbitrage opportunities enforces the expected returns to approximately line up linearly with common factor covariances, even in a world in which belief distortions affect asset prices. Our study focuses on the deviations of expected returns from this approximate linear relation and how statistical limits to arbitrage allow bigger deviations. A closely related paper to ours is **?**, which proposes a characteristics-based factor model to construct arbitrage portfolios. Their model does not preclude arbitrage opportunities with a theoretically infinite Sharpe ratio, whereas our framework rules out such a possibility. Relatedly, **?** propose a methodology to construct robust portfolios that can be decomposed into alpha (arbitrage) portfolios and beta (factor) portfolios. Our setting is considerably different in that alphas cannot possibly be recovered with certainty even when the sample size is large.

Our paper is also related 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., **?**, **?**, and **?**, as well as on testing for alphas, e.g., **?**, **?**, and **?**. The first literature focus on testing a null that all alphas are equal to zero. This is certainly an interesting null hypothesis, but as we emphasize in this paper, the APT does allow for alphas as long as they do not induce an explosive feasible Sharpe ratio. The second literature focuses on detecting strong alphas, in which widely used multiple testing methods, such as the BH method by **?**, or its extensions can be applied to control the false discovery rate (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.[1] 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 undetected by multiple testing methods, they may lead to a portfolio with a considerable Sharpe ratio.

There has been a long-standing critique of rational expectation models in macroeconomics and finance in which economic agents are not confronted with statistical uncertainty over structure parameters, see **?**. Bayesian learning is one way to expose model agents to statistical uncertainty. **?** survey the literature on learning in financial markets. In many settings, e.g., **?**, learning can be sufficiently slow such that its effects persist in empirically realistic sample sizes, even though convergence to rational expectations takes place in the long-run. An exception is **?** where learning effects persist because investors face a high-dimensional inference problem about the process

[1] **?** adopt the so-called higher criticism approach, dating back to **?**, to detect rare and weak signals in a stylized multiple testing problem.

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.[2]

Last but not least, our paper is related to **?** and **?** 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 expectation to bound the set of subjective probabilities, and seek robust inference with this form of misspecification. We develop an optimal arbitrage strategy that is optimal with respect to a large class of DGPs. This Sharpe ratio is optimal for arbitrageurs inside the economic model, which is in contrast with the (infeasible) optimal Sharpe ratio from an outside econometrician's point of view, in the spirit of **?**. We provide estimators for both and compare them empirically.

Our paper proceeds as follows. Section 2 develops a statistical limit of arbitrage. Section 2.5 constructs a feasible trading strategy that achieves the optimal Sharpe ratio. Section 2.7 analyzes two alternative trading strategies. Section 3 provides simulation evidence, followed by an empirical analysis in Section 4. Section 5 concludes. The appendix provides technical details.

2 Statistical Limit of Arbitrage

We start by revisiting the arbitrage pricing framework developed by **?**. This setting is ideal for explaining the statistical limit of arbitrage because the arbitrage pricing theory is largely developed based on a reduced-form statistical model for asset returns. This stylized model is sufficiently sophisticated to deliver theoretical insight, and is sufficiently relevant to guide empirical investment decisions.

2.1 Factor Model Setup

To be more concrete, 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, for $t = 1, 2, \ldots, T$:

$$r_t = \alpha + \beta\gamma + \beta v_t + u_t, \tag{1}$$

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 factor innovations with covariance matrix Σ_v, γ is a $K \times 1$ vector of risk premia (and zero beta rate), and u_t is a vector of idiosyncratic returns, independent of v_t, with a diagonal covariance matrix Σ_u. While it is customary to consider the approximate factor model, allowing for off-diagonal entries in the covariance matrix Σ_u would intro-

[2] Our analysis is related to a large literature in econometrics and statistics that discuss uniform validity of asymptotic approximations, see, e.g., **?**, **?**, **?**, **?**.

duce additional statistical obstacles due to the estimation of large covariance matrix for inference on alpha and for building optimal portfolios. For the purpose of illustrating the economic insight of limits on arbitrage, we instead focus on the statistical learning problem of alpha signals, which is sufficiently involved and is of first-order importance for building arbitrage portfolios, as will become clear below.

Throughout we will consider asymptotic limits as N and T increase while K is fixed. To facilitate our asymptotic analysis along 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, since their dimensions are fixed. We assume that α has mean zero, and is cross-sectionally independent of β, and that β has full column rank and is pervasive. These conditions are essential for identification of γ in a model that allows for pricing errors.

We formalize the conditions below.

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

(a) $\|\beta\|_{\text{MAX}} \lesssim_{\text{P}} 1$, *and* $\lambda_{\min}(\beta^{\intercal}\beta) \gtrsim_{\text{P}} N$.[3]

(b) v_t *is i.i.d. across* t, $\mathrm{E}(v_t) = 0$, *and its covariance matrix* Σ_v *satisfies* $1 \lesssim \lambda_{\min}(\Sigma_v) \leq \lambda_{\max}(\Sigma_v) \lesssim 1$.

(c) α_i *is i.i.d. across* t, *with density* $p_\alpha(\cdot)$, $\mathrm{E}(\alpha) = 0$, *and* $\mathrm{E}\big(\max_{i:i\leq N} \alpha_i^2\big) = o(1)$.

(d) u_t *is i.i.d. across* t, $\mathrm{E}(u_t) = 0$, Σ_u *is diagonal, and* $\sigma_i^2 = (\Sigma_u)_{i,i}$ *is i.i.d. across* i, *independent with* (α, v_t, β), *and satisfies* $1 \lesssim_{\text{P}} \lambda_{\min}(\Sigma_u) \leq \lambda_{\max}(\Sigma_u) \lesssim_{\text{P}} 1$.

(e) The pricing errors α, *factors* v_t, *factor loadings* β, *and idiosyncratic error* u_t *are mutually independent.*

Assumption 1 (a) and (b) are commonly seen in the literature of factor models. In particular, the assumption on $\lambda_{\min}(\beta^{\intercal}\beta)$ requires that all factors are pervasive.[4] (c) and (d) suggest that the signals in our model are weak, in that as N increases their magnitudes shrink towards 0, whereas volatilities are bounded from above and from below. Condition (c) on α implies that $\mathrm{Var}(\alpha_i) = \mathrm{E}(\alpha_i^2) = o(1)$. As will become clear (from footnote 8), a diminishing variance on α is necessary for precluding near-arbitrage opportunities in Ross' APT.

There are at least three variations of the factor model (1), depending on what econometricians assume to be observable. The most common setup in academic finance literature imposes that

[3]For a matrix A, we use $\|A\|$ and $\|A\|_{\text{MAX}} = \max_{i,j} |a_{ij}|$ to denote the operator norm (or $\mathbb{L}_2$ norm) and the $\mathbb{L}_\infty$ norm of A on the vector space. We use C to denote a generic constant that may change from line to line. We use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote the minimum and maximum eigenvalues of A. We also use the notation $a \lesssim b$ to denote $a \leq Cb$ for some constant $C > 0$ and $a \lesssim_{\text{P}} b$ to denote $a = O_{\text{P}}(b)$.

[4]See, e.g., Assumption I.1 of **?**. While our theoretical results may extend to certain weak factor settings, this is not our emphasis here.

factors are observable as in e.g., **?**.[5] The second setting, which has gained more popularity since its debut in **?**, assumes that factors are latent. The third setting, arguably most prevalent among practitioners, is the MSCI Barra model originally proposed by **?**, where factor exposures, i.e., characteristics, are assumed observable. The advantage of the last model lies in the fact that estimating a large number of (potentially) time-varying stock-level factor exposures is statistically inefficient and computationally expensive, as opposed to directly specifying risk exposures as (linear functions of) observable characteristics.[6]

Our core theoretical results below (e.g., Theorem 1) directly apply to all three cases aforementioned. In our empirical analysis we will adopt the third framework most convenient for modeling individual stocks. This makes our analysis highly relevant for practitioners.

2.2 Feasible Near-Arbitrage Opportunities

Building upon the insight of **?**, **?** and **?** 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 $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathrm{P})$, if it satisfies $w \in \mathcal{F}_t$, and along some diverging subsequence,[7] with probability approaching one,*

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

Intuitively, no near-arbitrage means there exist no sequence of portfolios that earn positive expected returns with vanishing risks. Under conditions similar to those in Assumption 1, **?** established that a sufficient and necessary condition for the absence of near-arbitrage is that

$$S^\star = \sqrt{\alpha^\intercal \Sigma_u^{-1} \alpha} \lesssim_{\mathrm{P}} 1. \tag{3}$$

$S^\star$ is the theoretically optimal Sharpe ratio arbitrageurs can achieve in this economy without exposure to factor risks. 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

[5] This 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.

[6] Strictly speaking, the MSCI Barra model is cast in a conditional version of (1):

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

where β_t is a vector of observed characteristics and γ_{t-1} is a vector of time-varying risk premia. Analyzing this conditional model will not yield additional economic insight relative to the unconditional model with respect to the theoretical limit of arbitrage. This model is overly parametrized that parameters are not identifiable without additional restrictions. Some examples of parsimonious conditional factor models include **?**, **?**, and **?**.

[7] We adopt the same subsequence definition as that used in **?**. 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.

alphas that are too large, to the extent that $S^\star$ explodes along some diverging sequence.[8]

To achieve this optimal Sharpe ratio, arbitrageurs should hold a portfolio with weights given by $w^\star = \Sigma_u^{-1}\alpha$, according to **?**.[9] Under the rational expectation assumption, arbitrageurs (agents in this model) know the true (population) parameters: α and Σ_u. In reality, however, the true parameters are blind to arbitrageurs as they can only learn these parameters from a finite sample of data. This learning effect is sometimes harmless since it can be expected that when the sample size is large enough, the true parameters are (asymptotically) revealed, and hence the predictions under rational expectation hold approximately. Fundamentally, this phenomenon is due to the assumption 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 set of investment opportunities, N. As N increases, it becomes increasingly difficult for arbitrageurs to determine which among all assets truly have nonzero alphas for a given sample size, T. If the learning problem remains difficult as N and T increase, the learning effect persists, which could lead to distinct limiting implications as opposed to the rational expectation case. It turns out that the rational expectation limit $S^\star$ 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 $S^\star$ because of their inability to make error-free inference. Therefore, the condition (3) 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:*

$$\alpha_i \overset{i.i.d.}{\sim} \begin{cases} \mu & \text{with prob. } \rho/2 \\ -\mu & \text{with prob. } \rho/2 \\ 0 & \text{with prob. } 1-\rho \end{cases}, \quad 1 \le i \le N, \tag{4}$$

where $\mu \ge 0$ and $0 \le \rho \le 1$, and they potentially vary with N and T. In addition, we also assume $\beta = 0$, $\Sigma_u = \sigma^2 \mathbb{I}_N$, for some $\sigma > 0$.

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 dimensions of the investment set, we can accurately characterize the difficulty of the finite sample problem arbitrageurs are facing.[10] 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.

[8] Using 3 and the fact that $\mathrm{E}(N\alpha_i^2) = \mathrm{E}(\alpha^\intercal\alpha) \lesssim \mathrm{E}(\alpha^\intercal\lambda_{\min}(\Sigma_u^{-1})\alpha) \lesssim \mathrm{E}(\alpha^\intercal\Sigma_u^{-1}\alpha) \lesssim 1$, implies $\mathrm{E}(\alpha_i^2) = o(1)$.

[9] In **?**, α is defined to be the cross-sectional projection of the expected returns onto β in the population model such that $\alpha^\intercal\Sigma_u^{-1}\beta = 0$. In this paper, we assume instead that α is random, satisfying $\mathrm{E}(\alpha^\intercal\beta) = 0$, and hence in our setting $w^\star = \mathbb{M}_\beta\Sigma_u^{-1}\alpha$, which achieves the Sharpe ratio $S^\star$ (asymptotically as N increases).

[10] Adopting a drifting sequence for parameters is a common trick in econometrics to provide more accurate finite sample approximations. As **?** 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."

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

$$\mu \sim T^{-1/2} \quad \text{and} \quad \rho \sim N^{-1/2}. \tag{5}$$

This condition (5) implies that the signal strength vanishes as the sample size increases and the signal count decays as the investment universe expands. That is, only a small portion of assets have a nonzero yet small alpha. We assume σ is a fixed constant, since in reality idiosyncratic risks never vanish, whereas alphas can be small driven by competition among arbitrageurs. This model rests on an uncommon territory in the existing literature of asset pricing: weak and rare alphas. In fact, the classical no near-arbitrage condition (3) imposes, implicitly, weakness or rareness on alphas; otherwise, if alphas are strong and dense, $\alpha^\intercal\alpha$ would explode rather rapidly. Even in the current setting, in light of the fact that $\mathrm{E}(\alpha^\intercal\alpha) = \rho\mu^2 N$, we still have $\alpha^\intercal\alpha \xrightarrow{p} \infty$ as long as $N^{1/2}/T \to \infty$. In this case, a near-arbitrage opportunity arises according to (3), with a strategy $w = \sigma^{-2}\alpha$.

However, the statistical obstacle 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 $T^{-1/2}$.[11] Since the true alpha is of the same order of magnitude as its level of statistical uncertainty by design, i.e., $\mu \sim T^{-1/2}$, it is impossible for arbitrageurs to determine precisely which assets among all have nonzero alpha.

For illustration purpose, suppose that arbitrageurs adopt the strategy $\widehat{w} = \sigma^{-2}\widehat{\alpha}$, replacing α in w with $\widehat{\alpha} = \bar{r} = \alpha + \bar{u}$.[12] Out of sample, this portfolio's conditional expected return and conditional variance can be written as:

$$\begin{aligned}
\mathrm{E}\left(\sigma^{-2}\left(\alpha+\bar{u}\right)^\intercal\left(\alpha+u_t\right)|\mathcal{F}_{t-1}\right) &= \sigma^{-2}(\alpha^\intercal\alpha + \bar{u}^\intercal\alpha), \\
\mathrm{Var}\left(\sigma^{-2}\left(\alpha+\bar{u}\right)^\intercal\left(\alpha+u_t\right)|\mathcal{F}_{t-1}\right) &= \sigma^{-2}(\alpha^\intercal\alpha + 2\alpha^\intercal\bar{u} + \bar{u}^\intercal\bar{u}),
\end{aligned}$$

where u_t denotes a future return at t, independent of $\bar{u}$ which is part of the information set up to $t-1$, $\mathcal{F}_{t-1}$, but share the same distribution as $\{u_s\}_{s\le t-1}$. The resulting squared Sharpe ratio is given by:

$$S^2 = \frac{\sigma^{-2}(\alpha^\intercal\alpha + \bar{u}^\intercal\alpha)}{\sigma^{-2}(\alpha^\intercal\alpha + 2\alpha^\intercal\bar{u} + \bar{u}^\intercal\bar{u})} \lesssim_{\mathrm{P}} T^{-1} \to 0. \tag{7}$$

Indeed, as we will show later, in this setting, the optimal Sharpe ratio, denoted by S^{OPT},

[11] **?** 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 $1 \le i \le N$,

$$\sqrt{T}(\widehat{\alpha}_i - \alpha_i) \xrightarrow{d} \mathcal{N}(0, \sigma_i^2(1 + \gamma^\intercal(\Sigma_v)^{-1}\gamma)), \tag{6}$$

where σ_i^2 is the ith 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 $(1 + \gamma^\intercal(\Sigma_v)^{-1}\gamma)$ disappears.

[12] For 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 $\widehat{w}$, which we will denote by $\widehat{w}^{\mathrm{CSR}}$, fails to achieve the optimal Sharpe ratio in all scenarios. We will discuss the optimal strategy in Section 2.5.

among all *feasible* trading strategies arbitrageurs adopt, vanishes asymptotically as $N, T \to \infty$, even though the *infeasible* optimal Sharpe ratio $S^\star \to \infty$. The gap between S^{OPT} and $S^\star$, as shown by this example, is enormous.

We say a strategy is *feasible* if it only uses observable data, combined with necessary statistical inference. We formalize the definition of a feasible portfolio strategy below:

Definition 2. *A portfolio strategy $\widehat{w}$ is said to be feasible at time t, if it is a function of observables from $t - T + 1$ to t, where T is the sample size.*

This Sharpe ratio gap is driven by the fact that the performance of a feasible portfolio depends on the difficulty of the learning problem. The learning problem in the above example remains difficult as N and T increase, to the extent that the learning effect does not diminish in the limit and the rational expectation limit becomes irrelevant. As we show below, the learning problem in practice is rather difficult, hence the optimal arbitrage Sharpe ratio achievable is expected to be much smaller than $S^\star$.

2.3 Upper Bound on Feasible Sharpe Ratios

We now demonstrate the impact of the feasibility constraint on the optimal arbitrage portfolio. For any feasible strategy $\widehat{w}$, its (conditional) Sharpe ratio can be written as:

$$S(\widehat{w}) \coloneqq \mathrm{E}(\widehat{w}^\intercal r_{t+1}|\mathcal{F}_t)/\mathrm{Var}(\widehat{w}^\intercal r_{t+1}|\mathcal{F}_t)^{1/2}.$$

Arbitrageurs, in our setting, strive to find a feasible strategy that maximizes $S(\widehat{w})$.

The next theorem provides an upper bound on $S(\widehat{w})$:

Theorem 1. *Suppose that r_t follows (1) and that Assumption 1 holds. For any feasible portfolio weight $\widehat{w}$, its Sharpe ratio, $S(\widehat{w})$, satisfies, as $N \to \infty$:*

$$S(\widehat{w}) \leq \left(S(\mathcal{G})^2 + \gamma^\intercal \Sigma_v^{-1} \gamma\right)^{1/2} + o_{\mathrm{P}}(1), \quad \textit{with} \quad S(\mathcal{G})^2 \coloneqq \mathrm{E}(\alpha|\mathcal{G})^\intercal \Sigma_u^{-1} \mathrm{E}(\alpha|\mathcal{G}), \tag{8}$$

where $\mathcal{G}$ is the information set (i.e., σ-algebra) generated by $\{(r_s, \beta, v_s, \Sigma_u) : t - T + 1 \leq s \leq t\}$. If, in addition, that, $\widehat{w}$ satisfies that $\widehat{w}^\intercal \beta = 0$, that is, the portfolio is factor-neutral, then

$$S(\widehat{w}) \leq S(\mathcal{G}) + o_{\mathrm{P}}(1). \tag{9}$$

Intuitively, $S(\mathcal{G})^2$ and $\gamma^\intercal \Sigma_v^{-1} \gamma$ are squared Sharpe ratios earned from arbitrage portfolios and factor portfolios, respectively. Naturally, $S(\mathcal{G})$ is an upper bound for Sharpe ratios of all feasible portfolio strategies that have no factor exposures.

Theorem 1 shows that it is $\mathrm{E}(\alpha|\mathcal{G})$, the posterior estimate of the pricing errors, α, that dictates the optimal feasible Sharpe ratio for arbitrageurs, rather than α themselves. In fact, it holds by the definition of $S(\mathcal{G})$ that

$$\mathrm{E}\big(S(\mathcal{G})^2\big) \leq \mathrm{E}\big(\alpha^\intercal \Sigma_u^{-1} \alpha\big),$$

with the equality holds only when $\mathrm{E}(\alpha|\mathcal{G}) = \alpha$ almost surely, where the right-hand side corresponds to the infeasible scenario in which arbitrageurs can learn α perfectly using their information set, which echoes (3), the result given by **?**.[13] In light of 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 1 hold. For any given return-generating process satisfying (1), there exists no feasible strategy $\widehat{w}$ that leads to a near-arbitrage, if*

$$S(\mathcal{G}) \lesssim_{\mathrm{P}} 1, \quad as \quad N \to \infty. \tag{10}$$

The form of $S(\mathcal{G})$ in Theorem 1 appears 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 (diagonal) covariance matrix of idiosyncratic errors, Σ_u. In fact, the upper bound in (8) 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 $\mathrm{E}(S(\mathcal{G}')^2) \leq \mathrm{E}(S(\mathcal{G})^2)$. Furthermore, we will show in Section 2.5 that $S(\mathcal{G})$ is in fact achievable by a feasible strategy we construct, which only assumes knowledge of β and r_t – the setting in which factor exposures are observable. This implies that the no near-arbitrage bound in (10) is sufficient and necessary.

The reason that Σ_u plays no significant role is that in our model idiosyncratic variances do not vanish as N and T increase, unlike alphas. This assumption makes sense empirically, because alphas are (potentially) small and rare, driven by competition among arbitrageurs, whereas idiosyncratic risks never diminish. Consequently, detecting alphas is more challenging as opposed to estimating idiosyncratic variances, and hence the latter plays a secondary (and negligible) role as opposed to the former in the limit of arbitrage.

To avoid technical details of the main text, we simplify model (1) by imposing that $\Sigma_u = \sigma^2 \mathbb{I}_N$ for some fixed constant $\sigma^2 > 0$ in the following-up sections, while leaving discussion of the general inhomogeneous variances case to the appendix, where we have Assumption A2 in place.

2.4 Explicit Formula of the Sharpe Ratio Bound

To gain more insight on $S(\mathcal{G})$, we seek a more explicit expression in this section. For that purpose, we need impose a normality assumption on the idiosyncratic error.

Proposition 1. *Suppose that r_t follows (1) and Assumption 1 holds. We also assume that for all (i,t), $u_{i,t} \sim \mathcal{N}(0, \sigma^2)$. Then it holds that*

$$\mathrm{E}(\alpha_i|\mathcal{G}) = \psi\big(\widehat{z}_i\big),$$

[13] For ease of discussion, we assume alpha is random. This difference with **?** by itself does not affect any economic or statistical conclusions we draw in this paper.

where $\widehat{z}_i = T^{1/2}(\alpha_i + \bar{u}_i)/\sigma$, $\bar{u}$ *is the average based on a sample of size* T,

$$\psi(a) = \frac{\mathrm{E}\left(\alpha_i \phi(a - T^{1/2}\alpha_i/\sigma)\right)}{\mathrm{E}\left(\phi(a - T^{1/2}\alpha_i/\sigma)\right)},$$

$\phi(\cdot)$ *is the normal pdf function, and* $\mathrm{E}(\cdot)$ *is the expectation taken with respect to the cross-sectional distributions of* α_i. *If it further holds that* $\mathrm{E}(\alpha_i^2 \mathbb{1}_{\{|\alpha_i| \geq c_N\}}) \leq c_N N^{-1}$ *for some sequence* $c_N \to 0$, *then we have*

$$S(\mathcal{G}) = S^{\mathrm{OPT}} + o_{\mathrm{P}}(1), \quad \textit{with} \quad S^{\mathrm{OPT}} = \left(N\sigma^{-2} \int \psi(a)^2 p(a) da\right)^{1/2},$$

where $p(a) = \mathrm{E}\left(\phi(a - T^{1/2}\alpha_i/\sigma)\right)$ *is the probability distribution function of* $\widehat{z}_i$.

The first part of Proposition 1 provides a closed-form formula for $\mathrm{E}(\alpha_i|\mathcal{G})$. Under the stated conditions on the independence across i, $\widehat{\alpha}_i = \alpha_i + \bar{u}$, and in turn $\widehat{z}_i$, are sufficient statistics, so that $\mathrm{E}(\alpha_i|\mathcal{G}) = \mathrm{E}(\alpha_i|\widehat{z}_i)$. The latter expectation can then be evaluated with the help of the normality assumption on $u_{i,t}$. The second part of this proposition aims to simplify $S(\mathcal{G})$ using $\mathrm{E}(\alpha_i|\mathcal{G})$, according to Theorem 1, which needs an additional condition on the tail behavior of the cross-sectional distribution of α_i. The tail condition is stronger than Assumption 1(c).[14] It is a technical condition used to simplify the expression of $S(\mathcal{G})$.

To shed more light on this result, we compare this optimal Sharpe ratio S^{OPT} with $S^\star$ of **?** using Example 1.

Corollary 2. *Suppose that the same assumptions as in Proposition 1 hold. In addition, we assume alpha follows (4) as in Example 1. Then we have* $S^\star = \sigma^{-1}\mu(\rho N)^{1/2} + o_{\mathrm{P}}(1)$. *Further, assuming that* $\sigma^{-1}\mu(\rho N)^{1/2}$ *does not vanish, then it holds that* $S^{\mathrm{OPT}} \leq (1-\epsilon)\sigma^{-1}\mu(\rho N)^{1/2}$ *for some* $\epsilon > 0$, *if and only if*

$$T^{1/2}\mu/\sigma - \sqrt{-2\log\rho} \lesssim 1. \tag{11}$$

Corollary 2 suggests that when $T^{1/2}\mu/\sigma$ is large, the constraint (11) is more likely violated, in which case $S^{\mathrm{OPT}} \approx \mathrm{E}(S^\star)$, that is, in the limit, the learning effect does not play any role, so that arbitrageurs in this scenario achieve the same optimal Sharpe ratio as in **?**. 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 to the extent 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 hence S^{OPT} is dominated by $\mathrm{E}(S^\star)$ in the limit, resulting in a strictly smaller Sharpe ratio than the infeasible Sharpe ratio in the classical case.

To give a concrete example of Corollary 2, consider an alternative DGP assumption as opposed

[14] Obviously, we have $\mathrm{E}(\max_i \alpha_i^2) \leq \mathrm{E}\left(\max_i \alpha_i^2 \mathbb{1}_{\{|\alpha_i| \geq c_N\}} + \max_i \alpha_i^2 \mathbb{1}_{\{|\alpha_i| < c_N\}}\right) \leq \mathrm{E}\left(\sum_i \alpha_i^2 \mathbb{1}_{\{|\alpha_i| \geq c_N\}}\right) + c_N^2 = o(1)$.

to (5):[15]

$$\mu \sim N^{-\lambda} \quad \text{and} \quad \rho > 0 \quad \text{is fixed.} \tag{12}$$

In this scenario, $\mathrm{E}\left((S^\star)^2\right) = \sigma^{-2}\mu^2\rho N \sim N^{1-2\lambda}$, which explodes unless $\lambda > 1/2$. If further assuming that $N/T \to \psi > 0$, then the left-hand-side of condition (11) is of order $N^{1/2-\lambda} \vee 1$, so that (11) holds if and only if $\lambda \geq 1/2$. Therefore, $\lambda < 1/2$ is not consistent with absence of (feasible) near arbitrage in that the infeasible Sharpe ratio explodes, while in the mean time the feasible Sharpe ratio approximately equals the infeasible Sharpe ratio (by Corollary 2) and hence explodes. If $\lambda > 1/2$, the infeasible Sharpe ratio (and hence the feasible one) vanishes, which does not seem like an economically plausible case. If we think that arbitrageur activity is required to prevent substantial mispricing, then a setting where mispricing disappears asymptotically even if the frictions faced by arbitrageurs are very large is not plausible. This suggests that under this DGP, the only economically plausible case with absence of near-arbitrage is $\lambda = 1/2$. That is, λ can be thought as determined in equilibrium, in which there are substantial asset demand distortions such that mispricing in the absence of arbitrageur action would be non-negligible asymptotically, and arbitrageurs are aggressive enough so that near-arbitrage opportunities do not exist asymptotically.

2.5 Constructing the Optimal Arbitrage Portfolio

In our previous discussion, we have shown in Theorem 1 that the optimal Sharpe ratio for any feasible strategy is bounded by $S(\mathcal{G})$. In Proposition 1, we have shown that $S(\mathcal{G}) \approx S^{\mathrm{OPT}}$ under additional assumptions. Corollary 2 further demonstrates that the optimal Sharpe ratio can vary with sequences of DGPs. In light of this, the optimal strategy should depend on the unobserved DGP as well.

In fact, we now show that arbitrageurs can construct a uniformly valid estimator of the optimal portfolio weights, which achieves S^{OPT} over a large class of data generating precesses, without perfect knowledge of the true DGP. This further implies that the no feasible near-arbitrage bound we develop is sharp.

We demonstrate this in the setting where factors are latent but factor exposures are observable, since this is the case we analyze empirically. Moreover, we maintain the assumption that $\Sigma_u = \sigma^2\mathbb{I}_N$ for simplicity of the main text, leaving the general case to the appendix (Algorithm 3). Since all stocks share the common volatility σ^2, it does not affect the relative weights of different stocks in our portfolio. As a result, in this simplified setup, σ^2 becomes a nuisance parameter that does not influence the Sharpe ratio of the constructed portfolio.

Algorithm 1 (Constructing the Optimal Arbitrage Portfolio)**.**
Inputs: r_t, $t \in \mathcal{T} = \{t - T + 1, \ldots, t\}$ *and* β.

[15] It is easy to show that the setup (12) satisfies all assumptions of Proposition 1 for all fixed $\lambda > 0$.

S1. We split the observed sample $\mathcal{T}$ into:

$$S' = \{t - \lfloor T^{1/2} \rfloor + 1, \ldots, t\} \quad \textit{and} \quad S = \mathcal{T} - S',$$

and we construct cross-sectional regression estimates of alpha $\breve{\alpha}$ and $\breve{\alpha}'$, and the t-statistics $\breve{z}_i = |S|^{1/2}\breve{\alpha}_i$ for each $i = 1, 2, \ldots, N$, using subsamples S and S':[16]

$$\breve{\alpha} = |S|^{-1} \sum_{s \in S} \mathbb{M}_\beta r_s, \quad \textit{and} \quad \breve{\alpha}' = |S'|^{-1} \sum_{s \in S'} \mathbb{M}_\beta r_s,$$

where $\mathbb{M}_\beta = \mathbb{I}_N - \beta(\beta^\intercal\beta)^{-1}\beta^\intercal$ and $\mathbb{I}_N$ denotes the $N \times N$ identity matrix.

S2. We choose the arbitrage portfolio weights as

$$\widehat{w}^{\mathrm{OPT}} = \mathbb{M}_\beta \breve{w}, \quad \textit{with} \quad \breve{w}_i = \begin{cases} \widehat{f}(\lfloor \breve{z}_i / k_N \rfloor), & |\breve{z}_i| \leq k_N^{-2/3}, \\ \breve{\alpha}_i, & |\breve{z}_i| > k_N^{-2/3}. \end{cases}$$

For any set of integers (l), we choose $k_N \sim (\log N)^{-1}$ and define

$$\widehat{f}(l) = \frac{1}{|B(l)|} \sum_{i \in B(l)} \breve{\alpha}'_i, \quad \textit{where} \quad B(l) = \{i \leq N : l \leq \breve{z}_i / k_N < l + 1\}.$$

Outputs: $\widehat{w}^{\mathrm{OPT}}$.

As we have discussed in footnote 9, the optimal strategy in the case that arbitrageurs know the true DGP is given by

$$w^\star = \mathbb{M}_\beta \Sigma_u^{-1} \alpha. \tag{13}$$

Intuitively, part of the construction in (13), $\Sigma_u^{-1}\alpha$, is the optimal allocation to the ex-factor returns, $\alpha + u_t = r_t - \beta(\gamma + v_t)$, in a simple mean-variance analysis. Multiplying by $\mathbb{M}_\beta$ in (13) simply eliminates factor exposures in r_t. Correspondingly, Step S1 of Algorithm 1 provides estimates of α_i in two separate samples. Sample splitting is a convenient approach to avoiding dependence of these alpha estimates. Step S2 first constructs a nonparametric estimate of the function $f(\alpha) = \mathrm{E}(\alpha|\{r_s, \beta\}_{s \in \mathcal{T}})$, with which the optimal weights on ex-factor returns are constructed as $\breve{w}$. This, in turn, leads to the optimal weight estimates, $\widehat{w}^{\mathrm{OPT}}$, on original input asset returns.

An essential step towards uniform optimality is the construction of $\breve{w}$, in which we deal with strong and weak signals separately. The strong signals (those with t-statistics greater than $k_N^{-2/3}$) are singled out, for which we can obtain relatively precise estimates of their optimal weights. For weaker signals, we consolidate information therein to obtain an estimate of the conditional expectation of their signal strength, using which we obtain their optimal portfolio weights. This strategy outperforms the alternatives, which either directly use estimated alphas as if these estimates are

[16] For any set S, we use $|S|$ to denote the number of elements in S.

not susceptible to errors even when they are rather weak, or simply ignore the contribution of all weaker signals.

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

Theorem 2. *Let $\mathbb{P}$ denote the collection of all data-generating processes under which r_t follows (1), and Assumption 1 holds. In addition, suppose that $N^d \leq T$ for some $d > 0$. We denote the Sharpe ratio generated by the portfolio strategy $\widehat{w}^{\mathrm{OPT}}$ as $\widehat{S}^{\mathrm{OPT}} := \mathrm{E}(r_{t+1}^\intercal \widehat{w}^{\mathrm{OPT}} | \mathcal{F}_t)/\mathrm{Var}(r_{t+1}^\intercal \widehat{w}^{\mathrm{OPT}} | \mathcal{F}_t)^{1/2}$. Then it holds that $\widehat{w}^{\mathrm{OPT}}$ achieves, asymptotically, the upper bound S^{OPT} uniformly over all sequences of data-generating processes. That is, for any $\epsilon > 0$,*

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

Theorem 2 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. Its gap to $\mathrm{E}(S^\star)$, the Sharpe ratio under rational expectation, is determined by the difficulty of the learning problem.

With Theorem 2, we establish the necessity for the no near-arbitrage condition given by (10).

Corollary 3. *Suppose the same assumptions as in Theorem 2 hold. The portfolio weights by $\widehat{w}^{\mathrm{OPT}}$ yields a near-arbitrage strategy under any sequences of data-generating processes for which condition (10) does not hold.*

We have shown that arbitrageurs can construct an optimal strategy that realizes S^{OPT}. Now suppose that the equilibrium "cost" of implementing an arbitrage is C in an economy with statistical limit of arbitrage. In equilibrium, $S^{\mathrm{OPT}} = C$, otherwise arbitrageurs can trade until it is no longer profitable to do so. We can thereby interpret $\widehat{S}^{\mathrm{OPT}}$ as an empirical estimate of the arbitrage cost, which we will estimate empirically.

2.6 Estimating Optimal Infeasible Sharpe Ratio

We are also interested in estimating the optimal infeasible Sharpe ratio, $S^\star$, which can be perceived as the optimal Sharpe ratio from an outside econometrician's point of view. Existing literature on testing APT often construct test statistics in the spirit of **?**, which are effectively based on $S^\star$, see, e.g., **?** and **?**. While such tests are powerful and may lead to discoveries of alpha signals, they are not relevant for arbitrageurs in that arbitrageurs may not construct a feasible portfolio to capture these statistical discoveries.

We now construct an estimator for $(S^\star)^2$. The most natural estimator of the infeasible Sharpe ratio is given by:

$$(\widetilde{S}^\star)^2 = \bar{r}^\intercal \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \mathbb{M}_\beta \bar{r}, \tag{14}$$

where $\bar{r} = T^{-1}\sum_{t\in\mathcal{T}} r_t$, $\widehat{\sigma}_i^2 = T^{-1}\sum_{t\in\mathcal{T}}(r_{i,t} - \bar{r}_i)^2$, and $\widehat{\Sigma}_u = \text{diag}(\widehat{\sigma}_1^2, \widehat{\sigma}_2^2, \ldots, \widehat{\sigma}_N^2)$.

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

$$(\widehat{S}^\star)^2 = \bar{r}^\intercal \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \mathbb{M}_\beta \bar{r} - T^{-2}\sum_{t\in\mathcal{T}} r_t^\intercal \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \mathbb{M}_\beta r_t. \tag{15}$$

The second estimator again takes the form of a summation over individual squared Sharpe ratios, but it eliminates the term that will be dominated by the estimation bias under some data generating processes. The next proposition summarizes the asymptotic properties of both estimators.

Proposition 2. *Suppose r_t follows (1), and Assumptions 1 and A2 hold. Then we have*

$$\left|\widehat{S}^\star - S^\star\right| \Big/ \left(1 + S^\star\right) = o_{\text{P}}\left(T^{-1/4}\right), \quad \left|\widetilde{S}^\star - \left((S^\star)^2 + NT^{-1}\right)^{1/2}\right| \Big/ \left(1 + S^\star\right) = o_{\text{P}}\left(T^{-1/4}\right).$$

As shown by this proposition, the estimation error is relative when $S^\star$ dominates one asymptotically, and absolute if $S^\star$ is dominated by one. This is necessary because we simultaneously consider a large class of models, some of which have an exploding or a shrinking $S^\star$, and because $S^\star$ also influences the convergence rate.

2.7 Alternative Strategies for Arbitrage Portfolios

Algorithm 1 suggests a relatively sophisticated procedure that distinguishes weaker and strong signals using t-statistics before constructing, separately, the optimal weights for these signals. In this section, we study several alternative methods, neither of which can achieve optimality uniformly across all DGPs we consider, but they are simpler and somewhat prevalent in practice. The contrast among these strategies helps illustrate their pros and cons in different scenarios.

2.7.1 Cross-Sectional Regression

The conventional approach to estimating alphas is through the cross-sectional regression:

$$\widehat{\alpha} = \left(\beta^\intercal \beta\right)^{-1} \beta^\intercal \bar{r},$$

with which the arbitrage portfolio weights can be constructed directly as:

$$\widehat{w}^{\text{CSR}} = \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \widehat{\alpha}. \tag{16}$$

This choice of portfolio weight is the sample analog of the optimal weight given by (13). Proposition 3 below compares the asymptotic behavior of the expected Sharpe ratio of this arbitrage portfolio with the optimal Sharpe ratio, S^{OPT}, in Example 1. For convenience, we adopt a simplified volatility

estimator: $\widehat{\Sigma}_u = \widehat{\sigma}^2 \mathbb{I}_N$, where $\widehat{\sigma}^2$ is averaged over all volatility estimates, because in this example, all assets share the same volatility. This further simplifies the analysis because the scaling factor, $\widehat{\sigma}^2$, is cancelled out, and hence does not play any role in this portfolio's Sharpe ratio.

Proposition 3. *Suppose that r_t follows (1) and Assumption 1 holds. In addition, we assume alpha follows (4) as in Example 1. The Sharpe ratio of the arbitrage portfolio, whose weights are given by $\widehat{w}^{\mathrm{CSR}} = \widehat{\sigma}^{-2}\mathbb{M}_\beta \widehat{\alpha}$, satisfies $\widehat{S}^{\mathrm{CSR}} - S^{\mathrm{CSR}} = o_{\mathrm{P}}(1)$, where*

$$\widehat{S}^{\mathrm{CSR}} = \mathrm{E}(r_{t+1}^\intercal \widehat{w}^{\mathrm{CSR}} | \mathcal{F}_t) / \mathrm{Var}(r_{t+1}^\intercal \widehat{w}^{\mathrm{CSR}} | \mathcal{F}_t)^{1/2}, \qquad S^{\mathrm{CSR}} = \frac{N^{1/2} \rho \mu^2 \sigma^{-2}}{(T^{-1} + \rho \mu^2 \sigma^{-2})^{1/2}}.$$

Further, assuming S^{OPT} does not vanish, then as $N, T \to \infty$, we have $S^{\mathrm{CSR}} \leq (1-\epsilon) S^{\mathrm{OPT}}$ for some fixed $\epsilon > 0$, if and only if

$$C \leq T\mu^2\sigma^{-2} \leq C' \rho^{-1}. \tag{17}$$

for some constants C and C'.

Proposition 3 suggests that arbitrageurs using this cross-sectional regression strategy cannot always achieve the optimal feasible Sharpe ratio. In fact, this strategy is dominated by the optimal strategy when signals are both very strong ($C \leq T\mu^2\sigma^{-2}$) and there are not too many strong signals ($T\mu^2\sigma^{-2} \leq C'\rho^{-1}$). Intuitively, the CSR approach treats all signals equally, without distinguishing insignificant signals from the significant ones. Consequently, even fake signals could get non-zero weights allocated. This strategy thereby works well when the strong signals are abundant (i.e., ρ is relatively large) or when all signals are weak (so that they do not differ too much from fake ones). The latter case is more interesting, as it also suggests that simply ignoring weaker signals is not optimal.

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.

2.7.2 False Discovery Rate Control

From the statistical point of view, we can formalize the search for alpha as a multiple testing problem. Say, there are N assets potentially with nonzero α, and for each i, we can define a null hypothesis: $\mathbb{H}_0^i : \alpha_i = 0$, so that alpha detection becomes a multiple testing problem. With multiple testing comes the concern of data snooping, meaning that a large fraction of tests that appear positive are in fact due to chance. One sensible approach is to control the false discovery rate (FDR), instead of the size of individual tests, a proposal advocated by **?**, **?**, and **?** in different asset pricing contexts.

The B-H procedure proposed by **?** is often adopted to control FDR in multiple testing problems. **?** have proved its validity in a general factor model setting for alpha detection. Below we describe

the algorithm for constructing alpha estimates, which will be used as inputs to the construction of an arbitrage portfolio.

Algorithm 2 (The B-H based Alpha Selection)**.** *Let $\widehat{\alpha}$ be the estimator of α via the cross-sectional regression, and $\{p_i : i = 1, \ldots, N\}$ be the p-values of the corresponding t-test statistics.*

S1. Sort in ascending order the collection of p-values, with the sorted p-values given by $p_{(1)} \leq \cdots \leq p_{(N)}$.

S2. For $i = 1, \ldots, N$, reject $\mathbb{H}_0^i : \alpha_i = 0$, if $p_i \leq p_{(\widehat{k})}$, where $\widehat{k} = \max\{i \leq N : p_{(i)} \leq \tau i/N\}$, for any pre-determined level τ, say, 5%.

We can then adjust our alpha estimates using

$$\widehat{\alpha}_i^{\mathrm{BH}}(\tau) = \widehat{\alpha}_i \mathbb{1}_{\{p_i \leq p_{(\widehat{k})}\}}. \tag{18}$$

The B-H procedure guarantees (in expectation) that at least a fraction $(1-\tau)$ of selected assets have nonzero alphas, regardless of the actual percentage of alphas in the data generating process. Effectively, it imposes a hard-thresholding procedure on the alpha estimates, replacing less significant alphas by zero. Similar to (16), the optimal portfolio weights are thus given by:

$$\widehat{w}^{\mathrm{BH}} = \mathbb{M}_\beta \widehat{\Sigma}_u^{-1} \widehat{\alpha}^{\mathrm{BH}}(\tau). \tag{19}$$

Controlling the false discovery rate on top of the CSR estimates is intuitively appealing, but doing so incurs a potential loss of power and hence less investment opportunities. Our focus is on optimal portfolio construction instead of false discovery control. The next proposition shows that in the context of Example 1, arbitrageurs who adopt the B-H based alpha estimator cannot achieve optimal portfolio for a large class of DGP sequences. In fact, we can prove a richer result. Even if arbitrageurs knew that all assets with non-zero alphas have an equally strong alpha with the same idiosyncratic volatility (which is true in this example), and that they adopted the following estimates for alpha instead of (18) to exploit this fact:

$$\overline{\alpha}_i^{\mathrm{BH}}(\tau) = sgn(\widehat{\alpha}_i)\overline{\alpha}\mathbb{1}_{\{p_i \leq p_{(\widehat{k})}\}}, \quad \overline{\alpha} = \sum_i |\widehat{\alpha}_i| \mathbb{1}_{\{p_i \leq p_{(\widehat{k})}\}} / \sum_i \mathbb{1}_{\{p_i \leq p_{(\widehat{k})}\}}, \tag{20}$$

still they would not be able to achieve the optimal performance.

Proposition 4. *Suppose that r_t follows (1) and Assumption 1 holds. In addition, we assume alpha follows (4) as in Example 1. The Sharpe ratio of the arbitrage portfolio with weights given by $\widehat{w}^{\mathrm{BH}} = \widehat{\sigma}^{-2}\mathbb{M}_\beta \widehat{\alpha}^{\mathrm{BH}}(\tau)$ and $\overline{w}^{\mathrm{BH}} = \widehat{\sigma}^{-2}\mathbb{M}_\beta \overline{\alpha}^{\mathrm{BH}}(\tau)$ satisfies $\overline{S}^{\mathrm{BH}} - \sqrt{1-\tau} S^{\mathrm{BH}} = o_{\mathrm{P}}(1)$ and $\widehat{S}^{\mathrm{BH}} \leq S^{\mathrm{BH}} + o_{\mathrm{P}}(1)$, where*[17]

$$\widehat{S}^{\mathrm{BH}} = \mathrm{E}(r_{t+1}^\intercal \widehat{w}^{\mathrm{BH}} | \mathcal{F}_t)/\mathrm{Var}(r_{t+1}^\intercal \widehat{w}^{\mathrm{BH}} | \mathcal{F}_t)^{1/2}, \quad \overline{S}^{\mathrm{BH}} = \mathrm{E}(r_{t+1}^\intercal \overline{w}^{\mathrm{BH}} | \mathcal{F}_t)/\mathrm{Var}(r_{t+1}^\intercal \overline{w}^{\mathrm{BH}} | \mathcal{F}_t)^{1/2}$$

[17] If $\widehat{w}^{\mathrm{BH}} = 0$, i.e., no asset is selected, we set $\widehat{S}^{\mathrm{BH}} = 0$ by convention.

are corresponding Sharpe ratios, and

$$S^{\mathrm{BH}} = \mu\sigma^{-1}\sqrt{\rho N\Phi(T^{1/2}\mu/\sigma - z^*)},$$

where $\Phi(\cdot)$ *is the normal cumulative distribution function, and* z^* *is the positive solution of the equation*

$$2(1-\tau(1-\rho))\Phi(-z) = \tau\rho\Phi(T^{1/2}\mu/\sigma - z). \tag{21}$$

Suppose further that S^{OPT} *does not vanish, and that* $CN^{-1+d} \le \rho \le CN^{-d}$ *for some fixed* $d > 0$. *Then it follows that, as* $N, T \to \infty$, $S^{\mathrm{BH}} \le (1-\epsilon)S^{\mathrm{OPT}}$ *for some fixed* $\epsilon > 0$, *if and only if, for some fixed* $\epsilon' > 0$,

$$T^{1/2}\mu/\sigma \le (1-\epsilon')\sqrt{-\lambda_\tau \log\rho}, \tag{22}$$

where $\lambda_\tau \in (2/3, 2)$ *only depends on* τ *and* $\lambda_\tau \to 2$ *as* $\tau \to 0$.

As Proposition 4 shows, (22) indicates that if the signal-to-noise ratio is not sufficiently strong, the B-H procedure is unlikely to reach S^{OPT}. This is because it ignores many individually impotent signals, which would hurt the portfolio performance, even though B-H remains a preferable approach to selecting truly significant alphas and controlling false discoveries. In contrast, the optimal arbitrage portfolio exploits information embedded in all alpha estimates, including false positives, beyond the set of significant ones selected via B-H procedure. This result also demonstrates a clear distinction between two objectives: alpha testing and portfolio construction, the objectives of which 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

The analysis above suggests that we can construct the optimal portfolio out of the ex-factor returns, while imposing regularization on portfolio weights, before rewriting the regularized portfolio weights in terms of raw returns (multiplying by $\mathbb{M}_\beta$). Regularizing portfolio weights amounts to imposing priors directly on the alpha estimates. To see this, we adopt a shrinkage approach, when constructing arbitrage portfolios on residual returns:

$$\arg\max_w \{w^\intercal\widehat{\alpha} - \frac{1}{2}w^\intercal\widehat{\Sigma}_u w - p_\lambda(w)\},$$

where $p_\lambda(w) = \lambda\|w\|_1$ or $\lambda\|w\|_2^2$, for some $\lambda > 0$. Since $\widehat{\Sigma}_u$ is diagonal, the closed-form solution is $\psi_q(\widehat{\alpha}, \widehat{\Sigma}_u, \lambda)$, where $q = 1$ corresponds to the LASSO penalty and $q = 2$ the ridge, and for $i = 1, 2, \ldots, N$,

$$\Big(\psi_1(\widehat{\alpha}, \widehat{\Sigma}_u, \lambda)\Big)_i = (\widehat{\sigma}_i)^{-2} sgn(\widehat{\alpha}_i)(|\widehat{\alpha}_i| - \lambda)_+, \quad \Big(\psi_2(\widehat{\alpha}, \widehat{\Sigma}_u, \lambda)\Big)_i = \big((\widehat{\sigma}_i)^2 + \lambda\big)^{-1}\widehat{\alpha}_i.$$

This leads to the optimal portfolio weight on r_t:[18]

$$\widehat{w}^q = \mathbb{M}_\beta \psi_q(\widehat{\alpha}, \widehat{\Sigma}_u, \lambda), \quad q = 1, 2.$$

Depending on the magnitude of λ, the LASSO approach replaces all smaller signals by zero and shrinks the larger signals 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 with a shrinkage factor depending on $\widehat{\sigma}_i^2$. Like the above analysis, when specialized to example (1), we can adopt $\widehat{\Sigma}_u = \widehat{\sigma}^2 \mathbb{I}_N$, in which case ridge becomes equivalent to CSR! This "embedded" shrinkage effect of CSR explains why it performs well in the case of small signals. Proposition 5, along with Proposition 3, demonstrate that neither LASSO nor ridge can achieve optimal Sharpe ratio in all DGPs even with the optimal tuning parameter λ.

Proposition 5. *Suppose that r_t follows (1) and Assumption 1 holds. In addition, we assume alpha follows (4) as in Example 1. The Sharpe ratio of the arbitrage portfolio with weights given by $\widehat{w}^q$, denoted as $\widehat{S}^q$ for $q = 1, 2$, satisfies $\widehat{S}^1 - S^{\text{LASSO}} = o_{\text{P}}(1)$ and $\widehat{S}^2 - S^{\text{CSR}} = o_{\text{P}}(1)$, where*

$$S^{\text{LASSO}} = \rho\mu\sigma^{-1}N^{1/2} \frac{\int_{-\infty}^{\infty} sgn(x)(T^{-1/2}\sigma|x| - \lambda)_+ \phi(T^{1/2}\sigma^{-1}\mu - x)dx}{\sqrt{\int_{-\infty}^{\infty} \big((T^{-1/2}\sigma|x| - \lambda)_+\big)^2 \big((1-\rho)\phi(x) + \rho\phi(T^{1/2}\sigma^{-1}\mu - x)\big)\big)dx}},$$

and S^{CSR} is defined in Proposition 3.

Suppose further that S^{OPT} does not vanish, and that $CN^{-1+d} \leq \rho \leq CN^{-d}$ for some fixed $d > 0$. Then it follows that, as $N, T \to \infty$, $S^{\text{LASSO}} \leq (1-\epsilon)S^{\text{OPT}}$ for some fixed $\epsilon > 0$ under all sequences of λ, if and only if

$$T^{1/2}\mu/\sigma \geq C, \quad \text{and} \quad \frac{T\mu^2/\sigma^2 + 2\log\rho}{\sqrt{-\log\rho}} \leq C.$$

3 Simulation Evidence

This section demonstrates the empirical relevance of our theory via simulations and examines the finite sample performance of the proposed portfolio strategies.

3.1 Numerical Illustration of Theoretical Predictions

We start by examining the theoretical predictions of Corollary 2 and Propositions 3 - 5. For simplicity and clarity, we simulate a one-factor (CAPM) model of returns given by (1). We choose the factor risk premium as 5% per year and set the annualized volatility at 25%. We model the cross-section of betas using a normal distribution with mean 1 and variance 1. Since we focus on the arbitrage portfolio, the parameters about the factor component (including the number of

[18] An 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.

factors) are inconsequential, because factors, if any, are eliminated by $\mathbb{M}_\beta$ in the first step when constructing these trading strategies. In addition, we adopt model (4) in Example 1 for the cross-sectional distribution of alpha, and fix the idiosyncratic volatilities of all assets at σ, since it is α/σ that determines the signal strength and that there is no need of varying both α and σ in the cross section.

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 (4), a ρ percentage 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 landscape of Sharpe ratios on signal weakness and rareness, despite that some of the resulting portfolio Sharpe ratios (the top left conner 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 the easiest when signal is strong and count is large (top left conner), and the most challenging towards the right bottom corner, where the optimal Sharpe ratios drop to near 0.

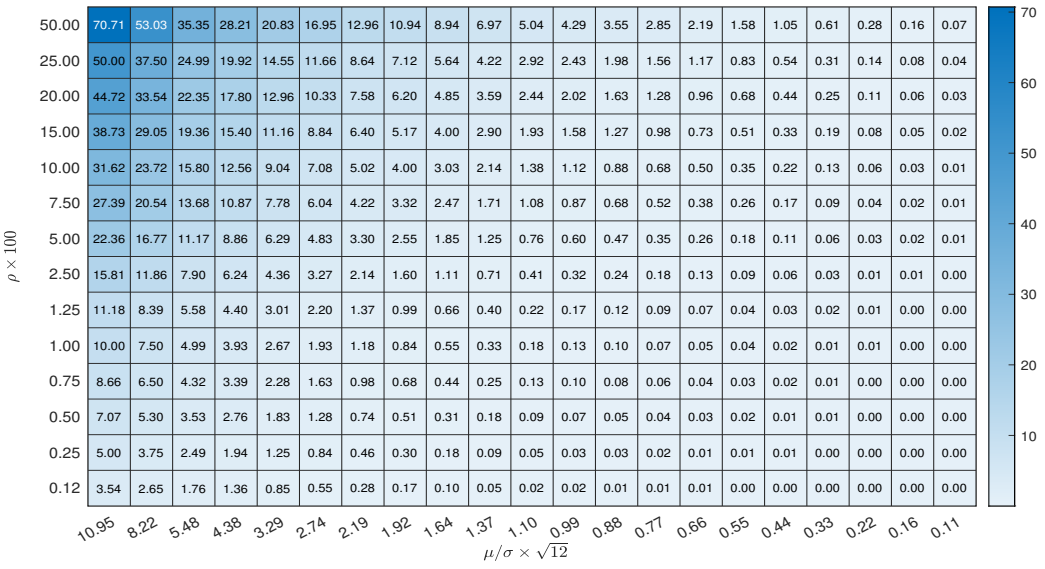


Figure 1: Optimal Sharpe Ratios (S^{OPT}) of Feasible Arbitrage Portfolios

Note: The figure reports optimal Sharpe ratios of feasible arbitrage portfolios in model (4), 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 on Figure 1 are only a fraction of the corresponding (infeasible)

Sharpe ratios, $S^\star = \sqrt{\alpha^\intercal(\Sigma_u)^{-1}\alpha} = \mu/\sigma\sqrt{\rho N}$, as shown by Figure 2. The pattern we see from Figure 2 agrees with theoretical predictions of Corollary 2. 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 $S^{\text{OPT}}/S^\star$ 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 $S^\star$ and S^{OPT} widens. In almost all empirically relevant scenarios, $S^\star$ is largely exaggerated.

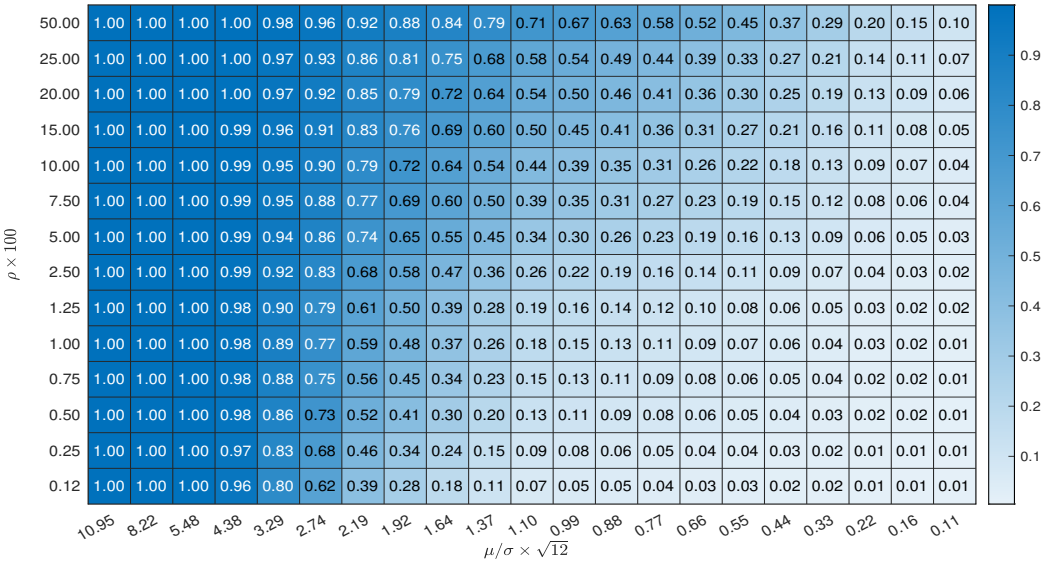


Figure 2: Ratios between S^{OPT} and $S^\star$

Note: The figure reports the ratios of optimal Sharpe ratios between feasible and infeasible arbitrage portfolios. The simulation setting is based on model (4), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

We now turn to the comparison of Sharpe ratios of optimal feasible arbitrage portfolios with those achieved by alternative strategies. Figure 3 compares with the cross-section regression approach in Section 2.7.1, Figure 4 with the B-H based procedure given by Section 2.7.2, and Figure 5 with LASSO given by Section 2.7.3, respectively.

According to Proposition 3, the optimal portfolio dominates the cross-sectional regression based portfolio if (17) holds. This dominance regime is bounded by a vertical line (as implied by the first inequality) and a cubic curve (as implied by the second inequality), which is visible from Figure 3 (black numbers on the heatmap). As $\mu/\sigma \times \sqrt{12}$ approaches 1.0 (a vertical line) from the right or the upper left corner, the gap between the two Sharpe ratios shrinks. Intuitively, when a large number of signals are clearly separable from the null (top left corner), the statistical inference becomes simpler so that the cross-sectional regression estimator of α is sufficient for building optimal portfolios. On the other hand, as the signal strength vanishes (the right vertical boundary), the

relative performance of the regression approach improves because it is equivalent to a ridge penalized regression that works well when all signals are weak and almost indistinguishable from noise. Figure 1 shows that the DGPs with respect to parameters for which the cross-sectional regression approach is strongly dominated by our optimal strategy are associated with realistic Sharpe ratios.

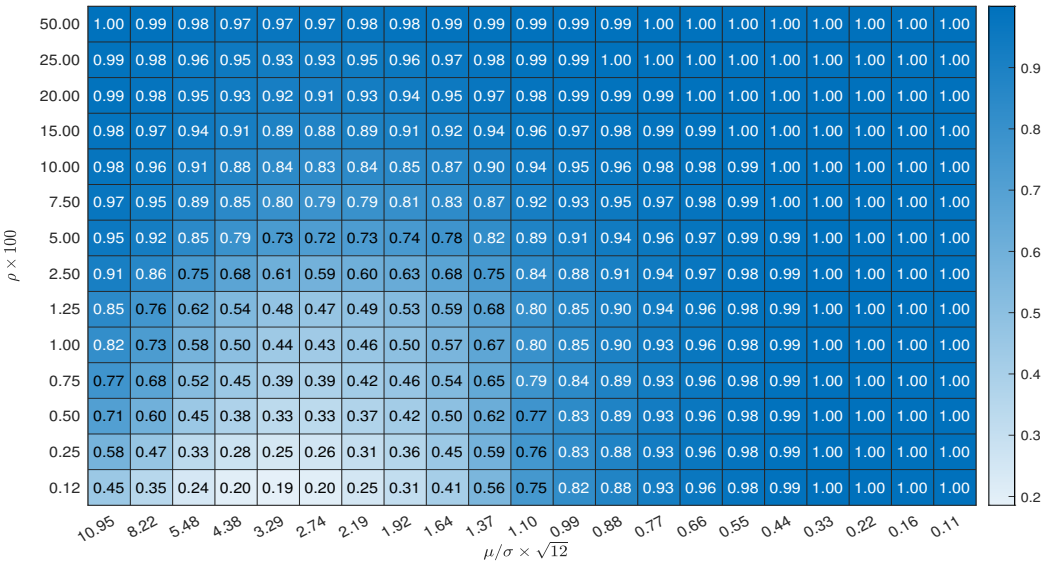


Figure 3: Ratios between S^{CSR} and S^{OPT}

Note: The figure reports the ratios between the Sharpe ratios of the OLS based portfolio and the feasible optimal arbitrage portfolio. The simulation setting is based on model (4), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

Similarly, the B-H procedure cannot achieve the optimal Sharpe ratio, as shown by Figure 4. According to Proposition 4, the gap between the optimal Sharpe ratio and the B-H approach largely depends on the signal strength. As long as $T\mu^2\sigma^{-2} \to 0$, the inequality (22) holds (since $\rho < 1$), the B-H procedure achieves the optimality. These scenarios correspond to the white values on Figure 4, where the border of the dominant region is located near the vertical line at $\mu/\sigma\sqrt{12} = 2.19$. Intuitively, the B-H is effective in singling out strong signals, so it leads to almost optimal portfolios as long as all signals are strong. However, when signals are weak, the B-H procedure, which amounts to hard-thresholding, performs worse than the cross-sectional regression, since in this case the embedded ridge regularization in the latter is more appropriate than hard-thresholding. As shown by Figure 1, even if alphas are individually weak, their empirical relevance should not be ignored because their collective contribution to the portfolio's Sharpe ratio can be highly non-trivial. The B-H approach is overly conservative compared to alternatives in this parameter regime.

Last but not least, Figure 5 presents the result for LASSO. This approach involves a tuning parameter, which calls for a cross-validation procedure. We adopt an infeasible and theoretically optimal tuning parameter, λ, which maximizes S^{LASSO}, making this approach a stronger competitor. Even though Proposition 5 suggests that LASSO is not uniformly optimal, it performs quite well,

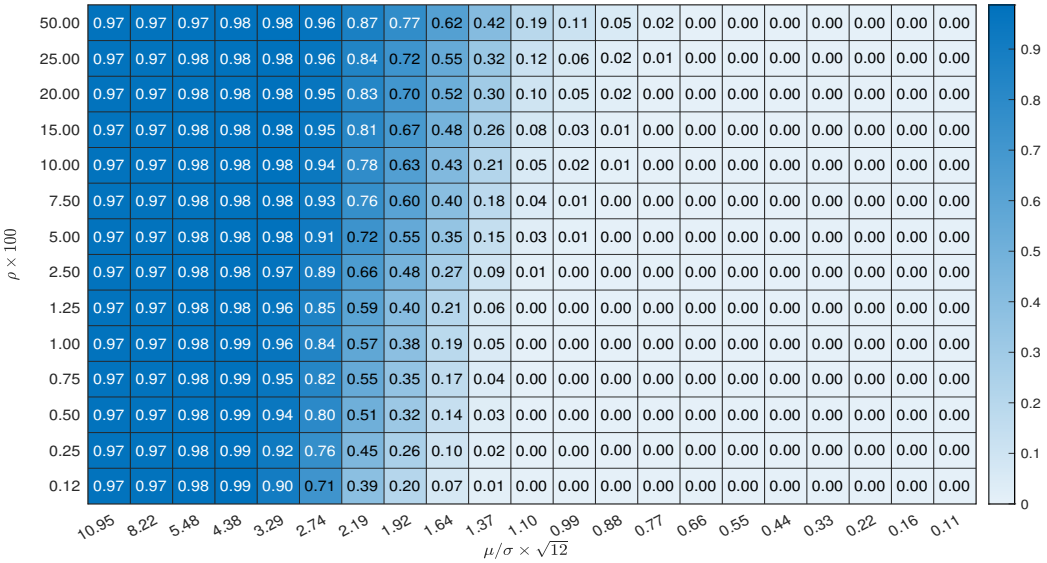


Figure 4: Ratios between S^{BH} and S^{OPT}

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 simulation setting is based on model (4), in which a $100 \times \rho\%$ of assets have alphas that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$.

achieving the optimal Sharpe ratio in almost all regimes. Intuitively, when signals are very strong, LASSO behaves like a hard-thresholding selector, as shrinkage does not play too much a role. When signals are rather weak, LASSO behaves like Ridge, because shrinking these signals does not change the fact that they are almost indistinguishable from noise.

3.2 Comparison of Portfolio Strategies in Finite Sample

We now compare the finite sample performance of our portfolio estimators over different DGPs. For any given parameter value $(\mu/\sigma, \rho)$ in a DGP, we estimate the portfolio weights, $\widehat{w}^{\text{OPT}}$, using our Algorithm 1, and calculate the resulting (theoretical) Sharpe ratio: $\widehat{w}^{\text{OPT}\intercal}\mu/\sqrt{\widehat{w}^{\text{OPT}\intercal}\Sigma_u^{-1}\widehat{w}^{\text{OPT}}}$. We then calculate the average Sharpe ratio over all Monte Carlo repetitions. Our approach requires a tuning parameter k_n. For robustness, we report results based on three parameter values $(0.5k_n, k_n, 2k_n)$ with $k_n = 0.25$. We repeat this exercise for the CSR, B-H, and LASSO methods for comparison.

In light of Theorem 2, a sensible choice of the estimation error can be written as:

$$\text{Err}^{\text{A}}(\mu/\sigma, \rho) = |\widehat{S}^{\text{A}} - S^{\text{OPT}}|/(1 + S^{\text{OPT}}),$$

where A denotes OPT, CSR, BH, or LASSO, and the dependence of $\widehat{S}^{\text{A}}$ and S^{OPT} on μ/σ and ρ is omitted. When S^{OPT} is large (i.e., $>> 1$), this error is in percentages relative to S^{OPT}; when S^{OPT} is small (i.e., $o_{\text{P}}(1)$), the error is measured in terms of the absolute difference. The error is

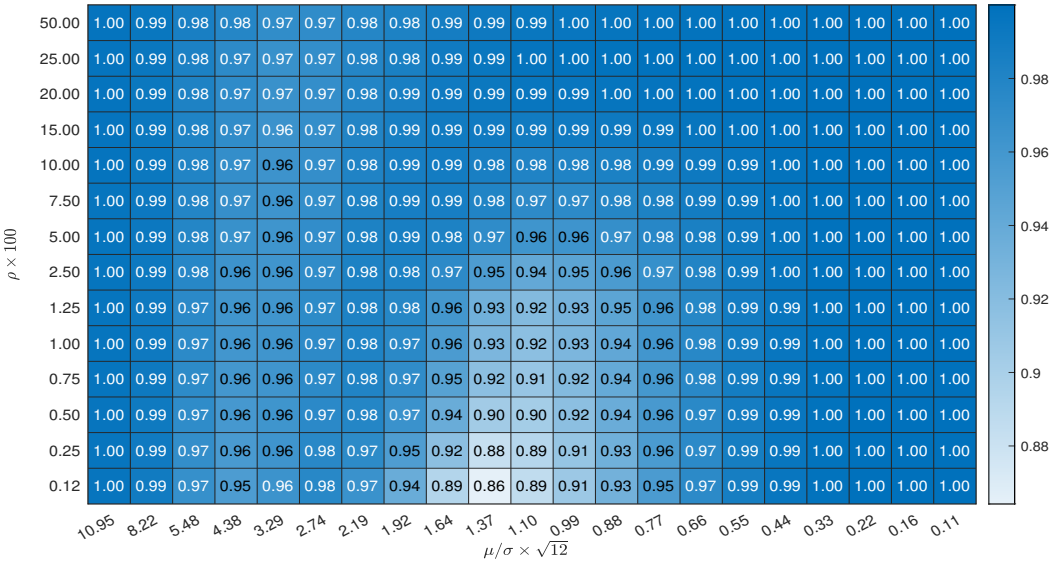


Figure 5: Ratios between S^{LASSO} and S^{OPT}

Note: The figure reports the ratios between the Sharpe ratios of the LASSO based portfolio and the feasible optimal arbitrage portfolio. The simulation setting is based on model (4), 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^{LASSO}.

defined this way because S^{OPT} itself can diverge or diminish depending on different parameters in the simulated DGPs.

Table 1 reports the maximal error over all values of μ/σ and ρ given in Section 3.1. The results show that OPT has a smaller error in almost all cases for all tuning parameters than CSR, BH, or LASSO. As T increases from 10 years to 40 years, the maximum error drops from 0.377 to 0.263 in the case of $N = 1,000$ for $k_n = 0.25$, whereas CSR, BH and LASSO stay above 0.44. The maximal error for CSR is achieved at the lower left conner of Figure 1, where signals are strong but rare; for BH, the worst performance occurs around the upper right corner, where many weak signals exist; for LASSO, the worse is near the bottom but in the middle, where signals are neither too strong nor too weak.

3.3 Finite Sample Performance of the Infeasible Sharpe Ratio Estimator

Finally, Figure 6 reports the estimation error $\left|\widehat{S}^{\star} - S^{\star}\right| / \left(1 + S^{\star}\right)$ in simulations. The result confirms the consistency result given by Proposition 2. The error is relative when $S^{\star}$ is large or moderate ($>> 1$). We find the relative error is around 1% towards the left top corner. For DGPs near the bottom right corner of Figure 6, $S^{\star}$ vanishes as shown by Figures 1 and 2, the error becomes absolute ($S^{\star} << 1$) and is moderately small given the sample size and the cross-sectional dimension.

	$N = 1,000$, Monthly			$N = 3,000$, Monthly			$N = 1,000$, Daily		
	$T = 10$	$T = 20$	$T = 40$	$T = 10$	$T = 20$	$T = 40$	$T = 10$	$T = 20$	$T = 40$
	0.385	0.332	0.289	0.442	0.367	0.320	0.449	0.440	0.408
OPT	0.377	0.309	0.263	0.437	0.333	0.282	0.411	0.382	0.356
	0.381	0.282	0.233	0.446	0.318	0.247	0.370	0.334	0.303
CSR	0.540	0.489	0.441	0.618	0.570	0.515	0.537	0.485	0.427
BH	0.742	0.703	0.651	0.814	0.789	0.748	0.760	0.715	0.657
LASSO	0.537	0.488	0.440	0.615	0.568	0.512	0.536	0.483	0.426

Table 1: Sharpe Ratio Comparison in Simulations

Note: This table reports the maximum error, defined by $\sup_{\mu/\sigma,\rho} \mathrm{Err}^{\mathrm{A}}(\mu/\sigma, \rho)$, where A denotes either OPT, or CSR, or BH, over all values of μ/σ and ρ in Figure 1, for several choices of N, T (in years), and data frequencies. The first three rows correspond to the OPT approach with three different values of tuning parameters, $0.5k_n$, k_n, and $2k_n$, respectively, where $k_n = 0.25$. The BH approach controls false discovery rate at a level 5%. The LASSO approach uses the optimal (infeasible) tuning parameter that optimizes S^{LASSO}.

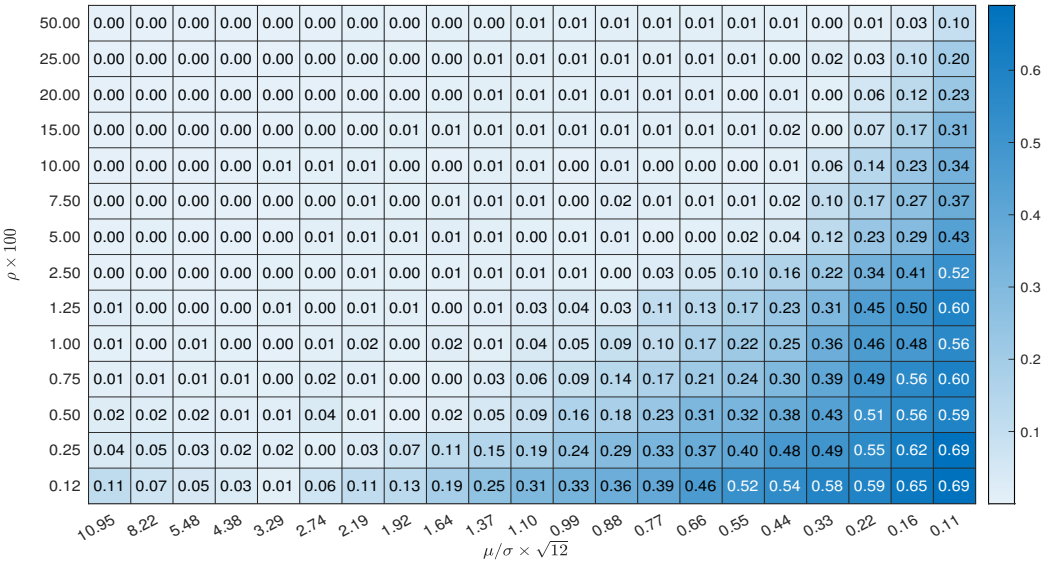


Figure 6: Comparison between $\widehat{S}^\star$ and $S^\star$

Note: The figure reports the error between $\widehat{S}^\star$ and $S^\star$ defined as $\left|\widehat{S}^\star - S^\star\right| / \left(1 + S^\star\right)$. The simulation setting is based on model (4), in which a $100 \times \rho\%$ of assets have αs that correspond to an annualized Sharpe ratio $\mu/\sigma \times \sqrt{12}$. In this experiment, $N = 1,000$ and $T = 20$ years.

4 Empirical Analysis of US Equities

To demonstrate the empirical relevance of the statistical limit of arbitrage, we study US monthly equity returns from January 1965 to December 2020.

4.1 Data Preprocessing

We adopt a multi-factor model with 16 characteristics and 11 GICS sectors, which are selected to incorporate empirical insight 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, see **?** for construction details.

We download the monthly return data for individual equities from CRSP. We take a number of steps to preprocess the data. First, we single out delisted stocks, and attach delisting returns as their last returns (on the delisting months). Next, we merge the returns data with the aforementioned characteristics database using permnos. The total number of unique permnos on average per month is 6,536. We then apply the usual filters (share codes 10 and 11 and exchange codes 1, 2, and 3) to the database, to eliminate (part of) the sampling periods for stocks that fail to meet these criteria. The remaining average number of stocks per month is 4,756. For stocks whose returns are missing for more than 3 months, we eliminate the missing periods, otherwise we fill the missing returns by zeros.

We now deal with missing characteristics. We start by removing all characteristics data for any stocks since their delisting months. We then fill missing GICS codes with the corresponding stocks' most recent records prior to their missing dates. Stocks without any GICS codes over the entire sample period are eliminated. If the GICS codes become available later in the sample for some stocks, their sample prior to the first dates when GICS become available are eliminated, which mainly occurs prior to 1990. With GICS information, we adopt a two-step procedure to fill in other missing characteristics. For any missing value in a stock's characteristic, we fill it with the sector-wise median of this characteristic each month. If a characteristic's values are not available for an entire sector in a certain month, we fill them with this characteristic's cross-sectional median over all stocks in this month. After data preprocessing, the final average number of stocks per month is reduced to 4,067.

The resulting panel is not balanced, because we do not fill in missing data before a stock's IPO or after its delisting. Our approach to filling missing data thereby avoids forward-looking bias. An empirical issue we have not discussed so far is time-varying exposures and missing data. We adapt Algorithm 1 and estimators (14) and (15) to this case in Appendix B.

4.2 Model Performance

At the end of each month, we run cross-sectional regressions of next month returns onto the 27 cross-sectional predictors (including the intercept). We do so using all stocks in the current month's cross sections. Following **?**, the 16 characteristics are rank-normalized within each cross-section, alleviating the impact of extreme outliers in characteristics, though this barely changes any follow-up results.

Figure 7 plots the time series of the cross-sectional regression R^2s over time. The R^2 has been on the decline since the beginning of the sample till 1990s. This coincides with the period when the number of stocks in the US equity markets increases. The R^2s are moderately low, with an average of 8.25%. The low R^2s suggest that a substantial portion of cross-sectional variation of individual equity returns is idiosyncratic noise. Therefore, learning alphas from residuals of the factor model is an incredibly difficult statistical task.

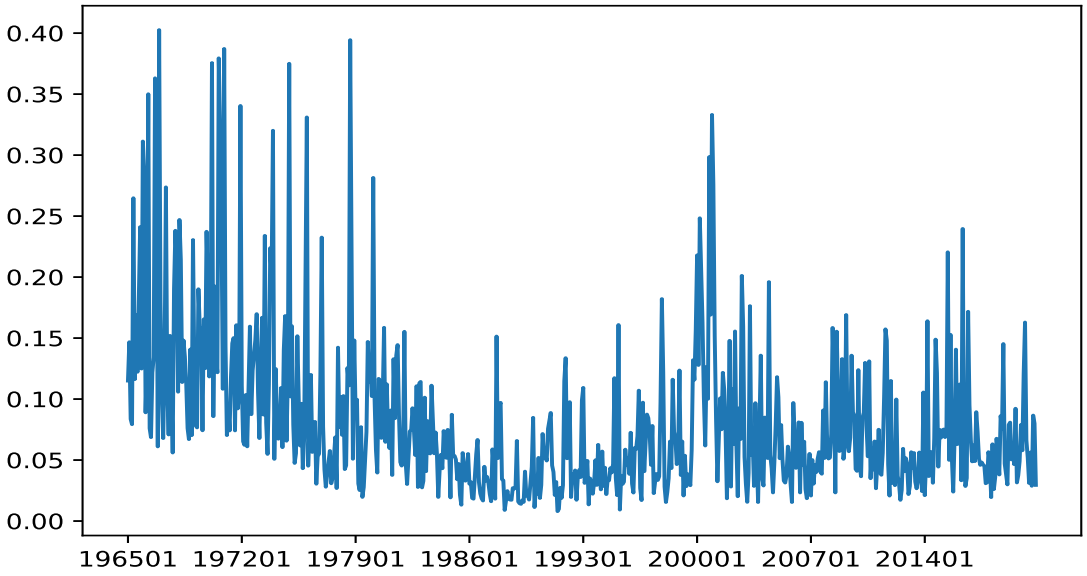


Figure 7: Time-series of the Cross-sectional R^2s

4.3 Rare and Weak Alphas

We now study the statistical properties of 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 impose that all residuals have at least 60 observations. This ensures enough sample size for inference on alpha, although the distribution of alphas' t-statistics turns out not sensitive to this requirement. Figure 8 provides histograms of the t-statistics and Sharpe ratios for alphas of all 12,415 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 6.35% of the t-statistics exceed 2.0 in magnitude, and more than 0.63% exceed 3.0. This suggests that truly significant alphas are extremely rare. Moreover, the largest Sharpe ratio of all individual stocks' alphas is rather modest, about 1.699. Only 0.505% 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.

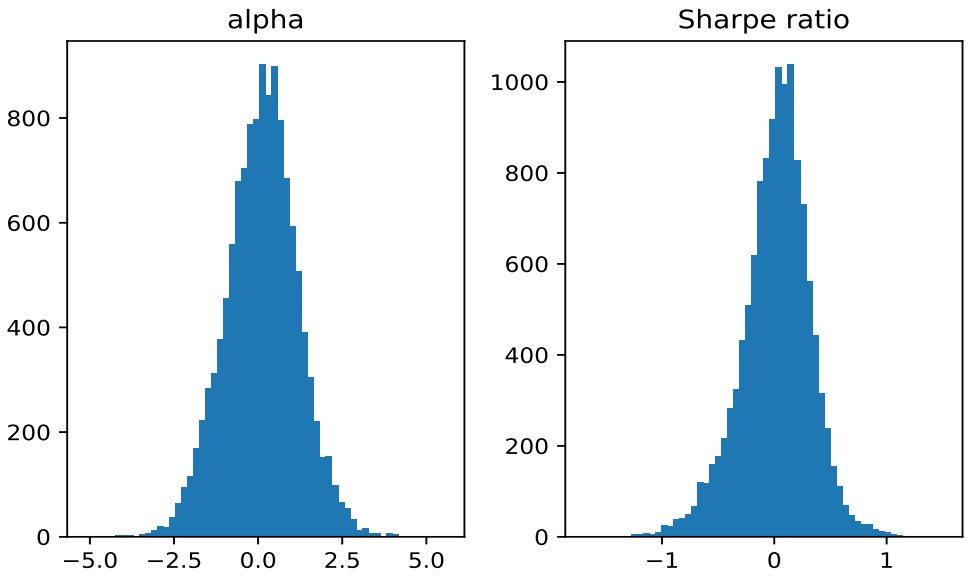


Figure 8: Histograms of the t-Statistics and Sharpe Ratios of Estimated Alphas

Note: The figure provides the histograms of the 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 12,482.

4.4 Performance of Arbitrage Portfolios

Throughout we assume alphas do not vary over time. If alphas are driven by some observable characteristics, then it is possible to construct a factor using these characteristics via cross-sectional regressions, which turns "alpha" into risk premia. In this regard, alphas are meaningless without reference to a specific factor model. Extracting more "factors" out of alphas would lead to even smaller arbitrage profits.

We now compare arbitrage portfolios based on various strategies, including the optimal strategy, the cross-sectional regression (CSR) approach, the multiple-testing based procedure (BH), and LASSO approach. The ridge approach is omitted, since it is equivalent to the CSR.

Specifically, at the end of each month, we build optimal portfolio weights using these strategies. We only invest in stocks with a continuous record for at least 96 months. We rebalance these portfolios at the end of each month, with weights recalculated using a 120-month rolling window. Both Lasso and the optimal strategy require a tuning parameter. Out of the 10-year rolling window, we leave the last 2 years as the validation sample for tuning parameter selection. As expected, optimal tuning parameter is difficult to select, which undermines the performance of both strategies.

All these strategies yield similar Sharpe ratios. BH and OPT tie for the top of the chart, yielding 0.497 and 0.496, respectively, followed by CSR that scores 0.450. The LASSO approach only obtains 0.384. The Sharpe ratios of different strategies are not influenced by risk aversion, though the cumulative returns are. 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 9.

Closely examining these strategies reveals more insight. BH is highly conservative. Out of 46 years of out-of sample trading months (1975/01 - 2020/12), 289 months have no trading activities. The largest number of stocks selected for trading in a month is 10, and the average over all non-zero periods is 2.43. In contrast, CSR trade all stocks that meet our trading criteria, with an average of 2,366 stocks per month. OPT almost does so, with an average of 2,359. The number of stocks traded by LASSO is rather volatile, varying between none and all stocks from month to month, with an average of 757.6 per month. This is likely caused by the noise in the tuning process.

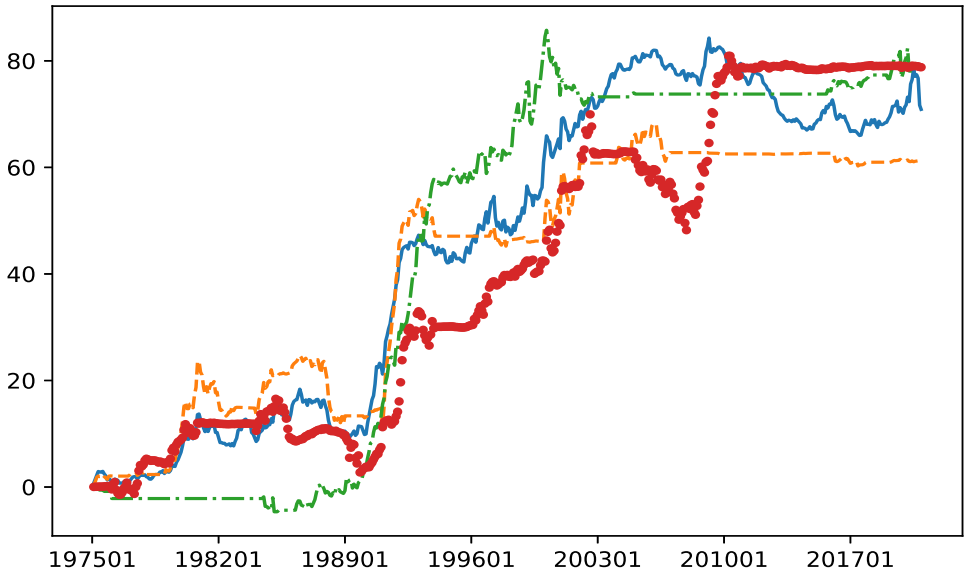


Figure 9: Normalized Cumulative Returns of Arbitrage Portfolios

Note: This figure compares the cumulative returns of OPT (red dotted), CSR (blue solid), BH (green dot-dashed), and LASSO (orange 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.

We also calculate the perceived Sharpe ratios using (14), and provide a time-series plot of $\widehat{S}^\star$ in Figure 10. We also compare it with the biased estimates $\widetilde{S}^\star$ using (15). We observe a huge gap between the estimated perceived Sharpe ratios using these formulae. As predicted by Proposition 2, $\widetilde{S}^\star$ overestimates $S^\star$, though it guarantees positive values. Our estimate $\widetilde{S}^\star$ is averaged around 2.55 (we truncate negative estimates by 0), but can sometimes exceed 7.5. These estimates are far greater than the feasible Sharpe ratios we obtain for any of these strategies. That said, even the infeasible Sharpe ratios can be as low as 0 for certain periods of the sample. The feasible portfolio returns seem in agreement with the prediction. For instance, OPT, LASSO, and BH's cumulative returns are almost flat post 2010, whereas CSR has negative returns.

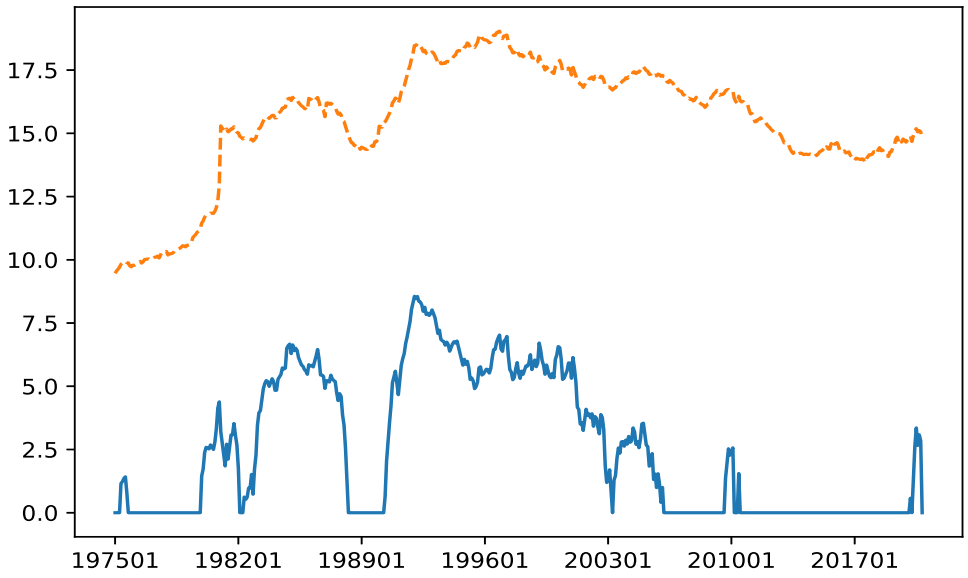


Figure 10: Time Series of Sharpe Ratios

Note: The figure compares the optimal feasible (blue solid) and infeasible (orange dashed) Sharpe ratios, estimated based on a rolling window of 120 months.

5 Conclusion

Taking stock, our paper provides a new theoretical framework to understanding the implications of statistical learning in asset pricing. In the age of big data, rational expectations assumption often fails to retain its relevance in practice, and hence understanding its limitation and the role of statistical learning is vitally important. We introduce new econometric tools in the spirit of nonparametric empirical Bayes, which could be adopted in other contexts.

The empirical message should be confined within the context of monthly rebalancing strategies via linear factor models. The gap between feasible and infeasible Sharpe ratios will further increase if arbitrageurs face additional statistical challenges, e.g., model misspecification, omitted factors, weak factors, large non-sparse idiosyncratic covariance matrix, etc. Consequently, the empirical gap should remain for any arbitrageurs, including those who engage in higher frequency trading or use more complex nonlinear models.

More broadly, existing literature have documented impressive Sharpe ratios on various machine learning based trading strategies. 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 in investment. Our paper provides a theoretical framework to understand the "upper bound" on the performance of any strategy in the specific context of arbitrage pricing theory, tying together this statistical limit with economic rationale.

On a side note, our theoretical and empirical analyses also have implications on the econometric

analysis in asset pricing. Examining the economic performance of asset pricing models is as important as and complementary to statistical tests. The criteria of a good statistical test are primarily statistical in nature, such as Type I and Type II errors, false discovery rate, etc; whereas in practice, it is the economic performance that agents in the economy fundamentally care about. There is often a wedge between these two objectives. For instance, a statistical procedure that guards against false discovery rate may be overly conservative for investment purpose; rejection by a powerful test statistic may not necessarily lead to the practical irrelevance of an economic theory. While the asset pricing literature has seen an explosion of statistical machine learning tools imported from other areas, we caution against their use without guidance of economics.