

NBER WORKING PAPER SERIES

TEST ASSETS AND WEAK FACTORS

Stefano Giglio
Dacheng Xiu
Dake Zhang

Working Paper 29002
<http://www.nber.org/papers/w29002>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

We benefited tremendously from discussions with seminar and conference participants at the University of California, Los Angeles, University of North Carolina at Chapel Hill, University of Miami, Temple University, University of Michigan, London School of Economics, Tilburg University, INSEAD Business School, Yale School of Management, Chicago Booth School of Business, Rotterdam, Rutgers University, Cornell University, Duke University, University of Texas at Dallas, ESSEC Business School, Chinese University of Hong Kong, University of Connecticut, Federal Reserve Board, ITAM Business School, Vienna Graduate School Finance, Paris Dauphine University, Goethe University, and 13th Annual SoFiE Conference. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Stefano Giglio, Dacheng Xiu, and Dake Zhang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Test Assets and Weak Factors
Stefano Giglio, Dacheng Xiu, and Dake Zhang
NBER Working Paper No. 29002
July 2021
JEL No. C58,G12

ABSTRACT

Estimation and testing of factor models in asset pricing requires choosing a set of test assets. The choice of test assets determines how well different factor risk premia can be identified: if only few assets are exposed to a factor, that factor is weak, which makes standard estimation and inference incorrect. In other words, the strength of a factor is not an inherent property of the factor: it is a property of the cross-section used in the analysis. We propose a novel way to select assets from a universe of test assets and estimate the risk premium of a factor of interest, as well as the entire stochastic discount factor, that explicitly accounts for weak factors and test assets with highly correlated risk exposures. We refer to our methodology as supervised principal component analysis (SPCA), because it iterates an asset selection step and a principal-component estimation step. We provide the asymptotic properties of our estimator, and compare its limiting behavior with that of alternative estimators proposed in the recent literature, which rely on PCA, Ridge, Lasso, and Partial Least Squares (PLS). We find that the SPCA is superior in the presence of weak factors, both in theory and in finite samples. We illustrate the use of SPCA by applying it to estimate the risk premia of several tradable and nontradable factors, to evaluate asset managers' performance, and to de-noise asset pricing factors.

Stefano Giglio
Yale School of Management
165 Whitney Avenue
New Haven, CT 06520
and NBER
stefano.giglio@yale.edu

Dake Zhang
Booth School of Business
University of Chicago
5807 S Woodlawn Ave
Chicago, IL 60637
dkzhang@chicagobooth.edu

Dacheng Xiu
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
dachxiu@chicagobooth.edu

A supplement is available at <http://www.nber.org/data-appendix/w29002>

1 Introduction

Inference on factor risk premia is a central element of empirical work in asset pricing. An essential role in this exercise is played by the set of test assets used in the estimation, yet little work has been dedicated to investigating rigorously and systematically how they should be chosen. In this paper, we show that there is an important connection between the selection of test assets and the long-standing problem of weak factors in asset pricing – factors to which the test assets have little or no exposure, resulting in a well-known failure in risk premia inference.

Central to understanding this connection is an alternative perspective on the issue of weak factors. We argue the strength or weakness of a factor should not be viewed as a property of the factor itself, as typical in the asset pricing literature; rather, it should be viewed as a property of the set of test assets used in the estimation. As an example, a liquidity factor may be weak in a cross-section of portfolios sorted by, say, size and value, but may be strong in a cross-section of assets sorted by characteristics that capture well exposure to liquidity. By exploiting this insight, we propose a new methodology for risk premia estimation, *supervised principal component analysis* (SPCA), which tackles the issue of weak factors via supervised test asset selection.

As discussed in the literature (e.g., Jagannathan and Wang (1998) and Giglio and Xiu (2021)), estimating and testing the risk premia of some factors requires properly controlling for all the other factors relevant to investors (whether they are observed or latent), in order to avoid an omitted variable bias. Importantly, the choice of test assets determines the strength not only of the factor of interest (e.g., liquidity), but also of *all* the other factors that drive the stochastic discount factor. We design the SPCA procedure using an iterative algorithm that uses the factor of interest to guide the selection of test assets. At the same time, the algorithm uses PCA to recover the relevant latent factors iteratively, thus controlling for potentially omitted factors. The integration of supervised selection and PCA yields a general methodology that is robust to the omission of factors, even when these omitted factors are weak.

In a nutshell, the procedure estimates the risk premium of a factor g_t as follows. We start from a large universe of potential test assets. In a first step of the procedure (selection step), we compute the univariate correlation of each asset's return with g_t . We select a relatively small portion of assets, only keeping those with sufficiently high correlation (in absolute value): these are assets that are particularly informative about the factor of interest g_t . We then compute the first principal component of these portfolios (PCA step), which will be our first estimated latent factor. Next, we remove via linear projection from both g_t and all the returns of the test assets the part explained by this first latent factor (projection). We then go back to the selection step, computing the univariate correlation of the *residuals* of the factor and the *residuals* of the assets from the projection step. Again, we select from the universe of test assets a subset for which this correlation is especially high, and compute the principal component of these residuals. This will be our second estimated latent factor. We then further remove (from g_t and the test assets) the part explained by this second

estimated factor as well, and iterate again on the residuals. We repeat this procedure p times, where p can be either a prior estimate of the number of factors in the data or can be regarded as a tuning parameter to be determined by some validation step. This procedure recovers from the data p latent factors that are informative about the factor of interest g_t . Importantly, the fact that at each iteration only test assets that are sufficiently correlated with the factor g_t are selected ensures that not only strong, but also weak factors (relative to the entire cross-section) are captured by the procedure – contrary to standard PCA that uses *all* assets at all steps to extract latent factors. Finally, a time-series regression of g_t on the p latent factors allows us to estimate the risk premium for g_t by linking it to the risk premia for these latent factors, yielding a consistent estimator of the risk premium of g_t .

The choice of test assets in the literature has mainly followed one of three approaches. The vast majority of the literature has adopted a “standard” set of portfolios sorted by a few characteristics, such as size and value, following the seminal work by [Fama and French \(1993\)](#). A second approach, taken more recently, e.g., [Kozak et al. \(2020\)](#), has been to expand this cross-section to include portfolios sorted by a much larger set of characteristics discovered in the last decades, on the order of hundreds of portfolios. Finally, a third approach, see, e.g., [Ang et al. \(2006\)](#), has been more “targeted” around the specific factor of interest: sorting assets into portfolios by their estimated exposure to the factor, and then estimating risk premia using only these sorted portfolios, that is, using a small cross-section expected to be particularly informative about that factor.

It is useful to contrast the asset selection procedure of SPCA with the three standard approaches to choose test assets summarized above. Using a standard, small cross-section (like the size- and value-sorted portfolios) to estimate risk premia has the problem that except for size and value, which are strong factors in this cross-section, many other factors are weak: the test assets do not contain sufficient information to identify their risk premia. Using a large cross-section of test assets (the second approach) may appear, on the surface, to address this issue: these assets contain returns that are exposed to a large number of underlying factors. However, and importantly, if only a *few* of those many assets are exposed to some factor, whereas most others are not, that factor will, again, be weak in this large cross-section, disrupting inference on the risk premium. Finally, the third approach – building targeted portfolios of assets sorted by the exposure to the factor of interest – is affected by the omitted factor problem, since it considers univariate exposures only (exposures with the factor of interest may also capture correlated exposures to other risks in the economy); in general, it will fail in a multi-factor context.

In the paper, we derive the asymptotic properties of SPCA, in a setting that allows for weak factors and test assets with highly correlated risk exposures. The latter scenario potentially involves the same (asymptotically) rank-deficiency issue as weak factors. We also analyze in this setting alternative estimators that have been proposed in the recent literature, which rely on PCA, Ridge, Lasso, and Partial Least Squares (PLS). We show that the PCA (and some other variations of it), Ridge, and PLS are inconsistent in the presence of weak factors, that the Lasso approach is consistent

for the estimation of the stochastic discount factor (SDF), and hence risk premia estimation, but is not as efficient as SPCA in general. Additionally, we perform an extensive set of simulations to study the performance of SPCA in different scenarios. These simulations isolate issues with the standard two-pass regressions, so that we can easily compare SPCA with other estimators. The simulations confirm the robustness of SPCA to both omitted factors and weak factors, as well as measurement error, which SPCA also tackles.

Finally, we illustrate the use of SPCA for estimating risk premia of a variety of tradable and nontradable factors proposed in the asset pricing literature. We use the large cross-section of test portfolios produced by [Chen and Zimmermann \(2020\)](#) and [Hou et al. \(2020\)](#), covering more than 900 and 1600 portfolios, respectively, for the period 1976-2020. We apply SPCA to estimate the factor risk premia, and study the out-of-sample performance of SPCA. We also explore the robustness of SPCA to the weakness of factors, by artificially changing the set of test assets used in the estimation: for example, we show that SPCA is able to recover the risk premium for momentum even when momentum assets are removed from the original set of test assets (and therefore the momentum factor is weak in the cross-section). In addition to estimating risk premia, we explore additional applications of SPCA, including the performance evaluation of money managers and the removal of measurement error (de-noising) of factors.

This paper builds on a large literature on risk premia and factor model estimation and their limits in the presence of weak and omitted factors. The seminal contribution of [Kan and Zhang \(1999\)](#) shows that the inference on risk premia estimates from Fama-MacBeth regressions becomes invalid when a “useless” factor – a factor to which test assets have zero exposures – is included in the model. [Kleibergen \(2009\)](#) further points out the failure of the standard inference if betas are relatively small.¹ This issue is quite relevant in practice because many test assets are not very sensitive to macroeconomic shocks. Moreover, the same rank-deficiency problem arises when betas are collinear, that is, some factors are redundant in terms of explaining the variation of expected returns. This is again a relevant issue in practice due to the existence of hundreds of factors discovered in the literature, see, e.g., [Harvey et al. \(2016\)](#), many of which are close cousins and do not add any explanatory power ([Feng et al. \(2020\)](#)). The weak factor problem appears to be caused by having seemingly more factors than necessary, which is why some suggest eliminating such factors ([Bryzgalova \(2015\)](#)) or shrinking their risk premia estimates ([Bryzgalova et al. \(2019\)](#)), so as to improve the estimates for strong factors. We instead argue that the weak factor problem is fundamentally an issue of test asset selection. Since weaker factors may still be priced, our solution is to accommodate them using an adapted procedure with carefully selected test assets.²

¹Also related is [Pesaran and Smith \(2019\)](#), who investigate the impact of factor strength and pricing error on risk premium estimation. They point out that the conventional two-pass risk premium estimator converges at a lower rate as the factors become weaker.

²It is worth noting that whereas some theories assume that only strong factors can be priced, this is not true in general for two reasons. First, many theoretical models – e.g., the consumption-CAPM – are silent on what assets are traded in equilibrium, and if markets are incomplete, it may very well be that some priced factors may not be reflected in many of the assets that are traded. Second, even if investors may have access to many assets exposed to a particular

Several recent papers have proposed different methodologies to deal with weak factors. [Lettau and Pelger \(2020\)](#) are among the first to study the issue of weak latent factors in a related problem, that of estimation of the SDF. They propose an estimator of the SDF in the presence of weak factors, which generalizes PCA with a penalty term that accounts for expected returns; they refer to the estimator as risk premium PCA, or rpPCA. Their objective is different from ours, but the SDF estimated using this procedure can still be used to estimate risk premia, since risk premia are covariances with the SDF. Whereas this estimator features desirable properties as explored by [Lettau and Pelger \(2020\)](#), we show that it is inconsistent for estimating risk premia in the weak-factor setting we consider.³ [Anatolyev and Mikusheva \(2021\)](#) propose an complementary approach to dealing with weak factors, based on sample-splitting and instrumental variables. This alternative procedure works well to address the weak factor bias, though it does not deal with omitted priced factors or with measurement error in the factors.

Our paper also relates to a literature that has explored different methods to form portfolios to test asset pricing models, like [Ahn et al. \(2009\)](#) or [Bryzgalova et al. \(2020\)](#). These methods are useful in helping to build or expand the starting cross-section for SPCA. In this paper, we use the simpler approach of working with an existing large cross-section of portfolios sorted by firm characteristics, as in [Chen and Zimmermann \(2020\)](#) and [Hou et al. \(2020\)](#). It also relates to a growing strand of econometrics literature on weak factor models, like [Bai and Ng \(2008\)](#) and [Huang et al. \(2021\)](#). Our SPCA approach shares the spirit of these approaches, but is more involved because we do not assume all factors are of the same strength, which thereby requires multiple selection steps. Also, our focus is on risk premia estimation instead of forecasting, for which we also provide inference. Also related are papers that propose estimators of factor count and strength, like [Freyaldenhoven \(2019\)](#) and [Bailey et al. \(2020\)](#).

The concept of supervised-PCA originated from a cancer diagnosis technique applied to DNA microarray data by [Bair and Tibshirani \(2004\)](#), and was later formalized by [Bair et al. \(2006\)](#) in a prediction framework, in which some predictors are not correlated with the latent factors that drive the outcome of interest. [Bair et al. \(2006\)](#) suggest a screening step using marginal correlations between predictors and the outcome variable to select the subset of useful predictors, before applying the standard PCA to this subset. They prove the consistency of this so-called SPCA procedure, but relying on a restrictive identification assumption that any important predictor must also have a substantial marginal correlation with the outcome. We provide several examples of multivariate

factor, the econometrician may not, making the factor weak for the set of test assets available to the econometrician.

³[Lettau and Pelger \(2020\)](#) focus their analysis on the case where factors are extremely weak – so much so that they are not statistically distinguishable from idiosyncratic noise. In that case, no estimator can be consistent for either risk premia or the SDF. They show that indeed, rpPCA does not recover consistently the SDF, but it correlates with the SDF more than the SDF estimator obtained from standard PCA. Rather than focusing on this extreme case of weak factors, our theory covers a range of factor weaknesses, which includes the cases from strong to very weak, and which permits consistent estimation of factors and risk premia. Formally, we study the case where the minimum eigenvalues of the factor component in the covariance matrix of returns diverges whereas the largest eigenvalue due to the idiosyncratic errors is bounded.

factor models in which this assumption fails. While the screening step of our SPCA procedure shares the spirit with theirs (in the sense that their outcome variable is our factor of interest, and their predictors are our test assets), our projection step and the subsequent iteration procedure are new, and are introduced precisely to eliminate the strong identification assumption used in the existing statistics literature. Also, our focus is not on prediction per se, but instead on inference on parameters (i.e., risk premia), which involves an additional step and more intricate analysis for the asymptotic theory.

The paper is organized as follows. Section 2 first sets up the notation and model (Sections 2.1 and 2.2), then discusses the inconsistency of existing estimators in the presence of weak factors (Section 2.3), provides our methodology (Sections 2.4 and 2.5) and finally the inference theory (Section 2.6). Section 3 provides simulation evidence, followed by an empirical study in Section 4. The appendix provides technical details.

2 Methodology

2.1 Notation

Throughout the paper, we use (A, B) to denote the concatenation (by columns) of two matrices A and B . e_i is a vector with 1 in the i th entry and 0 elsewhere, whose dimension depends on the context. ι_k denotes a k -dimensional vector with all entries being 1, and \mathbb{I}_d denotes the $d \times d$ identity matrix. For any time series of vectors $\{a_t\}_{t=1}^T$, we denote $\bar{a} = \frac{1}{T} \sum_{t=1}^T a_t$. In addition, we write $\bar{a}_t = a_t - \bar{a}$. We use the capital letter A to denote the matrix (a_1, a_2, \dots, a_T) , and write $\bar{A} = A - \bar{a}\iota_T^\top$ correspondingly. We denote $\mathbb{P}_A = A(A^\top A)^{-1}A^\top$ and $\mathbb{M}_A = \mathbb{I}_d - \mathbb{P}_A$, for some $d \times T$ matrix A . We use $a \vee b$ to denote the max of a and b , and $a \wedge b$ as their min for any scalars a and b . We also use the notation $a \lesssim b$ to denote $a \leq Kb$ for some constant $K > 0$ and $a \lesssim_p b$ to denote $a = O_p(b)$. If $a \lesssim b$ and $b \lesssim a$, we write $a \asymp b$ for short. Similarly, we use $a \asymp_p b$ if $a \lesssim_p b$ and $b \lesssim_p a$.

We use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote the minimum and maximum eigenvalues of A , and use $\lambda_i(A)$ to denote the i -th largest eigenvalue of A . Similarly, we use $\sigma_i(A)$ to denote the i th singular value of A . We use $\|A\|_1$, $\|A\|_\infty$, $\|A\|$, and $\|A\|_F$ to denote the \mathbb{L}_1 norm, the \mathbb{L}_∞ norm, the operator norm (or \mathbb{L}_2 norm), and the Frobenius norm of a matrix $A = (a_{ij})$, that is, $\max_j \sum_i |a_{ij}|$, $\max_i \sum_j |a_{ij}|$, $\sqrt{\lambda_{\max}(A^\top A)}$, and $\sqrt{\text{Tr}(A^\top A)}$, respectively. We also use $\|A\|_{\text{MAX}} = \max_{i,j} |a_{ij}|$ to denote the \mathbb{L}_∞ norm of A on the vector space. When a is a vector, we use $\|a\|_0$ to denote $\sum_i 1_{\{a_i \neq 0\}}$. We also denote $\text{Supp}(a) = \{i : a_i \neq 0\}$. Finally, we use $[N]$ to denote the set of integers: $\{1, 2, \dots, N\}$. For an index set $I \subset [N]$, we use $|I|$ to denote its cardinality. We use $A_{[I]}$ to denote a submatrix of A whose rows are indexed in I .

2.2 Model Setup

We study a standard linear factor model setup. Suppose that an $N \times 1$ vector of test asset excess returns, r_t , follows:

$$r_t = \beta\gamma + \beta v_t + u_t, \quad \mathbb{E}(v_t) = \mathbb{E}(u_t) = 0 \text{ and } \text{Cov}(v_t, u_t) = 0, \quad (1)$$

where β is an $N \times p$ matrix of factor exposures, v_t is a $p \times 1$ vector of factor innovations, and u_t is an $N \times 1$ vector of idiosyncratic errors.⁴ The v_t vector is unobservable, even though it may include factor innovations of observable factors, f_t , i.e., $v_t = f_t - \mu_f$, since μ_f is an unknown parameter.

In order to study the statistical properties of risk premia estimators in the presence of weak factors, we first define our asymptotic scheme. We will assume that both N and T go to ∞ , whereas p is fixed. The $p \times p$ factor covariance matrix Σ_v is asymptotically non-singular in the sense that $1 \lesssim \lambda_{\min}(\Sigma_v) \lesssim \lambda_{\max}(\Sigma_v) \lesssim 1$. This assumption is rather weak as it only rules out factors whose risks are (asymptotically) negligible or exploding. We also maintain the assumption that $\|\Sigma_u\| \lesssim 1$, so that there exists no factor structure in the residuals u_t . This condition is useful for identification purposes, and ensures that all factors must be distinguishable from the idiosyncratic errors, regardless of their strength, which we turn to next.

In this setting, a factor’s strength is entirely determined by test assets’ exposures to it, since all factors have non-negligible or non-exploding risks. In light of this, the strength of a factor is context specific — the selection of test assets dictates its strength. For instance, a momentum factor could be a strong factor for momentum-sorted portfolios, but this factor may be weak with portfolios sorted by size or value as test assets, because the latter portfolios may diversify the exposure to the momentum factor.

In the econometrics literature on factor models, the most prevalent assumption adopted by, e.g., [Bai and Ng \(2002\)](#), is that all factors are strong or pervasive, that is, $\lambda_i(\beta^\top \beta) \asymp N$ for $i = 1, 2, \dots, p$, which dominates the strength of the idiosyncratic component, as measured by $\|\Sigma_u\|$. Our focus is on the regime of *weak factors*, which covers a wide range of factor strength. In particular, the norm of columns of β is allowed to diverge at different and slower rates, which will be made more precise later.

The fact that weak factors are relevant in practice can be illustrated from a scree plot of eigenvalues of returns (for example, see [Figure 3](#) based on the large cross-section we use in our empirical analysis). Factors with a spectrum of strength, as indicated by various magnitudes of eigenvalues, are clearly present. Except for the first one or two eigenvalues, there is not a clear-cut gap between the next few eigenvalues (that would correspond to weaker factors) and the remaining eigenvalues that correspond to idiosyncratic components.

⁴Our model is set up for portfolios as test assets. To generalize this model for individual stocks, more structures should be imposed to address time-varying risk exposures, see, e.g., [Gagliardini et al. \(2016\)](#), [Kelly et al. \(2019\)](#), and [Kim et al. \(2020\)](#).

We develop our discussion of weak factors in the context of two standard asset pricing exercises: the estimation of risk premia and the recovery of the stochastic discount factor (SDF). In this model, an SDF can be defined in terms of asset pricing factors v_t as

$$m_t = 1 - \gamma^\top \Sigma_v^{-1} v_t, \quad (2)$$

where Σ_v is the covariance matrix of factor innovations. It also makes sense to consider the SDF represented in terms of the tradable test asset returns:

$$\tilde{m}_t = 1 - b^\top (r_t - E(r_t)), \quad (3)$$

where b is an $N \times 1$ vector of SDF loadings which satisfies $E(r_t) = \Sigma b$, where Σ is the covariance matrix of r_t . The relationship between the two SDFs depends on the degree of completeness of markets. As will be shown later, these two forms of the SDF are asymptotically equivalent in the asymptotic scheme we consider, with the number of assets N going to infinity, so that there is no ambiguity with respect to which estimand we consider.

In addition to the SDF, we are also interested in risk premia of some observable factors, summarized in a $d \times 1$ vector, g_t . Following [Giglio and Xiu \(2021\)](#), we do not impose that g_t is part of or is identical to v_t ; instead, we assume g_t and v_t are (potentially) correlated:

$$g_t = \xi + \eta v_t + z_t, \quad (4)$$

where $\xi = E(g_t)$, η is a $d \times p$ matrix, and z_t is measurement error orthogonal to v_t .⁵ The risk premia of the factors g_t are $\eta\gamma$, our parameter of interest in this paper. This model clearly nests the classic linear asset pricing model with observable factors only, in which case we can set $\eta = \mathbb{I}_p$ and $z_t = 0$.

Since the true factors in v_t are potentially weak, the observable factors in g_t may therefore also be weak because the exposure of r_t to g_t is partially determined by that to v_t . The risk exposure of g_t (to v_t), η , and risk premia, γ , are not necessarily diminishing (asymptotically). Specifically, $\eta\gamma$ could be a fixed parameter that does not vary with sample size.

2.3 Inconsistency of Existing Estimators

While the literature has proposed many different estimators of the SDF and risk premia, their properties in the weak factor setting have not been studied. In what follows, we revisit a number of existing procedures for estimating risk premia, and show that they are inconsistent in the presence of weak factors using a simple model with a single weak factor.

⁵When g_t is nontradable, measurement error could arise as the econometrician is implementing an empirical counterpart of some theory-predicted factor; when g_t is tradable, it captures the non-diversified errors in the portfolio.

2.3.1 PCA

Giglio and Xiu (2021) suggest a three-pass procedure to estimate $\eta\gamma$: 1) apply PCA to the sample covariance matrix of returns to obtain estimates of the latent factors, \hat{v}_t ;⁶ 2) use Fama-MacBeth regressions to recover the risk premia of \hat{v}_t , $\hat{\gamma}$; 3) use time series regressions of g_t on \hat{v}_t to estimate $\hat{\eta}$. The product of the estimates at steps 2 and 3 yields $\hat{\eta}\hat{\gamma}$, the estimate of risk premia. We summarize this procedure in the following algorithm:

Algorithm 1 (PCA-based Estimator of Risk Premia). *The estimator proceeds as follows:*

Inputs: \bar{R} and \bar{G} .

S1. Apply SVD on \bar{R} , and write the first p right singular vector as ξ . The estimated factors are given by $\hat{V} = \sqrt{T}\xi^\top$.

S2. Estimate the risk premia of \hat{V} by $\hat{\gamma} = (\hat{\beta}^\top \hat{\beta})^{-1} \hat{\beta}^\top \bar{r}$ where $\hat{\beta} = \bar{R} \hat{V}^\top (\hat{V} \hat{V}^\top)^{-1}$.

S3. Estimate the factor loading of g_t on v_t by $\hat{\eta} = \bar{G} \hat{V}^\top (\hat{V} \hat{V}^\top)^{-1}$.

Outputs: \hat{V} , $\hat{\eta}$, $\hat{\gamma}$, and $\hat{\gamma}_g^{PCA} = \hat{\eta}\hat{\gamma}$.

Giglio and Xiu (2021) establish the consistency of this estimator and derive its asymptotic inference, in the case that all latent factors are pervasive, whereas g_t can be either strong or weak (depending on the magnitude of η). This risk-premia estimator is appealing for its simplicity, efficiency, and robustness to missing factors. Unfortunately, it fails when some latent factors are weak, which we will show next.

To explain the intuition behind the failure of PCA, it is sufficient to consider a one-factor model with $p = d = 1$ and $\Sigma_v = 1$, in which case the covariance matrix of returns satisfies: $\Sigma = \beta\beta^\top + \Sigma_u$. This matrix has a noisy low rank structure in that $\beta\beta^\top$ has rank 1 whereas Σ_u is a full-rank covariance matrix. To make it simple, we also assume that the factor of interest g_t has no measurement error, i.e., $z_t = 0$ and $g_t = \eta v_t$.

A successful recovery of β via PCA of realized returns requires a favorable signal-to-noise ratio. If the “signal” as measured by $\|\beta\|$, dominates “noise”, which arises from the idiosyncratic component Σ_u and the estimation error in the sample covariance matrix $\hat{\Sigma} - \Sigma$, the first sample eigenvector of $\hat{\Sigma}$ would (approximately) span the same space spanned by the true β . Thus using $\hat{\beta}$, effectively the eigenvector of $\hat{\Sigma}$, in the cross-sectional regression would yield a consistent estimator of the risk premium of the estimated latent factor, which in turn leads to a consistent estimator of the risk premium of g_t . Otherwise, if signal $\|\beta\|$ is so weak that the estimation error in $\hat{\beta}$ dominates, there would be a non-vanishing angle between the space spanned by $\hat{\beta}$ and that by β , which eventually results in an inconsistent estimate of the risk premium $\eta\gamma$. Proposition 1 below shows that the PCA-based risk premium estimator is consistent only if $N/(\|\beta\|^2 T) \rightarrow 0$.

⁶Equivalently, one can directly apply the singular value decomposition (SVD) on \bar{R} .

Proposition 1. *Suppose that test asset returns follow a single-factor model in the form of (1) with $p = 1$, g_t satisfies (4) with $d = 1$, and u_t and v_t i.i.d. normally distributed and independent from each other and $z_t = 0$. In addition, suppose that β satisfies $N/(\|\beta\|^2 T) \rightarrow B \geq 0$ and $\|\beta\| \rightarrow \infty$. Then we have $\hat{\gamma}_g^{PCA} \xrightarrow{P} (1 + B)^{-1} \eta \gamma$.*

In the presence of strong factors, $\|\beta\| \asymp \sqrt{N}$, which leads to $B = 0$ as $T \rightarrow \infty$, so there is no bias. In general, the consistency depends on the relative magnitude of N , T , and $\|\beta\|$. When N and T are of the same order, $\|\beta\| \rightarrow \infty$ is sufficient for the consistency of risk premia estimation. This makes sense in that the eigenvalue of returns corresponding to this factor is proportional to $\|\beta\|^2$, whereas the eigenvalues for the idiosyncratic errors are bounded, so that $\|\beta\| \rightarrow \infty$ guarantees the separation between factors and errors and hence the identification of factors.

This example also shows that the risk premium estimator could be biased even if we have consistent estimator of the factors. In fact, the estimated factors in \hat{V} are consistent under the assumptions of Proposition 1 in the sense that $|\text{Corr}(\hat{V}, V)| \xrightarrow{P} 1$.⁷ However, estimating a large-dimensional vector β given \hat{V} remains a challenging problem, which requires this additional condition, $B = 0$, to achieve consistency.

2.3.2 PLS

Giglio and Xiu (2021) show that the PCA-based estimation procedure effectively constructs a mimicking portfolio for g_t via a principal component regression (PCR) on r_t , which amounts to a projection of g_t onto the first few PCs of the sample covariance matrix of r_t . This is an unsupervised approach, in that the PCs are obtained without any information from g_t . Therefore, PCA might be misled by large idiosyncratic errors in r_t when the signal is not sufficiently strong. In contrast with PCA, partial least squares (PLS) is a supervised procedure, which has been shown to work better than PCA in other settings, see, e.g., Kelly and Pruitt (2013). In the same spirit, we now propose a PLS-based approach for risk premia estimation, exploiting variation of returns that is relevant to the target factor of interest. The key difference is that PCA seeks linear combinations of r_t that maximize variation, ignoring information from the target g_t , whereas PLS seeks linear combinations that have the largest covariance with g_t . We formulate a general PLS-based algorithm for a $d \times 1$ vector of g_t below:

Algorithm 2 (PLS-based Estimator of Risk Premia). *The estimator proceeds as follows:*

Inputs: $\bar{R}_{(1)} := \bar{R}$, $\bar{r}_{(1)} := \bar{r}$ and \bar{G} , a $d \times T$ matrix.

S1. For $k = 1, 2, \dots, p$, repeat the following steps using $\bar{R}_{(k)}$, $\bar{r}_{(k)}$ and \bar{G} .

a. Obtain the weight vector w from the largest left singular vector of $\bar{R}_{(k)} \bar{G}^\top$.

⁷We can further establish that a sufficient condition for consistent recovery of factors is $N/(\|\beta\|^4 T) \rightarrow 0$, which clearly holds in the setup of Proposition 1.

- b. Estimate the k th factor as $\widehat{V}_{(k)} = \sqrt{T}w^\top \bar{R}_{(k)} / \|w^\top \bar{R}_{(k)}\|$. Here, $\widehat{V}_{(k)}$ is normalized to have norm \sqrt{T} .
- c. Estimate the risk premium of $\widehat{V}_{(k)}$ by $\widehat{\gamma}_{(k)} = \sqrt{T}w^\top \bar{r}_{(k)} / \|w^\top \bar{R}_{(k)}\|$.
- d. Estimate the k th factor loading of r_t by $\widehat{\beta}_{(k)} = T^{-1} \bar{R}_{(k)} \widehat{V}_{(k)}^\top$.
- e. Remove $\widehat{V}_{(k)}$ to obtain residuals for the next step: $\bar{R}_{(k+1)} = \bar{R}_{(k)} - \widehat{\beta}_{(k)} \widehat{V}_{(k)}$ and $\bar{r}_{(k+1)} = \bar{r}_{(k)} - \widehat{\beta}_{(k)} \widehat{\gamma}_{(k)}$.

S2. Estimate the factor loading of g_t on v_t by $\widehat{\eta} = T^{-1} \bar{G} \widehat{V}^\top$, where $\widehat{V} = (\widehat{V}_{(1)}^\top, \dots, \widehat{V}_{(p)}^\top)^\top$, and denote their risk premia estimated above as $\widehat{\gamma} = (\widehat{\gamma}_{(1)}, \dots, \widehat{\gamma}_{(p)})^\top$.

Output: $\widehat{\gamma}_g^{PLS} = \widehat{\eta} \widehat{\gamma}$.

The PLS estimator has a closed-form formula if \bar{G} is a $1 \times T$ vector and a single-factor is extracted ($p = 1$):

$$\widehat{\gamma}_g^{PLS} = \|\bar{G} \bar{R}^\top \bar{R}\|^{-2} \bar{G} \bar{R}^\top \bar{R} \bar{G}^\top \bar{G} \bar{R}^\top \bar{r}.$$

While the PLS procedure seems appealing, the next proposition shows that this approach is asymptotically equivalent to the PCA-based procedure, hence it fails in exactly the same weak factor setting as PCA.

Proposition 2. *Suppose that test asset returns follow a single-factor model in the form of (1) with $p = 1$, g_t satisfies (4) with $d = 1$, u_t and v_t i.i.d. normally distributed and independent from each other, and $z_t = 0$. In addition, suppose that β satisfies $N/(\|\beta\|^2 T) \rightarrow B \geq 0$ and $\|\beta\| \rightarrow \infty$. Then we have $\widehat{\gamma}_g^{PLS} \xrightarrow{p} (1 + B)^{-1} \eta \gamma$.*

Intuitively, the covariance information embedded in the objective function of PLS is dominated by its variance component, hence PLS yields the same asymptotic behavior as PCA with respect to estimating β , and therefore risk premia.

2.3.3 Ridge

Next, we consider an alternative ridge regression approach to the construction of mimicking portfolios, and the resulting risk premia estimator can be written as:

$$\widehat{\gamma}_g^{Ridge} = \bar{G} \bar{R}^\top (\bar{R} \bar{R}^\top + \mu \mathbb{I}_N)^{-1} \bar{r}, \quad (5)$$

where $\mu > 0$ is some tuning parameter. In the case of pervasive factors, [Giglio and Xiu \(2021\)](#) show that the ridge estimator yields consistent estimate of $\eta \gamma$. However, the ridge estimator also fails in the presence of weak factors:

Proposition 3. *Suppose that test asset returns follow a single-factor model in the form of (1) with $p = 1$, g_t satisfies (4) with $d = 1$, u_t and v_t i.i.d. normally distributed and independent from each other, and $z_t = 0$. In addition, suppose that β satisfies $N/(\|\beta\|^2 T) \rightarrow B \geq 0$ and $\|\beta\| \rightarrow \infty$, and the tuning parameter μ satisfies $\mu/(\|\beta\|^2 T) \rightarrow D$ for some constant $D \geq 0$ such that $B + D > 0$. Then we have $\hat{\gamma}_g^{\text{Ridge}} \xrightarrow{p} (1 + B + D)^{-1} \eta\gamma$.*

Even though the ridge-based risk premia estimator seemingly accounts for the impact of all eigenvectors as factors instead of only the first p of them, the resulting estimator remains inadequate for consistency. Intuitively, the tuning parameter μ in the ridge procedure serves as a threshold that impedes the influence of eigenvectors corresponding to small eigenvalues just like in PCA and PLS, which explains the appearance of B in the limit. The presence of μ also leads to a shrinkage bias to the first few eigenvectors (i.e., factors), which is why an extra term D appears in the limit as well.

2.3.4 Risk Premium PCA

Finally, we consider an estimator of $\eta\gamma$ based on the risk premium PCA (rpPCA) estimator proposed by Lettau and Pelger (2020) in the context of SDF estimation.

Algorithm 3 (rpPCA-based Estimator of Risk Premia). *The estimator proceeds as follows:*
Inputs: \bar{R} and \bar{G} .

- S1. *Apply PCA on $T^{-1}RR^\top + \mu\bar{r}\bar{r}^\top$, where μ is a tuning parameter, and write the first p eigenvectors as ς . The estimated factors are given by $\hat{V} = \varsigma^\top \bar{R}$.*
- S2. *Estimate the risk premia of \hat{V} by $\hat{\gamma} = \varsigma^\top \bar{r}$.*
- S3. *Estimate the factor loading of g_t on v_t by $\hat{\eta} = \bar{G}\hat{V}^\top(\hat{V}\hat{V}^\top)^{-1}$.*

Outputs: $\hat{\gamma}_g^{\text{rpPCA}} = \hat{\eta}\hat{\gamma}$.

The standard PCA is applied to the covariance matrix of returns, that is $T^{-1}RR^\top - \bar{r}\bar{r}^\top$. Lettau and Pelger (2020) show that assigning a larger weight $\mu > -1$ to the term related to average returns improves the Sharpe ratio of the estimated SDF.⁸ While this estimator was originally proposed for estimating the SDF, it can be used to estimate risk premia as well (since risk premia are just covariances with the SDF). We discuss here this risk premium estimator, in a setting where a single factor can be weak yet its strength is of a distinct order relative to idiosyncratic components asymptotically. This setting is more informative for comparing different approaches, because in this setting a consistent estimation procedure exists.

⁸They derive asymptotic properties of rpPCA in a setting where all factors are weak and N and T increase to infinity at the same rate. The setting they analyze is one where all factors are so weak that they cannot be recovered – specifically, the strength of weak factors remains indistinguishable from that of idiosyncratic errors as N and T increase. Under this assumption, consistent estimation of the SDF is impossible, including rpPCA, which, despite being more correlated with the SDF than PCA, is also inconsistent.

Proposition 4. *Suppose that test asset returns follow a single-factor model in the form of (1) with $p = 1$, g_t satisfies (4) with $d = 1$, u_t and v_t i.i.d. normally distributed and independent from each other, and $z_t = 0$. In addition, suppose that β satisfies $N/(\|\beta\|^2 T) \rightarrow B \geq 0$ and $\|\beta\| \rightarrow \infty$, that the factor has a non-zero risk premia, i.e., $\gamma \neq 0$. Then for some tuning parameter $\mu > -1$, we have*

$$\widehat{\gamma}_g^{rpPCA} \xrightarrow{p} w(1+B)^{-1}\eta\gamma + (1-w)\eta(\gamma + \gamma^{-1}B),$$

where

$$w = \frac{2 + 2B}{1 + 2B + \sqrt{(1-a)^2 + 4(1+\mu)\gamma + a}}, \quad a = (1 + \mu)(\gamma^2 + B) - B.$$

Proposition 4 suggests that this rpPCA estimator is inconsistent in the presence of a weak factor, with a more involved bias term compared to the above estimators. Like PCA and PLS, this estimator is consistent when all factors are strong ($B = 0$). When $B > 0$, we may design a different asymptotic setting, in which the tuning parameter $\mu \rightarrow \infty$, under which the rpPCA estimator converges to $\eta(\gamma + \gamma^{-1}B)$. If we further assume $\gamma \rightarrow \infty$ (while keeping $\eta\gamma$ constant), this estimator can be consistent as long as $\eta\gamma^{-1}B \xrightarrow{p} 0$. This suggests that rpPCA can be robust to weak factors if the information about β from the expected return dominates the information from return covariances (in which case factors have a diverging Sharpe ratio.)

An alternative approach to Algorithm 3, based on rpPCA, is to adapt Algorithm 1 by replacing its step S1 by S1 of Algorithm 3. It turns out that this approach yields the same asymptotic behavior as the PCA estimator of Algorithm 3, which is characterized by Proposition 1.⁹ Because its performance is essentially identical with that of PCA, we omit the discussion of this version of rpPCA from the rest of the paper.

2.4 Our Solution: Test Asset Selection

The results in the previous section shed light on the limitation of dimension reduction or shrinkage estimators, when factors are not pervasive.¹⁰ One potential solution is to screen test assets and only keep those that have nontrivial exposure to the factor of interest. Then, if the factor is strong *within this smaller set of test assets*, it is possible to apply PCA or any of the above procedures to recover its risk premium, as long as there remains a sufficient number of test assets.

This strategy echoes some of the practice in the empirical asset pricing literature. Very often,

⁹As shown by Giglio and Xiu (2021), using either left or right singular vectors of \bar{R} as factors yields asymptotically equivalent PCA-based estimators of risk premia. This is, however, not true for rpPCA, because its estimated “eigenvectors” do not correspond to any singular vectors of \bar{R} . This is the reason why using a rpPCA adapted Algorithm 1 would lead to a different asymptotic result (equivalent to Proposition 1), as opposed to Proposition 4 based on Algorithm 3.

¹⁰These results should not be regarded as evidence against the use of above estimators in all scenarios. Rather, we only establish that for data generating processes in the regime of weak factors we define, none of these estimators are consistent. It is however possible that, for some alternative sequences of data generating processes, or for purposes other than risk premia estimation, these estimators may perform well.

test assets are formulated using the exact characteristics-sorted portfolios that the factor of interest is generated from. For instance, [Fama and French \(1993\)](#) use size and value double-sorted portfolios as test assets when estimating a factor model that includes size and value as factors. In other cases, for nontradable factors, portfolios are sorted based on individual stock betas with respect to the factor of interest. These choices of test assets indeed help address the weak factor problem, though, as discussed in the introduction, they do not address the other issue that is relevant in practical applications – omitted factors. Our methodology formalizes the insight behind these traditional procedures and combines it with the use of PCA to address the omitted factor bias.

We start with a simple one factor setting as discussed in the previous propositions, which helps illustrate the intuition behind our proposal and facilitates the comparison with existing estimators (the next section is devoted to the general case). To ensure sufficient test assets after screening, we assume that there exists a subset $I_0 \subset [N]$ such that $\|\beta_{[I_0]}\| \asymp \sqrt{N_0}$, where $N_0 = |I_0| \rightarrow \infty$. Consequently, as long as we locate this subset of assets, within which there exists a strong factor structure, we can recover risk premia consistently. In practice, it is the researcher who decides which test assets to employ in an empirical study. Assuming that a strong factor structure exists at least within a subset of test assets seems practical and plausible.

We next formally present our SPCA procedure for test assets selection and risk premia estimation.

Algorithm 4 (SPCA-based Estimator of Risk Premia for a Single Factor Model ($p = 1$)). *The procedure is as follows:*

Inputs: \bar{R} and \bar{G} , a $1 \times T$ vector.¹¹

S1. Select a subset $\hat{I} \subset [N]$: $\hat{I} = \left\{ i \mid T^{-1} |\bar{R}_{[i]} \bar{G}^\top| \geq c_q \right\}$, where c_q is the $(1 - q)$ -quantile of $\left\{ T^{-1} |\bar{R}_{[i]} \bar{G}^\top| \right\}_{i \in [N]}$.

S2. Repeat S1. – S3. of [Algorithm 1](#) with selected return matrix $\bar{R}_{[\hat{I}]}$ and \bar{G} , and $p = 1$.

Outputs: $\hat{\gamma}_g^{SPCA} := \hat{\eta} \hat{\gamma}$, \hat{V} , $\hat{\eta}$, and $\hat{\gamma}$.

We establish the consistency of the SPCA estimator in the following proposition:

Proposition 5. *Suppose that $\log N/T \rightarrow 0$ and test asset returns follow a single-factor model in the form of [\(1\)](#) and that g_t satisfies [\(4\)](#), with u_t , v_t , and z_t i.i.d. normally distributed and independent from each other. The loading matrix β satisfies $\|\beta\|_{\text{MAX}} \lesssim 1$ and there exists a subset $I_0 \subset [N]$ such that $\|\beta_{[I_0]}\| \asymp \sqrt{N_0}$ where $N_0 = |I_0| \rightarrow \infty$. Then, for any choice of q in [Algorithm 4](#) such that $qN/N_0 \rightarrow 0$ and $qN \rightarrow \infty$, and that $|\beta|_{\{qN+1\}} \leq (1 + \delta)^{-1} |\beta|_{\{qN\}}$ for some $\delta > 0$, where $|\beta|_{\{k\}}$ denotes the k th largest value in $\{|\beta_{[i]}|\}_{i \in [N]}$, we have $\hat{\gamma}_g^{SPCA} \xrightarrow{p} \eta \gamma$.*

[Algorithm 4](#) involves a single tuning parameter q that determines how many assets we use to extract the factor. We select the first qN assets sorted by their covariances with the target variable

¹¹We discuss the case of a multivariate ($d \times T$) \bar{G} in [Section 2.6](#).

\bar{G} . The fact that \hat{I} incorporates information from the target reflects the distinctive nature of a supervised procedure. The technical condition on $|\beta|_{\{qN+1\}}$ simply states that these test assets should have (asymptotically) distinct risk exposure, which is a rather mild assumption used in the proof.

Propositions 2 - 4 show that in the single factor case, the consistency of PCA, PLS, and rpPCA requires $B = 0$. Suppose $\|\beta\|^2 = N^v$, for some $v > 0$, then $B = 0$ is equivalent to $N^{1-v}/T \rightarrow 0$. The consistency of SPCA, as shown by Proposition 5, nonetheless, only requires $\log N/T \rightarrow 0$.¹² That said, the condition $\|\beta\|^2 \gtrsim N_0 \rightarrow \infty$ rules out the case that the factor strength is of the same magnitude as that of idiosyncratic errors.¹³

2.5 The General Case: Selection and Projection

Propositions 1 - 5 focus on a perhaps unrealistic single-factor model since they are meant to illustrate the intuition behind our procedure as well as the failure of existing approaches due to the presence of a weak factor. In general, the DGP of returns is likely driven by more than one factors, some of which may be weak. In the same spirit of Proposition 1, we can show that a more general necessary condition for the consistency of PCA in a multi-factor model is that

$$N/(\lambda_{\min}(\beta^\top \beta)T) \rightarrow 0. \quad (6)$$

Intuitively, this condition requires that the weakest one among all p factors in (1) is sufficiently strong that it can be recovered by PCA. Once again, we consider below more challenging regimes in which the condition (6) fails.

In a multi-factor model, even if all factors are strong by themselves, a related problem arises when some of the factors' exposures are highly correlated. Consider, for example, a two-factor model where the beta matrix has the following form:

$$\beta = \left[\begin{array}{c|c} \beta_{11} & \beta_{12} \\ \hline \beta_{21} & \beta_{22} \end{array} \right], \quad (7)$$

¹²Another idea that shares this spirit is the scaled-PCA proposed by Huang et al. (2021), which uses regression coefficients of \bar{G} on \bar{R} to weight \bar{R} before feeding it into the PCA procedure. An advantage of the scaled PCA approach is that it does not involve any tuning parameter. Nonetheless, the scaled PCA still assigns weights of $1/\sqrt{T}$ magnitude to assets that have zero-correlations with the target variable, whereas our approach assigns zero weights to such assets. As a result, our procedure only requires $\log N$ to be small relative to T , whereas both the scaled PCA and PCA require N to grow no faster than a certain polynomial rate relative to T .

¹³Throughout this paper, an extremely weak factor is referred to as a factor whose strength is of the same order of magnitude as that of idiosyncratic errors. We preclude this extreme case from our discussion because no estimators under consideration could achieve consistency and a harmless modeling choice would be to treat these extremely weak factors as noise: their risk premia effectively become alpha. The weak-factor setting we investigate permits consistency, and allows for asymptotic comparison of different estimators.

where β_{11} and β_{12} are $N_0 \times 1$ vectors, β_{21} and β_{22} are $(N - N_0) \times 1$ vectors, and N_0 is small relative to N . Suppose that $\beta_{21} = \beta_{22}$. Then we can show that $\lambda_{\min}(\beta^\top \beta) \leq \|\beta_{11} - \beta_{12}\|^2 / 2 \lesssim N_0$. As a result, $N / (\lambda_{\min}(\beta^\top \beta) T) \gtrsim N / (N_0 T)$, which does not necessarily converge to 0 if N_0 and T are small, so that the condition (6) could fail. In this example, while either factor could be strong, the same “rank deficiency” issue may arise, since these factors could have highly correlated exposures.

Another important consideration is that applying the screening approach only once would in general not work in a multi-factor model. Take (7) again as an example. Suppose that $\beta_{21} \neq \beta_{22} = 0$, then it is easy to show that $\lambda_{\min}(\beta^\top \beta) \leq \|\beta_{12}\|^2 \lesssim N_0$, thus in light of the above discussion, the weak factor problem could occur in this example. In this case, it is the second factor that is weak since most of test assets’ exposure to it is zero. Now suppose that $\eta = (1, 1)$: the observed factor g is correlated with both factors and hence with all test assets. But in that case, the screening would not eliminate any test asset – and yet PCA with all test assets would not recover the weak factor, should $N / (N_0 T)$ not vanish. This example demonstrates that even though screening assets ensures that the *first* principal component after screening is strong, there is no guarantee that this procedure can solve the weak factor issue in one step if additional factors are weak.

It is worth pointing out that the two aforementioned cases are in fact equivalent, because we can rotate the beta matrix in the second case into the form of the first case. Thanks to the rotation invariance property illustrated in Giglio and Xiu (2021), both the risk premia and the SDF estimands remain unchanged after rotation, and hence the equivalence.

In the examples above, the problem was that the first screening step did not eliminate any assets, and therefore could not solve the weak factor issue. We provide next another example, that shows that in some situations screening can sometimes eliminate *too many* assets, making a strong factor model become weak or even rank-deficient. Suppose β has the following form:

$$\beta = \left[\begin{array}{c|c} \beta_{11} & \beta_{11} \\ \hline 0 & \beta_{22} \end{array} \right], \quad (8)$$

where β_{11} and β_{22} are $N/2 \times 1$ non-zero vectors satisfying $\|\beta_{11}\| \asymp \|\beta_{22}\| \asymp \sqrt{N}$. Clearly, β is full-rank and both factors are strong. Therefore, a standard PCA procedure should work smoothly. Suppose in addition that $\eta = (1, 0)$ (i.e., $g_t = v_{1t}$) and that v_{1t} and v_{2t} are uncorrelated. Then it implies that g_t is uncorrelated with the second half of test assets in r_t , so only the first half would remain, should screening be applied with g_t before extracting the principal components. In this example, however, the remaining test assets have perfectly correlated exposures to both factors, so that only one factor, $v_{1t} + v_{2t}$, is left. This example shows that the one-step supervised procedure (screening plus PCA) proposed by Bair et al. (2006), may be counterproductive for factor extraction in a multi-factor setting.

To resolve the issue of weak factors and avoid these screening traps, we propose a multi-step procedure that iteratively conducts selection and projection. The projection step eliminates the influence of the estimated factor, which ensures the success of the screening steps that occur over the following iterations. More specifically, Step S1 of Algorithm 4 can help identify one strong factor from a selected subset of test assets. Once we have estimated this factor, we project the returns of *all* test assets r_t (not just those selected at the first step) and g_t onto this factor, so that their residuals will not be correlated with this factor. Then we can repeat the same selection procedure with these residuals. This approach enables a continued discovery of factors, and guarantees that each new factor is orthogonal to the estimated factors in the previous steps, similar to the factors extracted by standard PCA. It is easy to check that this iterative screening and projection approach successfully addresses the problems of all three examples above. Formally, the algorithm is given by:

Algorithm 5 (Selection and Projection). *The selection and projection based procedure for risk premium estimation is as follows:*

Inputs: $\bar{R}_{(1)} := \bar{R}$, $\bar{r}_{(1)} := \bar{r}$, and $\bar{G}_{(1)} := \bar{G}$, a $d \times T$ vector.

S1. For $k = 1, 2, \dots$ iterate the following steps using $\bar{R}_{(k)}$, $\bar{r}_{(k)}$, and $\bar{G}_{(k)}$:

- a. Select an appropriate subset $\hat{I}_k \subset [N]$.
- b. Repeat S1. – S3. of Algorithm 1 with selected return matrix $(\bar{R}_{(k)})_{[\hat{I}_k]}$ and $\bar{G}_{(k)}$. Denote the estimates as $\hat{\lambda}_{(k)}$, $\hat{V}_{(k)}$, $\hat{\eta}_{(k)}$, $\hat{\gamma}_{(k)}$.
- c. Estimate the exposure of $\bar{R}_{(k)}$ on $\hat{V}_{(k)}$ by $\hat{\beta}_{(k)} = T^{-1} \bar{R}_{(k)} \hat{V}_{(k)}^\top$.
- d. Obtain $\bar{R}_{(k+1)} = \bar{R}_{(k)} - \hat{\beta}_{(k)} \hat{V}_{(k)}$, $\bar{r}_{(k+1)} = \bar{r}_{(k)} - \hat{\beta}_{(k)} \hat{\gamma}_{(k)}$, and $\bar{G}_{(k+1)} = \bar{G}_{(k)} - \hat{\eta}_{(k)} \hat{V}_{(k)}$.

Stop at $k = \hat{p}$, where \hat{p} is chosen based on some proper stopping rule.

S2. Estimate the risk premium by $\hat{\gamma}_g^{SPCA} = \sum_{k=1}^{\hat{p}} \hat{\eta}_{(k)} \hat{\gamma}_{(k)}$.

Outputs: $\hat{\gamma}_g^{SPCA}$, $\hat{\eta} = (\hat{\eta}_{(1)}^\top, \dots, \hat{\eta}_{(\hat{p})}^\top)^\top$, $\hat{\gamma} = (\hat{\gamma}_{(1)}, \dots, \hat{\gamma}_{(\hat{p})})^\top$, $\hat{V} = (\hat{V}_{(1)}^\top, \dots, \hat{V}_{(\hat{p})}^\top)^\top$ and $\hat{\beta} = (\hat{\beta}_{(1)}, \dots, \hat{\beta}_{(\hat{p})})$.

In Algorithm 5, we recover one latent factor and obtain its risk premium at each stage of S1. Both the factor and its risk premium are estimated using a subset of rows in the stage- k return residual matrix $\bar{R}_{(k)}$, within which this factor is strong. We then project all observables onto this factor and proceed again with residuals. Because each row of $\bar{R}_{(k+1)}$ is orthogonal to $\hat{V}_{(j)}$ for $j \leq k$ the factors we obtain are orthogonal with each other, as is the case with PCA.

Algorithm 5 yields a consistent estimator of γ_g as long as an appropriate choice of \hat{I}_k and a

stopping rule are adopted. One possible choice for \widehat{I}_k is:¹⁴

$$\widehat{I}_k = \left\{ i \mid T^{-1} \left\| (\bar{R}_{(k)})_{[i]} \bar{G}_{(k)}^\top \right\|_{\text{MAX}} \geq c_q^{(k)} \right\},$$

where $c_q^{(k)}$ is the $(1 - q)$ th-quantile of $\left\{ T^{-1} \left\| (\bar{R}_{(k)})_{[i]} \bar{G}_{(k)}^\top \right\|_{\text{MAX}} \right\}_{i \in [N]}$.

(9)

Correspondingly, we set the stopping criterion as:

$$c_q^{(k)} < c, \quad \text{for some threshold } c. \tag{10}$$

In other words, we select test assets that have predictive power for at least one variable in g_t and stop when most test assets are uncorrelated with all variables in g_t . With a good choice of tuning parameters, q and c , the iteration stops as soon as most of the rows of the projected residuals of returns appear uncorrelated with the projected residuals of g_t , which implies that all factors that are correlated with g_t are successfully recovered.

To establish the consistency of this estimator, we need a subset of assets, indexed by I_0 , such that within this subset all factors are strong, that is, $\lambda_{\min}(\beta_{[I_0]}^\top \beta_{[I_0]}) \asymp N_0$, where $N_0 = |I_0| \rightarrow \infty$. Because the number of factors, p , is finite, such a subset I_0 always exists as long as for each factor we can locate a sufficiently large subset, respectively, within which this factor is strong.¹⁵ With this identification assumption, along with moment conditions given in the appendix, the following theorem establishes the consistency of the SPCA estimator:

Theorem 1. *Suppose that test asset returns in r_t follow (1), the factor proxies in g_t satisfy (4), and that Assumptions A.1-A.8 hold. If $\log(NT)(N_0^{-1} + T^{-1}) \rightarrow 0$ then for any tuning parameters c and q that satisfy*

$$c \rightarrow 0, \quad c^{-1}(\log NT)^{1/2}(q^{-1/2}N^{-1/2} + T^{-1/2}) \rightarrow 0, \quad qN/N_0 \rightarrow 0,$$

we have $\widehat{\gamma}_g^{\text{SPCA}} \xrightarrow{p} \eta\gamma$.

The consistency result in Theorem 1 does not require a full recovery of all factors that drive the SDF. In fact, only factors correlated with g_t will be recovered. Missing any uncorrelated factors in the SDF does not affect the consistency of the risk premium of g_t because such factors do not help price g_t .

Moreover, this result does not rely on Gaussian error assumptions nor on an assumption that all factors have the same strength with respect to all test assets. The assumption on the relative size of

¹⁴Using covariance for screening allows us to replace all $\bar{G}_{(k)}$ in the definition of \widehat{I}_k and Algorithm 5 by \bar{G} , that is, only the projections of $\bar{R}_{(k)}$ and $\bar{r}_{(k)}$ are needed, because this replacement would not affect the covariance between $\bar{G}_{(k)}$ and $\bar{R}_{(k)}$, and in turn, the test assets after screening and the estimates of $\widehat{\eta}_{(k)}$. We use this fact in the proofs, which simplifies the notation. We can also use correlation instead of covariance in constructing \widehat{I}_k . Despite this does not affect the asymptotic analysis, we find correlation screening performs slightly better in finite samples.

¹⁵This assumption is weak in that it does not imply all factors should have identical strength with respect to the entire cross-section of assets in r_t . A detailed discussion on this point follows Assumption A.3 in the appendix.

N and T is also quite flexible, in contrast with existing results in the literature in which N cannot grow faster than a certain polynomial rate of T .

2.6 Asymptotic Inference on Risk Premia

In this section we develop the asymptotic distribution of the risk premium estimator from Algorithm 5. Not surprisingly, the conditions in Theorem 1 do not guarantee that $\hat{\gamma}_g^{SPCA}$ converges to $\eta\gamma$ at the desirable rate $T^{-1/2}$. The major obstacle lies in the recovery of factors, which we can explain with the previous single-factor example.

Recall that we use the sample correlation/covariance between r_t and g_t to screen test assets. Even if g_t is independent with respect to the test assets, their sample correlation can be as large as $T^{-1/2}$. Therefore, the threshold needs no smaller than $T^{-1/2}$. However, for any given threshold, say, $T^{-1/4}$, if it happens that $\eta \asymp T^{-1/3} < T^{-1/4}$, then it suggests that g_t is not too different from random noise, so that screening based on its correlation with r_t will likely not select any assets, which in turn leads to no discovery of factors. Our procedure thereby gives a risk premium estimate of 0, which is certainly consistent, but the estimation error is of an order $T^{-1/3}$, so that the usual central limit theorem (CLT) fails.

Generally speaking, this issue arises because of the potential failure to identify all factors in the DGP. Once all factors are identified, the central limit theorem holds regardless of the magnitude of η . So to make inference we need a stronger assumption that rules out cases like this, in order to insure against a higher order omitted factor bias that impedes the CLT even though it does not affect consistency. It turns out that so long as $\eta \in \mathbb{R}^{d \times p}$ satisfies $\lambda_{\min}(\eta^\top \eta) \gtrsim 1$, we can rule out the possibility of missing factors. On the other hand, our algorithm will not select more factors than needed, if we stop the iteration as soon as $c_q^{(k)}$ is sufficiently small. Of course, in a finite sample, a perfect recovery of the factor space is a stretch, but the assumptions here are substantially weaker than the pervasive factor assumption adopted in the literature, e.g., Bai (2003). The inference theory on factor models also relies on a perfect recovery of the count of (strong) factors, e.g., Bai and Ng (2002). We provide below the consistency result on the number of factors and the CLT result on risk premium, and investigate the finite sample behavior of SPCA in Section 3.

Theorem 2. *Under the same assumptions as Theorem 1, if we further have $T^{-1/2}N_0 \rightarrow \infty$, Assumption A.9 and $\lambda_{\min}(\eta^\top \eta) \gtrsim 1$, then for any tuning parameters c and q in (9) and (10) satisfying*

$$c \rightarrow 0, \quad c^{-1}(\log NT)^{1/2}(q^{-1/2}N^{-1/2} + T^{-1/2}) \rightarrow 0, \quad qN/N_0 \rightarrow 0, \quad q^{-1}N^{-1}T^{1/2} \rightarrow 0,$$

we have that \hat{p} defined in Algorithm 5 satisfies: $\hat{p} \xrightarrow{P} p$, and that the estimator constructed via Algorithm 5 satisfies

$$\sqrt{T}(\hat{\gamma}_g^{SPCA} - \eta\gamma) \xrightarrow{d} \mathcal{N}(0, \Phi),$$

where Φ is given by

$$\Phi = (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) \Pi_{11} (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) + (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) \Pi_{12} \eta^\top + \eta \Pi_{12}^\top (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) + \eta \Pi_{22} \eta^\top,$$

and Π_{11} , Π_{12} , and Π_{22} are specified by Assumption A.9.

We can adopt the same Newey-West-type estimator for Φ as in Section 4.5 of Giglio and Xiu (2021), since each component of Φ can be estimated from the outputs of the SPCA algorithm. These estimates are consistent up to some rotation matrices which will cancel each other and yield a consistent estimate of Φ .

The condition $\lambda_{\min}(\eta^\top \eta) \gtrsim 1$ implies $d \geq p$, that is, we need g_t to have at least equal number of variables as the true number of factors. Moreover, the condition also implies that for each factor in v_t , there is at least one variable in g_t with a non-vanishing exposure to it.

2.7 The Case of Observable Factors

The previous discussion does not assume any knowledge of the identities of the factors v_t in (1). If v_t corresponds to innovations of observable factors, denoted by f_t , which were known (by assumption), say, the Fama-French five factors, our procedure can be greatly simplified. It is meaningful to study this case, because it is most common in the empirical literature, albeit this is a (rather) strong assumption.

Suppose factors in f_t are tradable. If g_t is part of them, then we can estimate the risk premium of g_t by simply taking its time-series average. If g_t is either spanned by f_t or not tradable, then a simple time series regression of g_t onto the factors f_t can recover its loading, η , which along with the risk premia estimates of f_t by their averages, give rise to the risk premium estimate of g . These scenarios are simple, and do not require cross-sectional regressions.

If some of the observed factors in f_t are not tradable, say, GDP growth, then a cross-sectional regression is necessary, which effectively constructs their mimicking portfolios. In this setting, a weak factor problem potentially arises as documented in the literature, see, e.g., Kan and Zhang (1999), Kleibergen (2009). To tackle this issue, one could adopt a simplified version of Algorithm 5, to supervise the construction of mimicking portfolios for each of the observed non-tradable factors (in this case GDP growth), while using residuals from the projection of test asset returns onto tradable factors as new test assets.

2.8 Asymptotic Inference on Alpha

As a by-product, we can also make inference on the pricing error, α_g , defined as $E(g_t) - \gamma_g$, when g_t is tradable. In practice, this exercise is most relevant for inferring the “skill” of a fund manager; we explore this application in section 4.2.1. Using the SPCA estimator $\hat{\gamma}_g^{SPCA}$, we can directly construct $\hat{\alpha}_g = \bar{g} - \hat{\gamma}_g^{SPCA}$. We now provide its corresponding CLT result.

Theorem 3. *Suppose the same assumptions as those in Theorem 2 hold. If we further have Assumption A.10, then the estimator $\hat{\alpha}_g$ satisfies*

$$\sqrt{T}(\hat{\alpha}_g - \alpha_g) \xrightarrow{d} \mathcal{N}(0, \tilde{\Phi}),$$

where $\tilde{\Phi}$ is given by

$$\tilde{\Phi} = (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) \Pi_{11} (\Sigma_v^{-1} \gamma \otimes \mathbb{I}_d) - (\gamma^\top \Sigma_v^{-1} \otimes \mathbb{I}_d) \Pi_{13} - \Pi_{13}^\top (\Sigma_v^{-1} \gamma \otimes \mathbb{I}_d) + \Pi_{33}.$$

It is straightforward to construct a Newey-West-type estimator of the asymptotic variance $\tilde{\Phi}$ via its sample analog.

2.9 Recovery of the Stochastic Discount Factor

The main focus of the previous sections is on risk premia, whose consistency does not require a consistent recovery of the SDF, since some of these factors driving SDF might be uncorrelated with the factors of interest, and will therefore not play any role in the consistency of risk premia. Nonetheless, we have pointed out that constructing valid asymptotic inference requires the recovery of all factors that drive the SDF. In this case, we can also reconstruct the SDF. More specifically, from the outputs of Algorithm 5, we can estimate the SDF by:

$$\hat{m}_t^{SPCA} = 1 - \hat{\gamma}^\top \hat{v}_t, \quad \text{where } \hat{v}_1, \dots, \hat{v}_T \text{ are the columns of } \hat{V}. \quad (11)$$

Theorem 4. *Suppose the same assumptions as in Theorem 2 hold. In addition, we have Assumption A.11. Then the estimator (11) satisfies*

$$\frac{1}{T} \sum_{t=1}^T |\hat{m}_t^{SPCA} - m_t|^2 \lesssim_p \frac{1}{T} + \frac{\log N_0}{N_0}. \quad (12)$$

There are a number of alternative approaches for SDF estimation proposed in the literature, e.g., the selection/shrinkage approach by [Kozak et al. \(2020\)](#) and the risk premia PCA by [Lettau and Pelger \(2020\)](#). In what follows, we provide a theoretical comparison of Lasso and Ridge based estimators in our general framework where factors can potentially be weak. The ridge estimator shares the same spirit of PCA-based estimators as shown by [Giglio and Xiu \(2021\)](#) and propositions in previous sections. Examining the asymptotic behavior of these two approaches will provide useful insights that may guide their applications in practice.

[Kozak et al. \(2020\)](#) consider an SDF in the form of (3), whereas we represent it as in (2). Prior to the asymptotic analysis of their estimators, we first establish the asymptotic equivalence of these two definitions in our large- N setting:

Proposition 6. *Suppose that test asset returns in r_t follow (1), and Assumption A.11 holds. Then as $N \rightarrow \infty$, we have*

$$\frac{1}{T} \sum_{t=1}^T |m_t - \tilde{m}_t|^2 \lesssim_p \frac{1}{\lambda_{\min}(\beta^\top \beta)}.$$

Effectively, Proposition 6 proves that there is no ambiguity with respect to the definition of the estimand, since the two estimands are asymptotically equivalent as long as $\lambda_{\min}(\beta^\top \beta) \rightarrow \infty$. Given that this exact assumption is necessary for Theorem 4, and that $\lambda_{\min}(\beta^\top \beta) \gtrsim N_0$, we can replace m_t in the left-hand side of (12) by \tilde{m}_t .

Kozak et al. (2020) suggest estimating the SDF by solving an optimization problem:

$$\hat{b} = \arg \min_b \left\{ (\bar{r} - \hat{\Sigma}b)^\top \hat{\Sigma}^{-1} (\bar{r} - \hat{\Sigma}b) + p_\mu(b) \right\}, \quad (13)$$

with which the estimated pricing kernel is given by

$$\hat{m}_t = 1 - \hat{b}^\top (r_t - \bar{r}). \quad (14)$$

In the above, $\hat{\Sigma}$ is the sample covariance matrix of r_t and $p_\mu(b)$ is a penalty term through which economic priors are imposed. Depending on the penalty function, we will denote the resulting estimator of m by \hat{m}_t^{Ridge} or \hat{m}_t^{Lasso} .

The objective function in (13) appears to require the inverse of the sample covariance matrix $\hat{\Sigma}^{-1}$, which is not well-defined when $N > T$. Instead, we suggest optimizing an equivalent but different form of (13):

$$\hat{b} = \arg \min_b \left\{ b^\top \hat{\Sigma} b - 2b^\top \bar{r} + b^\top \hat{\Sigma} b + p_\mu(b) \right\}, \quad (15)$$

which avoids the calculation of $\hat{\Sigma}^{-1}$.

The following result sheds light on the asymptotic properties of this estimator in the cases of $p_\mu(b) = \mu \|b\|_1$ and $p_\mu(b) = \mu \|b\|^2$, respectively.

Theorem 5. *We investigate two distinct scenarios.*

- (a) *Suppose that r_t is driven by p latent factors as in (1). With $p_\mu(b) = \mu \|b\|^2$, if $(N+T)/(\lambda_p T) \rightarrow 0$ and Assumptions A.4-A.7, A.11-A.13 hold, we have*

$$\frac{1}{T} \sum_{t=1}^T |\hat{m}_t^{Ridge} - m_t|^2 \lesssim_p \frac{1}{T} + \frac{N+T}{\lambda_p T},$$

where λ_p is the p -th largest eigenvalue of $\beta \Sigma_v \beta^\top$. Since $\lambda_p \asymp \lambda_{\min}(\beta^\top \beta)$, we can replace m_t in the above equation by \tilde{m}_t .

(b) Suppose that the true SDF satisfies $E(\tilde{m}_t^2) \lesssim 1$. With $p_\mu(b) = \mu \|b\|_1$, if Assumptions A.11, A.12 hold, we have

$$\frac{1}{T} \sum_{t=1}^T |\hat{m}_t^{Lasso} - \tilde{m}_t|^2 \lesssim_p \|b\|_1 \sqrt{\frac{\log N}{T}}. \quad (16)$$

If, in addition, we assume that $\lambda_{\min}(\Sigma) \gtrsim 1$, and $\|b\|_0^2 \log N/T \rightarrow 0$, then we have a stronger result

$$\frac{1}{T} \sum_{t=1}^T |\hat{m}_t^{Lasso} - \tilde{m}_t|^2 \lesssim_p \|b\|_0 \frac{\log N}{T}. \quad (17)$$

Interestingly, both the Ridge and Lasso approaches deliver consistent estimates of the SDF, though under rather different sets of assumptions. First of all, the convergence rate of the Ridge approach depends critically on the strength of the weakest factor. If condition (6) fails, then the SDF is not consistent. The failure of this condition is precisely a symptom of weak factors which our SPCA estimator is designed for.

Second, with respect to the estimator using the Lasso penalty, the explicit factor model assumption on r_t is replaced by the sparsity assumption on b . The latter assumption requires that the SDF is spanned by a sparse linear combination of test assets, but place no explicit assumptions on the DGP of these test assets. This suggests that the Lasso estimator remains consistent regardless of the factor strength, but converges at a rather slow rate, $\|b\|_1 \sqrt{\log N/T}$ as shown in (16), so it is not as efficient as our SPCA estimator that exploits the factor structure. Nonetheless, under a much stronger sparsity assumption that $\|b\|_0^2 \log N/T \rightarrow 0$, the Lasso estimator can achieve a comparable rate to that of the SPCA. This stronger notion of sparsity effectively says that the set of true factors must be part of the test assets. In contrast, our SPCA estimator allows for idiosyncratic components in any of the test assets, which is a more acceptable assumption in practice.

Just like for the risk premia estimator based on rpPCA, we can adapt any SDF estimator to obtain an estimator of risk premia, because $-\text{Cov}(m_t, g_t) = \eta\gamma$. Naturally we have a Lasso-based risk premia estimator:¹⁶

$$\hat{\gamma}_g^{Lasso} = -\frac{1}{T} \sum_{t=1}^T \hat{m}_t^{Lasso} \times (g_t - \bar{g}).$$

Furthermore, the consistency of the SDF estimator translates to the consistency of the resulting risk premia estimator.¹⁷ Deriving a valid inference procedure is possible for Lasso, if we employ an

¹⁶The SDF-induced Ridge estimator is numerically equivalent to (5), so we do not mention it again.

¹⁷By Assumption A.12(1), Cauchy-Schwartz and triangle inequalities, we have

$$\|\hat{\gamma}_g^{Lasso} - \gamma_g\|_{\text{MAX}} \lesssim_p \sqrt{\frac{1}{T} \sum_{t=1}^T |\hat{m}_t^{Lasso} - \tilde{m}_t|^2} + \sqrt{\frac{\log N}{T}}.$$

additional de-biasing step, see, [Feng et al. \(2020\)](#), which is beyond the scope of the current paper.

As a side note, the SPCA estimator given by equation (11) can also be rewritten in the form of (14), so that it can yield an estimate of b in the definition of SDF given by equation (3). The reason is that \hat{v}_t is in fact a linear combination of r_t . Given that b is invariant to rotations of factors, we can use any rotation of \hat{v}_t to reconstruct an estimate of b . We can exploit this invariance property to construct a convenient estimator \hat{b} . In fact, in S1.b of Algorithm 5, we can construct an $N \times p$ matrix B such that the k th column of B is defined as: $B_{[I_k],k} = \varsigma_{(k)}$ and $B_{[I_k^c],k} = 0$, where $\varsigma_{(k)}$ is the left singular vector of $(\bar{R}_{(k)})_{[I_k]}$. It turns out the SPCA estimates of \hat{V} can be written as a rotation of $B^\top \bar{R}$, so to estimate \hat{b} we can use $B^\top \bar{R}$ as factors, denoted by, \tilde{V} , whose risk premia and covariance are denoted by $\tilde{\gamma}$ and $\tilde{\Sigma}$. Indeed, since the SDF is $m_t = 1 - \hat{\gamma}^\top (\hat{\Sigma}_v)^{-1} \hat{v}_t = 1 - \tilde{\gamma}^\top (\tilde{\Sigma}_v)^{-1} \tilde{v}_t = 1 - \tilde{\gamma}^\top (\tilde{\Sigma}_v)^{-1} B^\top (r_t - \bar{r})$, it follows that the SPCA-based estimate of b is given by

$$\hat{b} = B(\tilde{\Sigma}_v)^{-1} \tilde{\gamma} = TB (B^\top \bar{R} \bar{R}^\top B^\top)^{-1} B^\top \bar{r}.$$

Similarly, we can construct estimates of b using PCA, PLS, and rpPCA. With \hat{b} it is convenient to build out-of-sample SDF (optimal portfolios).

3 Simulations

In this section, we study the finite sample performance of our SPCA procedure using Monte Carlo simulations. We also implement a number of alternative estimators for comparison, some of which are robust to omitted or weak factors, including PCA and its related estimators (Ridge, PLS, and rpPCA), Lasso, as well as the four-split estimator by [Anatolyev and Mikusheva \(2021\)](#).¹⁸ Both the standard two-pass and four-split methods directly use g_t as if they were the true factors in their regressions. The PCA, rpPCA, Ridge, and Lasso effectively construct the SDF first without knowledge of g_t , then estimate the risk premia of g_t factor by factor, using the covariance between each factor and the resulting SDF. PLS and SPCA use all variables in g_t to supervise the estimation procedure.

To implement the SPCA estimator, we select the tuning parameters p and qN (or equivalently q) by cross-validation using the time series R^2 of the hedging portfolio for g_t built by SPCA as the criterion.¹⁹ Recall that any estimator of risk premia for a nontradable factor explicitly or implicitly builds a hedging portfolio exposed to g_t and not exposed to the other factors. We therefore use as a criterion for the choice of the tuning parameters the ability of this portfolio to hedge g_t in the validation sample. In order to produce a conservative comparison, except for SPCA, all the

¹⁸The four-split estimator, which does not rely on dimension reduction, selection, or shrinkage techniques, is valid in the presence of weak observable factors and strong omitted factors that are *not* priced. However, it does not have asymptotic guarantees against omitted and priced strong/weak factors, or measurement error in the observed factors.

¹⁹In finite samples, we find it more effective and more convenient to tune p and q than q and c . This is because the former are direct input to SPCA, which only take values from integers, so that multiple choices of the latter lead to the same integer values of the former.

remaining methods use optimal (even if infeasible) tuning parameters. Specifically, for PCA, PLS and rpPCA, we make use of the true number of factors, $p = 4$, even though it is difficult to obtain a consistent estimator of p in the regime of weak factors. The tuning parameter μ of Ridge estimator is determined via maximum likelihood estimation, with perfect knowledge of Σ_r and $E(r)$. The second tuning parameter of rpPCA is selected by maximizing the theoretical Sharpe ratio of the estimated SDF, using, again, perfect knowledge of Σ_r and $E(r)$. Due to limited sample size, estimating the sample mean and sample covariances in a separate validation sample is rather challenging, which would further deteriorate their performance.

To demonstrate and compare the performance of different estimators, we consider various DGP of returns and/or the observed variables in g_t .

We start with the benchmark case (a), in which all factors are strong and observed. Specifically, we consider a 4-factor DGP as given by equation (1), where the first three factors are calibrated to match the three Fama-French factors (RmRf, SMB, HML) as in Giglio and Xiu (2021), and the last one is a potentially weak factor, denoted by V . We calibrate the parameters such that the monthly Sharpe ratio for the optimal portfolio out of these factors is about 0.25. The realizations of u_t are generated independently from a Gaussian distribution with mean 0 and standard deviation σ_u calibrated such that the time-series R^2 ranges from 50-90%. The loadings of RmRf are generated independently from $\mathcal{N}(1, 1)$ and the loadings of SMB and HML are generated independently from $\mathcal{N}(0, 1)$. We generate the exposure to the fourth factor V , $\beta_{i,V}$, independently from a Gaussian mixture distribution, with probability a from $\mathcal{N}(0, 1)$ and $1 - a$ from $\mathcal{N}(0, 0.1^2)$. Based on our calibration, we choose $a = 0.5$, so that the factor V is sufficiently strong with respect to the cross-section of assets in simulations. g_t includes exactly these four factors in the DGP (RmRF, SMB, HML, and V), so that $\eta = \mathbb{I}_4$, and measurement error is absent.

In scenario b), we choose $a = 0.05$ so that V is weak in that for almost all test assets their factor loadings to V are tiny. In scenario c), the DGP is the same as that of the benchmark case, except that we add Gaussian measurement error, z_t , to each of the factors in g_t . In scenario d), we simulate β for V according to $\beta_{i,V} = -\beta_{i,HML} + e_i$ instead, where e_i s are generated independently from the same mixture Gaussian distribution as above with $a = 0.05$. In this case, the loading matrices of V and HML are very similar, which (almost) leads to a rank deficient factor loading matrix due to highly correlated exposures. The variable g_t contains all four factors with no measurement error. In scenario e), we consider the same DGP of returns as in scenario d), but in g_t we omit the HML factor. Finally, in scenario f), we further add measurement error to scenario d).

For each of these six scenarios (including the benchmark), we plot in Figure 1 the histograms of the estimated risk premium of V (one entry in g_t) for all estimators. If an estimator is consistent, then the histogram is expected to be centered around the true risk premium of V , whose value is represented by a vertical dashed line. This is indeed the case for SPCA in *all* scenarios. It is also the case for almost all estimators in the benchmark scenario, a), when factors are strong (except for Lasso and Ridge, which have a large shrinkage bias). This suggests that the latter two estimators

are not suitable for *inference* on risk premia. Furthermore, in scenario b), when weak factors are present, only SPCA and four-split are consistent. The same is true for scenario d) in which a similar rank-deficiency issue arises. In scenario c) the four-split estimator becomes inconsistent due to measurement error, and it is also ill-behaved in scenario e) because the omitted variable, HML, is priced. The PCA and PLS estimators are consistent in scenario c) but also fail in e), because they are robust to measurement error but not to omitted weak factors. The standard two-pass estimator is only consistent in the benchmark scenario.²⁰ Overall, the simulation evidence is in agreement with our theoretical predictions.

Next, we focus on the last scenario f), which includes the case of weak and omitted factors as well as measurement error. For this case, we report in Table 1 the bias and the RMSE (root-mean-square error) of all estimators for various sample size T . The four rows in each panel provide the results of risk premia estimation for RmRf, SMB, HML, and the weak factor V , respectively. We find that our SPCA approach has smaller biases for the weak factors, whereas the remaining estimators have larger biases and RMSEs, which agrees with our theoretical analysis and Figure 1.

We then investigate the finite sample performance of the inference result developed in Theorem 2. Figure 2 plots histograms of the standardized risk premia estimators using the estimated asymptotic standard errors for SPCA and PCA, respectively, using the DGP in scenario f) as an example. The histograms of PCA deviate from the standard normal distribution for the two highly correlated factors, V and HML. In contrast, the histograms corresponding to the SPCA match the normal distribution well, which verifies our central limit results.

Finally, we study the finite sample behavior of the SDF estimators. We compare the performance of SPCA, PCA, rpPCA, Lasso and Ridge estimators in scenario f). We report in Table 2 the MSE of the SDF estimators where the true SDF is defined by equation (3). The estimated number of factors from our SPCA approach is also reported. We also report in Table 3 the out-of-sample Sharpe ratios of different methods, given by $\hat{b}^\top E(r) / \sqrt{\hat{b}^\top \widehat{\Sigma} \hat{b}}$, where $E(r)$ and Σ are the true mean and covariance of all test assets and \hat{b} is the estimated SDF loading using each method. We find that in terms of the RMSE, SPCA outperforms all other methods, and that rpPCA performs the worst. That said, rpPCA performs the best in terms of the out-of-sample Sharpe ratio, followed by the SPCA. Last but not least, SPCA produces a decent estimator of p when T is large.

4 Empirical Analysis

In this section we perform different empirical exercises to illustrate the use of SPCA. First, we apply it to estimate the risk premia of a number of tradable and nontradable factors proposed in the literature, and we evaluate its out-of-sample performance. Second, we evaluate the robustness of SPCA as we change the universe of test assets to make the factors stronger or weaker. Third, we

²⁰The standard two-pass estimator appears to have a small bias in scenario f), but this happens to be true only for V since all sources of bias happen to balance out, as we show in Table 1.

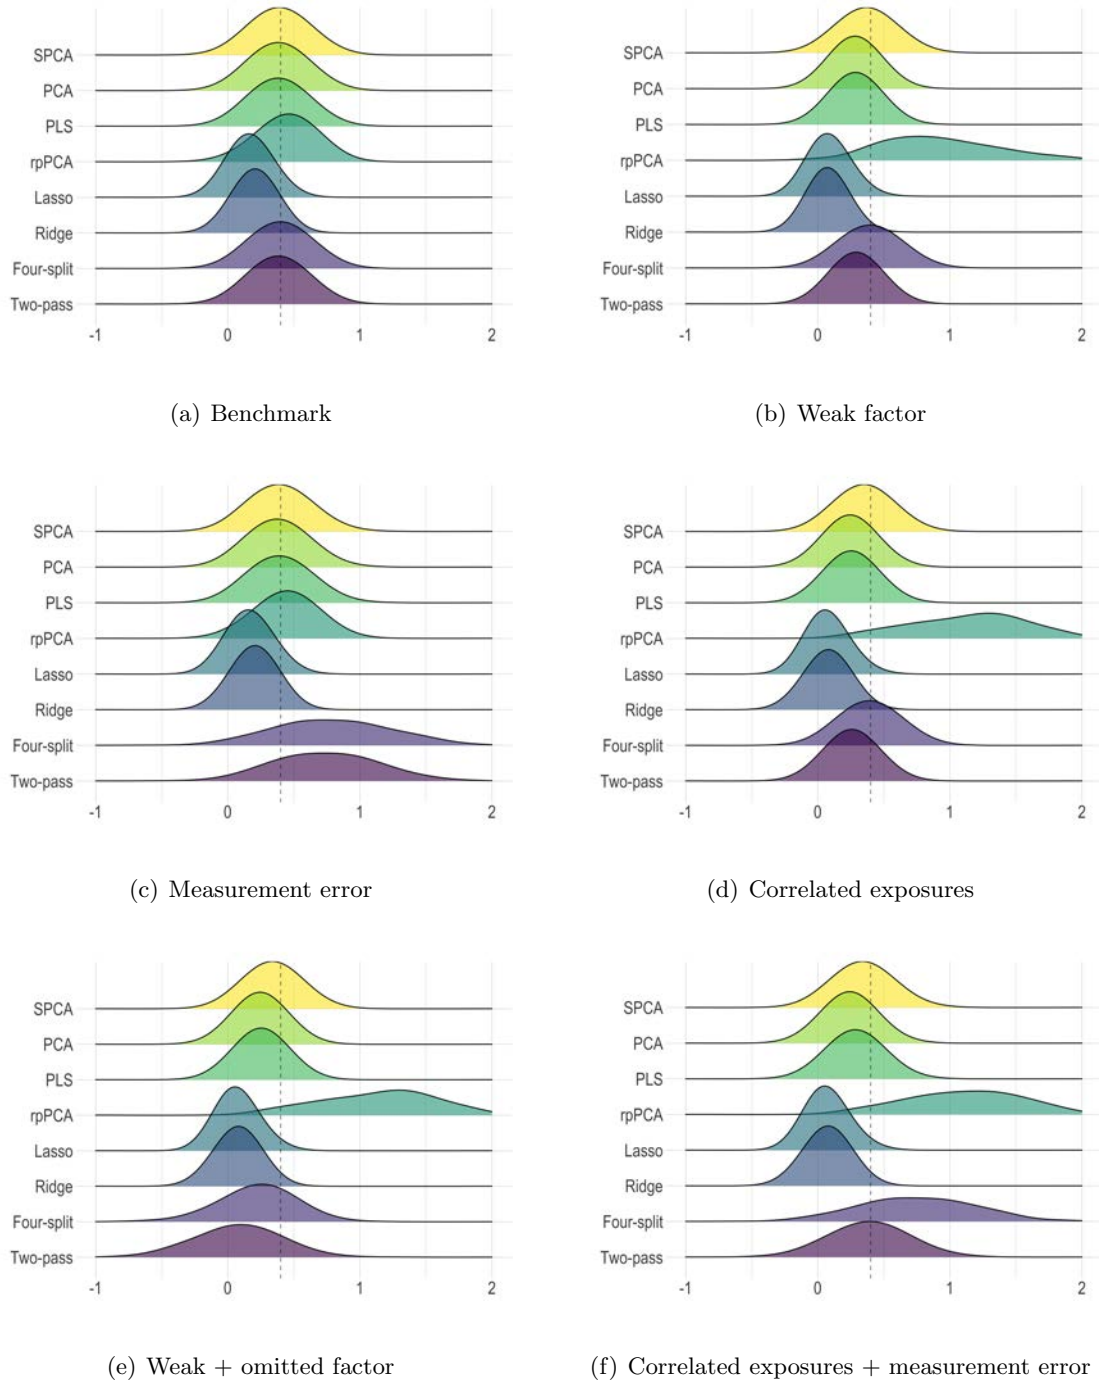


Figure 1: Histogram of Risk Premium Estimates of V

Note: The figure provides histograms of the risk premium estimates in six scenarios for eight estimators we compare, including SPCA, PCA, PLS, rpPCA, Lasso, Ridge, four-split, and the standard two-pass estimator. We simulate the models with $N = 2,000$ and $T = 120$. The number of Monte Carlo repetitions is 1,000.

T	Param	True	SPCA		PCA		rpPCA		PLS	
			Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
120	RmRf	53.7	0.9	39.2	1.0	38.9	0.8	42.4	1.0	39.1
	SMB	21.7	0.2	29.0	0.3	28.5	-0.1	32.7	0.4	28.7
	HML	25.4	-4.8	26.6	-15.3	28.1	76.4	93.5	-10.9	26.8
	V	40.0	-5.5	20.9	-15.7	23.1	74.5	89.7	-11.3	21.4
180	RmRf	53.7	0.8	33.9	0.8	33.8	0.8	34.9	0.9	33.8
	SMB	21.7	0.4	23.1	0.4	22.7	0.6	24.9	0.3	22.9
	HML	25.4	-3.7	21.5	-11.9	22.9	49.0	62.4	-7.4	21.7
	V	40.0	-3.6	17.0	-11.7	18.6	48.6	60.5	-7.1	17.1
240	RmRf	53.7	0.7	29.6	0.8	29.5	0.7	30.0	0.8	29.6
	SMB	21.7	0.5	20.2	0.5	20.0	0.3	21.4	0.5	20.1
	HML	25.4	-2.8	18.3	-9.4	19.3	35.5	45.7	-5.0	18.4
	V	40.0	-3.6	14.3	-10.1	16.0	33.9	42.7	-5.8	14.5
T	Param	True	Lasso		Ridge		Four-split		Two-pass	
			Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
120	RmRf	53.7	-16.3	28.8	-2.9	35.2	16.0	53.8	14.7	51.2
	SMB	21.7	-8.1	15.3	-3.1	20.4	7.1	48.1	7.1	44.7
	HML	25.4	-28.8	31.2	-31.1	35.9	19.2	50.2	-12.1	39.4
	V	40.0	-32.8	34.5	-32.7	34.8	36.3	57.8	-1.3	29.4
180	RmRf	53.7	-11.7	26.9	-1.5	31.3	16.9	46.7	15.0	44.9
	SMB	21.7	-6.4	14.2	-2.1	17.8	6.4	37.7	6.9	35.9
	HML	25.4	-29.7	31.6	-28.3	32.6	20.1	42.5	-6.0	31.8
	V	40.0	-31.6	33.0	-28.8	30.9	39.1	53.2	7.6	26.4
240	RmRf	53.7	-6.7	24.7	-0.6	28.1	16.4	41.7	14.9	39.9
	SMB	21.7	-3.7	14.8	-1.1	16.9	7.3	33.6	7.3	31.9
	HML	25.4	-24.9	28.0	-25.3	29.3	21.4	38.4	-0.8	27.0
	V	40.0	-26.5	28.6	-26.2	28.1	38.8	49.1	12.4	25.6

Table 1: Simulation Results for Risk Premia Estimators

Note: In this table, we report the bias (Column “Bias”) and the root-mean-square error (Column “RMSE”) of the risk premia estimates using SPCA, PCA, rpPCA, Lasso, PLS, Ridge, four-split, and the standard two-pass regression approaches, respectively. The true data-generating process, given by scenario f), has four factors, driven by RmRf, SMB, HML, and V , whereas we estimate the risk premia for noisy versions of these four factors. Their true risk premia are provided in Column “True.” We fix $N = 2,000$ while varying $T = 120, 180$, and 240 in this experiment. All values are in basis points.

T	SPCA		PCA	rpPCA	PLS	Lasso	Ridge
	\hat{p}	MSE	MSE	MSE	MSE	MSE	MSE
120	4.080	0.036	0.037	0.387	0.040	0.044	0.050
	(0.339)	(0.026)	(0.025)	(0.505)	(0.025)	(0.012)	(0.018)
180	4.000	0.024	0.025	0.163	0.027	0.041	0.041
	(0.000)	(0.017)	(0.017)	(0.209)	(0.017)	(0.011)	(0.015)
240	4.000	0.018	0.019	0.085	0.020	0.035	0.035
	(0.000)	(0.013)	(0.013)	(0.088)	(0.013)	(0.011)	(0.013)

Table 2: Simulation Results for SDF estimators

Note: In this table, we report the mean-squared errors (Column “MSE”) defined by $\frac{1}{T} \sum_{t=1}^T |\hat{m}_t - \tilde{m}_t|^2$ for various SDF estimates using SPCA, PCA, rpPCA, PLS, Lasso, and Ridge approaches, respectively. The reported MSEs are the sample average over 1,000 Monte Carlo repetitions and their standard errors are reported in the brackets. We also report the mean and standard deviation of the estimated number of factors \hat{p} using the SPCA approach. The true data-generating process, given by scenario f), has four factors, driven by RmRf, SMB, HML, and a weak factor V , whereas we estimate the SDF using a vector of factor proxies, g_t , that includes noisy versions of the four factors. We compare three scenarios with $T = 120, 180$, and 240 , where $N = 2,000$ is fixed.

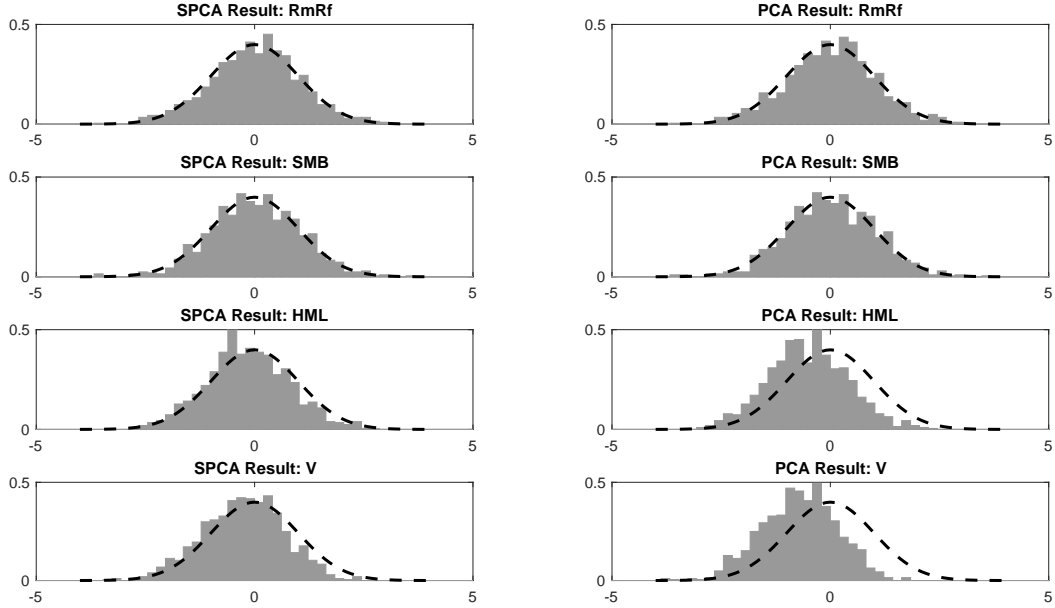


Figure 2: Histogram of the Standardized Estimates in Simulations

Note: The left panels provide the histograms of the standardized SPCA estimates as in Algorithm 5 with asymptotic standard errors given by Theorem 2, whereas the right panels provide those of the standardized PCA-based risk premia estimates as in Algorithm 1. We simulate the model in scenario f) with $N = 2,000$ and $T = 240$. The number of Monte Carlo repetitions is 1,000.

T	SPCA	PCA	rpPCA	PLS	Lasso	Ridge	Theoretical Value
120	0.186 (0.042)	0.159 (0.045)	0.214 (0.025)	0.155 (0.045)	0.133 (0.035)	0.126 (0.044)	0.245
180	0.204 (0.032)	0.186 (0.037)	0.224 (0.017)	0.183 (0.037)	0.144 (0.035)	0.148 (0.041)	0.245
240	0.214 (0.025)	0.202 (0.029)	0.229 (0.014)	0.201 (0.029)	0.160 (0.033)	0.164 (0.035)	0.245

Table 3: Simulation Results for Out-of-Sample Sharpe Ratios of Optimal Portfolios

Note: In this table, we report the mean and standard deviation of the out-of-sample Sharpe ratios for various optimal portfolios constructed by SPCA, PCA, rpPCA, PLS, Lasso, and Ridge approaches, respectively. The true data-generating process, given by scenario f), has four factors, driven by RmRf, SMB, HML, and a weak factor V , whereas we estimate the SDF using a vector of factor proxies, g_t , that includes noisy versions of the four factors. The reported Sharpe ratios are the sample average over 1,000 Monte Carlo repetitions and their standard errors are reported in the brackets. Column “Theoretical Value” provides the benchmark Sharpe ratio calculated by $b^T E(r) / \sqrt{b^T \Sigma b}$ using true parameter values. We compare three scenarios with $T = 120, 180,$ and 240 , where $N = 2,000$ is fixed.

propose two additional examples of applications of the SPCA methodology: the estimation of fund alpha, and the de-noising of observable factors (similar in spirit to [Daniel et al. \(2020\)](#)).

4.1 Estimation of Risk Premia using SPCA

4.1.1 Data

Our main dataset is the [Chen and Zimmermann \(2020\)](#) data, which includes a large number of equity portfolios sorted by characteristics. Specifically, we employ the April 2021 release of the data. For each characteristic considered, [Chen and Zimmermann \(2020\)](#) construct a variable number of portfolios (as many as are used in the original papers that introduced the anomaly in the literature: typically 2, 5, or 10). Not all test assets are available for the entire time period; for our analysis, we study the time period 1976m3 to 2020m12, for which 901 test portfolios are available without missing values. To these sorted portfolios, we add 49 industry portfolios from Ken French’s website. All of our results are at the monthly frequency.

We also consider an alternative dataset, proposed by [Hou et al. \(2020\)](#), that includes for the same period 1672 portfolios sorted by characteristics without missing values. [Hou et al. \(2020\)](#) classify their portfolios in six groups: momentum, value, investment, profitability, intangibles, frictions. The two datasets of [Hou et al. \(2020\)](#) and [Chen and Zimmermann \(2020\)](#) are similar and yield comparable results. Rather than producing two versions of each result using the two datasets, we choose [Chen and Zimmermann \(2020\)](#) to be our main dataset and report the robustness of the main results using the [Hou et al. \(2020\)](#) data (e.g., see section [4.1.6](#)).

We study the risk premium of both tradable and nontradable factors, focusing on the best-known ones from the literature. The tradable factors are: the market (in excess of the risk-free rate); size (SMB); value (HML); profitability (RMW); investment (CMA); momentum (MOM); betting-against-beta (BAB, from [Frazzini and Pedersen \(2014\)](#)); and quality-minus-junk (QMJ, from [Asness et al. \(2013\)](#)). The nontradable factors are: the liquidity factor from [Pástor and Stambaugh \(2003\)](#); the intermediary capital factor from [He et al. \(2017\)](#); AR(1) innovations in industrial production growth (IP); VAR(1) innovations in the first three principal components of 279 macro-finance variables from [Ludvigson and Ng \(2010\)](#); AR(1) innovations in the three uncertainty indexes of [Jurado et al. \(2015\)](#), representing financial uncertainty, macroeconomic uncertainty, and real uncertainty; AR(1) innovations in the term spread, the credit spread, and the unemployment rate; AR(1) innovations in two sentiment indexes, one from [Huang et al. \(2015\)](#) and one from [Baker and Wurgler \(2006\)](#); oil price growth AR(1) innovations; and consumption growth AR(1) innovations.²¹

²¹The market factor, SMB, HML, RMW, CMA and MOM are from Ken French’s website. BAB and QMJ are from AQR’s website. The liquidity factor is from Lubos Pastor’s website. The intermediary capital factor is from Asaf Manela’s website. The macro principal components and the uncertainty indexes are from Sydney Ludvigson’s website. Industrial production, the credit spread, unemployment rate, the term spread, and oil price are from Fred-MD. The [Huang et al. \(2015\)](#) sentiment index is from Huang’s webpage. The [Baker and Wurgler \(2006\)](#) sentiment index is from Wurgler’s website. The consumption factor was built from NIPA data using the methodology of [Schorfheide et al. \(2018\)](#).

4.1.2 Choice of Tuning Parameters and Implementation Details

To apply SPCA to the estimation of the risk premia and to evaluate its out-of-sample performance, we split the sample period into two equal-sized subsamples. The first half of the sample (training period) is used to choose the tuning parameters and produce the risk premium estimate. The second half of the sample (evaluation period) is used to evaluate the out-of-sample performance of the estimator.

For ease of presentation, we choose to select only one tuning parameter, q , for each plausible choice of p in our analysis. This approach reduces the number of tuning parameters to only one, and also conveniently serves as a robustness check.

To determine reasonable candidates of p , we examine the factor structure of the panel of test asset returns. Figure 3 provides the scree plot of the log of the first 25 eigenvalues. There appear to be at least three strong factors. In addition, it appears that factors 4-11 might also be relevant, though weaker. Motivated by the scree plot, in the empirical study below we highlight results for p equal to 3, 5, 7, and 11, therefore showing the robustness of our results to a wide range of model dimensions.




Figure 3: Logarithm of the First 25 Eigenvalues in the Chen-Zimmerman data

Note: The figure plots the logarithm of the first 25 eigenvalues of the data, obtained from [Chen and Zimmermann \(2020\)](#) plus 49 industry portfolios, covering the period 1976-2020.

To choose the tuning parameter q , we adopt the same criterion as in simulations to evaluate the estimator's out-of-sample performance, namely, the hedging ability of the portfolio built by SPCA for g_t . More concretely, recall that all estimators of risk premia (e.g., the standard Fama-MacBeth estimator, the PCA-based estimator of [Giglio and Xiu \(2021\)](#), and SPCA) recover the risk premium of a factor g_t by, implicitly or explicitly, building a tradable portfolio that isolates exposure to g_t . We

therefore compute the weights of the hedging portfolio built by SPCA using the training data only, and calculate the mean-squared-error of hedging g_t over the validation period using that portfolio (effectively, the out-of-sample R^2). We apply this criterion to pick q using cross-validation (CV) within the training sample.

Our empirical analysis proceeds as follows. We first choose the number of factors p in the model, based on the scree plot. Then, working exclusively in the training sample, we run 3-fold cross-validation 100 times. In each cross-validation run, the tuning parameter q is chosen to maximize the R^2 of the hedging portfolio described above (requiring a minimum of 100 assets selected). Our choice for q is the median across the 100 cross-validation runs, and the risk-premium estimate is the one obtained under that choice of q ; we also compute the weights of the hedging portfolio for the factor. All the analysis described so far uses only data from the training period. Next, we evaluate the out-of-sample performance of the estimator. Using the test asset returns from the evaluation period, together with the portfolio weights estimated previously, we compute the return of the hedging portfolio in the evaluation sample, and we calculate the fraction of the variance of g_t hedged by that portfolio, i.e., R^2 . This calculation does not re-estimate any parameter or coefficient in the evaluation data, and is thus a fully out-of-sample R^2 .

4.1.3 Results: Estimation of Risk Premia and Out-of-sample Evaluation

We report the main empirical results in Table 4 and Figures 4 and 5. Each row of Table 4 corresponds to one factor; the first 8 are tradable, the rest are nontradable. For tradable factors, the first two columns show the average excess return of the factor, in the training sample and in the evaluation samples, respectively; these numbers correspond to model-free estimates of the risk premia of tradable factors, and can be directly compared with the model-based estimate obtained from SPCA.

The next columns of the table show the SPCA results in 4 groups of columns, corresponding to the number of latent factors $p = 3, 5, 7,$ and 11 , respectively. For each choice of p , we report the risk-premium estimate (obtained in the training sample, in bp per month), the number of assets selected by SPCA (determined by q), and the out-of-sample R^2 obtained in the evaluation period. These estimates are obtained factor by factor: that is, in each case, g_t contains one factor, and the asset selection is driven by that factor only. In the last two columns of the table, we repeat the exercise (with $p = 11$) but estimate all risk premia simultaneously: g_t contains all the factors and the selection of the assets is based on all of them simultaneously (so that $d \geq p$ as opposed to $d = 1$). As discussed above, in this case, assets are sorted by the maximum of the correlation with any of the factors in g_t for the purpose of the selection step. In theory, both approaches are consistent. In practice, estimating risk premia factor by factor has the advantage that the latent factors zoom in immediately on the assets relevant for each factor. On the other hand, the joint estimation is required to satisfy the more stringent assumptions for the CLT of Section 2.6 (because for the CLT to work, we need to assume that g_t has exposure to the entire SDF, while this is not required for the

consistency of the estimates for an individual factor).

Consider first the market portfolio (first row of the table), a strong factor in this dataset. The average return of the market in the training sample is 74bp per month, and 62bp in the evaluation period. The SPCA estimates of the market risk premium, for the four chosen values of p , are 68, 70, 72, and 74bp per month, respectively, all close to the average excess return. To obtain these estimates, SPCA estimates the latent factors picking, in each iteration, 100 assets out of the total of 950. Finally, the portfolio that SPCA builds to hedge the market achieves, not surprisingly, a very high out-of-sample R^2 , above 0.98 for all p .

To better understand the out-of-sample performance of the estimator, we can examine the heatmap in Figure 4, panel (a), which focuses on the market factor. In the heatmap, the x axis reports the number of factors p ; the y axis reports the number of test assets selected by SPCA (in turn determined by q); for each combination of p and q , the heatmap reports the out-of-sample R^2 of SPCA, that is, the ability of the SPCA hedging portfolio to hedge g_t in the evaluation sample. The heatmap gives a complete description of the out-of-sample properties of SPCA as a function of the two parameters p and q . Panel (a) shows that for all combinations of p and q (that is, throughout the heatmap), out-of-sample R^2 s are overall very high for the market portfolio, above 85%. However, there appears to be a subset of the parameter space where performance is especially good: combinations with high p and low q .

The red marks in the heatmap correspond to the values of q chosen by cross-validation (CV) *in the training sample* (one for each value of p considered in the table: 3, 5, 7, 11). Ideally, the values of q chosen by CV in the training sample would perform well out of sample: that is, the marks should lie in areas in the heatmap with high out-of-sample R^2 s. This is indeed the case, as the figure shows, indicating good out-of-sample performance of the SPCA estimator and of the tuning parameter selection procedure.

Consider now another tradable factor, CMA, in the 5th row of Table 4. Like for the market, the estimated risk premium for CMA is not statistically significantly different from the average excess return of the factor. The number of assets selected by SPCA ranges between 100 and 350, and the out-of-sample R^2 is above 50%, indicating that our latent factor model is able to capture the majority of the variation on CMA out of sample.²²

The heatmap of the out-of-sample R^2 for this factor is panel (e) of Figure 4. The figure shows that for the case of CMA, different combinations of p and q yield very different hedging performance out-of-sample, with R^2 s ranging from above 50% to below 0. It is especially important then that our tuning parameter selection procedure yields good results out-of-sample. The red marks in the figure

²²Given that the universe of test assets includes portfolios sorted by the same characteristics used to construct the tradable factors like CMA, one may wonder why an out-of-sample R^2 of 100% is not always obtained for tradable factors. The reason is that SPCA is trying to build a hedging portfolio for the target g_t with factors that must also explain covariation among the universe of test assets. An advantage of our approach is that the hedging portfolio is able to avoid fitting the “measurement error” component in g_t , which, as discussed above, can be thought of as non-diversified idiosyncratic error for tradable factors, or more literally measurement error for nontradables. We come back to this point in section 4.2.2.

Table 4: Risk premia estimates

	3 Latent Factors		5 Latent Factors		7 Latent Factors		11 Latent Factors		Joint estim, 11 factors							
	Avg. ret. (train.)	Avg. ret. (eval.)	RP	# Assets	R^2	RP	# Assets	R^2	RP	# Assets	R^2	RP	Stderr			
Market	74	62	68	100	0.98	70	100	0.98	72	100	0.99	74	100	0.99	73	26
HML	39	-7	50	100	0.70	37	100	0.79	39	150	0.78	44	250	0.79	54	18
SMB	12	25	15	100	0.82	5	100	0.85	10	100	0.85	10	100	0.85	7	18
RMW	37	28	-8	100	-0.18	40	100	0.56	33	100	0.61	27	150	0.66	23	9
CMA	26	19	36	250	0.41	40	100	0.55	27	200	0.55	31	350	0.53	34	11
Momentum	91	30	67	100	0.79	86	100	0.87	102	100	0.87	101	100	0.88	96	23
BAB	126	56	112	100	0.43	120	100	0.38	112	150	0.35	128	150	0.45	93	20
QMJ	41	39	-9	100	0.43	28	100	0.81	31	100	0.80	36	150	0.78	20	10
Liquidity			70	550	0.01	85	650	0.02	83	700	0.04	95	900	0.03	105	25
Intermed. Cap.			112	100	0.59	101	100	0.56	121	150	0.55	116	350	0.52	109	41
IP growth			-4	950	-0.01	-4	950	-0.02	-5	950	-0.03	-2	950	0.00	-2	3
LN 1			225	550	-0.28	202	650	-0.19	150	700	-0.11	54	950	-0.12	35	146
LN 2			-70	950	-0.05	-79	950	-0.12	-24	950	-0.16	-29	950	-0.17	-53	82
LN 3			96	400	0.03	86	650	0.06	16	700	0.06	-21	850	0.05	-92	78
Consumption			2	950	-0.01	3	950	0.00	3	950	-0.01	2	950	-0.01	2	2
Fin. Unc.			-61	350	-0.08	-48	750	0.00	-40	850	0.09	-41	950	0.10	-46	17
Real Unc.			-6	950	0.05	-7	950	0.04	-9	950	0.04	-11	950	0.06	-17	12
Macro Unc.			-7	950	0.08	-10	950	0.08	-10	950	0.08	-16	950	0.09	-19	10
Term			229	950	-0.11	81	950	-0.36	-57	950	-0.54	262	950	-0.59	384	372
Credit			41	950	-0.03	62	950	-0.03	41	950	-0.02	-43	950	-0.03	-32	77
Unempl.			65	950	0.00	109	950	-0.01	112	950	-0.01	110	950	0.00	45	108
Sentiment HJTZ			-24	950	0.01	-27	950	-0.03	-18	950	-0.06	-40	950	-0.07	-34	76
Sentiment BW			57	950	0.00	64	950	0.00	50	950	0.01	16	950	-0.02	44	71
Oil			-37	950	-0.05	-62	950	-0.02	-42	950	-0.03	-20	950	-0.02	-9	41

Note: In this table, we report the estimation results for tradable and nontradable factors using SPCA. The first two columns report the average excess returns for tradable factors, in the training sample (first half of the sample period) and in the evaluation sample (second half of the sample period). The remaining columns report, for different values of the number of factors p , the risk premia estimates (in basis points per month, computed in the training period), the number of assets selected by SPCA (governed by the parameter q), and the out-of-sample R^2 of the implied hedging portfolio. The last two columns report risk premia estimates and standard errors including all factors in g_t simultaneously, with $p = 11$. Sample is the Chen and Zimmermann (2020) test portfolios plus 49 industry portfolios, over the period 1976-2020.

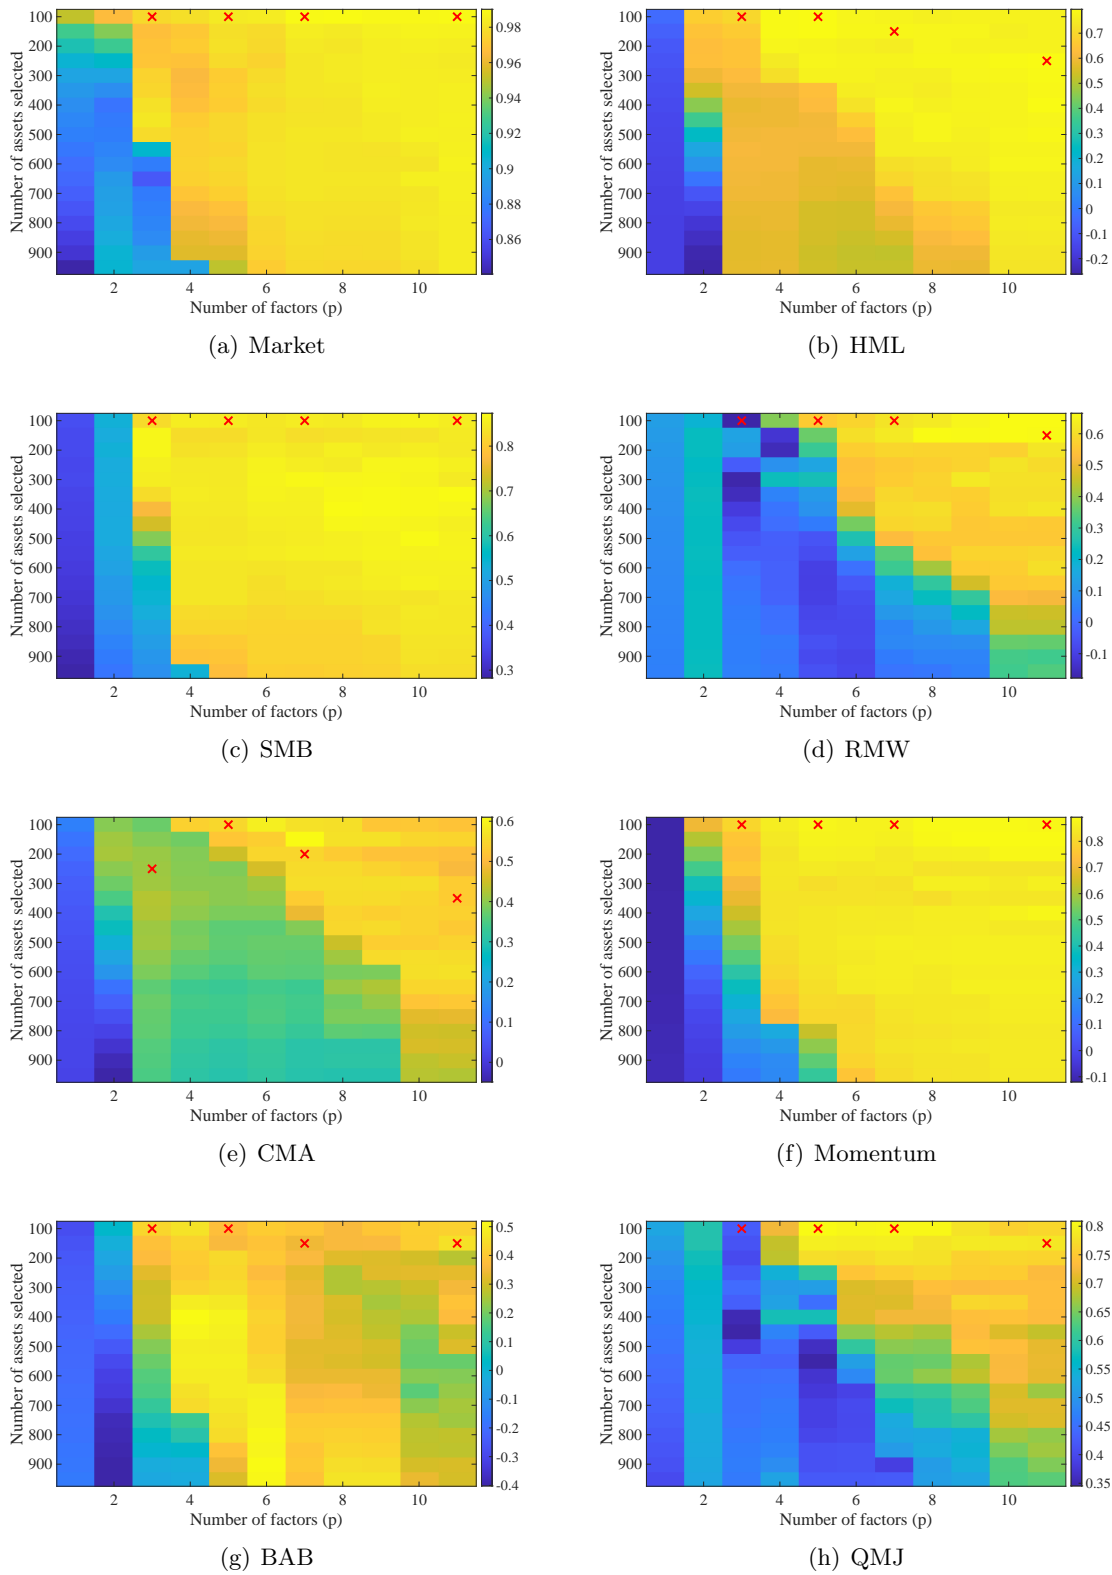
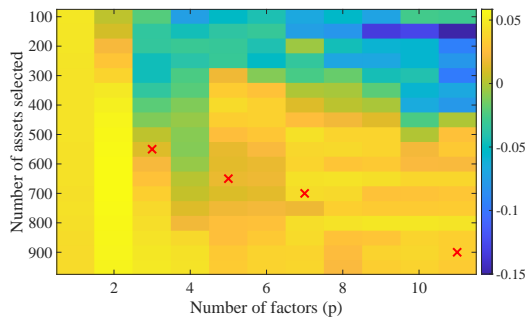



Figure 4: Out-of-sample R^2 Heatmaps, Tradable Factors


Note: Each panel reports the out-of-sample R^2 heatmap for a different factor. X-axis reports p . Y-axis reports the number of assets selected, governed by q . The colors in the heatmap correspond to the out-of-sample R^2 of the SPCA-implied hedging portfolio for the factor g_t ; this R^2 is computed entirely in the evaluation period. The red marks are the points chosen by CV within the training sample.




(a) Liquidity




(b) Intermediary Capital




(c) IP Growth




(d) LN #1




(e) LN #3



(f) Consumption



(g) Fin. Unc.



(h) Oil

Figure 5: Out-of-sample R^2 Heatmaps, Nontradable Factors

Note: Same as figure 4, but for a subset of nontradable factors.

show that this is indeed the case, especially for $p = 5$ and above.

In addition to showing the performance of SPCA, these heatmaps also allow us to compare the results with the PCA-based estimator of Giglio and Xiu (2021). This is because the last row of the heatmap corresponds to the case $q = 1$, that is, all assets are used to estimate the factors; but in that case, SPCA coincides with PCA. Looking across the various panels of Figure 4, it is clear that while for some factors (like the market) similar R^2 can be obtained (for appropriate choices of p) by PCA and SPCA, for other factors (like CMA and RMW) the out-of-sample R^2 s obtained by SPCA are substantially higher than those obtainable by PCA, for any choice of p (graphically: the area with the highest R^2 s is concentrated in the upper part of the heatmap, where $q < 1$). This shows that the case of weak factors studied in this paper is relevant in empirical applications.

One additional advantage of SPCA that is clearly visible in the heatmaps is that SPCA often manages to achieve the same (or better) R^2 than PCA, while estimating a much smaller number of factors. For example, consider the momentum factor in panel (f). The last row of the heatmap shows that extracting factors via PCA achieves an R^2 above 70% only once at least 6 factors are included; SPCA gets there even with 3 factors. The reason is intuitive: SPCA focuses on the test assets that are most informative about g_t , and therefore can zoom in quickly on the most relevant latent factors. This robustness to the number of factors is another advantage of SPCA that is relevant in practical applications.

For nontradable factors, we cannot compare the risk premium estimate from SPCA with the average excess return; the out-of-sample R^2 therefore plays an even more important role in evaluating the performance of the estimator. Note that it is well known in the literature that it is difficult to hedge nontradable factors, like consumption or IP growth, in equity markets. We will however show that SPCA gives a hedging portfolio that successfully hedges at least a part of the variation in many nontradable factors.

Consider first the liquidity factor of Pástor and Stambaugh (2003), in row 9 of Table 4 and panel (a) of Figure 5. The out-of-sample R^2 achieved by SPCA is above 0 (up to 4%), and the estimated risk premium appears to be high (between 70 and 95bp per month). Panel (a) of Figure 5 shows how strongly this R^2 depends on p and q . Among all combinations of parameters, a large fraction actually delivers a negative out-of-sample R^2 . This simply stresses how difficult it is to hedge this factor (like most macro factors) using equity markets, and emphasizes again the good performance of SPCA.

The remainder of the table and of the two figures shows the results for all the other factors (for reasons of space, the heatmaps only report a subset of the factors, while the table reports them all). A few interesting patterns emerge. First, for tradable factors, SPCA gives risk premia estimates that are always close to the model-free estimates obtained from average excess returns: the two are never statistically different at the 5% level (with the only exception of QMJ with $p = 3$). Second, confirming previous literature, nontradable factors are much harder to hedge than tradable factors; in fact, for several factors – like the first two JLN macro factors – we do not get positive R^2 at

all. For those factors, there is so little exposure in equity returns that SPCA cannot build a proper hedging portfolio. However, SPCA is able to hedge out of sample at least a part of the variation of many factors, like the third LN factor, the three uncertainty measures, the liquidity factor and the intermediary capital factor (for which it achieves an R^2 above 50%). Third, the risk premia estimated by SPCA – for those factors where SPCA can actually hedge some of the variation – make economic sense: for example, the liquidity and intermediary factors command significantly positive risk premia, whereas the three uncertainty measures command negative risk premia.

4.1.4 Asset Selection

To better understand how SPCA estimates the risk premium, we can study which assets are selected when extracting the latent factors. Table 5 shows, for four representative factors (two tradables, Momentum and RMW, and two nontradables, liquidity and intermediary capital), the top 10 test assets (by absolute value of correlation) selected at each step. The names of the portfolios follow [Chen and Zimmermann \(2020\)](#), with the numbers indicating the quintile or decile of the characteristic.

Consider Momentum, in the first set of rows of the table. To extract the first latent factor, SPCA selects the assets with the highest correlation with the momentum factor. As the table shows, the highest correlation is achieved by IntMom09 (an intermediate momentum portfolio). The correlation is 0.44. The other assets with high correlation are all momentum-related, not surprisingly. In the next columns, the table shows the assets selected at the second iteration of SPCA, after orthogonalizing g_t and the test assets to the first factor. Interestingly, the correlations among these residuals are even higher, up to 0.79 for a different momentum sort (Mom12mOffSeason, momentum without the seasonal component). This suggests that the first factor captures some of the asset variation that is not exclusively specific to momentum (for example, part of the market factor), which the projection step of SPCA removes. The residuals of the factor and the portfolio are then *more* correlated than the original factors and portfolios after the influence of the first factor is eliminated. In any case, momentum portfolios appear again at the second iteration, and, in part, at the third iteration.

The remainder of the table shows which assets are selected at the different iterations for RMW, Liquidity, and Intermediary Capital. For RMW (a profitability factor), the assets selected are often based on accounting measures, like asset growth, accruals, leverage, and operating profits. For liquidity, portfolios sorted by payout yield and beta seem to play an important role in hedging the risk. Finally, for intermediary capital, the portfolios selected by SPCA relate to idiosyncratic volatility, liquidity, as well as two industry portfolios (not surprisingly, banking and financials).

The selection of particularly informative assets is the central mechanism through which SPCA addresses the issue of weak factors. It is also responsible for the robustness of SPCA to the number of factors used, as we noted above: given that SPCA zooms in on the most informative assets, it can build a good hedging portfolio (and therefore a good estimate of the risk premium) even with a small number of factors.

Table 5: Assets Selected by SPCA

	Factor #1		Factor #2		Factor #3	
	Asset	Corr	Asset	Corr	Asset	Corr
Mom	IntMom09	0.44	Mom12mOffSeason02	0.79	Mom12m08	0.64
	IntMom10	0.4	Mom12mOffSeason03	0.76	BMdec05	0.63
	MomVol10	0.37	Size01	0.74	IntMom03	0.63
	MomVol09	0.36	ResidualMomentum01	0.73	SP05	0.62
	IntMom08	0.36	ResidualMomentum02	0.73	ShareIss5Y05	0.62
	Mom12m10	0.36	NumEarnIncrease01	0.72	BookLeverage02	0.62
	FirmAgeMom05	0.35	ShareIss5Y01	0.7	cfp05	0.61
	Mom12mOffSeason10	0.34	MomVol03	0.69	BMdec04	0.61
	Mom12mOffSeason09	0.33	CompEquIss01	0.68	ShareIss1Y05	0.6
	Mom12m09	0.33	Mom12m03	0.68	LReversal04	0.6
RMW	Industry:Gold	0.27	OperProf05	0.54	OperProfRD01	0.53
	MomOffSeason10	0.27	OperProfRD09	0.53	RoE01	0.47
	AccrualsBM02	0.27	CBOperProf09	0.5	GP01	0.45
	DelEqu05	0.27	RoE05	0.49	CBOperProf02	0.45
	LReversal05	0.27	CBOperProf10	0.49	DoIVol01	0.44
	roaq01	0.26	Leverage02	0.49	OperProfRD02	0.44
	AssetGrowth10	0.26	OperProfRD08	0.49	CBOperProf01	0.43
	DoIVol05	0.25	realestate03	0.49	OperProf01	0.41
	ChEQ05	0.25	GP05	0.49	RoE02	0.4
	Price05	0.25	GP04	0.48	VolMkt02	0.4
Liq.	InvGrowth06	0.47	InvGrowth06	0.28	InvGrowth06	0.3
	NetPayoutYield07	0.47	BetaFP09	0.26	DoIVol01	0.27
	PayoutYield05	0.46	EntMult06	0.25	XFIN08	0.26
	PayoutYield07	0.46	NetPayoutYield07	0.24	MeanRankRevGrowth01	0.26
	BetaFP03	0.46	PayoutYield07	0.24	BetaFP03	0.25
	DelLTI02	0.46	PayoutYield05	0.24	ShortInterest01	0.25
	IntanBM03	0.46	cfp04	0.23	BetaFP09	0.24
	EntMult06	0.46	BetaFP10	0.23	EntMult06	0.24
	VolMkt04	0.46	XFIN08	0.23	PayoutYield07	0.24
	PayoutYield06	0.46	ShortInterest01	0.22	ChEQ04	0.23
Interm.	Industry:Banks	0.9	Industry:banks	0.76	Industry:banks	0.7
	Industry:Fin	0.84	Industry:Fin	0.56	Industry:Fin	0.47
	IntMom05	0.8	DelEqu02	0.46	DebtIssuance02	0.38
	EquityDuration04	0.8	grcapx3y02	0.44	NOA10	0.36
	IdioVolAHT05	0.8	OScore02	0.43	ChAssetTurnover04	0.35
	IdioVol3F05	0.79	GrLTNOA10	0.43	HerfAsset05	0.35
	MaxRet08	0.79	ChAssetTurnover04	0.43	ShareRepurchase01	0.35
	Illiquidity01	0.79	IntMom05	0.43	HerfBE05	0.35
	IdioRisk05	0.79	IdioVolAHT05	0.42	DelEqu05	0.32
	CBOperProf03	0.78	Tax01	0.42	Beta05	0.32

Note: For each factor (one per panel) the table shows the top-10 assets selected by SPCA in extracting the latent factors. Assets are sorted by absolute value of the correlation. For each factor from 1 to 3, the table reports the names of the portfolios selected, and the absolute value of the correlation with g_t . Naming convention for the portfolios follows [Chen and Zimmermann \(2020\)](#).

4.1.5 SPCA and the Universe of Test Assets

The fact that SPCA estimates the latent factors using the most informative assets also makes it particularly robust to the universe of test assets used in the estimation. We explore this here in detail by considering three factors, value, momentum, and profitability, for which we can easily identify test assets that are informative about them. Specifically, we consider (for this section only) the dataset from [Hou et al. \(2020\)](#), which, as discussed in section 4.1.1, collects test portfolios by characteristics in six groups, among which one is labeled “value vs. growth”, one “momentum”, and one “profitability”. We can then ask: how does SPCA perform in estimating the value risk premium if we exclude the value and growth sorts from the universe? Similarly, how does it perform in estimating the momentum and profitability risk premium if momentum and profitability test assets, respectively, are removed? When the corresponding sorted portfolios are removed, the factors naturally become weaker. However, we should expect SPCA to still perform well, at least as long as sufficient exposure to the factor is present in the remaining test assets. On the contrary, we expect PCA’s performance to deteriorate more sharply.

Figure 6 reports the out-of-sample time-series R^2 heatmap for the three factors: value, momentum and profitability. On the left of each row we can see the R^2 obtained using all assets from the [Hou et al. \(2020\)](#) dataset; on the right we can see the results excluding the test assets corresponding to each factor. By looking at the last row of each heatmap, which corresponds to the PCA estimate with no selection, it is clear that the performance of PCA deteriorates significantly when the most informative assets are removed. Consider for example the case $p = 9$. For value, the PCA estimator’s out-of-sample R^2 decreases from 64% to 47%, as value and growth assets are removed; SPCA’s R^2 decreases by substantially less, from 74% to 62%. In the case of momentum, the R^2 decreases from 76% to 48% for PCA, but only from 86% to 77% for SPCA. Finally, for profitability, the R^2 decreases from 41% to 14% for PCA, but only from 71% to 60% for SPCA. In all cases, the SPCA estimator deteriorates little when the relative sorts are removed and the factor is made weaker, whereas the deterioration in performance is much larger for PCA.

To conclude, these empirical results mirror the simulations in section 3, that show SPCA performing well even when the factor of interest is weak in the universe of test assets considered. This is important in practical applications: given a certain factor g_t , we do not know ex ante if many or just a few assets are exposed to that factor. SPCA builds a good hedging portfolio and provides a consistent estimate of the risk premium in either case.

4.1.6 Robustness

We conclude by reporting in Table 6 a version of Table 4 obtained using the [Hou et al. \(2020\)](#) dataset instead of the [Chen and Zimmermann \(2020\)](#) data. The results are qualitatively similar to the ones obtained using the [Chen and Zimmermann \(2020\)](#) data, and, with a few exceptions, not statistically different. This confirms that, broadly, the results do not depend on using one particular universe of

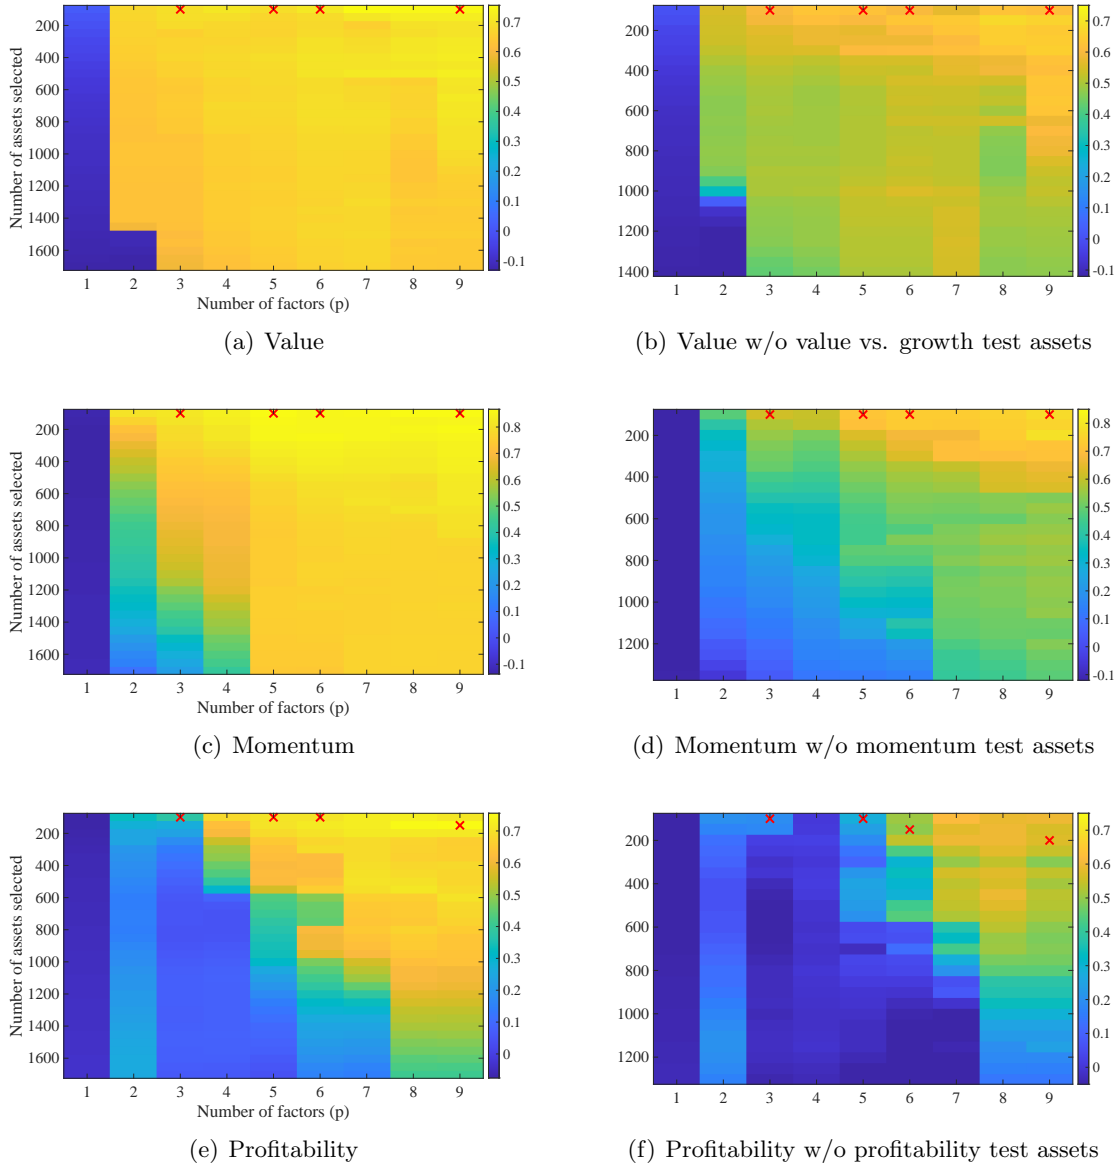


Figure 6: Varying the universe of test assets

Note: For value, momentum and RMW (profitability), the figure shows the out-of-sample R^2 heatmaps when all the test assets from Hou et al. (2020) are used in the estimation (left), and when value portfolios, momentum portfolios, or profitability portfolios, respectively, are excluded (right).

test assets.

4.2 Other Applications

In this section we study two additional applications of SPCA: the estimation of the alpha of a fund using latent factors to capture risk exposures, and using it to de-noise factor returns.

Table 6: Risk premia estimates, Hou et al. (2020) data

	Avg. ret. (train.)		3 Latent Factors		5 Latent Factors		6 Latent Factors		9 Latent Factors		Joint estim, 9 factors		
		(eval.)	RP	R^2	RP	R^2	RP	R^2	RP	R^2	RP	Stderr	
Market	74	62	72	100	0.98	74	100	0.99	74	100	0.99	71	26
HML	39	-7	22	100	0.69	20	100	0.69	16	100	0.71	18	16
SMB	12	25	-12	100	0.74	-13	100	0.72	-16	100	0.74	-15	18
RMW	37	28	12	100	0.38	26	100	0.71	25	100	0.71	36	9
CMA	26	19	8	100	0.70	11	100	0.65	12	100	0.66	4	11
Momentum	91	30	68	100	0.81	60	100	0.86	57	100	0.85	55	20
BAB	126	56	47	100	0.08	37	100	0.07	31	100	0.08	27	12
QMJ	41	39	-3	150	0.68	15	100	0.82	15	100	0.83	17	10
Liquidity			28	1700	0.05	35	1700	0.06	42	1700	0.06	35	18
Intermed. Cap.			107	100	0.49	98	100	0.46	91	100	0.51	63	37
IP growth			-2	1700	0.01	-4	1700	-0.02	-3	1700	-0.01	-3	2
LN 1			171	1200	-0.11	215	1650	-0.15	151	1700	-0.11	169	93
LN 2			-19	1700	-0.08	-17	1700	-0.08	-13	1700	-0.08	-4	55
LN 3			16	1000	0.03	69	1550	0.04	26	1700	0.02	15	62
Consumption			0	1700	0.00	0	1700	0.00	1	1700	0.00	0	1
Fin. Unc.			-5	1600	0.18	-15	1700	0.16	-15	1700	0.16	-9	11
Real Unc.			-4	1700	0.02	-5	1700	0.02	-8	1700	0.02	-6	6
Macro Unc.			-2	1700	0.05	-4	1700	0.05	-6	1700	0.05	-4	5
Term			-11	1700	-0.11	24	1700	-0.10	77	1700	-0.08	24	240
Credit			24	1700	-0.02	29	1700	-0.03	0	1700	-0.06	8	40
Unempl.			42	1700	0.00	116	1700	-0.01	112	1700	-0.01	101	61
Sentiment HJTZ			-44	1700	0.01	-39	1700	0.01	-22	1700	0.02	-20	44
Sentiment BW			-29	1700	0.03	-31	1700	0.02	-21	1700	0.02	-25	43
Oil			-8	1600	-0.03	-39	1500	0.00	-35	1600	0.00	-26	28

Note: Same as Figure 4, but using the characteristic-sorted portfolios from Hou et al. (2020) instead of those from Chen and Zimmermann (2020).

4.2.1 Estimating Buffett’s alpha using Latent Factors

When evaluating the performance of money managers, a crucial step is the choice of the benchmark against which the alpha of the fund is calculated. The benchmark model is often a standard tradable factor model like the Fama-French 3 or 5 factor model. How this benchmark should be selected is not entirely clear a priori; different benchmarks lead to different conclusions about the ability of managers to generate alpha. The asset pricing literature has proposed two approaches to address this issue that do not require arbitrarily choosing a benchmark: one approach is to use machine learning methods to select an optimal, parsimonious benchmark, chosen from a large set of candidate benchmarks in the “factor zoo” (Feng et al. (2020)). Another approach, proposed by Connor and Korajczyk (1986), uses latent factors to extract the relevant benchmarks, thus avoiding taking a stand on the identity of the factors (see also Giglio et al. (2020)).

SPCA offers a natural way to expand the second approach: by using a fund return as g_t , SPCA allows us to extract from the panel of returns all and only those factors that are informative about the fund’s risk exposures (and therefore compute the alpha after accounting for all risk exposures, including those to weak factors). In this section, we illustrate this possibility by applying SPCA to understand the alpha of Berkshire Hathaway, similar to Frazzini et al. (2013).

One of the headline results in Frazzini et al. (2013) is that Berkshire Hathaway’s returns display large alpha (13.4% annualized, statistically significant) when the fund’s return is benchmarked to the market factor alone. However, the alpha becomes much smaller (5.7%) and statistically insignificant when the benchmark model also includes SMB, HML, Momentum, BAB, and QMJ. Note that all the factor exposures together capture 29% of the time-series variation of the fund return.

We apply SPCA using the same data as in our main analysis (Chen-Zimmerman data plus 49 industry portfolios). The first factor in g_t will be Buffett’s return; we add to the vector of factors g_t all the other factors in our dataset, in order to get correct asymptotic standard errors. All the result here are in-sample, and we select the tuning parameter q by 3-fold cross-validation in the full sample.

The results are remarkably similar to those of Frazzini et al. (2013), with the main difference that we do not need to specify the identity of the factors. Specifically, when only one latent factor is extracted ($p = 1$), we obtain an alpha of 13%, with a t-stat of 4. The one-factor benchmark only explains 19% of the variation in the fund’s return. As more latent factors are included by SPCA, the R^2 increases and the alpha decreases. Once $p = 6$, the results effectively coincide with the ones of Frazzini et al. (2013): the time-series R^2 reaches 30% and the alpha drops to 6.1%, statistically insignificant (t-stat of 1.82).

The assets selected by SPCA give us some insights on what are the main risk exposures that determine Berkshire Hathaway’s risk premium. Specifically, among the portfolios most correlated with it, a large role is played by idiosyncratic volatility sorts, followed by a variety of portfolios related to profitability and leverage.

The results confirm the main insight of [Frazzini et al. \(2013\)](#): that Berkshire Hathaway’s return can be attributed in large part to exposure to priced factors; we however obtain this result without having to take a stand on the entire benchmark model, and relying only on SPCA to extract the latent factors. We can however gain some insights on the drivers of the risk premium by studying the assets selected by SPCA to build the hedging portfolio.

4.2.2 De-noising Factor Models

In a recent paper, [Daniel et al. \(2020\)](#) argue that the way standard factors are constructed based on characteristics sorts could be suboptimal as the portfolios used as factors may be contaminated by exposure to unpriced factors. They propose a procedure to remove the unpriced risk from observed factors, and produce a version of the Fama-French 5-factor model that achieves a higher Sharpe ratio and has better pricing ability for a cross-section of specifically-sorted test portfolios.

As discussed in the previous sections, SPCA constructs a hedging portfolio for any (tradable or nontradable) factor g_t that is built to capture the fundamental factors in the panel of test assets. By eliminating “measurement error” from g_t , SPCA effectively helps de-noise the factor from idiosyncratic (and therefore plausibly unpriced) risk. Similar in spirit to [Daniel et al. \(2020\)](#), we can use SPCA to strip out measurement error and build a de-noised version of the 5 factors in the FF5 model.

A natural exercise is then to compare the pricing ability for the panel of test assets of the original Fama-French 5-factor model, the de-noised model of [Daniel et al. \(2020\)](#), and the Fama-French 5 factors de-noised via SPCA. [Table 7](#) reports the average absolute alphas of these models (for SPCA, we consider different values of p to de-noise FF5 factors). The left set of columns reports the results using the [Chen and Zimmermann \(2020\)](#) data (CZ), the right column using the [Hou et al. \(2020\)](#) data (HXZ). For each set of results, we consider two versions of each model, restricting the zero-beta rate to be equal to the Tbill rate (left column), or with a free zero-beta rate (right column).

The table shows that for both the version with and without the zero-beta rate, SPCA produces an improvement over both the Fama-French 5-factor model and the [Daniel et al. \(2020\)](#) model, suggesting that removing the measurement error from the factor helps isolate the priced component of the factors.²³ The magnitude of the improvement of SPCA is as large as that by the [Daniel et al. \(2020\)](#) model.

5 Conclusions

The choice of test assets plays a fundamental role in empirical asset pricing tests. The recent explosion of anomaly discoveries and related characteristics in the empirical literature has provided researchers with a large universe of potential test assets to choose from. On the one hand, the availability of so

²³The assumptions on the zero-beta rate have a first-order effect only on the [Daniel et al. \(2020\)](#) results.

Table 7: Average absolute alphas across models

	CZ		HXZ	
	no zero-beta	w/ zero-beta	no zero-beta	w/ zero-beta
FF5	19	19	12	12
Daniel et al. (2020)	39	17	35	13
SPCA (5 factors)	20	20	10	10
SPCA (7 factors)	17	17	11	11
SPCA (11 factors)	17	17	11	11

Note: The table reports the average absolute alpha, in basis points per month, among the [Chen and Zimmermann \(2020\)](#) test assets (left) and the [Hou et al. \(2020\)](#) test assets (right), for different models: the Fama-French 5-factor model, the model of [Daniel et al. \(2020\)](#), and different versions of SPCA with different p . We consider two versions of each model, restricting the zero-beta rate to be equal to the Tbill rate (left), or with a free zero-beta rate (right).

many different characteristics gives us hope that the returns of these portfolios can help us uncover and identify the pricing of various dimensions of risk, including those that are not well captured by standard cross-sections. On the other hand, the large dimensionality goes hand in hand with the weak factor issue: a factor may well be captured by *some* assets within the large cross-section, but if most assets do not have exposure to that factor, it will be weak and inference will be incorrect.

Traditional methodologies to estimate risk premia take the cross-section of assets as given. In this paper, we present a new methodology, SPCA, that instead actively selects assets in order to estimate risk premia of factors of interest, whether they are strong or weak, and at the same time addresses the issue of potentially omitted factors, again regardless of whether they are strong or weak.

The paper confirms the performance of SPCA in a variety of simulations, and explores different empirical applications of SPCA to risk premia estimation, fund performance evaluation, and factor de-noising. Overall, the simulations and empirical analysis highlight a few important features of SPCA, that are particularly relevant for empirical applications: its robustness to the number of factors used, to the universe of test assets employed in the estimation, and, most importantly, to the strength of the factors in the data.

While the road to a full understanding of risk and risk premia in financial markets is still long, we believe that addressing systematically the issue of weak factors in empirical asset pricing is an important step forward, that opens the door to the study of factors that, while important to investors, may be not pervasive in the standard cross-sections.

References

- Ahn, D.-H., J. Conrad, and R. F. Dittmar (2009). Basis assets. *The Review of Financial Studies* 22(12), 5133–5174.
- Anatolyev, S. and A. Mikusheva (2021). Factor models with many assets: strong factors, weak factors, and the two-pass procedure. *Journal of Econometrics*, forthcoming.
- Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *Journal of Finance* 61, 259–299.
- Asness, C. S., A. Frazzini, and L. H. Pedersen (2013). Quality Minus Junk. Technical report, AQR.
- Bai, J. (2003). Inferential Theory for Factor Models of Large Dimensions. *Econometrica* 71(1), 135–171.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70, 191–221.
- Bai, J. and S. Ng (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics* 146(2), 304–317.
- Bailey, N., G. Kapetanios, and M. H. Pesaran (2020). Measurement of factor strength: Theory and practice.
- Bair, E., T. Hastie, D. Paul, and R. Tibshirani (2006). Prediction by supervised principal components. *Journal of the American Statistical Association* 101(473), 119–137.
- Bair, E. and R. Tibshirani (2004). Semi-supervised methods to predict patient survival from gene expression data. *PLoS Biology* 2(4), 511–522.
- Baker, M. and J. Wurgler (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance* 61(4), 1645–1680.
- Bryzgalova, S. (2015). Spurious Factors in Linear Asset Pricing Models. Technical report, Stanford University.
- Bryzgalova, S., J. Huang, and C. Julliard (2019). Bayesian solutions for the factor zoo: We just ran two quadrillion models. Available at SSRN 3481736.
- Bryzgalova, S., M. Pelger, and J. Zhu (2020). Forest through the trees: Building cross-sections of asset returns. Technical report, London School of Business and Stanford University.
- Chen, A. Y. and T. Zimmermann (2020). Open source cross-sectional asset pricing. Available at SSRN.
- Connor, G. and R. A. Korajczyk (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics* 15(3), 373–394.
- Daniel, K., L. Mota, S. Rottke, and T. Santos (2020). The cross-section of risk and returns. *The Review of Financial Studies* 33(5), 1927–1979.
- Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds.

- Journal of Financial Economics* 33(1), 3–56.
- Feng, G., S. Giglio, and D. Xiu (2020). Taming the factor zoo: A test of new factors. *Journal of Finance* 75(3), 1327–1370.
- Frazzini, A., D. Kabiller, and L. H. Pedersen (2013). Buffett’s alpha. Technical report, National Bureau of Economic Research.
- Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics* 111(1), 1–25.
- Freyaldenhoven, S. (2019). A generalized factor model with local factors.
- Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity datasets. *Econometrica* 84(3), 985–1046.
- Giglio, S., Y. Liao, and D. Xiu (2020). Thousands of alpha tests. *Chicago Booth Research Paper* (18-09), 2018–16.
- Giglio, S. W. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy* 129(7), 1947–1990.
- Harvey, C. R., Y. Liu, and H. Zhu (2016). ...and the Cross-Section of Expected Returns. *The Review of Financial Studies* 29(1), 5–68.
- He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics* 126(1), 1–35.
- Hou, K., C. Xue, and L. Zhang (2020). Replicating anomalies. *Review of Financial Studies* 33(5), 2019–2133.
- Huang, D., F. Jiang, K. Li, G. Tong, and G. Zhou (2021). Scaled pca: A new approach to dimension reduction. *Management Science*, forthcoming.
- Huang, D., F. Jiang, J. Tu, and G. Zhou (2015). Investor sentiment aligned: A powerful predictor of stock returns. *The Review of Financial Studies* 28(3), 791–837.
- Jagannathan, R. and Z. Wang (1998). An asymptotic theory for estimating beta-pricing models using cross-sectional regression. *The Journal of Finance* 53(4), 1285–1309.
- Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *The American Economic Review* 105(3), 1177–1216.
- Kan, R. and C. Zhang (1999). Two-Pass Tests of Asset Pricing Models with Useless Factors. *The Journal of Finance* 54(1), 203–235.
- Kelly, B. and S. Pruitt (2013). Market expectations in the cross-section of present values. *The Journal of Finance* 68(5), 1721–1756.
- Kelly, B., S. Pruitt, and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134(3), 501–524.
- Kim, S., R. A. Korajczyk, and A. Neuhierl (2020). Arbitrage portfolios. *Review of Financial Studies*,

Forthcoming.

- Kleibergen, F. (2009). Tests of risk premia in linear factor models. *Journal of Econometrics* 149(2), 149–173.
- Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics* 135(2), 271–292.
- Lettau, M. and M. Pelger (2020). Estimating latent asset-pricing factors. *Journal of Econometrics* 218, 1–31.
- Ludvigson, S. C. and S. Ng (2010). A factor analysis of bond risk premia. In A. Ulah and D. E. A. Giles (Eds.), *Handbook of empirical economics and finance*, Volume 1, Chapter 12, pp. 313–372. Chapman and Hall, Boca Raton, FL.
- Pástor, L. and R. F. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy* 111(3), 642–685.
- Pesaran, M. H. and R. Smith (2019). The role of factor strength and pricing errors for estimation and inference in asset pricing models.
- Schorfheide, F., D. Song, and A. Yaron (2018). Identifying long-run risks: A bayesian mixed-frequency approach. *Econometrica* 86(2), 617–654.