

Testing Beta-Pricing Models Using Large Cross-Sections

Valentina Raponi

Imperial College London

Cesare Robotti

University of Warwick

Paolo Zaffaroni

Imperial College London

We propose a methodology for estimating and testing beta-pricing models when a large number of assets is available for investment but the number of time-series observations is fixed. We first consider the case of correctly specified models with constant risk premia, and then extend our framework to deal with time-varying risk premia, potentially misspecified models, firm characteristics, and unbalanced panels. We show that our large cross-sectional framework poses a serious challenge to common empirical findings regarding the validity of beta-pricing models. In the context of pricing models with Fama-French factors, firm characteristics are found to explain a much larger proportion of variation in estimated expected returns than betas. (*JEL* G12, C12, C52)

We gratefully acknowledge comments and suggestions from three anonymous referees, Adrian Buss, Fernando Chague, Victor DeMiguel, Francisco Gomes, Cam Harvey, Andrew Karolyi (the editor), Ralph Koijen, Ľuboš Pástor, Tarun Ramadorai, Krishna Ramaswamy, Olivier Scaillet, Jay Shanken, Pietro Veronesi, Grigory Vilkov, Guofu Zhou, and especially Raman Uppal, and participants at CORE, Imperial College London, Luxembourg School of Finance, University of Georgia, University of Southampton, Toulouse School of Economics, Tinbergen Institute, University of Warwick, the 2015 Meetings of the Brazilian Finance Society, the CFE 2015, and the 2016 NBER/NSF Time Series Conference. Send correspondence to Paolo Zaffaroni, Imperial College Business School, Imperial College London, Tanaka Building, South Kensington Campus, London SW7 2AZ, United Kingdom; telephone: +44 20 7594 9186. E-mail: p.zaffaroni@imperial.ac.uk. Zaffaroni is also affiliated with Sapienza University of Rome.

Traditional econometric methodologies for estimating risk premia and testing beta-pricing models hinge on a large time-series sample size, T , and a small number of securities, N . At the same time, the thousands of stocks that are traded on a daily basis in financial markets provide a rich investment universe and an interesting laboratory for risk premia and cost of capital determination.¹ Moreover, although we have approximately a hundred years of U.S. equity data, much shorter time series are typically used in empirical work to mitigate concerns of structural breaks and to bypass the difficult issue of modeling explicitly the time variation in risk premia. Finally, when considering non-U.S. financial markets, only short time series are typically available.² Importantly, when N is large and T is small, the asymptotic distribution of any traditional risk premium estimator provides a poor approximation to its finite-sample distribution, thus rendering the statistical inference problematic.³

The main contribution of this paper is that it provides a methodology built on the large- N estimator of Shanken (1992), which allows us to perform valid inference on risk premia and assess the validity of the beta-pricing relation when N is large and T is fixed, possibly very small.⁴ Our novel methods are first illustrated for correctly specified models with constant risk premia and then extended to deal with time variation in risk premia, potential model misspecification, firm characteristics in the risk-return relation, and unbalanced panels. We also demonstrate that methodologies specifically designed for a large T and fixed N environment are no longer applicable when a large number of assets is used. Proposition 3 reveals the perils of inadvertently using the Fama and MacBeth (1973) t -ratios with the Shanken (1992) correction in our large N setting.

As emphasized by Shanken (1992), when T is fixed, one cannot reasonably hope for a consistent estimate of the traditional ex ante risk premium. For this reason, we focus on the ex post risk premia, which equal the ex ante risk premia plus the unexpected factor outcomes.⁵

¹For example, one can download the returns on 18,474 U.S. stocks for December 2013 from the Center for Research in Security Prices (CRSP), half of which are actively traded.

²For example, Table 1 in Hou, Karolyi, and Kho (2011) shows that, at most, only about thirty years of equity return data is available for emerging economies in Latin America, Europe/Middle East/Africa, and Asia-Pacific regions.

³The alternative approach of increasing the time-series frequency, although appealing, can lead to complications and is not always implementable. Potential problems with this approach include nonsynchronous trading and market microstructure noise. Furthermore, for models that include nontraded (macroeconomic) risk factors, high-frequency data is not available.

⁴Our methodology offers an alternative to the common practice of employing a relatively small number of portfolios for the purpose of estimating and testing beta-pricing models. Although the use of portfolios is typically motivated by the attempt of reducing data noisiness, it can also cause loss of information and lead to misleading inference due to data aggregation. (See, for example, Brennan, Chordia, and Subrahmanyam 1998; Berk 2002; and Ang, Liu, and Schwarz forthcoming, among others.)

⁵The ex post risk premium is a parameter with several attractive properties. It is unbiased for the ex ante risk

We start by considering the baseline case of a correctly specified beta-pricing model with constant risk premia when a balanced panel of test asset returns is available. We show that the estimator of Shanken (1992) is free of any pre-testing biases and that no data has to be sacrificed for the preliminary estimation of the bias. (See Proposition 1.) Next, we establish the asymptotic properties of the estimator, namely its \sqrt{N} -consistency and asymptotic normality. We derive an explicit expression for the estimator’s asymptotic covariance matrix and show how this expression can be used to construct correctly sized confidence intervals for the risk premia. Our assumptions are relatively mild and easily verifiable. In particular, we allow for a substantial degree of cross-correlation among returns (conditional on the factors’ realizations), and our assumptions are even weaker than the ones behind the arbitrage pricing theory (APT) of Ross (1976).

In the first extension of the baseline methodology, we show that the estimator continues to exhibit attractive properties even when risk premia vary over time. In particular, it accurately describes the time averages of the (time-varying) risk premia over a fixed time interval. We also derive a suitably modified version of the estimator that permits valid inference on risk premia at any given point in time. Noticeably, in our analysis we do not need to take a stand on the form of time variation in risk premia. Our time-varying risk premium estimator can accommodate non-traded as well as traded factors. For the latter, the traditional estimator based on the factors’ rolling sample mean is asymptotically valid for the true risk premium at a given point in time only for specific sampling schemes, and it requires a very large T to work when time variation is allowed for. (See the Online Appendix for details.)

Next, we allow for the possibility that the beta-pricing model is misspecified. We provide a new test of the validity of the beta-pricing relation and derive its large- N distribution under the null hypothesis that the model is correctly specified.⁶ Moreover, we show that our test enjoys nice size and power properties. We then establish the statistical properties of the estimator when the beta-pricing model is misspecified. This extension is particularly relevant when we reject the model’s validity based on the outcome of the specification test, but we are still interested in estimating the

premium, and the beta-pricing model is still linear in the ex post risk premia under the assumptions of either correctly specified or misspecified models. Finally, the corresponding ex post pricing errors can be used to assess the validity of a given beta-pricing model when T is fixed. Naturally, when T becomes large, any discrepancy between the ex ante and ex post risk premia vanishes because the sample mean of the factors converges to its population mean.

⁶Since our test is specifically designed for scenarios in which N is large, it alleviates the concerns of Lewellen, Nagel, and Shanken (2010), Harvey, Liu, and Zhu (2016), and Barillas and Shanken (2017) about a particular choice of test assets in the econometric analysis.

risk premia of a model with a possibly incomplete set of factors. Finally, we study an important case of deviations from exact pricing, that is, the cross-sectional dependence of expected returns on firm characteristics. The asymptotic covariance matrix of the normally distributed characteristic premia estimator is derived in closed form, unlike most approaches in this literature that typically rely on simulation-based arguments for inference purposes. Our method can be used to determine whether the beta-pricing model is invalid and to quantify the economic importance of the characteristics when there are deviations from exact pricing. By employing a new measure, which is immune to the often-documented cross-correlation between estimated betas and characteristics, we are able to determine the relative contribution of betas and characteristics to the overall cross-sectional variation in expected returns.

In the last methodological extension of our baseline analysis, we consider the case of unbalanced panels. This is a useful extension because eliminating observations for the sole purpose of obtaining a balanced panel could result in unnecessarily large confidence intervals for the risk premia and loss of power of the specification test.

We demonstrate the usefulness of our methodology by means of several empirical analyses. The three prominent beta-pricing specifications that we consider are the capital asset pricing model (CAPM), the three-factor Fama and French (1993) model (FF3), and the recently proposed five-factor Fama and French (2015) model (FF5). We also consider variants of these models augmented with the nontraded liquidity factor of Pástor and Stambaugh (2003). Our proposed methods under potential model misspecification uncover a significant pricing ability for all the traded factors in each of the three models, even when using a relatively short time window of three years. In contrast, the risk premia estimates often appear to be statistically insignificant when using the traditional large- T approaches. Based on our methodology, the liquidity factor appears to be priced in only about one-fifth of the three-year rolling samples examined. We also document strong patterns of time variation in risk premia, for both traded and nontraded factors. In addition, our specification test often rejects all beta-pricing models (with and without the liquidity factor), even when a short time window is used. Alternative methodologies, such as the finite- N approach of Gibbons, Ross, and Shanken (1989) and the more recent test of Gungor and Luger (2016), seem to have substantially lower power in detecting model misspecification. Finally, our results indicate that five prominent firm characteristics (book-to-market ratio, asset growth, operating profitability,

market capitalization, and six-month momentum) are important determinants of the cross-section of expected returns on individual assets. Although the characteristic premia estimates are not always found to be statistically significant, it seems that these characteristics jointly explain a fraction of the overall cross-sectional dispersion in expected returns that is about 30 times larger than the fraction explained by the estimated factors' betas, regardless of the beta-pricing model under consideration.

Our paper is related to a large number of studies in empirical asset pricing and financial econometrics. The traditional two-pass cross-sectional regression (CSR) methodology for estimating beta-pricing models, developed by Black, Jensen, and Scholes (1972) and Fama and MacBeth (1973), is valid when T is large and N is fixed. Shanken (1992) shows how the asymptotic standard errors of the second-pass CSR risk premia estimators are affected by the estimation error in the first-pass betas and provides standard errors that are robust to the errors-in-variables (EIV) problem.⁷ Shanken and Zhou (2007) derive the large- T properties of the two-pass estimator in the presence of global model misspecification.⁸ A different form of misspecification, not explored in this paper, can also occur when some of the factors have zero, or almost zero, betas, a situation that is referred to as the spurious or “useless” factors problem.⁹ Lack of identification of the risk premia also arises when at least one of the betas is cross-sectionally quasi-constant, as documented by Ahn, Perez, and Gadarowski (2013) with respect to the market factor empirical betas, a case also ruled out here.

Building on Litzenberger and Ramaswamy (1979), Shanken (1992) (Section 6) proposes a large- N estimator of the ex post risk premium and shows that it is asymptotically unbiased when N diverges and T is fixed. However, Shanken (1992) does not prove the consistency and asymptotic normality of this risk premium estimator.¹⁰

⁷Jagannathan and Wang (1998) relax the conditional homoscedasticity assumption of Shanken (1992). For a review of the large- T literature on beta-pricing models, see Shanken (1996), Jagannathan, Skoulakis, and Wang (2010), and Kan and Robotti (2012).

⁸See also Hou and Kimmel (2006) and Kan, Robotti, and Shanken (2013).

⁹Several methods have been developed to deal with this particular form of model misspecification. See, for example, Jagannathan and Wang (1998), Kan and Zhang (1999a, 1999b), Kleibergen (2009), Ahn, Perez, and Gadarowski (2013), Gospodinov, Kan, and Robotti (2014), Burnside (2015), Bryzgalova (2016), Gospodinov, Kan, and Robotti (2017), Ahn, Horenstein, and Wang (2018), Kleibergen and Zhan (2018, 2019), and Gospodinov, Kan, and Robotti (2019), among others.

¹⁰In the same paper, Shanken (1992) provides the well-known standard errors correction for ordinary least squares (OLS) and generalized least squares (GLS) estimators of the ex post risk premia, but his correction is valid only when T is large and N is fixed. (See his Section 3.2.)

Following these seminal contributions, other methods have been recently proposed to take advantage of the increasing availability of large cross-sections of individual securities. Our paper is close to Gagliardini, Ossola, and Scaillet (2016) in the sense that both studies provide inferential methods for estimating and testing beta-pricing models. However, their work is developed in a joint-asymptotics setting, where both T and N need to diverge. Moreover, they focus on a slightly different parameter of interest (obtained as the difference between the ex ante risk premia and the factors' population mean), which can be derived from the ex post risk premium by netting out the sample mean of the factor. Like us, Gagliardini, Ossola, and Scaillet (2016) need a bias adjustment, because in their setting N is diverging at a much faster rate than T .¹¹ Moreover, while Gagliardini, Ossola, and Scaillet (2016) assume random betas, as a consequence of their sampling framework with a continuum of assets, in our analysis we prefer to keep the betas nonrandom. This is for us mostly a convenience assumption, since we show in the Online Appendix that allowing for randomness of the betas in a large- N environment leaves our theoretical results unchanged. Gagliardini, Ossola, and Scaillet (2016) characterize the time variation in risk premia by conditioning on observed state variables, whereas we leave the form of time variation unspecified. Like us, they show how to carry out inference when the beta-pricing model is globally misspecified. Finally, Gagliardini, Ossola, and Scaillet (2016) allow for a substantial degree of cross-sectional dependence of the returns' residuals. Although our setup and assumptions differ from theirs (mainly because in our framework only N diverges), we also allow for a similar form of cross-sectional dependence in the residuals' covariance matrix.

Bai and Zhou (2015) investigate the joint asymptotics of the modified ordinary least squares (OLS) and generalized least squares (GLS) CSR estimators of the ex ante risk premia. Although the CSR estimators are asymptotically unbiased when T diverges, they propose an adjustment to mitigate the finite-sample bias. Their bias adjustment differs from the one suggested by Litzenberger and Ramaswamy (1979) and Shanken (1992), and studied in this paper, because it relies on a large T for its validity. However, their simulation results suggest that their bias-adjusted estimator performs well for various values of N and T . Moreover, since T must be large in their setting, the bias adjustment of Bai and Zhou (2015) is asymptotically negligible, implying that the asymptotic distribution of their CSR estimators is identical to the asymptotic distribution of the traditional

¹¹In contrast, recall that in the traditional analysis of the CSR estimator (where T diverges and N is fixed), no bias adjustment is required.

OLS and GLS CSR estimators.¹² In contrast, we show that the asymptotic distribution of the risk premia estimator must necessarily change in the fixed- T case, where the traditional trade-off between bias and variance emerges. Moreover, consistent estimation of the asymptotic covariance matrix of our risk premia estimator requires a different analysis because only N is allowed to diverge. Bai and Zhou (2015) focus exclusively on the case of a balanced panel under the assumption of correctly specified models. Unlike us, they do not account for time variation in the risk premia and do not analyze model misspecification.

Giglio and Xiu (2017) propose a modification of the two-pass methodology based on principal components that is robust to omitted priced factors and mismeasured observed factors, and establish its validity under joint asymptotics. Kim and Skoulakis (2018) employ the so-called regression calibration approach used in EIV models to derive a \sqrt{N} -consistent estimator of the ex post risk premia in a two-pass CSR setting.¹³ Finally, Jegadeesh et al. (forthcoming) propose instrumental-variable estimators of the ex post risk premia, exploiting the assumed independence over time of the return data.¹⁴

As for specification testing, Pesaran and Yamagata (2012) extend the classical test of Gibbons, Ross, and Shanken (1989) to a large- N setting. Besides accommodating only traded factors, the feasible version of their tests requires joint asymptotics, and N needs to diverge at a faster rate than T . Gungor and Luger (2016) propose a nonparametric testing procedure for mean-variance efficiency and spanning hypotheses (with tests of the beta-pricing restriction as a special case), and they derive bounds on the null distribution of the test statistics using resampling techniques. Their procedure, which is designed for traded factors only, is valid for any N and T , even though they show that the power of their test increases when both N and T diverge. Gagliardini, Ossola, and Scaillet (2016) derive the asymptotic distribution of their specification test under joint asymptotics

¹²Gagliardini, Ossola, and Scaillet (2016) show that the bias adjustment in their framework is not asymptotically negligible when N diverges at a much faster rate than T , a case not explicitly studied in Bai and Zhou (2015).

¹³Building on Jagannathan, Skoulakis, and Wang (2010), the Kim and Skoulakis (2018) estimator can be seen as an alternative to the Shanken estimator, the only difference being that in Kim and Skoulakis (2018) the first- and second-pass regressions are evaluated on non-overlapping time periods.

¹⁴Besides the classical econometric challenges associated with the choice of potentially weak instruments, these instrumental-variable approaches require a relatively larger T in order to achieve the same statistical accuracy of the Shanken (1992) estimator. Moreover, the construction of the instruments in Jegadeesh et al. (forthcoming) hinges upon the assumption of stochastic independence over time of the return data. The same assumption is also required in Kim and Skoulakis (2018). In contrast, it can be shown that the Shanken (1992) estimator retains its asymptotic properties even when the data is not independent over time. In fact, an arbitrary degree of serial dependence of the return data can be allowed for.

and, like us, they allow for general factors. Finally, Gagliardini, Ossola, and Scaillet (2018) propose a diagnostic criterion for detecting the number of omitted factors from a given beta-pricing model and establish its statistical behavior under joint asymptotics.

1. The Two-Pass Methodology

This section introduces the notation and describes the two-pass OLS CSR methodology. We assume that the asset returns $R_t = [R_{1t}, \dots, R_{Nt}]'$ are governed by the following beta-pricing model:

$$R_{it} = \alpha_i + \beta_{i1}f_{1t} + \dots + \beta_{iK}f_{Kt} + \epsilon_{it} = \alpha_i + \beta_i'f_t + \epsilon_{it}, \quad (1)$$

where i denotes the i -th asset, with $i = 1, \dots, N$, t refers to time, with $t = 1, \dots, T$, α_i is a scalar parameter representing the asset specific intercept, $\beta_i = [\beta_{i1}, \dots, \beta_{iK}]'$ is a K -vector of multiple regression betas of asset i with respect to the K factors $f_t = [f_{1t}, \dots, f_{Kt}]'$, and ϵ_{it} is the i -th return's idiosyncratic component at time t . In matrix notation, we can write the model as

$$R_t = \alpha + Bf_t + \epsilon_t, \quad t = 1, \dots, T, \quad (2)$$

where $\alpha = [\alpha_1, \dots, \alpha_N]'$, $B = [\beta_1, \dots, \beta_N]'$, and $\epsilon_t = [\epsilon_{1t}, \dots, \epsilon_{Nt}]'$. Let $\Gamma = [\gamma_0, \gamma_1']'$, where γ_0 is the zero-beta rate and γ_1 is the K -vector of ex ante factor risk premia, and denote by $X = [1_N, B]$ the beta matrix augmented with 1_N , an N -vector of ones. The following assumption of exact pricing (correct model specification) is used at various points in the analysis.

Assumption 1

$$E[R_t] = X\Gamma. \quad (3)$$

Equation (3) follows, for example, from no-arbitrage (Condition A in Chamberlain 1983) and a well-diversified mean-variance frontier (Definition 4 in Chamberlain 1983).¹⁵

Averaging Equation (2) over time, where we set $\bar{R} = \frac{1}{T} \sum_{t=1}^T R_t = [\bar{R}_1, \dots, \bar{R}_N]'$, $\bar{\epsilon} = \frac{1}{T} \sum_{t=1}^T \epsilon_t$, and $\bar{f} = [\bar{f}_1, \dots, \bar{f}_K]'$ = $\frac{1}{T} \sum_{t=1}^T f_t$, imposing assumption 1, and noting that $E[R_t] = \alpha + BE[f_t]$ from Equation (2), yields

$$\bar{R} = X\Gamma^P + \bar{\epsilon}, \quad (4)$$

¹⁵It should be noted that the mere absence of arbitrage is not sufficient for exact pricing, that is, nonzero pricing errors can coexist with no-arbitrage, as in the case of the APT of Ross (1976).

where $\Gamma^P = [\gamma_0, \gamma_1^P]'$, and

$$\gamma_1^P = \gamma_1 + \bar{f} - E[f_t]. \quad (5)$$

From Equation (4), average returns are linear in the asset betas conditional on the factor outcomes through the quantity γ_1^P , which, in turn, depends on the factors' sample mean innovations, $\bar{f} - E[f_t]$. The random coefficient vector γ_1^P in Equation (5) is referred to as the vector of ex post risk premia.¹⁶

Equation (5) shows that Γ and Γ^P will coincide when $\bar{f} = E[f_t]$, which happens for $T \rightarrow \infty$. When T is small, ex ante and ex post risk premia can differ substantially, as emphasized in the empirical section of the paper, although γ_1^P remains an unbiased measure for the ex ante risk premia, γ_1 .¹⁷

Note that Equation (4) cannot be used to estimate the ex post risk premia Γ^P since X is not observed. For this reason, the popular two-pass OLS CSR method first obtains estimates of the betas by running the following multivariate regression for every i :

$$R_i = \alpha_i 1_T + F\beta_i + \epsilon_i, \quad (6)$$

where $R_i = [R_{i1}, \dots, R_{iT}]'$, $\epsilon_i = [\epsilon_{i1}, \dots, \epsilon_{iT}]'$, $F = [f_1, \dots, f_T]'$ is the $T \times K$ matrix of factors, and 1_T is a T -vector of ones. Then, the OLS estimates of B are given by

$$\hat{B} = R' \tilde{F} (\tilde{F}' \tilde{F})^{-1} = B + \epsilon' \mathcal{P}, \quad (7)$$

where $\hat{B} = [\hat{\beta}_1, \dots, \hat{\beta}_N]'$, $R = [R_1, \dots, R_N]$, $\epsilon = [\epsilon_1, \dots, \epsilon_N]$, and $\mathcal{P} = \tilde{F} (\tilde{F}' \tilde{F})^{-1}$ with $\tilde{F} = [\tilde{f}_1, \dots, \tilde{f}_T]' = \left(I_T - \frac{1_T 1_T'}{T} \right) F = F - 1_T \bar{f}'$, where I_T is the identity matrix of order T . The corresponding matrix of OLS residuals is given by $\hat{\epsilon} = [\hat{\epsilon}_1, \dots, \hat{\epsilon}_N] = R - 1_T \bar{R}' - \tilde{F} \hat{B}'$.

We then run a single CSR of the sample mean vector \bar{R} on $\hat{X} = [1_N, \hat{B}]$ to estimate the risk premia. Note that we have two alternative feasible representations of Equation (4), that is,

$$\bar{R} = \hat{X} \Gamma + \eta, \quad (8)$$

with residuals $\eta = \left[\bar{\epsilon} + B(\bar{f} - E[f_t]) - (\hat{X} - X)\Gamma \right]$, and

$$\bar{R} = \hat{X} \Gamma^P + \eta^P, \quad (9)$$

¹⁶For traded factors, Equation (5) reduces to $\gamma_1^P = \bar{f} - \gamma_0 1_K$, where 1_K is a K -vector of ones. (See Shanken 1992.)

¹⁷It should be noted that any valid estimator of γ_1^P provides, as a byproduct, a valid estimator of the population parameter $\nu = \gamma_1 - E[f_t] = \gamma_1^P - \bar{f}$, namely, the portion of the ex ante risk premia that is nonlinearly related to the factors. This is the quantity studied in Gagliardini, Ossola, and Scaillet (2016).

with residuals $\eta^P = [\bar{\epsilon} - (\hat{X} - X)\Gamma^P]$. The OLS CSR estimator applied to either Equation (8) or Equation (9) yields

$$\hat{\Gamma} = \begin{bmatrix} \hat{\gamma}_0 \\ \hat{\gamma}_1 \end{bmatrix} = (\hat{X}'\hat{X})^{-1}\hat{X}'\bar{R}. \quad (10)$$

However, when T is fixed, $\hat{\Gamma}$ is not a consistent estimator of the ex ante risk premia, Γ , in Equation (8) and of the ex post risk premia, Γ^P , in Equation (9). The reason is that neither \hat{B} converges to B , nor does \bar{f} converge to $E[f_t]$ unless $T \rightarrow \infty$. Focusing on the representation in Equation (9), the OLS CSR estimator can be corrected as follows. Denote by $\text{tr}(\cdot)$ the trace operator and by 0_K a K -vector of zeros. In addition, let

$$\hat{\sigma}^2 = \frac{1}{N(T - K - 1)}\text{tr}(\hat{\epsilon}'\hat{\epsilon}). \quad (11)$$

The bias-adjusted estimator of Shanken (1992) is then given by

$$\hat{\Gamma}^* = \begin{bmatrix} \hat{\gamma}_0^* \\ \hat{\gamma}_1^* \end{bmatrix} = (\hat{\Sigma}_X - \hat{\Lambda})^{-1} \frac{\hat{X}'\bar{R}}{N}, \quad (12)$$

where

$$\hat{\Sigma}_X = \frac{\hat{X}'\hat{X}}{N} \quad \text{and} \quad \hat{\Lambda} = \begin{bmatrix} 0 & 0'_K \\ 0_K & \hat{\sigma}^2(\tilde{F}'\tilde{F})^{-1} \end{bmatrix}. \quad (13)$$

The formula for the estimator $\hat{\Gamma}^*$ exhibits a multiplicative bias adjustment through the term $(\hat{\Sigma}_X - \hat{\Lambda})^{-1}$.¹⁸ This prompts us to explore the analogies of $\hat{\Gamma}^*$ with the more conventional class of additive bias-adjusted OLS CSR estimators. To this end, it is useful to consider the following expression for the OLS CSR estimator, $\hat{\Gamma}$, obtained from Theorem 1 of Bai and Zhou (2015):

$$\begin{aligned} \hat{\Gamma} &= \Gamma^P + \left(\frac{\hat{X}'\hat{X}}{N}\right)^{-1} \begin{bmatrix} 0 & 0'_K \\ 0_K & -\hat{\sigma}^2(\tilde{F}'\tilde{F})^{-1} \end{bmatrix} \Gamma^P + O_p\left(\frac{1}{\sqrt{N}}\right) \\ &= \Gamma^P - \left(\frac{\hat{X}'\hat{X}}{N}\right)^{-1} \hat{\Lambda}\Gamma^P + O_p\left(\frac{1}{\sqrt{N}}\right). \end{aligned} \quad (14)$$

Equation (14) suggests a simple way to construct an additive bias-adjusted estimator of Γ^P ; that is,

$$\hat{\Gamma}^{bias-adj} = \hat{\Gamma} + \left(\frac{\hat{X}'\hat{X}}{N}\right)^{-1} \hat{\Lambda}\hat{\Gamma}^{prelim}, \quad (15)$$

¹⁸Equation (15) in Shanken (1992) differs slightly from our Equation (12). The reason is that we do not impose the traded-factor restriction of Shanken (1992), that is, consistent with much of the extant literature on the two-pass methodology, throughout our analysis we do not constrain the risk premia on the traded factors to be equal to the corresponding factor means, even when the beta-pricing model is correctly specified.

where $\hat{\Gamma}^{prelim}$ is an arbitrary preliminary estimator of Γ^P .¹⁹ The next proposition shows that, by imposing that the preliminary estimator, $\hat{\Gamma}^{prelim}$, and the bias-adjusted estimator, $\hat{\Gamma}^{bias-adj}$, coincide, the unique solution to Equation (15) is the Shanken (1992) estimator $\hat{\Gamma}^*$ in Equation (12).

Proposition 1 Assume that $\hat{\Sigma}_X - \hat{\Lambda}$ is nonsingular. Then, the Shanken (1992) estimator $\hat{\Gamma}^*$ in Equation (12) is the unique solution to the linear system of equations:

$$\hat{\Gamma}^* = \hat{\Gamma} + \left(\frac{\hat{X}'\hat{X}}{N} \right)^{-1} \hat{\Lambda}\hat{\Gamma}^*. \quad (16)$$

Proof: See the Online Appendix.

Therefore, $\hat{\Gamma}^*$ is the unique additive bias-adjusted OLS CSR estimator that does not require the preliminary estimation of the risk premia. As a computational precaution, it is possible that the EIV correction in Equation (12) overshoots, making the matrix $(\hat{\Sigma}_X - \hat{\Lambda})$ almost singular for a given N and potentially leading to extreme values for the estimator.²⁰ To alleviate this risk, our suggestion is to multiply the matrix $\hat{\Lambda}$ by a scalar k ($0 \leq k \leq 1$) and to substitute $(\hat{\Sigma}_X - \hat{\Lambda})^{-1}$ with $(\hat{\Sigma}_X - k\hat{\Lambda})^{-1}$ in Equation (12), effectively yielding a shrinkage estimator.²¹ If k is zero, we obtain the OLS CSR estimator $\hat{\Gamma}$, whereas if k is one, we obtain the Shanken (1992) estimator $\hat{\Gamma}^*$.²² In our simulation experiments, we find that this shrinkage estimator is virtually unbiased, leading to $k = 1$. In contrast, in our empirical application in Section 4, shrinking is applied to roughly 75% of the cases (the average k is 0.58) when $T = 36$ and to 5% of the cases (the average k is 0.71) when $T = 120$. Our shrinkage adjustment can also alleviate the documented evidence of cross-sectional quasi-homogeneity for the loadings associated with certain risk factors, in particular for the market

¹⁹For example, Bai and Zhou (2015) propose using the OLS CSR $\hat{\Gamma}$ itself as the preliminary estimator, plugging it into the formula above in place of $\hat{\Gamma}^{prelim}$. However, this adjustment is justified only when $T \rightarrow \infty$. In general, the use of a preliminary estimator would decrease the precision of the bias-adjusted estimator and, in addition, it would make its properties harder to study.

²⁰This important point was first made by Chordia, Goyal, and Shanken (2015).

²¹Our asymptotic theory would require $k = k_N$ to converge to unity at a suitably slow rate as N increases. We omit the details to simplify the exposition.

²²The choice of the shrinkage parameter k can be based on the eigenvalues of the matrix $(\hat{\Sigma}_X - k\hat{\Lambda})$ as follows. Starting from $k = 1$, if the minimum eigenvalue of this matrix is negative and/or the condition number of this matrix is larger than 20 (as suggested by Greene 2003, 60), then we lower k by an arbitrarily small amount. In our empirical application we set this amount equal to 0.05 and perform shrinkage whenever the absolute value of the relative change between the Shanken (1992) and the OLS CSR estimators is greater than 100%. We iterate this procedure until the minimum eigenvalue is positive and the condition number becomes less than 20. Gagliardini, Ossola, and Scaillet (2016) rely on similar methods to implement their trimming conditions. Alternatively, one could use cross-validation to set the value of k .

factor (see Ahn, Perez, and Gadarowski 2013).²³

Before turning to the challenging task of deriving the large- N distribution of the Shanken (1992) estimator (and the associated standard errors), we discuss the perils of using the traditional t -ratios (specifically designed for a large- T environment) when N diverges. We first introduce the necessary assumptions and then present our results in Proposition 3.

Assumption 2 As $N \rightarrow \infty$,

$$\frac{1}{N} \sum_{i=1}^N \beta_i \rightarrow \mu_\beta \quad \text{and} \quad \frac{1}{N} \sum_{i=1}^N \beta_i \beta_i' \rightarrow \Sigma_\beta, \quad (17)$$

such that the matrix

$$\begin{bmatrix} 1 & \mu_\beta' \\ \mu_\beta & \Sigma_\beta \end{bmatrix} \quad (18)$$

is positive-definite.

Assumption 2 states that the limiting cross-sectional averages of the betas, and of the squared betas, exist. The second part of assumption 2 rules out the possibility of spurious factors and situations in which at least one of the elements of β_i is cross-sectionally constant. It implies that X has full (column) rank for N sufficiently large. To simplify the exposition, we assume that the β_i are nonrandom.²⁴

Assumption 3 The vector ϵ_t is independently and identically distributed (i.i.d.) over time with

$$E[\epsilon_t|F] = 0_N \quad (19)$$

and a positive-definite matrix,

$$\text{Var}[\epsilon_t|F] = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \cdots & \sigma_{1N} \\ \sigma_{21} & \sigma_2^2 & \cdots & \sigma_{2N} \\ \vdots & \vdots & \cdots & \vdots \\ \sigma_{N1} & \sigma_{N2} & \cdots & \sigma_N^2 \end{pmatrix} = \Sigma, \quad (20)$$

²³Ahn, Perez, and Gadarowski (2013) propose the so-called invariance beta (IB) coefficient as a measure of cross-sectional homogeneity. Applying their measure to our data on FF5, we find that the IB coefficient corresponding to the market factor equals 0.74 and 0.81 for rolling samples of size $T = 36$ and $T = 120$, respectively (averages across rolling samples). The IB coefficient is equal to 0.93 when considering the whole sample. According to Ahn, Perez, and Gadarowski (2013), these values signal a very moderate risk of multicollinearity due to cross-sectional homogeneity. Similar values of the IB coefficient associated with the loadings on the market factor are obtained when estimating CAPM and FF3.

²⁴See Gagliardini, Ossola, and Scaillet (2016) for a treatment of the beta-pricing model with random betas. In the Online Appendix, we discuss the consequences of relaxing the nonrandomness of the β_i .

where 0_N is a N -vector of zeros, and σ_{ij} denotes the (i, j) -th element of Σ , for every $i, j = 1, \dots, N$ with $\sigma_i^2 = \sigma_{ii}$.

The i.i.d. assumption over time is common to many studies, including Shanken (1992). However, our large- N asymptotic theory, in principle, permits the ϵ_{it} to be arbitrarily correlated over time, but the expressions would be more complicated. Conditions (19) and (20) are verified if the factors f_t and the innovations ϵ_s are mutually independent for any s, t . Noticeably, condition (20) is not imposing any specific structure on the elements of Σ . In particular, we are not assuming that the returns' innovations are uncorrelated across assets or exhibit the same variance. However, our large- N asymptotic theory needs to discipline the degree of cross-correlation among the innovations, although still allowing for a substantial degree of heterogeneity in the cross-section of asset returns. (See assumption 5 below.)

As for the factors, we impose minimal assumptions because our asymptotic analysis holds conditional on the factors' realizations.

Assumption 4 $E[f_t]$ does not vary over time. Moreover, $\tilde{F}'\tilde{F}$ is a positive-definite matrix for every $T \geq K$.

Assumption 5 As $N \rightarrow \infty$,

(i)

$$\frac{1}{N} \sum_{i=1}^N (\sigma_i^2 - \sigma^2) = o\left(\frac{1}{\sqrt{N}}\right), \quad (21)$$

for some $0 < \sigma^2 < \infty$.

(ii)

$$\sum_{i,j=1}^N |\sigma_{ij}| \mathbb{1}_{\{i \neq j\}} = o(N), \quad (22)$$

where $\mathbb{1}_{\{\cdot\}}$ denotes the indicator function.

(iii)

$$\frac{1}{N} \sum_{i=1}^N \mu_{4i} \rightarrow \mu_4, \quad (23)$$

for some $0 < \mu_4 < \infty$ where $\mu_{4i} = E[\epsilon_{it}^4]$.

(iv)

$$\frac{1}{N} \sum_{i=1}^N \sigma_i^4 \rightarrow \sigma_4, \quad (24)$$

for some $0 < \sigma_4 < \infty$.

(v)

$$\sup_i \mu_{4i} \leq C < \infty, \quad (25)$$

for a generic constant C .

(vi)

$$E[\epsilon_{it}^3] = 0. \quad (26)$$

(vii)

$$\frac{1}{N} \sum_{i=1}^N \kappa_{4,iiii} \rightarrow \kappa_4, \quad (27)$$

for some $0 \leq |\kappa_4| < \infty$, where $\kappa_{4,iiii} = \kappa_4(\epsilon_{it}, \epsilon_{it}, \epsilon_{it}, \epsilon_{it})$ denotes the fourth-order cumulant of the residuals $\{\epsilon_{it}, \epsilon_{it}, \epsilon_{it}, \epsilon_{it}\}$.

(viii) For every $3 \leq h \leq 8$, all the mixed cumulants of order h satisfy

$$\sup_{i_1} \sum_{i_2, \dots, i_h=1}^N |\kappa_{h, i_1 i_2 \dots i_h}| = o(N), \quad (28)$$

for at least one i_j ($2 \leq j \leq h$) different from i_1 .

Assumption 5 essentially describes the cross-sectional behavior of the model disturbances. In particular, assumption 5(i) limits the cross-sectional heterogeneity of the return conditional variance. Assumption 5(ii) implies that the conditional correlation among asset returns is sufficiently weak. Assumptions 5(i) and 5(ii) allow for many forms of strong cross-sectional dependence, as emphasized by the following proposition, which considers the case in which the ϵ_{it} obey a factor structure.

Proposition 2 Assume that

$$\epsilon_{i,t} = \lambda_i u_t + \eta_{i,t}, \quad (29)$$

where

$$\sum_{i=1}^N |\lambda_i| = O(N^\delta), \quad 0 \leq \delta < 1/4, \quad (30)$$

and (without loss of generality) for some fixed $q < N$ and some constant C ,

$$\lambda_1 + \dots + \lambda_q \sim CN^\delta, \quad (31)$$

with u_t i.i.d. $(0, 1)$ and $\eta_{i,t}$ i.i.d. $(0, \sigma_\eta^2)$ over time and across units, where the u_t and the $\eta_{i,s}$ are mutually independent for every i, s, t . Then,

- (i) Assumptions 5(i) and 5(ii) are satisfied with $\sigma^2 = \sigma_\eta^2$.
- (ii) The maximum eigenvalue of Σ diverges as $N \rightarrow \infty$.²⁵

Proof: See the Online Appendix.

Note that the boundedness of the maximum eigenvalue is the most common assumption on the covariance matrix of the disturbances in beta-pricing models. (See, e.g., the generalization of the APT by Chamberlain and Rothschild 1983.) Our assumptions are weaker than the ones for the APT because the maximum eigenvalue can now diverge. This implies that the row-column norm of Σ , $\sup_{1 \leq i \leq N} \sum_{j=1}^N |\sigma_{ij}|$, diverges.²⁶ Equation (29) is adopted in our Monte Carlo experiments reported in the Online Appendix. Other special cases nested by assumption 5 for which the cross-covariances σ_{ij} are nonzero are network and spatial measures of cross-dependence and a suitably modified version of the block-dependence structure of Gagliardini, Ossola, and Scaillet (2016).²⁷

In assumption 5(iii), we simply assume the existence of the limit of the conditional fourth-moment, averaged across assets. In assumption 5(iv), the magnitude of σ_4 reflects the degree of cross-sectional heterogeneity of the conditional variance of the asset returns. Assumption 5(v) is a bounded fourth-moment condition uniform across assets, which implies that $\sup_i \sigma_i^2 \leq C < \infty$. Assumption 5(vi) is a convenient symmetry assumption, but it is not strictly necessary for our results. Without assumption 5(vi) the asymptotic distribution would be more involved, due to the presence of terms such as the third moment of the disturbance (averaged across assets). Assumption 5(vii) allows for non-Gaussianity of the asset returns when $|\kappa_4| > 0$. For example, this

²⁵ The maximum eigenvalue of Σ is given by $\sup_{z \text{ s.t. } \|z\|=1} z' \Sigma z$.

²⁶ Assumption 5 allows for the maximum eigenvalue of Σ to diverge at rate $o(\sqrt{N})$. (See the proof of Proposition 2 for details.) Gagliardini, Ossola, and Scaillet (2016) can allow for a faster rate, $o(N)$, of divergence of the maximum eigenvalue of Σ because both T and N diverge in their double-asymptotics setting.

²⁷ Gagliardini, Ossola, and Scaillet (2016)'s assumption BD.2 on block sizes and block numbers requires that the largest block size shrinks with N and that there are not too many large blocks; that is, the partition in independent blocks is sufficiently fine-grained asymptotically. They show formally that such block-dependence structure is compatible with the unboundedness of the maximum eigenvalue of Σ .

assumption is satisfied when the marginal distribution of asset returns is a Student t with degrees of freedom greater than four. However, when estimating the asymptotic covariance matrix of the Shanken (1992) estimator, one needs to set $\kappa_4 = 0$ merely for identification purposes, as explained in Lemma 6 in the Online Appendix. This said, higher-order cumulants are not constrained to be zero, implying that $\kappa_4 = 0$ is not equivalent to Gaussianity. We are now ready to state our Proposition 3.

Proposition 3 Under assumptions 1–5 and as $N \rightarrow \infty$, the Fama and MacBeth (1973) t -ratios for $\hat{\Gamma} = [\hat{\gamma}_0, \hat{\gamma}_{11}, \dots, \hat{\gamma}_{1k}, \dots, \hat{\gamma}_{1K}]'$ based on the correction of Shanken (1992) satisfy the following relations.

(i) For the ex ante risk premia $\Gamma = [\gamma_0, \gamma_{11}, \dots, \gamma_{1k}, \dots, \gamma_{1K}]'$, we have

$$|t_{FM}(\hat{\gamma}_0)| = \frac{|\hat{\gamma}_0 - \gamma_0|}{SE_0^{FM}} \rightarrow_p \infty \quad (32)$$

and

$$|t_{FM}(\hat{\gamma}_{1k})| = \frac{|\hat{\gamma}_{1k} - \gamma_{1k}|}{SE_k^{FM}} \rightarrow_p \left| \frac{\bar{f}_k - E[f_{kt}]}{\hat{\sigma}_k/\sqrt{T}} - \frac{v'_{k,K} A^{-1} C \gamma_1^P}{\hat{\sigma}_k/\sqrt{T}} \right| \text{ for } k \geq 1. \quad (33)$$

(ii) For the ex post risk premia $\Gamma^P = [\gamma_0, \gamma_{11}^P, \dots, \gamma_{1k}^P, \dots, \gamma_{1K}^P]'$, we have

$$|t_{FM,P}(\hat{\gamma}_0)| = \frac{|\hat{\gamma}_0 - \gamma_0|}{SE_0^{FM,P}} \rightarrow_p \infty \quad (34)$$

and

$$|t_{FM,P}(\hat{\gamma}_{1k})| = \frac{|\hat{\gamma}_{1k} - \gamma_{1k}^P|}{SE_k^{FM,P}} \rightarrow_p \infty \text{ for } k \geq 1, \quad (35)$$

where SE_k^{FM} and $SE_k^{FM,P}$ are the Fama and MacBeth (1973) standard errors with the Shanken (1992) correction corresponding to the ex ante and ex post risk premia, respectively (see the Online Appendix for details), and where $v_{k,K}$ is k -th column of the identity matrix I_K , $\hat{\sigma}_k^2$ is the (k, k) -th element of $\tilde{F}'\tilde{F}/T$, $A = \Sigma_\beta - \mu_\beta\mu_\beta' + C$, and $C = \sigma^2(\tilde{F}'\tilde{F})^{-1}$.

Proof: See the Online Appendix.

In summary, Proposition 3 shows that a methodology designed for a fixed N and a large T , such as the one based on the Fama and MacBeth (1973) standard errors with the Shanken's correction,

is likely to lead to severe over-rejections when N is large, thus rendering the inference on the beta-pricing model invalid.²⁸ Our Monte Carlo simulations corroborate this finding, as emphasized in the Online Appendix. Moreover, Proposition 3 shows that when N and T are large, there is no need to apply the correction of Shanken (1992) to the Fama and MacBeth (1973) standard errors.

2. Asymptotic Analysis Under Correctly Specified Models

In this section, we establish the limiting distribution of the Shanken (1992) bias-adjusted estimator, $\hat{\Gamma}^*$, and explain how its asymptotic covariance matrix can be consistently estimated.

2.1 Baseline case

Our baseline case assumes that the beta-pricing model is correctly specified, that the risk premia are constant, and that the panel is balanced. This corresponds to the setup of Shanken (1992).

Let $\Sigma_X = \begin{bmatrix} 1 & \mu'_\beta \\ \mu_\beta & \Sigma_\beta \end{bmatrix}$, $\sigma^2 = \lim \frac{1}{N} \sum_{i=1}^N \sigma_i^2$, $U_\epsilon = \lim \frac{1}{N} \sum_{i,j=1}^N E \left[\text{vec}(\epsilon_i \epsilon'_i - \sigma_i^2 I_T) \text{vec}(\epsilon_j \epsilon'_j - \sigma_j^2 I_T)' \right]$, $M = I_T - D(D'D)^{-1}D'$, where μ_β , Σ_β , and σ_i^2 are defined in our assumptions above, $D = [1_T, F]$, $Q = \frac{1_T}{T} - \mathcal{P}\gamma_1^P$, $Z = (Q \otimes \mathcal{P}) + \frac{\text{vec}(M)}{T-K-1} \gamma_1^{P'} \mathcal{P}' \mathcal{P}$, and \otimes and $\text{vec}(\cdot)$ denote the Kronecker product and vec operators, respectively.

We make the following further assumption to derive the large- N distribution of the Shanken (1992) estimator.

Assumption 6 As $N \rightarrow \infty$, we have

(i)

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N \epsilon_i \xrightarrow{d} \mathcal{N}(0_T, \sigma^2 I_T). \quad (36)$$

(ii)

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N \text{vec}(\epsilon_i \epsilon'_i - \sigma_i^2 I_T) \xrightarrow{d} \mathcal{N}(0_{T^2}, U_\epsilon). \quad (37)$$

²⁸In particular, the t -ratio of the OLS CSR estimator for a particular element of the ex ante risk premium vector, γ_1 , equals the standardized sample mean of the associated factor plus a bias term. When T is allowed to diverge, the convergence of this t -ratio to a standard normal is re-obtained, but, for any given T , the deviations from normality can be substantial.

(iii) For a generic T -vector C_T ,

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N \left(C'_T \otimes \begin{pmatrix} 1 \\ \beta_i \end{pmatrix} \right) \epsilon_i \xrightarrow{d} \mathcal{N}(0_{K+1}, V_c), \quad (38)$$

where $V_c = c\sigma^2\Sigma_X$ and $c = C'_T C_T$. In particular, $\frac{1}{\sqrt{N}} \sum_{i=1}^N (C'_T \otimes \beta_i) \epsilon_i \xrightarrow{d} \mathcal{N}(0_K, V_c^\dagger)$, where $V_c^\dagger = c\sigma^2\Sigma_\beta$.

Primitive conditions for assumption 6 can be derived but at the cost of raising the level of complexity of our proofs. For instance, when Equations (29)–(30) hold, then Equation (36) follows by Theorem 2 of Kuersteiner and Prucha (2013) when the η_{it} satisfy their martingale difference assumptions. (See their assumptions 1 and 2.) This result extends easily to Equations (37)–(38) under suitable additional assumptions. (Details are available upon request.) We are now ready to state our first theorem.

Theorem 1 As $N \rightarrow \infty$, we have

(i) Under assumptions 1–5,

$$\hat{\Gamma}^* - \Gamma^P = O_p\left(\frac{1}{\sqrt{N}}\right). \quad (39)$$

(ii) Under assumptions 1–6,

$$\sqrt{N}(\hat{\Gamma}^* - \Gamma^P) \rightarrow_d \mathcal{N}(0_{K+1}, V + \Sigma_X^{-1}W\Sigma_X^{-1}), \quad (40)$$

where

$$V = \frac{\sigma^2}{T} \left[1 + \gamma_1^{P'} \left(\frac{\tilde{F}'\tilde{F}}{T} \right)^{-1} \gamma_1^P \right] \Sigma_X^{-1} \quad (41)$$

and

$$W = \begin{bmatrix} 0 & 0'_K \\ 0_K & Z'U_\epsilon Z \end{bmatrix}. \quad (42)$$

Proof: See the Online Appendix.

The expression in Equation (40) is remarkably simple and has a neat interpretation. The first term of this asymptotic covariance, V , accounts for the estimation error in the betas, and it is essentially identical to the large- T expression of the asymptotic covariance matrix associated

with the OLS CSR estimator in Shanken (1992). (See his Theorem 1(ii).) The term $\frac{\sigma^2}{T}\Sigma_X^{-1}$ in Equation (41) is the classical OLS CSR covariance matrix, which one would obtain if the betas were observed. The term $c = \gamma_1^{P'} \left(\tilde{F}' \tilde{F} / T \right)^{-1} \gamma_1^P$ is an asymptotic EIV adjustment, with $c \frac{\sigma^2}{T} \Sigma_X^{-1}$ being the corresponding overall EIV contribution to the asymptotic covariance matrix. As Shanken (1992) points out, the EIV adjustment reflects the fact that the variability of the estimated betas is directly related to the residual variance, σ^2 , and inversely related to the factors' variability, $\left(\tilde{F}' \tilde{F} / T \right)^{-1}$. The last term of the asymptotic covariance, $\Sigma_X^{-1} W \Sigma_X^{-1}$ in Equation (40), arises because of the bias adjustment that characterizes $\hat{\Gamma}^*$. The W matrix in Equation (42) accounts for the cross-sectional variation in the residual variances of the asset returns through U_ϵ . This term will vanish when $T \rightarrow \infty$. In the Appendix, we provide an explicit expression for U_ϵ , and we show that U_ϵ only depends on the fourth-moment structure of the ϵ_{it} , that is, on κ_4 and σ_4 .²⁹ The \sqrt{N} -rate of convergence obtained in Theorem 1(i) coincides with the rate of convergence established by Gagliardini, Ossola, and Scaillet (2016) with respect to their \sqrt{NT} -consistent estimator of $\nu = \gamma_1^P - \bar{f}$ when T is fixed.

To conduct statistical inference, we need a consistent estimator of the asymptotic covariance matrix, which we present in the next theorem. Let $M^{(2)} = M \odot M$, where \odot denotes the Hadamard product operator. In addition, define

$$\hat{Z} = (\hat{Q} \otimes \mathcal{P}) + \frac{\text{vec}(M)}{T - K - 1} \hat{\gamma}_1^{*'} \mathcal{P}' \mathcal{P} \quad \text{with} \quad \hat{Q} = \frac{1_T}{T} - \mathcal{P} \hat{\gamma}_1^*. \quad (43)$$

Theorem 2 Under assumptions 1–5 and the identification condition $\kappa_4 = 0$, as $N \rightarrow \infty$, we have

$$\hat{V} + \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \hat{W} \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \rightarrow_p V + \Sigma_X^{-1} W \Sigma_X^{-1}, \quad (44)$$

where

$$\hat{V} = \frac{\hat{\sigma}^2}{T} \left[1 + \hat{\gamma}_1^{*'} \left(\frac{\tilde{F}' \tilde{F}}{T} \right)^{-1} \hat{\gamma}_1^* \right] (\hat{\Sigma}_X - \hat{\Lambda})^{-1}, \quad (45)$$

$$\hat{W} = \begin{bmatrix} 0 & 0'_K \\ 0_K & \hat{Z}' \hat{U}_\epsilon \hat{Z} \end{bmatrix}, \quad (46)$$

and \hat{U}_ϵ is a consistent estimator of U_ϵ (see the Online Appendix), obtained by replacing σ_4 with

$$\hat{\sigma}_4 = \frac{\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N \hat{\epsilon}_{it}^4}{3 \text{tr}(M^{(2)})}. \quad (47)$$

²⁹See assumption 5 for the definition of κ_4 (the cross-sectional average of the fourth-order cumulants of the ϵ_{it}) and σ_4 (the cross-sectional average of the σ_i^4).

Proof: See the Online Appendix.

A remarkable feature of the result above is that a consistent estimate of the asymptotic covariance matrix of $\hat{\Gamma}^*$ can be obtained while leaving the residual covariance matrix Σ unspecified. In fact, with Σ having in general $N(N+1)/2$ distinct elements and our asymptotic theory being valid only for $N \rightarrow \infty$, consistent estimation of Σ would be infeasible. A convenient feature of the Shanken (1992) estimator is that it depends on Σ only through the average of the σ_i^2 . Moreover, its asymptotic covariance matrix depends on the limits of $\sum_{i,j=1}^N \sigma_{ij}/N$ and $\sum_{i=1}^N \sigma_i^4/N$. Our large- N asymptotic theory shows how these quantities can be estimated consistently. In contrast, the individual covariances σ_{ij} cannot be consistently estimated due to the fixed T . The condition $\kappa_4 = 0$ is required as a consequence of the small- T and large- N framework.³⁰ However, $\kappa_4 = 0$ is not as restrictive as it may seem. A sufficiently large level of heterogeneity in the σ_i^2 generates a substantial level of volatility in the conditional distribution of assets' returns by inducing a mixture-distribution effect.³¹

Finally, it should be noted that the theorems accommodate traded as well as nontraded factors. However, for traded factors, the factors' sample means could in principle be used as risk premia estimators. Clearly, a sufficiently large T is required for the sample means to converge to their population counterparts. For nontraded factors, for example, macroeconomic variables, a panel of test asset returns is required to pin down the factors' risk premia, as the time series of the factors do not suffice. Mimicking portfolio excess returns could also be used in place of the nontraded factors, with the population means of the mimicking portfolio excess returns serving as the true risk premia.³² However, the mimicking portfolio projection requires $N < T$, which is violated under our reference sampling scheme.³³

³⁰As we show in Lemma 6 of the Online Appendix, the limit of $\hat{\sigma}_4$ in Equation (47) converges to a linear combination of k_4 and σ_4 . These two parameters could be identified and consistently estimated only under the stronger assumption of independence across assets, since, in this case, σ_4 would reduce to σ^4 (which could be easily estimated using the square of $\hat{\sigma}^2$). In contrast, allowing for some arbitrary degree of cross-correlation implies that k_4 and σ_4 cannot be separately identified. This is the reason for setting $k_4 = 0$.

³¹In our empirical applications our estimate σ_4 is about 10 times the estimate for σ^4 .

³²See Breeden, Gibbons, and Litzenberger (1989), Chan, Karceski, and Lakonishok (1998), and Lamont (2001), among others, for empirical studies based on the mimicking portfolio methodology. Balduzzi and Robotti (2008) demonstrate by means of Monte Carlo simulations the greater accuracy of the mimicking portfolio risk premia estimates relative to the CSR risk premia estimates associated with the corresponding nontraded factors.

³³When $N > T$, one could obtain the first \tilde{N} principal components from a large panel of test assets returns, and then construct the mimicking portfolio for the nontraded factor using these \tilde{N} assets (assuming that $\tilde{N} < T < N$). Although this approach is feasible and is used in our empirical application, the theoretical properties of this double-projection approach are difficult to derive; see Giglio and Xiu (2017) for a theoretical analysis of a similar approach.

2.2 Time-varying case

In this section, we study the behavior of the estimator $\hat{\Gamma}^*$ when the risk premia are allowed to be time-varying, again under the assumption of correct model specification. It turns out that $\hat{\Gamma}^*$ is suitable for time-varying risk premia estimation because it estimates accurately local averages (over the, possibly very short, time window of size $T > K + 1$) of the true time-varying risk premia, regardless of their form and degree of time variation. Noticeably, we are also able to derive a consistent estimator of the true t -th period risk premia and to characterize its asymptotic distribution.³⁴

Throughout this section, we substitute assumption 1 with

$$E_{t-1}[R_{it}] = \gamma_{0,t-1} + \beta_i' \gamma_{1,t-1}, \quad (48)$$

where $E_{t-1}[\cdot]$ denotes the conditional expectation with respect to all the available information up to time $t - 1$. Importantly, our theory does not need to restrict the type of time variation in $\Gamma_{t-1} = [\gamma_{0,t-1}, \gamma_{1,t-1}]'$. To simplify the treatment of time variation in the premia, without altering the estimation procedure developed in this paper, we maintain the β_i in Equation (48) constant over time.³⁵ Our results below easily extend to the case of $\beta_{i,t-1} = B_i z_{t-1}$, for some (vector of) predetermined state variables z_{t-1} and a suitable matrix of loadings B_i .

Under Equation (48), asset returns are now given by $R_{it} = [1, \beta_i'] \Gamma_{t-1}^P + \epsilon_{it}$, where Γ_{t-1}^P are the $(t-1)$ -th ex post risk premia:

$$\Gamma_{t-1}^P = \Gamma_{t-1} + f_t - E_{t-1}[f_t], \text{ with a sample average } \bar{\Gamma}^P = \frac{1}{T} \sum_{t=1}^T \Gamma_{t-1}^P. \quad (49)$$

By construction, the ex post time-varying risk premia Γ_{t-1}^P have a conditional mean that equals Γ_{t-1} , the ex ante time-varying risk premia.

We are grateful to an anonymous referee for suggesting this approach to us.

³⁴Our new estimator for the time-varying risk premia appears useful also for traded factors, and not just for nontraded factors, especially when T is small. It should be noted that the (rolling) sample mean of the excess return on the traded factor will capture, in general, the average, over T observations, of the true time-varying risk premium associated with the factor. Alternatively, one can adopt the sampling scheme typical of nonparametric methods, with the implication that now the (rolling) sample mean will capture the time-varying risk premium and not just its average. However, a large T would be necessary to obtain accurate estimates and a certain degree of smoothness, over time, of the true time-varying risk premium would be required. (See the Online Appendix for further details.) Our method for estimation of time-varying risk premia works for any T and makes no smoothness assumption.

³⁵See, e.g., Ferson and Harvey (1991), who argue that the time variation in expected returns is mainly due to time variation in the premia as opposed to time variation in the betas.

To estimate the $(t-1)$ -th risk premia, for $t = 1, \dots, T$, we introduce the following novel estimator:

$$\hat{\Gamma}_{t-1}^* = \begin{bmatrix} \hat{\gamma}_{0,t-1}^* \\ \hat{\gamma}_{1,t-1}^* \end{bmatrix} = \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \frac{\hat{X}' R_t}{N} - \hat{\sigma}^2 \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \begin{pmatrix} 0 \\ (\tilde{F}' \tilde{F})^{-1} \tilde{F}' \iota_{t,T} \end{pmatrix}, \quad (50)$$

where, as before, $\iota_{t,T}$ denotes the t -th column, for $t = 1, \dots, T$, of the identity matrix I_T .³⁶ The next theorem derives the large- N behavior of both $\hat{\Gamma}^*$ and $\hat{\Gamma}_{t-1}^*$.

Theorem 3 Under Equation (48) and assumptions 2–6, as $N \rightarrow \infty$, we have

(i) $\hat{\Gamma}^*$ and $\sqrt{N}(\hat{\Gamma}^* - \bar{\Gamma}^P)$ satisfy Theorem 1 with Γ^P replaced by $\bar{\Gamma}^P$.

(ii) $\hat{\Gamma}_{t-1}^* - \Gamma_{t-1}^P = O_p\left(\frac{1}{\sqrt{N}}\right)$ and

$$\sqrt{N} \left(\hat{\Gamma}_{t-1}^* - \Gamma_{t-1}^P \right) \rightarrow_d \mathcal{N} \left(0_{K+1}, V_{t-1} + \Sigma_X^{-1} W_{t-1} \Sigma_X^{-1} \right), \quad (51)$$

where $V_{t-1} = \sigma^2 Q'_{t-1} Q_{t-1} \Sigma_X^{-1}$, $W_{t-1} = \begin{bmatrix} 0 & 0'_K \\ 0_K & Z'_{t-1} U_\epsilon Z_{t-1} \end{bmatrix}$, $Q_{t-1} = \iota_{t,T} - \mathcal{P} \gamma_{1,t-1}^P$, and $Z_{t-1} = (Q_{t-1} \otimes \mathcal{P}) - \frac{\text{vec}(M)}{T-K-1} Q'_{t-1} \mathcal{P}$, with U_ϵ as in Theorem 1.

Proof: See the Online Appendix.

Theorem 3 states that, when Equation (48) holds, $\hat{\Gamma}^*$ consistently estimates the local average of the ex post time-varying risk premia over T periods, the only requirement being that $T > K + 1$. If one is interested in the ex post risk premia for a specific time period, Γ_{t-1}^P , then asymptotically correct inference can be carried out by using $\hat{\Gamma}_{t-1}^*$. Interestingly, $\hat{\Gamma}^*$ is numerically identical to the sample mean of $\hat{\Gamma}_{t-1}^*$, over $t = 1, \dots, T$, because the additive bias adjustment, on the right-hand side of Equation (50), vanishes due to the identity $\sum_{t=1}^T \tilde{F}' \iota_{t,T} = \tilde{F}' 1_T = 0$.

To better understand the importance of our large- N results, it is useful to consider the behavior of the OLS CSR estimator $\hat{\Gamma}$ when Equation (48) holds. In this case, we have

$$\hat{\Gamma} \rightarrow_p \Gamma_\infty \quad \text{as } T \rightarrow \infty, \quad (52)$$

where $\Gamma_\infty = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \Gamma_s ds$ denotes the integrated risk premia, namely the long-run average over the entire timeline.³⁷ Next, consider $\hat{\Gamma}_{t-1} = (\hat{X}' \hat{X})^{-1} \hat{X}' R_t$, which can be thought of as the

³⁶Note that $\hat{\Gamma}_{t-1}^*$ is a new estimator that successfully tackles the problem of estimating time-varying risk premia in a large- N setting. It should not be confused with the Shanken (1992) formula in his Theorem 5.

³⁷If one assumes, as in Ang and Kristensen (2012), that $\Gamma_t = \Gamma(t/T)$, $1 \leq t \leq T$, for a smooth function $\Gamma(\cdot)$, then the integrated risk premia Γ_∞ become $\int_0^1 \Gamma_s ds$.

OLS CSR estimator for the $(t-1)$ -th risk premia.³⁸ It follows that

$$\hat{\Gamma}_{t-1} \rightarrow_p \Gamma_{t-1}^P + \begin{pmatrix} N & 1'_N B \\ B' 1_N & B' B \end{pmatrix}^{-1} \begin{pmatrix} 1'_N \\ B' \end{pmatrix} \epsilon_t \quad \text{as } T \rightarrow \infty. \quad (53)$$

Hence, the limit of $\hat{\Gamma}_{t-1}$ is the sum of two components, that is, the $(t-1)$ -th ex post risk premia Γ_{t-1}^P and a random term that is a function of ϵ_t . This last term cannot be consistently estimated, thus making $\hat{\Gamma}_{t-1}$ an unreliable estimator of both Γ_{t-1} and Γ_{t-1}^P , even when $T \rightarrow \infty$. In contrast, in our large- N framework, $\left(\hat{\Sigma}_X - \hat{\Lambda}\right)^{-1} \frac{\hat{X}' R_t}{N} \rightarrow_p \Gamma_{t-1}^P + \sigma^2 \Sigma_X^{-1} \begin{pmatrix} 0 \\ (\tilde{F}' \tilde{F})^{-1} \tilde{F}' \nu_{t,T} \end{pmatrix}$ as $N \rightarrow \infty$, where the bias term $\sigma^2 \Sigma_X^{-1} \begin{pmatrix} 0 \\ (\tilde{F}' \tilde{F})^{-1} \tilde{F}' \nu_{t,T} \end{pmatrix}$ can now be consistently estimated, leading to the bias-adjusted estimator $\hat{\Gamma}_{t-1}^*$ in Equation (50). Finally, a consistent estimator of the asymptotic covariance matrix of $\hat{\Gamma}_{t-1}^*$ in Equation (51) can be easily obtained. (See Theorem 2 and its proof in the Online Appendix.)

3. Asymptotic Analysis Under Potentially Misspecified Models

In this section, we explore the implications of model misspecification for model and parameter testing. Under the full rank assumption on the X matrix, the focus of the analysis is on the fixed (global) type of misspecification considered in Shanken and Zhou (2007) and several follow-up papers. A beta-pricing model is misspecified if there exists no value of the risk premia Γ for which the associated vector of pricing errors is zero. This misspecification might be due, for example, to the omission of some relevant risk factor, imperfect measurement of the factors, or failure to incorporate some relevant aspect of the economic environment – taxes, transaction costs, irrational investors, and the like. Thus, misspecification of some sort seems inevitable, given the inherent limitations of asset pricing models.

3.1 Testing for model misspecification

When a beta-pricing model is correctly specified (see assumption 1),

$$H_0 : e_i = 0 \quad \text{for every } i = 1, 2, \dots, \quad (54)$$

³⁸The quantity $\hat{\Gamma}_{t-1}$ is well known in empirical finance because its sample variance is routinely used to compute the Fama and MacBeth (1973) standard errors of $\hat{\Gamma}$.

where $e_i = E[R_{it}] - \gamma_0 - \beta'_i \gamma_1$ is the population (ex ante) pricing error associated with asset i . Denoting the vector of sample ex post pricing errors by

$$\hat{e}^P = (\hat{e}_1^P, \dots, \hat{e}_N^P)' = \bar{R} - \hat{X} \hat{\Gamma}^*, \quad (55)$$

we have

$$\begin{aligned} \hat{e}_i^P &= \bar{R}_i - \hat{X}_i \hat{\Gamma}^* \\ &= e_i + Q' \epsilon_i - \hat{X}_i (\hat{\Gamma}^* - \Gamma^P). \end{aligned} \quad (56)$$

Theorem 1(i) implies that, for every i ,

$$\hat{e}_i^P \rightarrow_p e_i + Q' \epsilon_i \equiv e_i^P. \quad (57)$$

Equation (57) shows that even when the ex ante pricing errors, e_i , are zero, \hat{e}_i^P will not converge in probability to zero because T is fixed. Nonetheless, a test of H_0 with correct size and good power can be developed. Define the sum of the sample squared ex post pricing errors as

$$\hat{Q} = \frac{1}{N} \sum_{i=1}^N (\hat{e}_i^P)^2. \quad (58)$$

Consider the centered statistic

$$\mathcal{S} = \sqrt{N} \left(\hat{Q} - \frac{\hat{\sigma}^2}{T} \left(1 + \hat{\gamma}_1^{*'} (\tilde{F}' \tilde{F} / T)^{-1} \hat{\gamma}_1^* \right) \right). \quad (59)$$

The centering is needed because of Equation (57). To see this, from the population ex post pricing errors, e_i^P , we have

$$\frac{1}{N} \sum_{i=1}^N (e_i^P)^2 = \frac{1}{N} \sum_{i=1}^N e_i^2 + Q' \left(\frac{1}{N} \sum_{i=1}^N \epsilon_i \epsilon_i' \right) Q + o_p(1) = \frac{1}{N} \sum_{i=1}^N e_i^2 + \sigma^2 Q' Q + o_p(1). \quad (60)$$

Therefore, even under $H_0 : e_i = 0$ for all i , the average of the population squared ex post pricing errors will not converge to zero but rather to $\sigma^2 Q' Q = \sigma^2 (1 + \gamma_1^{*'} (\tilde{F}' \tilde{F} / T)^{-1} \gamma_1^*)$. This is the quantity whose consistent estimate we need to demean our test statistic by in order to obtain its limiting distribution. The following theorem provides the limiting distribution of \mathcal{S} under $H_0 : e_i = 0$ for every i .

Theorem 4 Under Equation (54) and assumptions 2–6, as $N \rightarrow \infty$, we have

$$\mathcal{S} \rightarrow_d \mathcal{N}(0, \mathcal{V}), \quad (61)$$

where $\mathcal{V} = Z_Q' U_\epsilon Z_Q$ and $Z_Q = (Q \otimes Q) - \frac{\text{vec}(M)}{T-K-1} Q' Q$.

Proof: See the Online Appendix.

Our novel specification test easily follows from Theorem 4. Under H_0 ,

$$\mathcal{S}^* = \frac{\mathcal{S}}{\hat{\mathcal{V}}^{\frac{1}{2}}} \rightarrow_d \mathcal{N}(0, 1), \quad (62)$$

where

$$\hat{\mathcal{V}} = \hat{Z}'_Q \hat{U}_\epsilon \hat{Z}_Q \quad (63)$$

is a consistent estimator of \mathcal{V} in Equation (61), $\hat{Z}_Q = (\hat{Q} \otimes \hat{Q}) - \frac{\text{vec}(M)}{T-K-1} \hat{Q}' \hat{Q}$, and \hat{U}_ϵ is defined in Theorem 2.³⁹

Our test statistic \mathcal{S}^* has power when the squared pricing errors, e_i^2 , are greater than zero for the majority of the test assets.⁴⁰ Moreover, it is straightforward to show that the distribution of our test under the null hypothesis is invariant to asset repackaging.

3.2 Estimation under potential model misspecification

If the null hypothesis of correct model specification, for the beta-pricing model under consideration, is rejected, one has two options. The first possibility is to conclude that the model is wrong, and to modify the model accordingly before proceeding with risk premia estimation. If one still wishes to conduct inference on risk premia with the same beta-pricing model, then the standard errors of the risk premia estimates need to be made robust against potential model misspecification. This is the approach we propose in this section. Suppose that assumption 1 is violated and assume that

$$E[R_t] = 1_N \tilde{\gamma}_0 + B \tilde{\gamma}_1 + e, \quad (64)$$

where, following Shanken and Zhou (2007), the pseudo-true values $\tilde{\Gamma} = [\tilde{\gamma}_0, \tilde{\gamma}'_1]'$ are given by

$$\tilde{\Gamma} = \underset{C}{\text{argmin}} \frac{(E[R_t] - XC)'(E[R_t] - XC)}{N}, \text{ for an arbitrary } (K+1)\text{-vector } C. \quad (65)$$

When the model is correctly specified, $\tilde{\Gamma} = \Gamma$, the vector of ex ante risk premia.⁴¹

³⁹ By Theorem 2 and Lemma 6 in the Online Appendix, we have $\hat{Z}'_Q \hat{U}_\epsilon \hat{Z}_Q \rightarrow_p Z'_Q U_\epsilon Z_Q$ as $N \rightarrow \infty$.

⁴⁰ Specifically, our test will reject H_0 when the pricing errors e_i are zero for only a number N_0 of assets, such that $N_0/N \rightarrow 0$ as $N \rightarrow \infty$. This condition allows N_0 to diverge, although not too fast. A formal power analysis can be developed by using the notion of local alternatives as in Gagliardini, Ossola, and Scaillet (2016). In the Online Appendix, we present a Monte Carlo simulation experiment calibrated to real data that shows the desirable size and power properties of our test.

⁴¹ Under the i.i.d. normality assumption and Equation (64), Shanken and Zhou (2007) establish the asymptotic distribution of the OLS and GLS CSR estimators of $\tilde{\Gamma}$ as $T \rightarrow \infty$. (See also Hou and Kimmel 2006.) Kan, Robotti, and Shanken (2013) generalize their results to the case of temporally dependent and nonnormal test asset returns and factors, and derive the large- T distribution of the OLS and GLS CSR R^2 .

We now introduce an additional assumption that governs the behavior of the population pricing errors in terms of cross-sectional moments with the returns' innovations.

Assumption 7 As $N \rightarrow \infty$, we have

(i)

$$\frac{1}{N} \sum_{i=1}^N \epsilon_i e_i \rightarrow_p 0. \quad (66)$$

(ii)

$$\frac{1}{N} \sum_{i=1}^N \epsilon_i \epsilon_i' e_i^2 \rightarrow_p \tau_\Omega I_T. \quad (67)$$

(iii)

$$\frac{1}{N} \sum_{i=1}^N \epsilon_i \epsilon_i' e_i \rightarrow_p \tau_\Phi I_T. \quad (68)$$

(iv)

$$\sum_{i,j=1}^N |\sigma_{ij} e_i e_j| \mathbb{1}_{\{i \neq j\}} = o(N), \quad (69)$$

for some constants $\tau_\Omega = \text{plim} \frac{1}{N} \sum_{i=1}^N \epsilon_{it}^2 e_i^2$ and $\tau_\Phi = \text{plim} \frac{1}{N} \sum_{i=1}^N \epsilon_{it}^2 e_i$.

Assumption 7(i) implies that the ϵ_{it} and the pricing errors are cross-sectionally uncorrelated, although, by assumptions 7(ii) and 7(iii), they could be cross-sectionally dependent in terms of second moments of the ϵ_{it} . Assumption 7(iv) implies that the pricing errors are not altering the degree of cross-sectional dependence of the ϵ_{it} .

Let $\tilde{\Gamma}^P = [\tilde{\gamma}_0, \tilde{\gamma}_1^P]' = \tilde{\Gamma} + \bar{f} - E[f_t]$. The following theorem extends Theorems 1 and 2 to the case of globally misspecified beta-pricing models.

Theorem 5 As $N \rightarrow \infty$, we have

(i) Under assumptions 2–5, assumption 7, and Equation (64),

$$\hat{\Gamma}^* - \tilde{\Gamma}^P = O_p \left(\frac{1}{\sqrt{N}} \right). \quad (70)$$

(ii) Under assumptions 2–7 and Equation (64),

$$\sqrt{N} \left(\hat{\Gamma}^* - \tilde{\Gamma}^P \right) \rightarrow_d \mathcal{N} \left(0_{K+1}, V + \Sigma_X^{-1} (W + \Omega + \Phi + \Phi') \Sigma_X^{-1} \right), \quad (71)$$

where V and W are defined in Theorem 1 by replacing γ_1^P with $\tilde{\gamma}_1^P$,

$$\Omega = \begin{bmatrix} 0 & 0'_K \\ 0_K & \tau_\Omega \mathcal{P}' \mathcal{P} \end{bmatrix} \quad \text{and} \quad \Phi = \begin{bmatrix} 0 & \tau_\Phi Q' \mathcal{P} \\ 0_K & \tau_\Phi (Q' \otimes \mu_\beta) \mathcal{P} \end{bmatrix}. \quad (72)$$

(iii) Under assumptions 2–5, assumption 7, Equation (64), and $\kappa_4 = 0$,

$$\hat{V} + \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \left(\hat{W} + \hat{\Omega} + \hat{\Phi} + \hat{\Phi}' \right) \left(\hat{\Sigma}_X - \hat{\Lambda} \right)^{-1} \rightarrow_p V + \Sigma_X^{-1} (W + \Omega + \Phi + \Phi') \Sigma_X^{-1}, \quad (73)$$

where \hat{V} and \hat{W} are defined in Theorem 2,

$$\hat{\Omega} = \begin{bmatrix} 0 & 0'_K \\ 0_K & \hat{\tau}_\Omega \mathcal{P}' \mathcal{P} \end{bmatrix} \quad \text{and} \quad \hat{\Phi} = \begin{bmatrix} 0 & \hat{\tau}_\Phi Q' \mathcal{P} \\ 0_K & \hat{\tau}_\Phi \left(Q' \otimes \frac{\hat{B}'_{1N}}{N} \right) \mathcal{P} \end{bmatrix}, \quad (74)$$

and $\hat{\tau}_\Phi$ and $\hat{\tau}_\Omega$ are defined in Lemmas 8 and 9 in the Online Appendix, respectively.

Proof: See the Online Appendix.

Similar to the expressions in Shanken and Zhou (2007) and Kan, Robotti, and Shanken (2013), the asymptotic covariance of $\hat{\Gamma}^*$ contains three additional terms, Ω , Φ , and Φ' . The contribution of the pricing errors to the overall asymptotic covariance increases when the variability of the residuals ϵ_{it} increases or, alternatively, when the variability of the pricing errors e_i increases, leading to a larger τ_Ω .

Notice that under model misspecification $\tilde{\Gamma}$ changes with N and, as a consequence, one can define the limit risk premia $\tilde{\Gamma}_\infty = \lim_{N \rightarrow \infty} \tilde{\Gamma}$. Theorem 3 of Ingersoll (1984) provides the conditions for the existence and the uniqueness of $\tilde{\Gamma}_\infty$.⁴² It follows that, by Theorem 5, $\hat{\Gamma}^*$ also converges to $\tilde{\Gamma}_\infty^P = \left[\tilde{\gamma}_{0,\infty}^P, \tilde{\gamma}_{1,\infty}^{P'} \right]' = \tilde{\Gamma}_\infty + \bar{f} - E[f_t]$. Moreover, if $\tilde{\Gamma} - \tilde{\Gamma}_\infty$ is $o\left(1/\sqrt{N}\right)$, then the asymptotic distribution of $\hat{\Gamma}^*$ around $\tilde{\Gamma}_\infty^P$ is the same as the one in Equation (71).⁴³ Interestingly, even under

⁴²In particular, asymptotic no-arbitrage (see Ingersoll 1984, Equation (7)), our assumption 2, and boundedness of the maximum eigenvalue of Σ imply Ingersoll's result.

⁴³It can be shown that (deterministic) convergence of $\tilde{\Gamma}$ to $\tilde{\Gamma}_\infty$ occurs at most at rate $O\left(1/\sqrt{\sum_{i=1}^N \beta'_i \beta_i}\right)$, which equals $O\left(1/\sqrt{N}\right)$ by assumption 2, although any faster rate is allowed for in principle. Notice that if $\tilde{\Gamma} - \tilde{\Gamma}_\infty$ is exactly $O\left(1/\sqrt{N}\right)$, then we need to modify our sampling scheme and select an arbitrary, slightly smaller, set of assets n such that $n/N \rightarrow 0$ as N diverges. When evaluating $\hat{\Gamma}^*$ using these n assets, then the slower $O(\sqrt{n})$ rate of convergence to $\tilde{\Gamma}_\infty^P$ is obtained.

model misspecification, there is no loss of speed of convergence. This differs from Gagliardini, Ossola, and Scaillet (2016), who obtain a slower rate of convergence, $O(\sqrt{N})$ instead of $O(\sqrt{NT})$, of their estimator to the true ex ante risk premia, $\tilde{\Gamma}_\infty$, when the model is misspecified.

3.3 Misspecification due to priced characteristics

We follow Section 3.3 of Shanken (1992) and allow for assumption 1 to be potentially violated because the cross-section of expected returns now satisfies

$$E[R_{it}] = \gamma_0 + \gamma_1' \beta_i + \delta' c_i, \quad (75)$$

where c_i denotes a K_c -vector of time-invariant firm characteristics and δ denotes the corresponding vector of characteristic premia. Our theory requires characteristics and loadings to be sufficiently heterogenous across assets although we allow them to be (almost) arbitrarily cross-sectionally correlated.⁴⁴ Since characteristics exhibit only modest changes over short time windows, Equation (75) would be a good approximation to the true data-generating process also in a time-varying setting with a small T .⁴⁵

Imposing Equation (75), averaging Equation (2) over time, and replacing X with \hat{X} , we obtain

$$\bar{R} = \hat{X}\Gamma^P + C\delta + \eta^P, \quad (76)$$

where $C = [c_1, \dots, c_N]'$ and $\eta^P = (\bar{\epsilon} - (\hat{X} - X)\Gamma^P)$. The estimates of Γ^P and δ are given by

$$\begin{bmatrix} \hat{\Gamma}^* \\ \hat{\delta}^* \end{bmatrix} = \begin{bmatrix} \hat{X}'\hat{X} - N\hat{\Lambda} & \hat{X}'C \\ C'\hat{X} & C'C \end{bmatrix}^{-1} \begin{bmatrix} \hat{X}'\bar{R} \\ C'\bar{R} \end{bmatrix}, \quad (77)$$

where $\hat{\Lambda}$ is the bias adjustment from Theorem 1. In line with the discussion around Theorem 3, $\hat{\Gamma}^*$ and $\hat{\delta}^*$ will also estimate (consistently) the local averages of the risk and characteristic premia if these are allowed to be time-varying.

In this setting with characteristics, we need to make the following additional assumption. Let $z_i = \epsilon_i \otimes c_i$ and $\Sigma_{zz,ij} = \text{Cov}(z_i, z_j') = \sigma_{ij} [I_T \otimes c_i c_j']$.

⁴⁴The case for (linear or nonlinear) dependence, whereby $\beta_i = \beta(c_i)$, has been forcefully made by both the empirical (see Connor, Hagmann, and Linton 2012; Chordia, Goyal, and Shanken 2015; and Kelly, Pruitt, and Su forthcoming, among others) and theoretical literature (see the survey in Kogan and Papanikolaou 2013) in order to resolve the debate on systematic risk versus characteristic-based stories of expected returns that was spurred from the influential empirical findings of Daniel and Titman (1997).

⁴⁵Chordia, Goyal, and Shanken (2015) highlight the challenges that arise when estimating time-varying characteristic premia and propose a bootstrap procedure to perform correct inference in this setting.

Assumption 8 As $N \rightarrow \infty$,

(i)

$$\hat{\mu}_C = \frac{C'1_N}{N} \rightarrow_p \mu_C = [\mu_{c1}, \dots, \mu_{cK}]', \text{ a finite } K_c\text{-vector,} \quad (78)$$

$$\hat{\Sigma}_{CC} = \frac{C'C}{N} \rightarrow_p \Sigma_{CC}, \text{ a finite positive-definite } (K_c \times K_c) \text{ matrix,} \quad (79)$$

$$\hat{\Sigma}_{CB} = \frac{C'B}{N} \rightarrow_p \Sigma_{CB}, \text{ a finite } (K_c \times K) \text{ matrix,} \quad (80)$$

with positive-definite matrices

$$\Sigma_{CC} - \mu_C \mu_C' \text{ and } \begin{bmatrix} \Sigma_{CC} & \Sigma_{CB} \\ \Sigma_{CB}' & \Sigma_{\beta} \end{bmatrix} - \begin{bmatrix} \mu_C \\ \mu_{\beta} \end{bmatrix} \begin{bmatrix} \mu_C \\ \mu_{\beta} \end{bmatrix}'. \quad (81)$$

(ii)

$$\frac{C'\epsilon'}{N} \rightarrow_p 0_{(K_c \times T)}. \quad (82)$$

(iii)

$$\frac{1}{N} \sum_{i=1}^N \Sigma_{zz,ii} \rightarrow \sigma^2(I_T \otimes \Sigma_{CC}) \text{ and } \sum_{i,j=1}^N \Sigma_{zz,ij} 1_{\{i \neq j\}} = o(N). \quad (83)$$

(iv)

$$\frac{1}{\sqrt{N}} \sum_{i=1}^N z_i \rightarrow_d \mathcal{N}(0_{K_c T}, \sigma^2(I_T \otimes \Sigma_{CC})). \quad (84)$$

Since $\begin{bmatrix} \Sigma_{CC} & \Sigma_{CB} \\ \Sigma_{CB}' & \Sigma_{\beta} \end{bmatrix} - \begin{bmatrix} \mu_C \\ \mu_{\beta} \end{bmatrix} \begin{bmatrix} \mu_C \\ \mu_{\beta} \end{bmatrix}'$ in assumption 8(i) is positive-definite, then $\begin{bmatrix} \Sigma_{CC} & \Sigma_{CB} \\ \Sigma_{CB}' & \Sigma_{\beta} \end{bmatrix}$ is also positive-definite, and this implies that the β_i and the c_i cannot be proportional.

In the next two theorems, we characterize the asymptotic properties of the estimators $\hat{\Gamma}^*$ and $\hat{\delta}^*$.

Theorem 6 As $N \rightarrow \infty$, we have

(i) Under assumptions 2–5 and 8, and Equation (75),

$$\hat{\Gamma}^* - \Gamma^P = O_p\left(\frac{1}{\sqrt{N}}\right), \hat{\delta}^* - \delta = O_p\left(\frac{1}{\sqrt{N}}\right). \quad (85)$$

(ii) Under assumptions 2–6 and 8, and Equation (75),

$$\sqrt{N} \begin{bmatrix} \hat{\Gamma}^* - \Gamma^P \\ \hat{\delta}^* - \delta \end{bmatrix} \rightarrow_d \mathcal{N}(0_{K+K_c+1}, \sigma^2(Q'Q)L^{-1} + L^{-1}OL^{-1}), \quad (86)$$

with

$$L = \begin{bmatrix} \Sigma_X & \begin{bmatrix} \mu'_C \\ \Sigma'_{CB} \end{bmatrix} \\ \begin{bmatrix} \mu_C \\ \Sigma_{CB} \end{bmatrix} & \Sigma_{CC} \end{bmatrix}, \quad O = \begin{bmatrix} \begin{bmatrix} 0 & 0'_K \\ 0_K & Z'U_\epsilon Z \end{bmatrix} & 0_{(K+1) \times K_c} \\ 0_{K_c \times (K+1)} & 0_{K_c \times K_c} \end{bmatrix}, \quad (87)$$

where Q , Z , and U_ϵ are defined in Theorem 1.

Proof: See the Online Appendix.

A consistent estimator of the asymptotic covariance matrix of $\hat{\Gamma}^*$ and $\hat{\delta}^*$ is provided in the next theorem.⁴⁶

Theorem 7 Under assumptions 2–5 and 8, Equation (75), and the identification condition $\kappa_4 = 0$, as $N \rightarrow \infty$, we have

$$\hat{\sigma}^2(\hat{Q}'\hat{Q})\hat{L}^{-1} + \hat{L}^{-1}\hat{O}\hat{L}^{-1} \rightarrow_p \sigma^2(Q'Q)L^{-1} + L^{-1}OL^{-1}, \quad (88)$$

with

$$\hat{L} = \begin{bmatrix} \hat{\Sigma}_X - \hat{\Lambda} & \begin{bmatrix} \hat{\mu}'_C \\ \hat{\Sigma}'_{CB} \end{bmatrix} \\ \begin{bmatrix} \hat{\mu}_C \\ \hat{\Sigma}_{CB} \end{bmatrix} & \hat{\Sigma}_{CC} \end{bmatrix}, \quad \hat{O} = \begin{bmatrix} \begin{bmatrix} 0 & 0'_K \\ 0_K & \hat{Z}'\hat{U}_\epsilon\hat{Z} \end{bmatrix} & 0_{(K+1) \times K_c} \\ 0_{K_c \times (K+1)} & 0_{K_c \times K_c} \end{bmatrix}, \quad (89)$$

where $\hat{\sigma}^2$ is defined in Equation (11), and \hat{Q} , \hat{Z} , and \hat{U}_ϵ are defined in Theorem 2.

4. Empirical Analysis

In this section, we show empirically that the results obtained with our fixed- T and large- N methodology can differ substantially from the results obtained with traditional large- T and fixed- N methods. Using a large number of individual equity returns from CRSP, we estimate and test FF5 and an extension of this model that includes the nontraded liquidity factor (LIQ) of Pástor and Stambaugh (2003).⁴⁷ The empirical success of FF5 in explaining the cross-sectional variation in expected equity returns is what motivates our interest in this model. In the second part of this section, we analyze the extent to which firm characteristics contribute to explaining the cross-section of expected equity returns. The risk and characteristic premia estimators, their confidence intervals, and the various test statistics employed are based on our theoretical analysis in Sections 2 and 3.

⁴⁶The proof of Theorem 7 follows the same steps of the proof of Theorem 2 and is therefore omitted.

⁴⁷The Online Appendix reports further empirical results for FF5, as well as results for CAPM and FF3.

4.1 Data

The monthly data on the traded factors of FF5 is available from Kenneth French’s website, and the nontraded liquidity factor of Pástor and Stambaugh (2003) is taken from Ľuboš Pástor’s website.⁴⁸ As for the test assets, we download monthly stock returns (from January 1966 to December 2013) from CRSP and apply two filters in the selection of stocks. First, we require that a stock has a Standard Industry Classification (SIC) code. (We adopt the 49 industry classifications listed on Kenneth French’s website.) Second, we keep a stock in our sample only for the months in which its price is at least three dollars. The resulting data set consists of 3,435 individual stocks. We perform the empirical analysis using balanced panels over fixed-time windows of three and 10 years (that is, $T = 36$ and 120), respectively. We obtain time series of estimated risk premia and test statistics by shifting the time window month by month over the 1966–2013 period. After filtering the data, we obtain an average number (over the overlapping time windows) of approximately 2,800 stocks when $T = 36$ and 1,200 stocks when $T = 120$.

4.2 Specification testing

For the analysis with traded factors only, we report the p -values of our specification test, \mathcal{S}^* , as well as the p -values of two alternative tests, the Gibbons, Ross, and Shanken (1989) (GRS) and Gungor and Luger (2016) (GL) tests. It should be noted that GRS requires N to be fixed, while the Gungor and Luger (2016) test is valid for any N and T . All three tests are tests of the same null hypothesis; that is, $H_0 : e_i = 0$, for every $i = 1, 2, \dots$.

4.2.1 \mathcal{S}^* test

We first assess the performance of FF5 using \mathcal{S}^* .

Figure 1 about here

The black line in Figure 1 denotes the time series of p -values associated with our test statistic

⁴⁸The five traded factors of FF5 are the market excess return (MKT), the return difference between portfolios of stocks with small and large market capitalizations (SMB), the return difference between portfolios of stocks with high and low book-to-market ratios (HML), the average return on two robust operating profitability portfolios minus the average return on two weak operating profitability portfolios (RMW), and the average return on two conservative investment portfolios minus the average return on two aggressive investment portfolios (CMA).

\mathcal{S}^* for time windows of three years (top panel) and 10 years (bottom panel), respectively. When the black line is below the 5% significance level (dotted red line), we reject FF5. Figure 1 shows that based on our test, we reject the validity of FF5 about 60% of the time when $T = 36$. As expected, the rejection of FF5 happens more frequently when we increase the time window from $T = 36$ to $T = 120$. The rejection of FF5 occurs in about 95% of the cases when the latter scenario is considered. Given the availability of a time series of p -values, one could cast the analysis in a multiple testing framework, as suggested by Barras, Scaillet, and Wermers (2010). Applying their methodology to \mathcal{S}^* , we reject the null of correct model specification in 61% and 95% of the cases for $T = 36$ and $T = 120$, respectively. In Figure 2, we perform the same analysis for the liquidity-augmented FF5.

Figure 2 about here

This variant of FF5 turns out to be strongly rejected, even when $T = 36$. The rejection frequencies are approximately equal to 82% and 92% for $T = 36$ and 120, respectively. Overall, the frequent and strong rejections of FF5 justify our use of confidence intervals that are robust to model misspecification in the subsequent analysis.

4.2.2 GRS and GL tests

Figure 3 reports the GRS p -values (blue line) as well as the GL p -values (green line).

Figure 3 about here

Unlike ours, these two tests are only applicable to beta-pricing models with traded factors. As a consequence, we consider only FF5 here. Since GRS is a GLS-based test, effectively, it is implementable only when N is substantially smaller than T . Therefore, we construct 25 equally weighted portfolio returns from our individual stock returns and analyze the performance of these two tests, using this smaller asset set.⁴⁹ Differently from our large- N test, we are much less likely to reject FF5 based on the GRS test. When considering time windows of $T = 36$, the average rejection rate for FF5 is only about 30%. In addition, FF5 is rejected almost always when $T = 120$. We obtain

⁴⁹The results in Figure 3 are obtained by randomly assigning the various stocks to 25 portfolios. For instance, when $T = 36$, each of the 25 portfolios contains approximately 110 randomly selected stocks. We also experimented with 25 portfolios formed on CAPM betas. The results of the analysis are qualitatively similar to those in Figure 3.

similar results when using the GL test, although it is harder to quantify the rejection rates in this case because the GL test often leads to an inconclusive outcome. Based on the GL test, FF5 is not rejected in about 70% of the cases when $T = 36$, but the test is inconclusive about 29% of the time. Moreover, FF5 is not rejected in only about 18% of the cases when $T = 120$, but the test is inconclusive about 76% of the time. The main message here is that using our test can lead to qualitatively different conclusions relative to existing methods.

4.3 Risk premia estimates

Since our test, \mathcal{S}^* , points to serious misspecification of the risk-return relation, in this section we mainly perform parameter testing by means of standard errors that are robust to model misspecification. Specifically, we use the large- N standard errors derived in Theorem 5. To highlight the differences between our approach and standard large- T methods, we also consider the OLS CSR estimator and the corresponding large- T standard errors from Theorem 1(ii) in Shanken (1992). In the Online Appendix, we report in all figures the rolling sample mean of the traded-factor returns, which is a valid risk premium estimator when T is large.⁵⁰ We briefly comment on the sample mean results in the following analysis.

4.3.1 FF5

Based on a time window of three years, the top panel of Figure 4 presents the rolling-window estimates of the risk premium on MKT and the corresponding 95% confidence intervals. (The results for the other four factors are in the Online Appendix.)

Figure 4 about here

In the figure, the bold black line and the dotted red line refer to the Shanken (1992) and OLS CSR estimators, respectively. The gray band represents the large- N 95% confidence intervals that are robust to model misspecification, whereas the striped orange band is for the large- T confidence intervals. Noticeably, the large- T confidence intervals include the zero value in about 60% of the cases. In contrast, our large- N confidence intervals include the zero value only about 30% of the

⁵⁰Similarly, when considering nontraded factors such as liquidity, we report the rolling sample mean of the corresponding mimicking portfolio return. (See note 33 above.)

time. Not surprisingly, the bottom panel of Figure 4 ($T = 120$ cases) shows that the risk premia estimates are smoother than in the $T = 36$ scenario. However, the large- T confidence intervals are still larger than the corresponding large- N confidence intervals, and they indicate that the OLS CSR and the Shanken (1992) estimates are statistically significant 30% and 80% of the time, respectively. The large- N estimates appear to be systematically larger than the corresponding large- T estimates for most dates, especially for the longer time window. This is the result of the systematic (negative) bias that affects the OLS CSR estimator when N is large. As emphasized in the Online Appendix, the relationship between the large- N and the rolling sample mean estimates (the latter are based on windows of $T = 36$ and $T = 120$ monthly data, respectively) is less stable. The two sets of estimates exhibit a correlation of about 0.5 when $T = 36$ and 0.7 when $T = 120$. Figure 4 shows that the large- T approach supports the hypothesis of constant risk premia, whereas our large- N results point toward a significant time variation in risk premia. Therefore, it seems plausible to interpret $\hat{\Gamma}^*$ as the estimator of the local average, over T periods, of the (time-varying) risk premia, $\bar{\Gamma}^P$, as explained in Section 2.2.

The top panel of Figure 5 reports the Shanken (1992) large- N estimates, expressed in terms of a single line (black line) and in terms of local averages (horizontal bars of length $T = 36$, blue lines), with the corresponding 95% confidence intervals for these local averages based on the large- N standard errors of Theorem 5 (gray band).

Figure 5 about here

The local average estimates appear to be significantly different from each other in most cases, which is a clear symptom of time variation in risk premia. In the Online Appendix, we also report the rolling sample mean (over fixed windows of six months of daily data) of the market excess return and the corresponding 95% confidence interval. As our results in the Online Appendix indicate, although the latter is a suitable nonparametric estimator of the time-varying risk premium, it requires a large number of observations (over a short time window) to produce sufficiently narrow confidence intervals. The correlation between the Shanken (1992) large- N estimator and the six-month rolling sample mean based on daily data is positive but small (the sample correlation coefficient is 0.14). In addition, differently from the Shanken (1992) large- N estimator, the six-month rolling sample mean based on daily data appears to be very noisy.

Given the pronounced time variation in risk premia, the bottom panel of Figure 5 reports our novel estimator $\hat{\gamma}_{1,t-1}^*$ (black line), formally defined in Equation (50), and the corresponding 95% confidence interval (gray band). Although noisier than $\hat{\gamma}_1^*$, the $\hat{\gamma}_{1,t-1}^*$ estimates are still statistically significant about 50% of the time. As the figure indicates, there is a sharp increase in risk premia volatility in correspondence and in the aftermath of major economic and financial crises and episodes such as the Black Monday of October 1987 and the U.S. savings and loan crisis of the 1980s and 1990s. Our empirical findings on risk premia countercyclicality confirm the results in Gagliardini, Ossola, and Scaillet (2016) and corroborate the predictions of many theoretical models. (See the discussion in Section 4.3 of Gagliardini, Ossola, and Scaillet 2016.)

4.3.2 Liquidity-augmented FF5

As for the liquidity-augmented FF5, Figure 6 presents the estimated liquidity risk premium in a time-invariant setting.

Figure 6 about here

The estimated liquidity risk premia in Figure 6 are positive 55% and 37% of the time for $T = 36$ and $T = 120$, respectively. However, the risk premia estimates are statistically significant at the 5% level only in the 21% and 32% of the cases, for $T = 36$ and $T = 120$, respectively. In the same figure, we also report the OLS CSR estimator. The OLS CSR estimates in this case are not too far from the Shanken (1992) estimates. In contrast, the rolling mimicking portfolio sample means (based on windows of $T = 36$ and $T = 120$ monthly data and reported in the Online Appendix) are now only mildly positively correlated with the $\hat{\Gamma}^*$ estimates. (The correlation coefficients are 0.15 and 0.27 for $T = 36$ and $T = 120$, respectively.)

As in the traded factor case, Figure 7 indicates that the time variation in risk premia is pronounced.

Figure 7 about here

The correlation between the mimicking portfolio six-month rolling sample mean (reported in the Online Appendix) and the Shanken (1992) large- N estimates is about 0.19. Similar to the FF5 case, the large- N estimator seems to exhibit a higher precision. Looking at the bottom panel of

Figure 7, the risk premia countercyclicality emerges again, especially around major economic and financial downturns.

4.3.3 Percentage difference between estimated risk premia

Finally, Table 1 reports the percentage difference (averaged over rolling time windows of size $T = 36$ and $T = 120$, respectively) between the Shanken (1992) estimator, $\hat{\Gamma}^*$, and the OLS CSR estimator, $\hat{\Gamma}$, for the various risk premia in the liquidity-augmented CAPM, FF3, and FF5 models.

Table 1 about here

For MKT in panel A, the percentage difference between estimators is quite large (about 64% when $T = 36$ and 27% when $T = 120$). As for FF3 in panel B, the discrepancy between the two estimators is sizeable for HML, ranging from 31% to 52%, and less pronounced for MKT and SMB. Moreover, for FF5 in panel C, the percentage difference between the two estimators is relatively large for CMA, ranging from 33% to about 43%. Finally, sizable differences between the two estimators exist for LIQ, especially in panel A.

In summary, we often find substantial differences between the results based on our large- N approach and the results based on traditional large- T methods. The difference mainly stems from the smaller standard errors of the Shanken (1992) estimator relative to the OLS CSR estimator and the nontrivial bias correction induced by the Shanken (1992) estimator when N is large. These differences are even more pronounced when comparing the results based on the Shanken (1992) estimator with those based on the rolling sample mean estimator. Finally, the estimated risk premium on the (nontraded) liquidity factor of Pástor and Stambaugh (2003) is often found to be statistically insignificant.

4.4 Characteristics

In this section, for ease of comparison with Chordia, Goyal, and Shanken (2015), we use balanced panel data from January 1980 to December 2015.⁵¹ In the data set we use, an average of 3,071 firms have return data in a particular month. Consistent with Daniel and Titman (1997) and Chordia,

⁵¹We thank Alberto Martín-Utrera for sharing his data with us and refer to DeMiguel et al. (forthcoming) for data details.

Goyal, and Shanken (2015), among others, we focus on five firm characteristics that have often been found to be related to the cross-section of expected returns: book-to-market ratio (B/M), asset growth (ASSGR), operating profitability (OPERPROF), market capitalization (MCAPIT), and six-month momentum (MOM6). As it is common in this literature, we cross-sectionally standardize the characteristics.

In the interest of space, we focus only on the $T = 36$ case. For each time window, we compute the average of the characteristics. In the first pass, we obtain beta estimates for CAPM, FF3, and FF5. We then estimate the ex post risk and characteristic premia using our second-pass CSR estimator in Equation (77). Figure 8 reports the time series of the characteristic premia estimates, $\hat{\delta}^*$, and the corresponding 95% confidence intervals for FF5. (The figures for CAPM and FF3 can be found in the Online Appendix.)

Figure 8 about here

Averaging across the three models, the estimated B/M premium is positive about 59% of the time, but it is only statistically significant at the 5% level in about 3% of the cases. The estimated ASSGR premium is almost always negative (in 81% of the cases) and significantly so about 16% of the time, whereas the estimated OPERPROF premium is positive in about 32% of the cases and statistically significant only about 19% of the time. For MCAPIT, the estimated premium is positive 32% of the time and statistically significant in about 12% of the cases, while the MOM6 estimate is almost always positive (99.6% of the time) and significant in 86% of the cases.

We now analyze the importance of the five characteristics in jointly explaining deviations from correct model specification; that is, we assess whether the expected returns on individual stocks represent a compensation for risk or firm characteristics. We consider two alternative approaches. First, we conduct formal tests of the two hypotheses, $H_0 : \gamma_1^P = 0_K$ and $H_0 : \delta = 0_{K_c}$ using the asymptotic distribution theory in Theorems 6 and 7. The results are in panel A of Table 2. The F -tests indicate that the characteristic premia estimates are statistically significant at any conventional level, with the average F -test (over rolling windows of size $T = 36$) for the null hypothesis $H_0 : \delta = 0_{K_c}$ being equal to 1,278.60, 1,108.41, and 927.04 for CAPM, FF3, and FF5, respectively. In contrast, the average F -test for the null hypothesis $H_0 : \gamma_1^P = 0_K$ equals 12.45, 17.19, and 57.18 for CAPM, FF3, and FF5, respectively, with rejection rates, in the order, of

25.70%, 25.90%, and 37.90%.

Next, panel B of Table 2 presents the cross-sectional variance contribution of betas and characteristics to the overall cross-sectional dispersion in the (sample) average returns, \bar{R}_i . Chordia, Goyal, and Shanken (2015) suggest to consider the ratios of the (cross-sectional) variance of the beta component (betas times the factor risk premia) and of the characteristics component (characteristics times the characteristic premia), with respect to the overall (cross-sectional) variance of average returns. However, since the beta and characteristics components are not orthogonal cross-sectionally, this can lead to a percentage of the cross-sectional variance explained by the betas and by the characteristics that is jointly greater than 100%.⁵² In addition, the estimated pricing errors based on our bias-adjusted estimator are not necessarily orthogonal to the regressors of the CSR, thus complicating the interpretation even further.

We modify the approach of Chordia, Goyal, and Shanken (2015) as follows. From the estimated CSR, we have $\bar{R} = \hat{X}\hat{\Gamma}^* + C\hat{\delta}^* + \hat{\eta}^P$, where $\hat{\eta}^P$ are the sample counterparts of η^P in Equation (76). Consider the orthogonalization of the estimated pricing errors, $\hat{\eta}^P$,

$$\begin{aligned}\bar{R} &= \hat{X}\hat{\Gamma}^* + C\hat{\delta}^* + P_{\hat{Z}}\hat{\eta}^P + (I_N - P_{\hat{Z}})\hat{\eta}^P \\ &\equiv \hat{X}\hat{\Gamma}^* + C\hat{\delta}^* + P_{\hat{Z}}\hat{\eta}^P + \hat{\eta}^{*P},\end{aligned}\tag{90}$$

where $P_{\hat{Z}} = \hat{Z}(\hat{Z}'\hat{Z})^{-1}\hat{Z}'$ with $\hat{Z} = [\hat{X}, C]$, and I_N denotes the identity matrix of order N . By construction, the orthogonalized estimated pricing errors, $\hat{\eta}^{*P} = (I_N - P_{\hat{Z}})\hat{\eta}^P$, satisfy $\hat{Z}'\hat{\eta}^{*P} = 0_{K+K_c+1}$. Setting $P_C = C(C'C)^{-1}C'$, rewrite the estimated CSR as

$$\begin{aligned}\bar{R} &= \left(\hat{X}\hat{\Gamma}^* + P_{\hat{Z}}\hat{\eta}^P\right) + C\hat{\delta}^* + \hat{\eta}^{*P} \\ &= \left[(I_N - P_C)\left(\hat{X}\hat{\Gamma}^* + P_{\hat{Z}}\hat{\eta}^P\right)\right] + \left[P_C\left(\hat{X}\hat{\Gamma}^* + P_{\hat{Z}}\hat{\eta}^P\right) + C\hat{\delta}^*\right] + \hat{\eta}^{*P} \\ &\equiv \bar{R}_{\perp C} + \bar{R}_C + \hat{\eta}^{*P},\end{aligned}\tag{91}$$

where $\bar{R}_{\perp C} \equiv (I_N - P_C)\left(\hat{X}\hat{\Gamma}^* + P_{\hat{Z}}\hat{\eta}^P\right)$ is the component of the average returns that is explained only by the estimated betas, and thus (perfectly) uncorrelated with C in sample, and $\bar{R}_C \equiv P_C\left(\hat{X}\hat{\Gamma}^* + P_{\hat{Z}}\hat{\eta}^P\right) + C\hat{\delta}^*$ is the component of the average returns due to C only. Since $\bar{R}_{\perp C}$ and \bar{R}_C are orthogonal to each other and to $\hat{\eta}^{*P}$, the sample variance of the average returns equals the sum of the sample variances of the beta component, of the characteristics component, and of the

⁵²This problem is acknowledged, although not solved, in Chordia, Goyal, and Shanken (2015).

orthogonalized pricing errors, that is,

$$\begin{aligned}
S_{\bar{R}}^2 &= \frac{\bar{R}'\bar{R}}{N} - \left(\frac{1'_N\bar{R}}{N}\right)^2 &= \left(\frac{\bar{R}'_{\perp C}\bar{R}_{\perp C}}{N} - \left(\frac{1'_N\bar{R}_{\perp C}}{N}\right)^2\right) + \left(\frac{\bar{R}'_C\bar{R}_C}{N} - \left(\frac{1'_N\bar{R}_C}{N}\right)^2\right) + \frac{\hat{\eta}^{*P'}\hat{\eta}^{*P}}{N} \\
&&\equiv S_{\bar{R}_{\perp C}}^2 + S_{\bar{R}_C}^2 + S_{\hat{\eta}^{*P}}^2.
\end{aligned} \tag{92}$$

Panel B of Table 2 reports the average, over rolling windows of size $T = 36$, of the variance ratios $100 \times S_{\bar{R}_C}^2/S_{\bar{R}}^2$ and $100 \times S_{\bar{R}_{\perp C}}^2/S_{\bar{R}}^2$.

Table 2 about here

The results are largely supportive of our findings based on the F -tests; that is, characteristics overwhelmingly dominate the cross-sectional variation in average individual stock returns. Averaging across the three beta-pricing models, the characteristic variance ratio, $100 \times S_{\bar{R}_C}^2/S_{\bar{R}}^2$, is about 76%, whereas the beta variance ratio, $100 \times S_{\bar{R}_{\perp C}}^2/S_{\bar{R}}^2$, is about 2.8%. The rest (about 21.5%) represents the unexplained portion of the average return cross-sectional variance.⁵³ Overall, our empirical findings support the conclusions of Chordia, Goyal, and Shanken (2015), who argue that in characteristic-augmented beta-pricing models with the Fama-French factors, it is mainly the characteristics that contribute to the cross-sectional variation in expected stock returns, regardless of whether the premia are allowed to be time-varying.

5. Conclusion

This paper is concerned with estimation of risk premia and testing of beta-pricing models when the data is available for a large cross-section of securities, N , but only for a fixed number of time periods, T . Since in this context the traditional OLS CSR estimator of the risk premia is asymptotically biased and inconsistent, we provide a new methodology built on the appealing bias-adjusted estimator of the ex post risk premia proposed by Shanken (1992). We establish its consistency and asymptotic normality for the baseline case of correctly specified beta-pricing models with constant risk premia, and then extend our setting to deal with time-varying risk premia. We also explore in detail the case of misspecified beta-pricing models by deriving a new specification test (and its large- N properties) and by showing how to robustify the asymptotic standard errors

⁵³Confidence intervals for these variance ratios could be computed based on our asymptotic results. The details are available upon request.

of the risk premia estimates when the beta-pricing relation is violated. The important case of misspecification due to priced firm characteristics is also considered. Finally, we analyze the case of unbalanced panels.

Compelling reasons for using our methodology arise when T is fairly small (and, in particular, smaller than N),⁵⁴ when considering models with nontraded factors (as in this case the risk premia on the nontraded factors cannot be inferred from the factors' means even when the model is correctly specified), and when interest lies in the time variation in risk premia. In addition, our misspecification-robust standard errors should be used when the researcher rejects the asset pricing specification but is still interested in performing inference on the model's risk premia.

We apply our large- N methodology to empirically investigate the performance of some prominent beta-pricing models using individual stock return data, that is, the monthly returns (from CRSP) on about 3,500 individual stocks for the January 1966–December 2013 period. We consider three empirical models: the CAPM, the three-factor model of Fama and French (1993), and the five-factor model of Fama and French (2015). We also augment these models with the (nontraded) liquidity factor of Pástor and Stambaugh (2003).

Our large- N test often rejects the Fama and French (2015) model, with and without the liquidity factor, at conventional significance levels even for short time windows of three years. In contrast, when using a suitable aggregation of the same data, in most cases we are unable to reject the Fama and French (2015) model using the traditional large- T methodologies. Similar conclusions hold when testing the validity of the CAPM and the Fama and French (1993) three-factor model, with and without the liquidity factor. The empirical rejection of these models suggests that the misspecification-robust standard errors derived in this paper should be employed when performing inference on risk premia.

Turning to estimation, our results indicate that all the traded-factor risk premia estimates are statistically significant most of the time, even over short time windows of three years. In contrast, the (nontraded) liquidity factor is often not priced. We also provide evidence of significant time variation in risk premia for both traded and nontraded factors. Our overall evidence of pricing is at

⁵⁴Based on numerous Monte Carlo experiments, previous studies have found that the large- T approximations of the CSR estimators are reliable only when five or more decades of data are used. (See Chen and Kan 2004 and Shanken and Zhou 2007, among others.) Therefore, our methodology could be useful also in scenarios where the time-series dimension is relatively large.

odds with the results obtained using the traditional approach based on the large- T Shanken (1992) standard errors.

Finally, allowing for characteristics in the risk-return relation, we find that the book-to-market ratio, asset growth, operating profitability, market capitalization, and six-month momentum explain most of the cross-sectional variation in estimated expected stock returns. Monte Carlo simulations (in the Online Appendix) corroborate our theoretical findings, both in terms of estimation and in terms of testing of the beta-pricing restriction.

Appendix: Explicit Form of U_ϵ

Denote by U_ϵ the $T^2 \times T^2$ matrix

$$U_\epsilon = \begin{bmatrix} U_{11} & \cdots & U_{1t} & \cdots & U_{1T} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ U_{t1} & \cdots & U_{tt} & \cdots & U_{tT} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ U_{T1} & \cdots & U_{Tt} & \cdots & U_{TT} \end{bmatrix}. \quad (\text{A.1})$$

Each block of U_ϵ is a $T \times T$ matrix. The blocks along the main diagonal, denoted by U_{tt} , $t = 1, 2, \dots, T$, are themselves diagonal matrices, with $(\kappa_4 + 2\sigma_4)$ in the (t, t) -th position and σ_4 in the (s, s) position for every $s \neq t$; that is,

$$U_{tt} = \begin{matrix} \downarrow \\ t\text{-th column} \\ \left[\begin{array}{ccccccc} \sigma_4 & \cdots & 0 & \cdots & \cdots & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \sigma_4 & 0 & \cdots & \cdots & 0 \\ 0 & \cdots & 0 & (\kappa_4 + 2\sigma_4) & 0 & \cdots & 0 \\ 0 & \cdots & \cdots & 0 & \sigma_4 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & \sigma_4 \end{array} \right] \end{matrix}. \quad (\text{A.2})$$

The blocks outside the main diagonal, denoted by U_{ts} , $s, t = 1, 2, \dots, T$ with $s \neq t$, are all made of zeros except for the (s, t) -th position that contains σ_4 , that is,

$$U_{ts} = \begin{matrix} \downarrow \\ t\text{-th column} \\ \left[\begin{array}{ccccccc} 0 & \cdots & 0 & \cdots & \cdots & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & \cdots & 0 \\ 0 & \cdots & 0 & \sigma_4 & 0 & \cdots & 0 \\ 0 & \cdots & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & 0 \end{array} \right] \end{matrix}. \quad (\text{A.3})$$

Under assumption 5, by Lemma 6 in the Online Appendix, it is easy to show that \hat{U}_ϵ in Theorem 2 is a consistent plug-in estimator of U_ϵ that only depends on $\hat{\sigma}_4$.

References

- Ahn, S. C., A. R. Horenstein, and N. Wang. 2018. Beta matrix and common factors in stock returns. *Journal of Financial and Quantitative Analysis* 53:1417–40.
- Ahn, S. C., M. F. Perez, and C. Gadarowski. 2013. Two-pass estimation of risk premiums with multicollinear and near-invariant betas. *Journal of Empirical Finance* 20:1–17.
- Ang, A., and D. Kristensen. 2012. Testing conditional factor models. *Journal of Financial Economics* 106:132–56.
- Ang, A., J. Liu, and K. Schwarz. Forthcoming. Using stocks or portfolios in tests of factor models. *Journal of Financial and Quantitative Analysis*.
- Bai, J., and G. Zhou. 2015. Fama-MacBeth two-pass regressions: Improving risk premia estimates. *Finance Research Letters* 15:31–40.
- Balduzzi, P., and C. Robotti. 2008. Mimicking portfolios, economic risk premia, and tests of multi-beta models. *Journal of Business & Economic Statistics* 26:354–68.
- Barillas, F., and J. Shanken. 2017. Which Alpha? *Review of Financial Studies* 30:1316–38.
- Barras, L., O. Scaillet, and R. Wermers. 2010. False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance* 65:179–216.
- Berk, J. B. 2002. Sorting out sorts. *Journal of Finance* 55:407–27.
- Black, F., M. C. Jensen, and M. Scholes. 1972. The capital asset pricing model: Some empirical tests. In Michael C. Jensen (Ed.), *Studies in the theory of capital markets*. New York: Praeger.
- Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger. 1989. Empirical tests of the consumption-oriented CAPM. *Journal of Finance* 44:231–62.
- Brennan, M., T. Chordia, and A. Subrahmanyam. 1998. Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of Financial Economics* 49:345–73.

- Bryzgalova, S. 2016. Spurious factors in linear asset pricing models. Working Paper. Stanford University.
- Burnside, C. 2015. Identification and inference in linear stochastic discount factor models with excess returns. *Journal of Financial Econometrics* 14:295–330.
- Chamberlain, G. 1983. Funds, factors, and diversification in arbitrage pricing models. *Econometrica* 51:1305–23.
- Chamberlain, G., and M. Rothschild. 1983. Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica* 51:1281–1304.
- Chan, L. K. C., J. Karceski, and J. Lakonishok. 1998. The risk and return from factors. *Journal of Financial and Quantitative Analysis* 33:159–88.
- Chen, R., and R. Kan. 2004. Finite sample analysis of two-pass cross-sectional regressions. Working Paper. University of Toronto.
- Chordia, T., A. Goyal, and J. Shanken. 2015. Cross-sectional asset pricing with individual stocks: Betas versus characteristics. Working Paper. Emory University.
- Connor, G., M. Hagmann, and O. Linton. 2012. Efficient semiparametric estimation of the Fama-French model and extensions. *Econometrica* 80:713–54.
- Daniel, K., and S. Titman. 1997. Evidence on the characteristics of cross sectional variation in stock returns. *Journal of Finance* 52:1–33.
- DeMiguel, V., A. Martín-Utrera, F. J. Nogales, and R. Uppal. Forthcoming. A transaction-cost perspective on the multitude of firm characteristics. *Review of Financial Studies*.
- Fama, E. F., and K. R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3–56.
- Fama, E. F., and K. R. French. 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116:1–22.
- Fama, E. F., and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81:607–36.

- Ferson, W. E., and C. R. Harvey. 1991. The variation of economic risk premiums. *Journal of Political Economy* 99:385–415.
- Gagliardini, P., E. Ossola, and O. Scaillet. 2016. Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84:985–1046.
- . 2018. A diagnostic criterion for approximate factor structure. Working Paper. Swiss Finance Institute.
- Gibbons, M. R., S. A. Ross, and J. Shanken. 1989. A test of the efficiency of a given portfolio. *Econometrica* 57:1121–52.
- Giglio, S., and D. Xiu. 2017. Inference on risk premia in the presence of omitted factors. Working Paper. National Bureau of Economic Research.
- Gospodinov, N., R. Kan, and C. Robotti. 2014. Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors. *Review of Financial Studies* 27:2139–70.
- . 2017. Spurious inference in reduced-rank asset-pricing models. *Econometrica* 85:1613–28.
- . 2019. Too good to be true? Fallacies in evaluating risk factor models. *Journal of Financial Economics* 132:451–71.
- Greene, W. H. 2003. *Econometric analysis*. Upper Saddle River, NJ: Pearson Education.
- Gungor, S., and R. Luger. 2016. Multivariate tests of mean-variance efficiency and spanning with a large number of assets and time-varying covariances. *Journal of Business & Economic Statistics* 34:161–75.
- Harvey, C. R., Y. Liu, and H. Zhu. 2016. ... and the cross-section of expected returns. *Review of Financial Studies* 29:5–68.
- Hou, K., G. A. Karolyi, and B. C. Kho. 2011. What factors drive global stock returns? *Review of Financial Studies* 24:2527–74.
- Hou, K., and R. Kimmel. 2006. On the estimation of risk premia in linear factor models. Working Paper. Ohio State University.

- Ingersoll, J. E. 1984. Some results in the theory of arbitrage pricing. *Journal of Finance* 39:1021–39.
- Jagannathan, R., G. Skoulakis, and Z. Wang. 2010. The analysis of the cross-section of security returns. In *Handbook of financial econometrics: Applications*, vol. 2 of *Handbooks in finance*, ed. Y. Aït-Sahalia and L. P. Hansen, 73–134. San Diego, CA: Elsevier.
- Jagannathan, R., and Z. Wang. 1998. An asymptotic theory for estimating beta-pricing models using cross-sectional regression. *Journal of Finance* 53:1285–1309.
- Jegadeesh, N., J. Noh, K. Pukthuanthong, R. Roll, and J. L. Wang. Forthcoming. Empirical tests of asset pricing models with individual assets: Resolving the errors-in-variables bias in risk premium estimation. *Journal of Financial Economics*.
- Kan, R., and C. Robotti. 2012. Evaluation of asset pricing models using two-pass cross-sectional regressions. In *Handbook of computational finance*, ed. J. C. Duan, J. C. Gentle, and W. Hardle, 223–51. Berlin: Springer.
- Kan, R., C. Robotti, and J. Shanken. 2013. Pricing model performance and the two-pass cross-sectional regression methodology. *Journal of Finance* 68:2617–49.
- Kan, R., and C. Zhang. 1999a. GMM tests of stochastic discount factor models with useless factors. *Journal of Financial Economics* 54:103–27.
- . 1999b. Two-pass tests of asset pricing models with useless factors. *Journal of Finance* 54:203–35.
- Kelly, B. T., S. Pruitt, and Y. Su. Forthcoming. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*.
- Kim, S., and G. Skoulakis. 2018. Ex-post risk premia estimation and asset pricing tests using large cross sections: The regression-calibration approach. *Journal of Econometrics* 204:159–88.
- Kleibergen, F. 2009. Tests of risk premia in linear factor models. *Journal of Econometrics* 149:149–73.
- Kleibergen, F., and Z. Zhan. 2018. Identification-robust inference on risk premia of mimicking portfolios of non-traded factors. *Journal of Financial Econometrics* 16:155–90.

- . 2019. Robust inference for consumption-based asset pricing. Working Paper. University of Amsterdam.
- Kogan, L., and D. Papanikolaou. 2013. Firm characteristics and stock returns: The role of investment-specific shocks. *Review of Financial Studies* 26:2718–59.
- Kuersteiner, G. M., and I. R. Prucha. 2013. Limit theory for panel data models with cross sectional dependence and sequential exogeneity. *Journal of Econometrics* 174:107–26.
- Lamont, O. A. 2001. Economic tracking portfolios. *Journal of Econometrics* 105:161–84.
- Lewellen, J., S. Nagel, and J. Shanken. 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96:175–94.
- Litzenberger, R. H., and K. Ramaswamy. 1979. The effect of personal taxes and dividends on capital asset prices: Theory and empirical evidence. *Journal of Financial Economics* 7:163–95.
- Pástor, Ľ., and R. F. Stambaugh. 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111:642–85.
- Pesaran, M. H., and T. Yamagata. 2012. Testing CAPM with a large number of assets. Working Paper. Cambridge University.
- Ross, S. A. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13:341–60.
- Shanken, J. 1992. On the estimation of beta-pricing models. *Review of Financial Studies* 5:1–33.
- . 1996. 23 statistical methods in tests of portfolio efficiency: A synthesis. In *Handbook of statistics*, vol. 14, ed. G. Maddala and C. Rao, 693–711. New York: Elsevier.
- Shanken, J., and G. Zhou. 2007. Estimating and testing beta pricing models: Alternative methods and their performance in simulations. *Journal of Financial Economics* 84:40–86.

Table 1
Percentage difference between estimated risk premia

Factor	$T = 36$	$T = 120$
Panel A: CAPM (with liquidity)		
MKT	64.3%	27.2%
LIQ	41.3%	54.2%
Panel B: FF3 (with liquidity)		
MKT	13.9%	7.3%
SMB	14.7%	12.3%
HML	51.6%	31.2%
LIQ	22.9%	46.1%
Panel C: FF5 (with liquidity)		
MKT	15.3%	11.1%
SMB	13.2%	9.7%
HML	14.1%	15.2%
RMW	13.3%	15.2%
CMA	43.3%	33.0%
LIQ	13.9%	38.7%

The table reports the percentage difference between the Shanken (1992) estimator, $\hat{\gamma}_1^*$, and the OLS CSR estimator, $\hat{\gamma}_1$, averaged over rolling windows of size $T = 36$ and $T = 120$, respectively. The three panels refer to the CAPM, Fama and French (1993) three-factor model (FF3), and Fama and French (2015) five-factor model (FF5). Each of these models has been augmented with the nontraded liquidity factor of Pástor and Stambaugh (2003). We use monthly return data on individual stocks from CRSP and factor data from Kenneth French's and Ľuboš Pástor's websites from January 1966 to December 2013.

Table 2
Betas versus characteristics

	CAPM	FF3	FF5
Panel A: F -tests and rejection frequencies			
$H_0 : \gamma_1^P = 0_K$			
F -tests	14.54	17.33	21.14
Rejection frequencies	25.84%	28.72%	29.91%
$H_0 : \delta = 0_{K_c}$			
F -tests	888.27	960.01	927.04
Rejection frequencies	100%	100%	100%
Panel B: Variance ratios			
$100 \times \frac{S_{RC}^2}{S_R^2}$	73.84%	76.36%	76.70%
$100 \times \frac{S_{R \perp C}^2}{S_R^2}$	2.21%	3.11%	3.19%

The top panel of the table reports the F -tests (average over rolling windows of size $T = 36$) for the null hypotheses $H_0 : \gamma_1^P = 0_K$ and $H_0 : \delta = 0_{K_c}$, respectively, and the rejection frequencies at the 95% confidence level (average over rolling windows of size $T = 36$). Each column refers to a different beta-pricing model, that is, the CAPM (first column), the Fama and French (1993) three-factor model (FF3) (second column), and the Fama and French (2015) five-factor model (FF5) (third column). The bottom panel reports the variance ratios $100 \times S_{RC}^2/S_R^2$ and $100 \times S_{R \perp C}^2/S_R^2$ defined in Section 4.4 (average over rolling windows of size $T = 36$). The data is from DeMiguel et al. (forthcoming) and Kenneth French's website (from January 1980 to December 2015).




Figure 1
Specification testing for the Fama and French (2015) (FF5) five-factor model

The figure presents the time series of p -values (black line) of S^* for FF5. Rolling time windows of three (top panel) and 10 years (bottom panel) are used. The dashed dotted red line denotes the 5% significance level of the test. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French's website from January 1966 to December 2013.




Figure 2
Specification testing for the liquidity-augmented Fama and French (2015) (FF5) five-factor model

The figure presents the time series of p -values (black line) of \mathcal{S}^* for the liquidity-augmented FF5 model. Rolling time windows of three (top panel) and 10 years (bottom panel) are used. The dashed dotted red line denotes the 5% significance level of the test. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French's and Ľuboř Pástor's websites from January 1966 to December 2013.




Figure 3
Specification testing for the Fama and French (2015) (FF5) five-factor model using the Gibbons, Ross, and Shanken (1989) and Gungor and Luger (2016) tests

The figure presents the time series of p -values of the GRS (blue line) and GL (green line) tests for FF5. Rolling time windows of three (top panel) and 10 years (bottom panel) are used. The dashed dotted red line denotes the 5% significance level of the tests. The gray bars are for the periods in which the GL test is inconclusive. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French's website from January 1966 to December 2013.

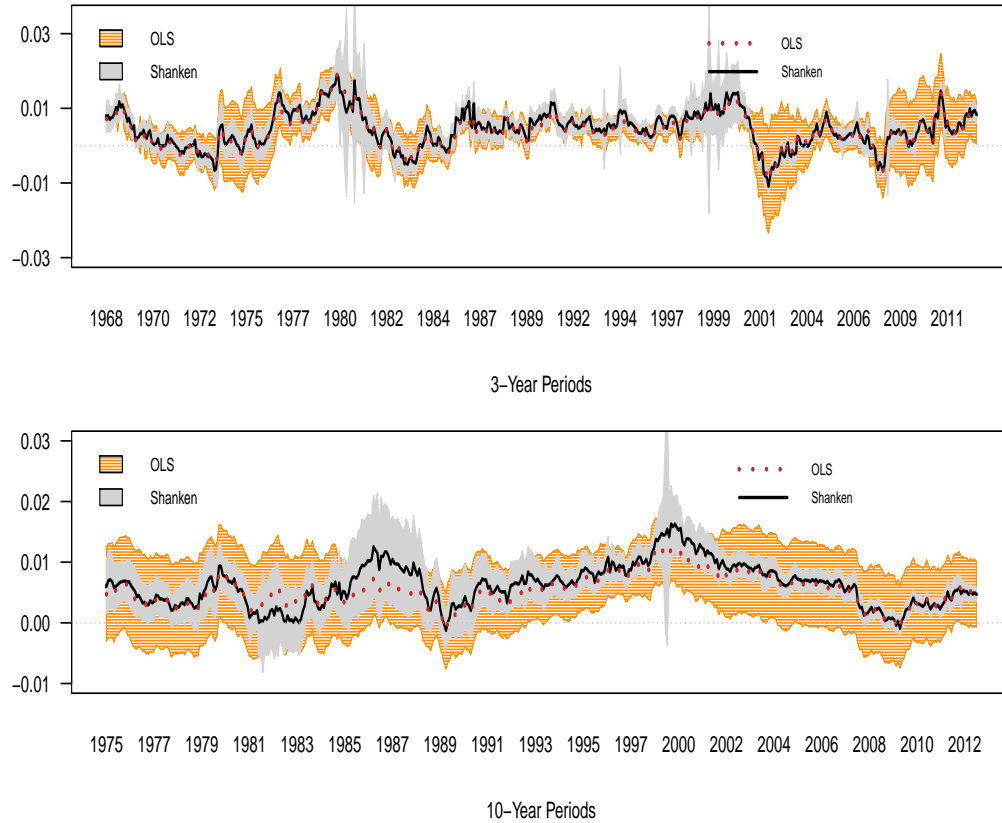


Figure 4
Estimates and confidence intervals for the market risk premium

The figure presents the estimates and the associated confidence intervals for the market risk premium in the Fama and French (2015) five-factor model. The bold black line is for the Shanken (1992) estimator. The corresponding gray band represents the 95% confidence intervals based on the large- N standard errors of Theorem 5. We also report the OLS CSR estimator (dotted red line) and the corresponding 95% confidence interval (striped orange band) based on the traditional large- T standard errors. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French’s website from January 1966 to December 2013.




Figure 5

Estimates and confidence intervals for the time-varying market risk premium

The figure presents the estimates and the associated confidence intervals for the time-varying market risk premium in the Fama and French (2015) five-factor model. The top panel reports the Shanken (1992) large- N estimates, expressed in terms of a single line (black line) and in terms of horizontal bars of length $T = 36$ observations (blue line), with the corresponding 95% confidence intervals based on the large- N standard errors of Theorem 5 (gray band). The bottom panel reports the modified Shanken (1992) estimator (black line) and the corresponding 95% confidence interval (gray band) based on the large- N standard errors of part (ii) of Theorem 3. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French’s website from January 1966 to December 2013. The light gray bands correspond to the NBER recession dates and to various economic and financial crises. They are numbered as follows: [1] 1969:10-1970:11, [2] 1973:11-1975:3, [3] 1977:8-1977:11, [4] 1980:1-1980:7, [5] 1981:7-1982:11, [6] 1986:10-1986:12, [7] 1987:9-1987:11, [8] 1989:9-1989:12, [9] 1990:7-1991:3, [10] 1991:8-1992:12, [11] 1994:7-1994:10, [12] 1997:5-1997:9, [13] 1998:8-1998:10, [14] 2000:2-2000:4, [15] 2001:3-2001:11, [16] 2005:8-2005:11, [17] 2007:12-2009:6, [18] 2010:8-2010:10, [19] 2012:5-2012:7.




Figure 6
Estimates and confidence intervals for the liquidity risk premium

The figure presents the estimates and the associated confidence intervals for the liquidity risk premium in the liquidity-augmented Fama and French (2015) five-factor model. The bold black line is for the Shanken (1992) estimator. The corresponding gray band represents the 95% confidence intervals based on the large- N standard errors of Theorem 5. We also report the OLS CSR estimator (dotted red line) and the corresponding 95% confidence interval (striped orange band) based on the traditional large- T standard errors. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French’s and Ľuboš Pástor’s websites from January 1966 to December 2013.

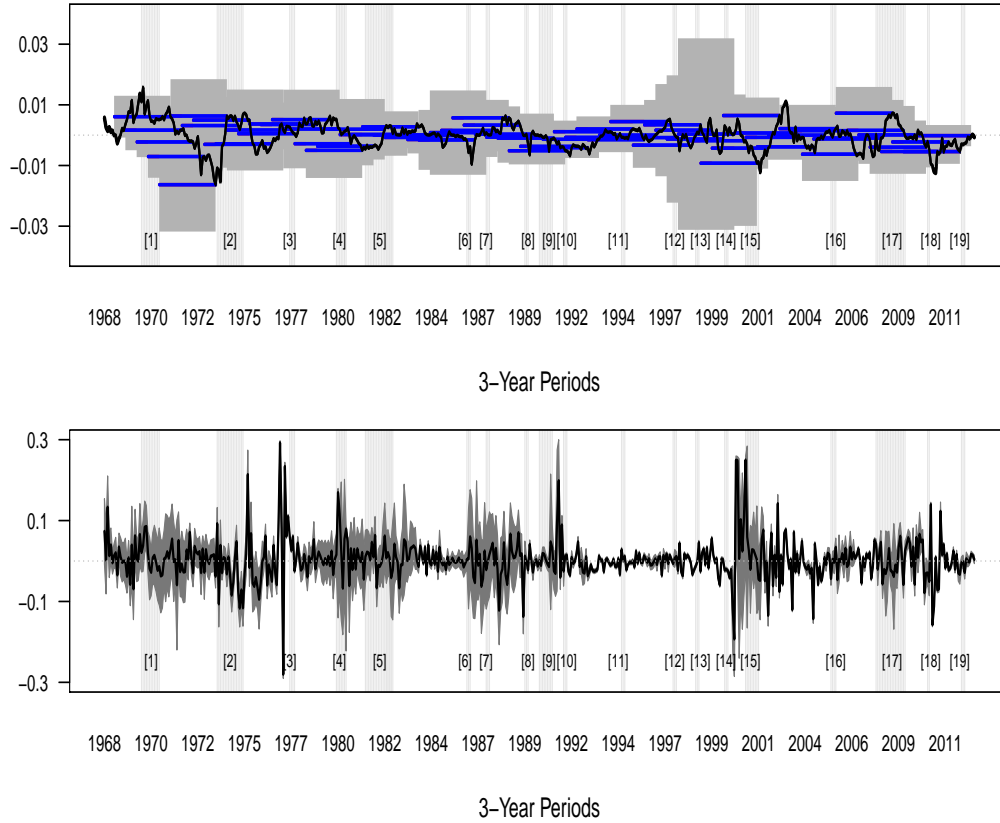


Figure 7

Estimates and confidence intervals for the time-varying liquidity risk premium

The figure presents the estimates and the associated confidence intervals for the time-varying liquidity risk premium in the liquidity-augmented Fama and French (2015) five-factor model. The top panel reports the Shanken (1992) large- N estimates, expressed in terms of a single line (black line) and in terms of horizontal bars of length $T = 36$ observations (blue line), with the corresponding 95% confidence intervals based on the large- N standard errors of Theorem 5 (gray band). The bottom panel reports the modified Shanken (1992) estimator (black line) and the corresponding 95% confidence interval (gray band) based on the large- N standard errors of part (ii) of Theorem 3. We use monthly return data on individual stocks from CRSP and factor data from Kenneth French’s and Ľuboř Pástor’s websites from January 1966 to December 2013. The light gray bands correspond to the NBER recession dates and to various economic and financial crises. They are numbered as follows: [1] 1969:10-1970:11, [2] 1973:11-1975:3, [3] 1977:8-1977:11, [4] 1980:1-1980:7, [5] 1981:7-1982:11, [6] 1986:10-1986:12, [7] 1987:9-1987:11, [8] 1989:9-1989:12, [9] 1990:7-1991:3, [10] 1991:8-1992:12, [11] 1994:7-1994:10, [12] 1997:5-1997:9, [13] 1998:8-1998:10, [14] 2000:2-2000:4, [15] 2001:3-2001:11, [16] 2005:8-2005:11, [17] 2007:12-2009:6, [18] 2010:8-2010:10, [19] 2012:5-2012:7.




Figure 8
Estimates and confidence intervals for the characteristic premia

The figure presents estimates (blue line) of the characteristic premia on the book-to-market ratio (B/M), asset growth (ASSGR), operating profitability (OPERPROF), market capitalization (MCAPIT), and six-month momentum (MOM6), and the associated confidence intervals based on Theorem 7 (light blue band), for the Fama and French (2015) five-factor model. The data is from DeMiguel et al. (forthcoming) and Kenneth French’s website (from January 1980 to December 2015).