

Changes in the Span of Systematic Risk Exposures*

Yuan Liao[†] Viktor Todorov[‡]

February 1, 2024

Abstract

We develop a test for deciding whether the linear spaces spanned by the factor exposures of a large cross-section of assets towards latent systematic risk factors at two distinct points in time are the same. The test uses a panel of asset returns in local windows around the two time points. The asymptotic setup is of joint type: the number of assets and the number of return observations per asset increase asymptotically while the length of both time windows shrinks. We estimate the factor exposures, up to rotation, over the two periods using classical principal component analysis and evaluate their projection discrepancy, which is rotation invariant. This projection discrepancy is then centered with one between factor exposures computed over a partition of the pooled return data into odd and even increments. We derive the limit distribution of the statistic under the null hypothesis and develop an easy-to-implement bootstrap for constructing the critical region of the test. The test is applied to intraday financial data to determine whether the linear span of assets' systematic risk exposures differ during a trading day or after a release of important economic information.

Keywords: asset pricing, high-frequency data, latent factor model, nonparametric test, PCA, systematic risk.

JEL classification: C51, C52, G12.

1 Introduction

Measuring assets' sensitivity towards systematic risk, or betas, plays a central role in asset pricing, see e.g., part II of the book of Cochrane (2009). Early asset pricing models, such as the classical CAPM of Sharpe (1964) and Lintner (1965a,b), are static and imply constant assets' exposure to systematic risk. However asset pricing models can hold only conditionally,

*We would like to thank the CoEditor (James Hamilton) and two anonymous referees for helpful comments and suggestions. We thank also Alexei Onatski and Torben G. Andersen for stimulating discussions and Huidi Lin for excellent research assistance.

[†]Department of Economics, Rutgers University; e-mail: yuan.liao@rutgers.edu.

[‡]Kellogg School of Management, Northwestern University; e-mail: v-todorov@northwestern.edu.

see e.g., Hansen and Richard (1987), and a changing investment opportunity set, see e.g., Merton (1973), can induce time-varying systematic risk exposures of assets. Incorporating this time variation is important for evaluation and testing of asset pricing models, see for example Shanken (1990), Jagannathan and Wang (1996) and Ferson and Harvey (1999). Existing work has either modeled time-varying betas explicitly using macro variables or firm specific quantities, see e.g., Connor et al. (2012), Fan et al. (2016), Gagliardini et al. (2016) and Kelly et al. (2019), among many others or has used a long window of low frequency returns to estimate assets' exposure to observable systematic risk factors following Fama and MacBeth (1973).

Sampling at high frequency allows to improve the measurement of betas. Indeed, quadratic covariation between two processes can be inferred from high-frequency observations of these processes via the so-called realized quadratic covariation, see e.g., Barndorff-Nielsen and Shephard (2004a) and also Mykland and Zhang (2006, 2009). If betas remain constant over short time intervals, then they are simple and known transforms of the quadratic covariation between the assets and the observable factors as well as the quadratic covariation between the factors. Therefore, high-frequency data allows implementing the approach of Fama and MacBeth (1973) but using much shorter time windows while at the same time maintaining high level of precision in the beta estimation. The improved precision in measuring betas can lead to nontrivial gains in asset pricing applications, see e.g., Bollerslev et al. (2016) and Ait-Sahalia et al. (2023), among others.

The key underlying assumption for utilizing the high-frequency data in the nonparametric estimation of betas discussed above is that the latter remain constant over the estimation window. The goal of this paper is to design a nonparametric test that allows us to decide if this is the case without taking a stand on what the systematic risk factors are. More specifically, we design a test for difference in the linear span of assets' exposure to latent systematic risk factors at two distinct points in time. If the assumption for constant assets' exposures to systematic risk factors (our null hypothesis) is violated, then one needs to take into account this time-variation via parametric or nonparametric methods.

Given the fact that variables that have been used to model variation in betas in prior work typically do not change over very short time intervals, one would expect that betas remain constant over days. However, Andersen et al. (2021) find that market betas exhibit pronounced intraday pattern, with monotonically declining cross-sectional beta dispersion during the trading day.¹ The result in Andersen et al. (2021) is for the case when the only observable systematic risk factor is the market portfolio.²

¹This pattern is distinct from the well-known U-shape intraday pattern of assets' volatility, see e.g., Andersen et al. (2023) and references therein.

²Andersen et al. (2023) extend this analysis by including the Fama-French factors and show that the

There are two major drawbacks, however, of using observable factors that our analysis can overcome. First, omitted factors that are correlated with observable ones can yield time-variation in estimated assets’ factor exposures to the observable factors even when the true exposures (the object of economic interest) remain unchanged.³ The reason for this is time variation in the standard omitted variable bias in the beta estimation over different parts of the trading day. This is a serious limitation of using only observable factors in the analysis given the many candidate asset pricing factors proposed in the finance literature (also referred to as factor zoo), see e.g., Cochrane (2011), Harvey et al. (2016) and Feng et al. (2020). Second, if the observable factors are merely noisy proxies of the true ones, this can also generate time variation in the exposure to observable factors when there is none towards the true latent factors. This can happen because the errors-in-variables bias induced by the measurement error in the observable factors can change over time due to time varying volatility in the latter.⁴

For these reasons, in this paper, we develop a general nonparametric test for deciding whether the *linear span* of assets’ exposures to systematic risk at two distinct points in time remains unchanged without assuming knowledge of the systematic risk factors. Our test can discriminate against two types of scenarios. One is a situation in which the number of factors at the two points in time are the same but nevertheless the linear spaces spanned by the factor loadings at the two time points differ. Another scenario that our test can discriminate against is one in which the number of “active” systematic factors at the two points in time differ, even though the factor loadings for the common factors present in the two periods are the same. This can be the case, for example, if some of the systematic risk factors are present only at one of the two time points but remain dormant at the other one.⁵

In our asymptotic setup the number of assets increases and so does the sampling frequency while the length of the two time periods shrinks to zero. Assets are exposed to a fixed number of latent factors and even though exposures can change, they remain constant in arbitrarily

intraday market beta pattern gets attenuated somewhat when accounting for these additional risk factors in the beta estimation. On a theoretical level, Andersen et al. (2021) and Andersen et al. (2023) focus only on changes in the cross-sectional distribution of betas while here we are interested in changing factor exposures for each of the assets in the sample. As a result, the rate of convergence of our statistics is faster and depends on the cross-sectional dimension of the return panel, which is not the case for the test statistics of Andersen et al. (2021) and Andersen et al. (2023).

³Such omitted factors can also impact significantly inference about risk premia of observable factors as shown recently by Giglio and Xiu (2021).

⁴For example, the feasible decomposition of market returns into cash flow shocks and discount-rate shocks (both of which are latent), proposed by Campbell and Vuolteenaho (2004), involves estimation of a model of expected returns. Similarly, issues with the correct measurement of the value factor have been recently discussed by Eisfeldt et al. (2022).

⁵For example, a release of important economic information, e.g., during an FOMC announcement, might lead to certain systematic shocks being active only during that period.

small local neighbourhoods of the two time points. We follow Connor and Korajczyk (1986); Bai and Ng (2002), Stock and Watson (2002) and Bai (2003), and use standard principal component analysis (PCA) to recover the factor exposures, up to a rotation, towards the active factors.⁶ Our test statistic is then formed on the basis of the Frobenius norm of the projection discrepancy between the estimated latent factor exposures over the two time periods. This statistic is invariant to rotation of the factor loadings and its limit should be zero under the null hypothesis that the linear span of the assets' systematic risk exposures does not differ across the two points in time that are being compared.

A methodological innovation of the current paper is the development of a new debiasing method. To remove higher-order biases, we introduce a *bias-mimicking* statistic, which is defined as the projection discrepancy between factor loadings estimated from the odd and even increments, respectively, of the pooled over the two periods return data. Our test statistic is then centered by the bias-mimicking one. The latter mimics the higher-order biases of the original test statistic, and is asymptotically zero both under the null and the alternative hypotheses. Unlike the usual analytical debiasing or the Jackknife, our new bias-mimicking statistic does not require us to know *a priori* either the convergence rate or the explicit form of the higher-order biases.

We derive a Central Limit Theorem (CLT) for our bias-corrected test statistic under the null hypothesis. The rate of convergence depends on both dimensions of the panel of the return observations. It is a product of the number of returns per asset and the square root of the number of assets used in the analysis. The limiting distribution is mixed Gaussian and is determined by the idiosyncratic risk in the asset prices. The limiting variance of the test statistic is random and adapted to the so-called common information set, using the terminology of Andrews (2005), that includes information about economy-wide variables such as the systematic risk factors, their volatility, etc. For feasible implementation of a test of fixed asymptotic size, we propose a simple cross-sectional bootstrap consisting of resampling with replacement of the available stocks.

The test developed in the current paper is related to two testing problems considered in earlier work. First, Ang and Kristensen (2012), Reiß et al. (2015) and Kalnina (2023) develop tests for the constancy of factor loadings of assets over fixed time interval using high-frequency data. On a theoretical level, the asymptotic setup of these tests is for a fixed cross-section. It is not clear whether and how these tests can be extended to a high-dimensional setting, which is the focus of this paper. In addition, the methods of the above-cited work are designed only for observable factors. As we discussed earlier, the observable

⁶Earlier work that applies PCA analysis to high-frequency data with large high cross-sectional dimension include Aït-Sahalia and Xiu (2017) and Pelger (2019, 2020).

factor setup has drawbacks related to possible spurious time variation in betas induced by omitted factors or measurement error in the factors. Second, our paper relates to existing work on testing for structural breaks in factor loadings in classical linear factor models, see e.g., Breitung and Eickmeier (2011), Chen et al. (2014), Corradi and Swanson (2014), Yamamoto and Tanaka (2015), Cheng et al. (2016), Baltagi et al. (2017), Su and Wang (2020) and Bai et al. (2022). The goals of the current paper and this strand of work are different: in our case we are interested in factor exposures at two fixed distinct points in time while in the above-cited papers the interest is in long-span properties of linear factor models. The implication of this difference in goals is that in our setup the rotation matrices of the factor loadings over the two time periods, up to which the factor loadings can be identified, are random and can differ even under the null hypothesis. This is not the case in the long-span tests for structural stability of linear factor models. As a result, one cannot adapt the test statistics of the structural break cited above to our problem. Instead, our statistic is based on the projection discrepancy in the estimated factor loadings over the two periods, which is rotation invariant. Another difference between the above strand of work and our paper is that the limit distribution here is mixed Gaussian and common shocks can impact the limiting variance of our test statistic.

The rest of the paper is organized as follows. We introduce the continuous-time factor model and state the assumptions in Section 2. Section 3 presents the sampling scheme and the discrete-time factor model of asset returns. The test and its asymptotic properties are given in Section 4. Section 5 contains a simulation study and Section 6 an empirical application. Proofs are given in the Supplementary Appendix.

Throughout the paper, we will use $\|A\|_F$ and $\|A\|$ to denote the Frobenius norm and operator norm of a matrix A , respectively. In addition, for a random sequence X_{np} , depending on n and p , and a nonrandom nonzero sequence a_{np} , we write $X_{np} = o_P(a_{np})$ if X_{np}/a_{np} converges to zero in probability, and $X_{np} = O_P(a_{np})$ if $X_{np}/a_{np} = O_P(1)$, when $n, p \rightarrow \infty$.

2 The continuous-time factor model and assumptions

Our interest in this paper is in the behavior of a $p \times 1$ vector of log asset prices, denoted by Y , at two fixed points in time. The vector of prices is defined on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_t)_{t \geq 0})$ and is assumed to obey the following time-varying factor model dynamics:

$$dY_t = \alpha_t dt + \beta_t df_t + dJ_t + d\epsilon_t, \quad (1)$$

where α_t , J_t and ϵ_t are $p \times 1$ vectors, f_t is $K \times 1$ vector of diffusive systematic risk factors, for some positive integer K , and β_t is a $p \times K$ matrix of factor loadings. The diffusive factors evolve according to

$$df_t = \Lambda_t dW_t, \quad (2)$$

where Λ_t is $K \times K$ matrix-valued stochastic process, and W_t is $K \times 1$ standard Brownian motion. The systematic risk factors f_t are latent and our interest in this paper is to design a test for the hypothesis that the linear span of the factor loadings at two distinct points in time, $0 < b < a$, is the same.

The diffusive idiosyncratic risk is given by

$$d\epsilon_{t,j} = \sigma_{t,j} dW_{t,j}, \quad j = 1, \dots, p, \quad (3)$$

where $\epsilon_{t,j}$ is the j -th element of ϵ_t , and where $\{W_{t,j}\}_{j=1,\dots,p}$ is a sequence of independent univariate Brownian motions, which are also independent from W_t that drives the systematic risk factors.

Finally, the process J is a pure-jump process, i.e., it is of the form $J_t = \sum_{s \leq t} \Delta Y_s$, where as usual $\Delta Y_s = Y_s - Y_{s-}$ and $Y_{s-} = \lim_{u \uparrow s} Y_u$. We are not interested in the jumps of the asset prices in this paper, and therefore we will make no assumptions regarding whether they arrive together in the asset prices or not. We note that the analysis of the cross-sectional dependence in J requires different techniques than the ones used here, see e.g., Jacod et al. (2022).

We need several assumptions for deriving our results. For stating these assumptions, we introduce the ‘‘common shocks’’ σ -algebra \mathcal{C} , using the terminology of Andrews (2005), which contains the information about the systematic risk and more generally about any economy-wide random variable. The processes Λ , W and f are all adapted to \mathcal{C} .

Our interest in this paper is in the factor loadings β_b and β_a , where $0 < b < a$ are two specific points in time of interest. For stating the assumptions about them, we partition them in the following way

$$\begin{aligned} \beta_b &= (\beta_b^{(1)}, \beta_b^{(2)}, \beta_b^{(3)}, \mathbf{0}_{p \times K_4}), \\ \beta_a &= (\beta_a^{(1)}, \beta_a^{(2)}, \mathbf{0}_{p \times K_3}, \beta_a^{(4)}), \end{aligned} \quad (4)$$

where $\beta_c^{(j)}$ are of dimensions $p \times K_j$, with $K_1 + K_2 + K_3 + K_4 = K$. In addition, we have $\beta_a^{(1)} = \beta_b^{(1)} H$ for some invertible $K_1 \times K_1$ matrix H .

With this decomposition of the two factor loadings, we allow for: (1) a common component in them up to a rotation, $\beta_b^{(1)}$ and $\beta_a^{(1)}$, (2) different nontrivial components, $\beta_b^{(2)}$ and

$\beta_a^{(2)}$, and (3) components of the factor loadings (or equivalently of the factors) which are present in one of the two periods only, $\beta_b^{(3)}$ and $\beta_a^{(4)}$. We combine the unique components of the factor loadings over the two periods into:

$$\beta_b^r = (\beta_b^{(1)}, \beta_b^{(2)}, \beta_b^{(3)}), \beta_a^r = (\beta_a^{(1)}, \beta_a^{(2)}, \beta_a^{(4)}) \text{ and } \beta_{mix} = (\beta_b^{(1)}, \beta_b^{(2)}, \beta_b^{(3)}, \beta_a^{(2)}, \beta_a^{(4)}), \quad (5)$$

which are of size $p \times (K_1 + K_2 + K_3)$, $p \times (K_1 + K_2 + K_4)$ and $p \times (K + K_2)$, respectively.

We next denote by Λ_c^r the sub-matrix of Λ_c corresponding to columns of β_c^r and by Λ_{mix} the sub-matrix of (Λ_a, Λ_b) corresponding to columns of β_{mix} . Formally, consider partitioning of the identity matrix $I_{K \times K}$ into $I_{K \times K} = (\iota'_1, \iota'_2, \iota'_3, \iota'_4)$, with ι_j being of size $K_j \times K$ for $j = 1, \dots, 4$. With this notation, we set

$$\Lambda_b^r = ((\iota_1 \Lambda_b)', (\iota_2 \Lambda_b)', (\iota_3 \Lambda_b)'), \Lambda_a^r = ((\iota_1 \Lambda_a)', (\iota_2 \Lambda_a)', (\iota_4 \Lambda_a)'), \quad (6)$$

$$\Lambda_{mix} = ((\iota_1 \Lambda_b)', (\iota_2 \Lambda_b)', (\iota_3 \Lambda_b)', (\iota_2 \Lambda_a)', (\iota_4 \Lambda_a)'). \quad (7)$$

Finally, we set

$$\Sigma_{f,c}^r = \Lambda_c^r (\Lambda_c^r)'. \Sigma_{f,mix} = \Lambda_{mix} \Lambda_{mix}', \quad c \in \{a, b\}. \quad (8)$$

Our assumptions are as follows:

A1. *There exists a sequence T_1, T_2, \dots of stopping times increasing to infinity, such that for $s, t \leq T_m$, we have:*

$$\sup_{j \geq 1} \mathbb{E} |\chi_{t,j}|^q + \sup_{j \geq 1} \mathbb{E} |J_{t,j}|^q < \infty, \quad \text{for any } q > 0, \quad (9)$$

$$\sup_{j \geq 1} \mathbb{E} |\chi_{t,j} - \chi_{s,j}|^2 + \sup_{j \geq 1} |\mathbb{E}(\chi_{t,j} - \chi_{s,j})| + \sup_{j \geq 1} \mathbb{P}(J_{t,j} - J_{s,j} \neq 0) \leq C_m |t - s|, \quad (10)$$

$$\mathbb{E} \|\Lambda_t - \Lambda_s\|_F^2 + \|\mathbb{E}(\Lambda_t - \Lambda_s)\|_F \leq C_m |t - s|, \quad (11)$$

for some sequence of positive constants C_m , and where $\chi_{t,j}$ is one of $\alpha_{t,j}$, $\beta_{t,j}$ and $\sigma_{t,j}$. Furthermore, β_t remains constant in local neighborhoods of $t = a, b$, i.e., we have $\beta_t = \beta_c$ for $t \in (c - \varepsilon, c + \varepsilon)$ with $\varepsilon > 0$ being an arbitrary small number and $c = a, b$.

A2. (i) *Conditionally on \mathcal{C} , the processes $W_{t,j}$, $\beta_{t,j}$, $\sigma_{t,j}$ and $J_{t,j}$ are independent across j . Furthermore, $\beta_{c,j}$ and $\sigma_{c,j}$ are identically distributed across j and $\mathbb{P}(\sigma_{c,1} \neq 0 | \mathcal{C}) > 0$ a.s., for $c = a, b$.*

(ii) *Denote $\Sigma_{\beta,c}^r = \text{plim}_{p \rightarrow \infty} \frac{1}{p} \beta_c^r \beta_c^r$ and $\Sigma_{\beta,mix} = \text{plim}_{p \rightarrow \infty} \frac{1}{p} \beta'_{mix} \beta_{mix}$, for $c \in \{a, b\}$. Then, the eigenvalues of the matrices $\Sigma_{\beta,c}^r$, $\Sigma_{f,c}^r$, $\Sigma_{\beta,mix}$ and $\Sigma_{f,mix}$ are all bounded away from zero and infinity, almost surely. In addition, the matrix H that satisfies $\beta_a^{(1)} = \beta_b^{(1)} H$ has eigenvalues*

that are bounded away from both zero and infinity.

(iii) In the case when $\beta_b^r = \beta_a^r H$, for an invertible matrix H , let $M_c = (\Sigma_{\beta,c}^r)^{1/2} \Sigma_{f,c}^r (\Sigma_{\beta,c}^r)^{1/2}$, for $c \in \{a, b\}$, and also let $M_{ab} := (\Sigma_{\beta,a}^r)^{1/2} (\Sigma_{f,a}^r + H \Sigma_{f,b}^r H') (\Sigma_{\beta,a}^r)^{1/2}$. Then, M_a, M_b and M_{ab} are of full rank and have distinct non-zero eigenvalues. That is, if the eigenvalues of any of these matrices are denoted with v_1, \dots, v_k , then there is $c_0 > 0$ so that $|v_i - v_j| > c_0 > 0$ for all $i \neq j$.

A3. We have

$$\max_{i=1, \dots, p} \sigma_{c,i}^2 = O_P(\zeta_p), \quad c = a, b, \quad \text{as } p \rightarrow \infty, \quad (12)$$

for some deterministic sequence ζ_p , with $\lim_{p \rightarrow \infty} \zeta_p$ being either finite or infinite.

We make several comments regarding the above assumptions. First, assumption A1 is a standard integrability and smoothness condition for the various processes entering the dynamics of Y . In particular, the second and third conditions in A1 are satisfied when the processes involved in them are Itô semimartingales. We note also that the moment conditions in A1 are for the stopped processes, so the values of the processes in A1 at a point in time might not have finite moments.

Assumption A2(i) is about cross-sectional dependence. We note that we are interested in such dependence only after conditioning on the common information set \mathcal{C} . Without conditioning on \mathcal{C} , the processes in A2(i) will typically have strong cross-sectional dependence due to their systematic risk exposure. The requirement for identical distribution across j of $\beta_{c,j}$ and $\sigma_{c,j}$, conditional on \mathcal{C} , is for ease of exposition. It can be replaced with a weaker one requiring convergence in probability of cross-sectional averages involving $\beta_{c,j}$ and $\sigma_{c,j}$. The bound from below for the eigenvalues of $\Sigma_{\beta,c}^r$ and $\Sigma_{f,c}^r$ in A2(ii) guarantees that the factors corresponding to β_c^r are active, i.e., have nontrivial variance, and none of the factor loadings are redundant, i.e., are not in the linear span of the other factors.

Finally, assumption A3 is a bound on the cross-sectional maxima of the idiosyncratic variance. This assumption will hold with finite $\lim_{p \rightarrow \infty} \zeta_p$ if, for example, the idiosyncratic variances are adapted to the common shock information set \mathcal{C} . This will be also the case if, after conditioning on \mathcal{C} , $\sigma_{c,j}$ are all bounded for $c = a, b$. If neither of the above two conditions hold, then the rate of growth of ζ_p as the cross-sectional dimension increases can be determined by use of extreme value theory, see e.g., chapter 3 of Embrechts et al. (2013). For example, if the \mathcal{C} -conditional law of $\sigma_{c,i}^2$ is exponential-like, then $\zeta_p = \log(p)$ while if the \mathcal{C} -conditional law of $\sigma_{c,i}^2$ is Gaussian-like, then $\zeta_p = \sqrt{\log(p)}$.

Overall, the assumptions A1-A3 are relatively weak and allow for many relevant features of assets' dynamics that have been documented empirically in earlier work. In particular, we allow for stochastic volatility in asset prices of general form. Idiosyncratic volatility of

individual assets can have a common component (adapted to \mathcal{C}). Shocks to asset prices and to stochastic volatility can be correlated, i.e., the so-called leverage effect is allowed for. Both idiosyncratic jumps that arrive at distinct times in asset prices and systematic jumps that arrive together in asset prices can be present. Systematic jumps do not need to obey a factor model. The factor exposures towards diffusive risk can change over time.

3 Sampling scheme and the discrete factor model

Prices are observed at times $0, \frac{T}{n}, \frac{2T}{n}, \dots, T$, for some fixed T and n asymptotically increasing. We have $0 < b < a < T$. The gap between observations is denoted with $\Delta_n = \frac{T}{n}$ and the increment of a generic process X with

$$\Delta_i^n X = X_{i\Delta_n} - X_{(i-1)\Delta_n}. \quad (13)$$

Since the factor loadings are allowed to vary over time, when forming the test, we will use local blocks of k_n increments around the times b and a , where k_n is a sequence of integers increasing to infinity and satisfying $k_n \Delta_n \rightarrow 0$, so that the local windows used in the estimation around $t = a$ and $t = b$ are both asymptotically shrinking. The indices of the price increments used in the construction of the test are given by

$$i_t^c = \lfloor c/\Delta_n \rfloor - k_n + t, \text{ for } c = a, b \text{ and } t = 1, \dots, k_n. \quad (14)$$

The price increments will be trimmed, following Mancini (2001), in order to eliminate the price jumps. More specifically, let $h(x, a) = (x \vee -a) \wedge a$, for some $x \in \mathbb{R}$ and $a \in \mathbb{R}_+$. We will then denote the trimmed price increments as

$$\bar{Y}_{c,tj} = \frac{h(\Delta_{i_t^c}^n Y_j, \gamma_j \Delta_n^\varpi)}{\sqrt{\Delta_n}}, \quad \bar{Y}_c = (\bar{Y}_{c,tj} : t = 1, \dots, k_n, j = 1, \dots, p), \text{ for } c = a, b, \quad (15)$$

and for some $\gamma_j > 0$ that is uniformly bounded in j and $\varpi \in (0, 1/2)$. We will introduce similar notation for the increments of the latent factors and of the idiosyncratic risk:

$$\bar{f}_{c,t} = \frac{\Delta_{i_t^c}^n f}{\sqrt{\Delta_n}}, \quad \bar{\epsilon}_{c,tj} = \frac{\Delta_{i_t^c}^n \epsilon_{tj}}{\sqrt{\Delta_n}}, \text{ for } c = a, b, t = 1, \dots, k_n, j = 1, \dots, p, \quad (16)$$

and

$$\bar{F}_c = (\bar{f}_{c,t} : t = 1, \dots, k_n), \quad \bar{U}_c = (\bar{\epsilon}_{c,tj} : t = 1, \dots, k_n, j = 1, \dots, p), \text{ for } c = a, b. \quad (17)$$

Finally, we denote the residual component of asset returns as follows

$$\begin{aligned}
r_{c,tj} &= \frac{1}{\sqrt{\Delta_n}} \left(\int_{(i_t^c-1)\Delta_n}^{i_t^c\Delta_n} \alpha_{s,j} ds + \Delta_{i_t^c}^n J_j \right) \mathbf{1}_{\{|\Delta_{i_t^c}^n Y_j| \leq \gamma_j \Delta_n^{\bar{\omega}}\}} \\
&\quad + (\gamma_j \Delta_n^{\bar{\omega}-1/2} - \iota_j' \beta_c \bar{F}_{c,t} - \bar{\epsilon}_{c,tj}) \mathbf{1}_{\{\Delta_{i_t^c}^n Y_j > \gamma_j \Delta_n^{\bar{\omega}}\}} \\
&\quad - (\gamma_j \Delta_n^{\bar{\omega}-1/2} + \iota_j' \beta_c \bar{F}_{c,t} + \bar{\epsilon}_{c,tj}) \mathbf{1}_{\{\Delta_{i_t^c}^n Y_j < -\gamma_j \Delta_n^{\bar{\omega}}\}},
\end{aligned} \tag{18}$$

where ι_j is a $p \times 1$ vector with j th element of 1 and rest of the elements being zero. The corresponding matrix of residuals is $R_c = (r_{c,tj} : t = 1, \dots, k_n, j = 1, \dots, p)$, for $c = a, b$.

With the above notation, we have the following discrete factor model for the truncated price increments over the two time windows:

$$\bar{Y}_c = \beta_c \bar{F}'_c + \bar{U}_c + R_c, \quad c = a, b, \tag{19}$$

where \bar{Y}_c , \bar{U}_c and R_c are $p \times k_n$ matrices, \bar{F}'_c is a $k_n \times K$ matrix and β_c is a $p \times K$ matrix, and provided $k_n \Delta_n < \varepsilon$ with ε being the constant in assumption A1. One can show also that in a certain asymptotic sense, see the proofs for formal results, the matrix of residual terms R_c is small. This is the reason why we refer to the above representation for the truncated increments of the asset prices as an discrete factor model.

4 Formulation and asymptotic properties of the test

Our goal is to test whether the number of systematic factors and the factor loadings, up to a rotation, are the same at two distinct points in time. That is, we are interested in testing the hypothesis:

$$\mathbb{H}_0 : \text{span}(\beta_b) = \text{span}(\beta_a), \tag{20}$$

where, for an arbitrary matrix β , $\text{span}(\beta)$ denotes the linear space spanned by the columns of β . The null hypothesis is equivalent to: $\beta_b = \beta_a H$ for some invertible matrix H . Using the notation in (4), under the null, $K = K_1$, and $K_2 = K_3 = K_4 = 0$, so $\beta_c = \beta_c^{(1)}$. First, we start with estimating the dimensions of β_b and β_a , K_a and K_b , respectively. Following that, we introduce our test statistic based on projection discrepancy of estimated factor loadings over the two periods and derive its limit distribution.

4.1 Formulation of the test

4.1.1 Step 0: Estimating the number of factors

We estimate separately the number of factors for the two time windows, by applying an information criterion as in Bai and Ng (2002). We extend their analysis here to our high-frequency setting. Alternative methods for estimating the number of factors include Ahn and Horenstein (2013); Hallin and Liška (2007); Onatski (2010); Kapetanios (2010), among many others.

Let $v_{c,1} \geq v_{c,2} \geq \dots$ be the sorted eigenvalues of

$$S_c = \frac{1}{k_n p} \bar{Y}_c \bar{Y}_c', \quad c \in \{a, b\}. \quad (21)$$

The number of factors K_c is estimated by

$$\hat{K}_c = \operatorname{argmin}_{0 \leq K \leq K_{\max}} \log \sum_{m > K} v_{c,m} + K \frac{k_n + p}{k_n p} g_{np}, \quad (22)$$

for some deterministic sequence $g_{np} > 0$ and a predetermined upper bound K_{\max} , both of which go asymptotically to infinity. In our implementation of the test, we use $g_{n,p} = \log \left(\frac{k_n p}{k_n + p} \right)$ as in Bai and Ng (2002). We also note that we allow for $K = 0$ in the search set for \hat{K}_c .

In addition to the dimensions of β_b and β_a , we also need the dimension of the non-redundant factor loadings across the two periods. We denote $K_{mix} = \dim(\beta_b, \beta_a)$, with $\dim(\beta)$ being the dimension of the linear space spanned by the columns of an arbitrary matrix β . Our estimator of K_{mix} is $\hat{K}_{mix} = \min\{\hat{K}_{mix,o}, \hat{K}_{mix,e}\}$, where $\hat{K}_{mix,o}$ and $\hat{K}_{mix,e}$ are estimates of the number of factors of two return panels that partition the pooled data $\frac{1}{k_n p} (\bar{Y}_b, \bar{Y}_a) (\bar{Y}_b, \bar{Y}_a)'$ into one from odd and even increments, respectively. Note that $\hat{K}_{mix,o}$ and $\hat{K}_{mix,e}$ are both valid estimators of K_{mix} . For this reason and in order to keep the factor specification more parsimonious in finite samples, we set $\hat{K}_{mix} = \min\{\hat{K}_{mix,o}, \hat{K}_{mix,e}\}$. The consistency of $(\hat{K}_a, \hat{K}_b, \hat{K}_{mix})$ is established in the proofs of our theoretical results.

Remark 4.1. We assume that the risk factors that assets are exposed to are *strong*, in the sense that they satisfy the well-known *pervasive condition*. In this setting the number of latent factors, while unknown, are consistently estimable. It is not difficult to extend the study to allow a mix of strong and so-called “semi-strong” factors, where the strength of the factors might also affect the rate of convergence of the test statistics. On the other hand, our assumption rules out the case of *weak factors* in the sense of Onatski (2012). Studying the span of betas corresponding to these types of weak factors is an open question.

4.1.2 Step 1: The projection discrepancy statistic

We proceed with formulating our test. Let $\widehat{\beta}_c$ denote the $p \times \widehat{K}_c$ estimate of β_c^r constructed as follows. Each column of $\widehat{\beta}_c/\sqrt{p}$ is a $p \times 1$ eigenvector of the sample covariance matrix S_c defined in (21). Our test statistic is then based on the projection discrepancy $\|P_{\widehat{\beta}_a} - P_{\widehat{\beta}_b}\|_F^2$, where for arbitrary matrix A , $P_A = A(A'A)^{-1}A'$, is $p \times p$ matrix and $\|A\|_F^2 = \text{trace}(A'A)$.

Using standard PCA analysis, it can be shown that, under the null hypothesis and for some normalization matrices $\widehat{A}_a, \widehat{A}_b$ and \widehat{G} of dimension $K \times K$, each with norm of asymptotic order $O_P(1)$,⁷, we have with probability approaching one:

$$\begin{aligned} \|P_{\widehat{\beta}_a} - P_{\widehat{\beta}_b}\|_F^2 &= \widehat{\mu}_a + \widehat{\mu}_b - \widehat{\mu}_{ab} + \Delta_5, \\ \widehat{\mu}_a &= \frac{2}{pk_n^2} \text{tr} \widehat{A}'_a [\overline{F}'_a \overline{U}'_a \overline{U}_a \overline{F}_a] \widehat{A}_a \\ \widehat{\mu}_b &= \frac{2}{pk_n^2} \text{tr} \widehat{A}'_b [\overline{F}'_b \overline{U}'_b \overline{U}_b \overline{F}_b] \widehat{A}_b \\ \widehat{\mu}_{ab} &= \frac{2}{pk_n^2} \text{tr} \widehat{A}'_a \overline{F}'_a \overline{U}'_a \overline{U}_b \overline{F}_b \widehat{A}_b \widehat{G}, \end{aligned} \quad (23)$$

where Δ_5 collects all higher-order terms. The leading term in the above expansion is $\widehat{\mu}_a + \widehat{\mu}_b - \widehat{\mu}_{ab}$. Its components, $\widehat{\mu}_a, \widehat{\mu}_b$ and $\widehat{\mu}_{ab}$, are jointly asymptotically normal but only $\widehat{\mu}_{ab}$ has zero mean. The other terms $\widehat{\mu}_a$ and $\widehat{\mu}_b$ are chi-square statistics with growing degrees of freedom p and whose means are nonzero. Therefore, we need to center them.

It can be shown that the asymptotic mean of $\widehat{\mu}_a$ is:

$$B_c = \frac{2}{k_n^2} \text{tr} \widehat{A}'_c \overline{F}'_c \overline{F}_c \widehat{A}_c \mathbb{E}(\sigma_{c,1}^2 | \mathcal{C}),$$

where $\mathbb{E}(\sigma_{c,1}^2 | \mathcal{C})$ denotes the expected idiosyncratic volatility conditional on the common information set \mathcal{C} (recall from assumption A2(i) that $\sigma_{c,i}^2$ are identically distributed across $i = 1, 2, \dots$, for $c = a, b$ conditional on \mathcal{C}). B_c is of order $O_P(k_n^{-1})$. After centering using B_c for $c \in \{a, b\}$, we have

$$\|P_{\widehat{\beta}_a} - P_{\widehat{\beta}_b}\|_F^2 - (B_a + B_b) = (\widehat{\mu}_a - B_a) + (\widehat{\mu}_b - B_b) - \widehat{\mu}_{ab} + \Delta_5. \quad (24)$$

All three leading terms $(\widehat{\mu}_a - B_a)$, $(\widehat{\mu}_b - B_b)$ and $\widehat{\mu}_{ab}$ are of order $O_P(k_n^{-1}p^{-1/2})$, which is the rate of convergence of our (unscaled) test statistic, and as we show in the Appendix, the higher-order term Δ_5 converges at a faster rate.

It remains to replace B_c by a consistent estimator \widehat{B}_c . Towards this end, define the

⁷Note that under the null hypothesis $K_a = K_b = K$ and $\mathbb{P}(\widehat{K}_a = K_a, \widehat{K}_b = K_b) \rightarrow 1$ from Theorem B.1 in the Appendix.

factor estimator $\widehat{F}_c = \overline{Y}'_c \widehat{\beta}_c / p$, and the estimated idiosyncratic terms $\widehat{U}_c = \overline{Y}_c - \widehat{\beta}_c \widehat{F}'_c$, whose elements are denoted by $\widehat{\epsilon}_{c,ti}$. Our estimate for the asymptotic bias of $\widehat{\mu}_c$, $c \in \{a, b\}$, is then given by

$$\widehat{B}_c = \frac{2}{k_n} \text{tr}(\widehat{Q}_c^{-1}) \mathbb{E}(\widehat{\sigma}_{c,1}^2 | \mathcal{C}), \quad c \in \{a, b\}, \quad (25)$$

where \widehat{Q}_c is a $\widehat{K}_c \times \widehat{K}_c$ diagonal matrix of the top eigenvalues of S_c . To understand the intuition behind this expression, note that B_c depends on the rescaled factors $\overline{F}_c \widehat{A}_c$. Although the scaling matrix \widehat{A}_c is not estimable, the rescaled factor can be directly estimated by $\widehat{F}_c \widehat{Q}_c^{-1}$, where the scaling matrix \widehat{Q}_c^{-1} is feasible. As such, the following term $\frac{1}{k_n} \widehat{A}'_c \overline{F}'_c \overline{F}_c \widehat{A}_c$ in B_c can be estimated by:

$$\widehat{Q}_c^{-1} \frac{1}{k_n} \widehat{F}'_c \widehat{F}_c \widehat{Q}_c^{-1} = \widehat{Q}_c^{-1}, \quad \text{given the identity } \frac{1}{k_n} \widehat{F}'_c \widehat{F}_c = \widehat{Q}_c.$$

The problem of estimating the conditional mean of idiosyncratic volatility $\mathbb{E}(\sigma_{c,1}^2 | \mathcal{C})$ is much more delicate. A natural candidate is $\frac{1}{k_n p} \|\widehat{U}_c\|_F^2$. Unfortunately, this estimator (although consistent) under-estimates its population counterpart $\frac{1}{k_n p} \|\overline{U}_c\|_F^2$ due to higher-order bias terms. This would lead to underestimation of $\mathbb{E}(\sigma_{c,1}^2 | \mathcal{C})$ in finite samples. Therefore, we employ a bias-adjusted estimator defined as

$$\mathbb{E}(\widehat{\sigma}_{c,1}^2 | \mathcal{C}) := \frac{1}{k_n p} \|\widehat{U}_c\|_F^2 + \delta, \quad (26)$$

where the extra term δ is given by

$$\delta = \frac{1}{k_n p} \|\widehat{U}_c\|_F^2 \widehat{K}_c / k_n + \frac{1}{p^2} \text{tr}(\widehat{\beta}'_c D_c \widehat{\beta}_c), \quad D_c = \text{diag} \left\{ \frac{1}{k_n} \sum_{t=1}^{k_n} \widehat{\epsilon}_{c,ti}^2, \quad i = 1, \dots, p \right\}. \quad (27)$$

Obviously, $\delta \geq 0$, and as we show in the proof, this term corrects for the downward bias in the naive estimator.

Altogether, the projection discrepancy statistic is then

$$\widehat{\mathcal{A}} := \|P_{\widehat{\beta}_a} - P_{\widehat{\beta}_b}\|_F^2 - (\widehat{B}_a + \widehat{B}_b). \quad (28)$$

4.1.3 Step 2: The bias-mimicking statistic

Although the expression in (28) is an asymptotically unbiased statistic under the null, higher-order biases might nevertheless lead to finite sample distortions of a test based on it. One approach to correct for those is to use a Jackknife method. Such a method, however, can eliminate only a bias of specific asymptotic order in terms of k_n and p . In our situation, the

higher-order bias terms can be multiple and with different asymptotic orders that depend on (k_n, p) in a rather complicated way. For this reason, we propose an alternative bias-reduction method that is more robust to the form of higher-order biases than the Jackknife.

The idea of our method is to center the statistic in (28) by a projection discrepancy of factor loadings estimated from two return panels of the same dimension as \bar{Y}_b and \bar{Y}_a and which have the same number factors both under the null and the alternative hypotheses. These two return panels are formed from pooling the return data from the two periods and splitting the combined data set into two: one formed from the odd increments and one from the even ones. Note that the number of factors of each of the two return panels is K_{mix} , and under the null hypothesis, $K_a = K_b = K_{mix}$.

We now provide the details. Let $\bar{Y}_{a,e}$ and $\bar{Y}_{b,e}$ denote the sub-matrices of \bar{Y}_a and \bar{Y}_b corresponding to *even* columns. We define in a similar way $\bar{Y}_{a,o}$ and $\bar{Y}_{b,o}$ from the *odd* columns. We then form the pooled data sets corresponding to odd and even increments:

$$\bar{Y}_{mix,k} = (\bar{Y}_{a,k}, \bar{Y}_{b,k}), \quad k \in \{o, e\}.$$

Let $\hat{\beta}_{mix,k}$ be the $p \times \hat{K}_{mix}$ matrix of PCA estimates for beta from the data matrix $\bar{Y}_{mix,k}$.

Now, consider the projection discrepancy:

$$\hat{\mathcal{A}}_{mix} := \|P_{\hat{\beta}_{mix,o}} - P_{\hat{\beta}_{mix,e}}\|_F^2 - (\hat{B}_{mix,o} + \hat{B}_{mix,e}), \quad (29)$$

which is a centered statistic similar to the one in (28) and the bias-correction terms are given by

$$\hat{B}_{mix,k} = \frac{2}{k_n^2} \text{tr} \hat{Q}_k^{-1} \hat{F}'_{a,k} \hat{F}_{a,k} \hat{Q}_k^{-1} \widehat{\mathbb{E}(\sigma_{a,1}^2 | \mathcal{C})} + \frac{2}{k_n^2} \text{tr} \hat{Q}_k^{-1} \hat{F}'_{b,k} \hat{F}_{b,k} \hat{Q}_k^{-1} \widehat{\mathbb{E}(\sigma_{b,1}^2 | \mathcal{C})},$$

where \hat{Q}_k is a $\hat{K}_{mix} \times \hat{K}_{mix}$ diagonal matrix of top eigenvalues of $\frac{1}{pk_n} \bar{Y}'_{mix,k} \bar{Y}_{mix,k}$; $\hat{F}_{c,k} = \bar{Y}'_{c,k} \hat{\beta}_{mix,k} / p$ for $c \in \{a, b\}$, and where we continue using $\widehat{\mathbb{E}(\sigma_{c,1}^2 | \mathcal{C})}$ as our estimate of the idiosyncratic variance.

The new projection discrepancy in (29) is negligible in the sense that $\hat{\mathcal{A}}_{mix} = o_P(1)$ under the null and alternative hypotheses. Its exact order of magnitude is $O_P(k_n^{-1} p^{-1/2})$.

Altogether, our final test statistic is given by:

$$\mathcal{S} := k_n \sqrt{p} \left(\hat{\mathcal{A}} - \hat{\mathcal{A}}_{mix} \right). \quad (30)$$

In Section 4.3, we will show formally that the centering by $\hat{\mathcal{A}}_{mix}$ can indeed lead to bias-reduction and in turn allow for weaker conditions on the relative size of the two dimensions

of the return panel used in the estimation.

4.1.4 Step 3: The cross-sectional bootstrap

For feasible inference, we propose a simple to implement bootstrap. Under the null hypothesis, the test statistic has the asymptotic expansion

$$\mathcal{S} = \frac{1}{\sqrt{p}} \sum_{i=1}^p z_{i,n} + o_P(1),$$

where $z_{i,n}$ is a zero-mean triangular array, whose variance depends on the asymptotic variances of $\|P_{\hat{\beta}_a} - P_{\hat{\beta}_b}\|_F^2$ and $\|P_{\hat{\beta}_{mix,o}} - P_{\hat{\beta}_{mix,e}}\|_F^2$ as well as their asymptotic covariance. To avoid computing estimates for these sophisticated expressions, we rely on a cross-sectional bootstrap by re-sampling the rows of the $p \times (2k_n)$ matrix (\bar{Y}_b, \bar{Y}_a) .

More specifically, let $\{(\bar{Y}_{b,i}, \bar{Y}_{a,i})^* : i = 1, \dots, p\}$ be a simple random sample with replacement from $\{(\bar{Y}_{b,i}, \bar{Y}_{a,i}) : i = 1, \dots, p\}$, where $\bar{Y}'_{c,i}$ denotes the $1 \times k_n$ vector of the i -th row of \bar{Y}_c . Let $(\bar{Y}_a^*, \bar{Y}_b^*)$ denote the resulting bootstrap matrix. We compute the test statistic on the new data as:

$$\mathcal{S}^* := k_n \sqrt{p} \left[\|P_{\hat{\beta}_a}^* - P_{\hat{\beta}_b}^*\|_F^2 - (\hat{B}_a^* + \hat{B}_b^*) \right] - k_n \sqrt{p} \mathcal{A}_{mix}^*,$$

where $(P_{\hat{\beta}_c}^*, \hat{B}_c^*, \mathcal{A}_{mix}^*)$ denotes the bootstrap counterpart of $(P_{\hat{\beta}_c}, \hat{B}_c, \hat{\mathcal{A}}_{mix})$. We repeat this procedure many times and compute $q_\tau^*\{\mathcal{S}^* - \mathcal{S}\}$, the τ th upper quantile of the bootstrap distribution of $\mathcal{S}^* - \mathcal{S}$, i.e., $\mathbb{P}(\mathcal{S}^* - \mathcal{S} > q_\tau^*\{\mathcal{S}^* - \mathcal{S}\}) = \tau$. The critical region for a test of size τ is then given by

$$\mathcal{S} > q_\tau^*\{\mathcal{S}^* - \mathcal{S}\}.$$

4.2 Limit behavior of the test

Our goal in this section is to derive the behavior of the test statistic \mathcal{S} , defined in (30), under the null hypothesis and a set of alternative hypotheses. We start with a CLT result under the null. Henceforth, $\mathcal{L}|\mathcal{C}$ and $\mathcal{L}|\mathcal{F}$ denote \mathcal{C} -conditional and \mathcal{F} -conditional, respectively, convergence in law, see Section VIII.5(b) in Jacod and Shiryaev (2003).

Theorem 4.1 (size). *Suppose Assumptions A1-A3 hold as well as \mathbb{H}_0 given in (20). Further, let*

$$K_{max} \rightarrow \infty, K_{max}/\sqrt{k_n} \rightarrow 0, g_{np} \rightarrow \infty, \left(\frac{1}{k_n} + \frac{1}{p} \right) g_{np} \rightarrow 0, \frac{\zeta_p}{g_{np}} \rightarrow 0, \quad (31)$$

$$\left(\frac{p}{k_n^2} + \frac{k_n^2}{p^3}\right) \zeta_p^8 \rightarrow 0 \quad \text{and} \quad pk_n \Delta_n^{2\tilde{\varpi}} \rightarrow 0, \quad (32)$$

as $n, p \rightarrow \infty$ and for some $\tilde{\varpi} \in (0, \varpi)$ that is arbitrary close to $\varpi \in (0, 1/2)$. Then, we have

$$\mathcal{S} \xrightarrow{\mathcal{L}|\mathcal{C}} Z_d, \quad (33)$$

where Z_d is defined on an extension of the original probability space and, conditional on \mathcal{C} , it is a mean-zero normal random variable with some \mathcal{C} -adapted variance. In addition, the bootstrap statistic satisfies

$$\mathcal{S}^* - \mathcal{S} \xrightarrow{\mathcal{L}|\mathcal{F}} Z_d^*, \quad (34)$$

where Z_d^* has the same \mathcal{F} -conditional law as that of the \mathcal{C} -conditional law of Z_d .

Hence, for any $\tau \in (0, 1)$, we have

$$\mathbb{P}(\mathcal{S} > q_\tau^* \{\mathcal{S}^* - \mathcal{S}\}) \rightarrow \tau. \quad (35)$$

The conditions in (31)-(32) put restrictions on the asymptotic size of the two dimensions of the return panel and on the tuning parameters for choosing the number of factors. The rate requirements for g_{np} are relatively weak. We make several comments about the conditions in (32). First, it is optimal to pick the jump truncation parameter ϖ arbitrary close to, but below, $1/2$. This is similar to other applications of jump truncation procedures, see e.g., Jacod and Protter (2012). The second condition in (32) is due to the effect of the discretization error on the estimation, i.e., due to the residual term R_c in the discrete factor model in (19). The second condition in (32) is also needed in order to guarantee that the error due to the time-variation in volatility over the local windows has only asymptotically negligible effect on the estimation. This condition puts an upper bound on the size of the cross-sectional dimension p relative to the length of the two local windows used in the estimation.

Second, the conditions for p and k_n in (32) allow for the standard case in which p and k_n grow at the same rate, provided of course ζ_p increases at a slower rate than $p^{1/8}$ (see the discussion after assumption A3 regarding the size of ζ_p). We note that the first of the two conditions in (32) allows for either of the two dimensions of the return panel to grow at a faster rate than the other one.

The rate of convergence of our statistic is determined by both dimensions of the return panel. This is unlike the rate of convergence of the estimators of the factor loadings of individual assets, i.e., of the rows of $\hat{\beta}_c$, which depends only on k_n , see e.g., Bai (2003). The bias-correction term $\hat{B}_a + \hat{B}_b$ is of asymptotic order $O_P(1/k_n)$. This term becomes non-

negligible when multiplied by $k_n\sqrt{p}$, and therefore such a bias-correction of the projection discrepancy statistic is necessary for the CLT result.

We finish this section with characterising the behavior of our statistic under a set of alternatives. This is formally done in the next theorem.

Theorem 4.2 (power). *Suppose Assumptions A1-A3 hold and the conditions in (31)-(32) are satisfied. In addition, let either of the following two alternatives hold:*

Alternative (i): $\beta_a = (\beta_b^{(1)}H, \mathbf{0}_{p \times K_3})$, and $\beta_b = (\beta_b^{(1)}, \beta_b^{(3)})$, so $K_2 + K_4 = 0$.

Alternative (ii): $\beta_a = \beta_a^{(2)}$, and $\beta_b = \beta_b^{(2)}$, so $K_1 + K_3 + K_4 = 0$. Also, there is $c > 0$ so that with probability approaching one,

$$\min_{H \in \mathbb{R}^{K \times K}} \frac{1}{\sqrt{p}} \|\beta_a H - \beta_b\|_F > c.$$

Then, under either alternative, for any $\tau \in (0, 1)$,

$$\mathbb{P}(\mathcal{S} > q_\tau^* \{\mathcal{S}^* - \mathcal{S}\}) \rightarrow 1.$$

The result of Theorem 4.2 shows that a one-sided test based on the statistic \mathcal{S} will have power against two types of alternatives.

The first alternative covers a situation in which the cross-section of assets at one of the two time points that are compared contains more systematic risk factors than the other. We can think of this scenario as a situation in which some systematic risk factors become dormant over certain periods of time or alternatively as a situation in which some risk factors appear only at some unique points of time (e.g., following the release of a macroeconomic announcement). We note that because of the invertibility of $\Sigma_{\beta,c}^r$ from Assumption A2(ii), the probability limit of the Schur complement $\frac{1}{p} \beta_b^{(3)'} (I - P_{\beta_b^{(1)}}) \beta_b^{(3)}$ is also invertible. This intuitively means that asymptotically $\beta_b^{(3)}$ is not in the linear space of $\beta_b^{(1)}$.

The second alternative that we consider in Theorem 4.2 is one in which the number of systematic risk factors at the two time points is the same but the linear spaces spanned by the factor loadings differ. Our test does not have power against situations in which the factor loadings of only a finite number of assets change. For the test to have power, the change in the factor loadings should be pervasive, i.e., it should affect an increasing number of assets. Note that in many applications in finance, one is interested in large portfolios of assets constructed on the basis of their factor loadings. For these types of applications, it is the pervasive changes in factor loadings that matter and our test is designed to detect those.

4.3 Higher-order bias correction using $\widehat{\mathcal{A}}_{mix}$

Recall that our test statistic is defined as $\mathcal{S} = k_n \sqrt{p}(\widehat{\mathcal{A}} - \widehat{\mathcal{A}}_{mix})$, where

$$\widehat{\mathcal{A}} = \|P_{\widehat{\beta}_a} - P_{\widehat{\beta}_b}\|_F^2 - (\widehat{B}_a + \widehat{B}_b),$$

and $\widehat{\mathcal{A}}_{mix}$ is the bias-mimicking statistic that is defined similarly to $\widehat{\mathcal{A}}$ but by return matrices whose columns are mixtures of odd and even observations within the two time periods. As we mentioned earlier, this term can provide higher-order bias corrections.

In this section we formalize this statement by showing that including $\widehat{\mathcal{A}}_{mix}$ in \mathcal{S} can allow for weaker condition for the asymptotic size of k_n relative to p . This result is rather complicated to establish in general. We will illustrate it here under the following stronger condition, which essentially requires volatilities to stay constant in the two periods:

A4. *Suppose that under the null $\beta_b = \beta_a$, $\Lambda_b = \Lambda_a$ and $\sigma_{b,i}^2 = \sigma_{a,i}^2$, for $i = 1, \dots$. In addition, suppose that Λ_t and $\{\sigma_{t,i}^2\}_{i \geq 1}$ remain constant for $t \in (c - \varepsilon, c + \varepsilon)$ with $\varepsilon > 0$ being an arbitrary small number and $c = a, b$.*

From our discussion in Sections 4.1.2, $\widehat{\mathcal{A}}$ can be expressed as:

$$k_n \sqrt{p} \widehat{\mathcal{A}} = \widehat{Z}_1 + \widehat{\mathcal{R}\mathcal{A}},$$

where \widehat{Z}_1 is the leading term that admits a CLT and the residual term

$$\widehat{\mathcal{R}\mathcal{A}} := \sqrt{p} k_n \Delta_5 - \sqrt{p} k_n \sum_{c \in \{a, b\}} (\widehat{B}_c - B_c),$$

contains higher-order terms. In the proof of Theorem 4.1 we show that \widehat{Z}_1 admits a CLT, provided only $k_n, p \rightarrow \infty$, $\zeta_p/p \rightarrow 0$ and $p k_n \Delta_n = O_p(1)$ hold. These conditions are much weaker than those in (32). The latter are needed to show that $\widehat{\mathcal{R}\mathcal{A}}$ is negligible under the null hypothesis. In particular, one of the conditions in (32) is $\frac{p}{k_n^2} \zeta_p^8 \rightarrow 0$. This condition limits how fast p can grow relative to k_n and is due to higher-order biases in $\widehat{\mathcal{R}\mathcal{A}}$. In fact, we have the following higher-order expansion:

$$\widehat{\mathcal{R}\mathcal{A}} = \frac{\sqrt{p}}{k_n} \mathcal{R}\mathcal{A} + o_P(1) + o_P \left(\frac{\sqrt{p}}{k_n} \right)^{1/2}, \quad (36)$$

where $\mathcal{R}\mathcal{A}$ is a random variable (depending on the factor loadings and the volatilities) given in the Appendix.

Using Assumption A4, we can make a similar expansion for $\widehat{\mathcal{A}}_{mix}$:

$$k_n \sqrt{p} \widehat{\mathcal{A}}_{mix} = \widehat{Z}_{1,mix} + \widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix},$$

where $\widehat{Z}_{1,mix}$ is the leading term that admits a CLT and the residual term $\widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix}$ yields a same higher order bias term:

$$\widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix} = \frac{\sqrt{p}}{k_n} \mathcal{R}\mathcal{A} + o_P(1) + o_P\left(\frac{\sqrt{p}}{k_n}\right)^{1/2}. \quad (37)$$

We note that $(\widehat{Z}_1, \widehat{Z}_{1,mix})$ admit a joint CLT and the limit of $\widehat{Z}_{1,mix}$ is non-trivial in general. This means that the centering with $\widehat{\mathcal{A}}_{mix}$ affects the limit distribution of our statistic.

Comparing (36) and (37), we get $\widehat{\mathcal{R}}\widehat{\mathcal{A}} - \widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix} = o_P(1) + o_P\left(\frac{\sqrt{p}}{k_n}\right)^{1/2}$, i.e., we have bias cancellation and this difference is negligible under a weaker condition for k_n and p . In particular, this argument will allow us to weaken the condition from $\frac{p}{k_n^2} \zeta_p^8 \rightarrow 0$ to allowing $\frac{p}{k_n^2} \zeta_p^8 \rightarrow \kappa$ for some $\kappa \geq 0$. Formally, we have the following result:

Theorem 4.3. *Suppose Assumptions A1-A4 hold as well as \mathbb{H}_0 given in (20). Then*

$$\mathcal{S} = \frac{1}{\sqrt{p}} \sum_{i=1}^p z_{i,n} + \widehat{\mathcal{R}}\widehat{\mathcal{A}} - \widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix},$$

where $\frac{1}{\sqrt{p}} \sum_{i=1}^p z_{i,n} \xrightarrow{\mathcal{L}^C} Z_d$, for Z_d being the limit variable in (33), and $\widehat{\mathcal{R}}\widehat{\mathcal{A}}$ and $\widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix}$ satisfy (36) and (37), with the expression for $\mathcal{R}\mathcal{A}$ given in equation (E.1) in the Appendix.

Hence, $\widehat{\mathcal{R}}\widehat{\mathcal{A}} - \widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix} = o_P(1)$ if the requirement $\frac{p}{k_n^2} \zeta_p^8 \rightarrow 0$ in (32) is replaced with the weaker one $\frac{p}{k_n^2} \zeta_p^8 \rightarrow \kappa$, for some finite $\kappa \geq 0$, and all other conditions in (31)-(32) remain true.

As is clear from the discussion above, the improvement in the above theorem is due to cancellation of biases in $\widehat{\mathcal{A}}$ and $\widehat{\mathcal{A}}_{mix}$. Naturally, such improvement will not be available if one was to use $k_n \sqrt{p} \widehat{\mathcal{A}}$ instead. We note that the improvements offered by centering with $\widehat{\mathcal{A}}_{mix}$ are likely much larger than what is implied by Theorem 4.3 because here we only made an expansion of $\widehat{\mathcal{R}}\widehat{\mathcal{A}}$ and $\widehat{\mathcal{R}}\widehat{\mathcal{A}}_{mix}$ with the leading term being of order \sqrt{p}/k_n . A full asymptotic expansion of these higher-order terms that can show further bias cancellations is much harder to prove and we do not do this here. In our Monte Carlo simulations, we will show the practical gains in realistic situations from using $\widehat{\mathcal{A}}_{mix}$.

4.4 Time aggregation

In order to reduce the uncertainty from estimating the unknown number of factors, we can consider pooling data in the estimation of the factor loadings over several local windows. This can be done, of course, under the assumption that the factor loadings remain constant (up to rotation) over these periods. The results of Theorems 4.1-4.2 are easily extended to cover this situation and in this section we only sketch such an extension.

We denote two finite sets of points in time with \mathbf{A} and \mathbf{B} . All of the elements of \mathbf{A} and \mathbf{B} are in the interior of $[0, T]$. The number of factors and factor loadings at each point in each of these two sets are the same (up to a multiplication by a rotation matrix) under the null hypothesis. The counterpart of \bar{Y}_c is denoted with $\bar{Y}_{\mathbf{C}}$, for $c = a, b$ and $\mathbf{C} = \mathbf{A}, \mathbf{B}$. Then we set $S_{\mathbf{C}} = \frac{1}{\mathbf{T}k_n p} \bar{Y}_{\mathbf{C}} \bar{Y}'_{\mathbf{C}}$, for $\mathbf{C} = \mathbf{A}, \mathbf{B}$ and where \mathbf{T} denotes the number of elements in \mathbf{C} (thus, this number is the same for \mathbf{A} and \mathbf{B}). The counterpart of $\hat{\beta}_c$, when S_c is replaced with $S_{\mathbf{C}}$, is denoted with $\hat{\beta}_{\mathbf{C}}$ and that of \hat{Q}_c with $\hat{Q}_{\mathbf{C}}$. With this notation, the bias correction term is now given by

$$\hat{B}_{\mathbf{C}} = \frac{2}{\mathbf{T}^2 k_n^2} \sum_{c \in \mathbf{C}} \left\{ \text{tr} \left(\hat{Q}_{\mathbf{C}}^{-1} \hat{F}'_c \hat{F}_c \hat{Q}_{\mathbf{C}}^{-1} \right) \mathbb{E}(\widehat{\sigma_{c,1}^2} | \mathcal{C}) \right\}. \quad (38)$$

Similarly, let $\hat{Q}_{mix,k}$ denote the diagonal matrix of the top $K_{mix,k}$ eigenvalues from the aggregated matrix $(\bar{Y}_{\mathbf{A},k}, \bar{Y}_{\mathbf{B},k})$ and let $\hat{F}_{c,k}$ denote the counterpart of \hat{F}_c for $k \in \{o, e\}$. The bias correction term for $\hat{\mathcal{A}}_{mix}$ is

$$\hat{B}_{mix,k} = \frac{2}{\mathbf{T}^2 k_n^2} \text{tr} \left(\sum_{a \in \mathbf{A}} \hat{Q}_{mix,k}^{-1} \hat{F}'_{a,k} \hat{F}_{a,k} \hat{Q}_{mix,k}^{-1} \mathbb{E}(\widehat{\sigma_{a,1}^2} | \mathcal{C}) + \sum_{b \in \mathbf{B}} \hat{Q}_{mix,k}^{-1} \hat{F}'_{b,k} \hat{F}_{b,k} \hat{Q}_{mix,k}^{-1} \mathbb{E}(\widehat{\sigma_{b,1}^2} | \mathcal{C}) \right).$$

We can then show, under the same conditions as those of Theorem 4.1, the following convergence result under \mathbb{H}_0 :

$$\mathbf{T}k_n \sqrt{p} \left[\|P_{\hat{\beta}_{\mathbf{A}}} - P_{\hat{\beta}_{\mathbf{B}}}\|_F^2 - (\hat{B}_{\mathbf{A}} + \hat{B}_{\mathbf{B}}) - \hat{\mathcal{A}}_{mix} \right] \xrightarrow{\mathcal{L}|\mathcal{C}} Z_d. \quad (39)$$

We can similarly establish the counterparts of the bootstrap result in Theorem 4.1 as well as the power result of Theorem 4.2 adapted to the current situation.

5 Simulation study

5.1 Setting

The model used in the Monte Carlo is given by

$$dY_{t,j} = \sigma_t \left(\beta_j' \Sigma_f^{1/2} dW_t + \bar{\sigma}_j dW_{t,j} \right), \quad j = 1, \dots, p, \quad (40)$$

where the univariate stochastic volatility process σ_t has the following square-root diffusion dynamics

$$d\sigma_t^2 = 8.3(1 - \sigma_t^2)dt + \sigma_t dB_t, \quad (41)$$

with B_t being a standard Brownian motion that is independent from W_t and $\{W_{t,j}\}_{j \geq 1}$. The process σ_t drives variation both in diffusive and idiosyncratic variances. The parameter specification of the dynamics of σ_t implies that the half-life of a shock to it is one month.

The specification of the factors and the factor loadings is calibrated to match estimates from five-minute frequency S&P500 returns over January 2021. We run PCA on the high-frequency returns for $K = 3$ factors, and calculate the covariance of the estimated factors Σ_f , as well as the mean μ_β and covariance Σ_β of the estimated loadings. We then generate factor and betas from $\beta_i \sim N(\mu_\beta, \Sigma_\beta)$, $i = 1, \dots, p$. The calibrated parameters are as follows:

$$\Sigma_f = \text{diag}(1.198, 0.377, 0.264) \times 10^{-6}, \quad (42)$$

$$\mu_\beta = \begin{pmatrix} 0.8763 \\ 0.2818 \\ -0.2965 \end{pmatrix}, \quad \Sigma_\beta = \begin{pmatrix} 0.2326 & -0.2474 & 0.2603 \\ -0.2474 & 0.9224 & 0.0837 \\ 0.2603 & 0.0837 & 0.9139 \end{pmatrix}. \quad (43)$$

The scale of the idiosyncratic variance $\bar{\sigma}_j$ is cross-sectionally i.i.d. and is drawn according to

$$\bar{\sigma}_j \sim \text{Uniform}([0.5, 1.5]) \times g, \quad j = 1, \dots, p. \quad (44)$$

where $g = 1.1 \times 10^{-3}$ is chosen such that the share of idiosyncratic risk in total asset variance is around 40% for the median stock in the cross-section.

The unit of time in the above specification of the model is one year. We adopt a business day time convention, which means that a period from market close on one day to market close on the next trading day has a length of $1/252$.

The specification for the process under the alternative hypothesis is the same as that under the null described above with one exception regarding the specification of the factor loadings. Mainly, under alternative $\mathbb{A}1$, the factor loadings at the two time points are given

by

$$\mathbb{A}1 : \beta_a^{(1)} = \beta_b^{(1)} + N \left(0, \frac{\tau}{p} \|\beta_b\|_F^2 \times I \right), \quad \beta_a^{(j)} = \beta_b^{(j)} \text{ for } j \neq 1, \quad (45)$$

where $\beta_c^{(j)}$ denotes the j -th column of the $N \times 3$ matrix β_c , for $c \in \{a, b\}$. The variance $\frac{\tau}{p} \|\beta_b\|_F^2$ of the perturbation is chosen such that $\|\beta_a - \beta_b\|_F^2 / \|\beta_b\|_F^2 \xrightarrow{\mathbb{P}} \tau$ as $p \rightarrow \infty$, so τ represents the amount of percentage change of the accumulated signals in betas. We experiment with $\tau \in \{0.01, 0.05\}$ in the simulation.

Under alternative $\mathbb{A}2$, we set one of the columns of β_a to be zero, and the other two columns to be the same as those of β_b . Formally,

$$\mathbb{A}2 : \beta_a^{(k)} = 0, \quad \beta_a^{(j)} = \beta_b^{(j)} \text{ for } j \neq k. \quad (46)$$

In the simulation, we experiment with $k \in \{2, 3\}$.

We turn next to the observation scheme. We consider a setup that is similar to the one in our empirical application. In particular, we use a cross-section of $p = 500$ assets and assume that we sample asset prices either 80 or 40 times during a trading day. This corresponds approximately to sampling every five or ten minutes in a 6.5 hour trading day. On each day, we consider a time window of 2 hours at the beginning and the end of the trading day. Under the five-minute frequency, there are $k_n = 24$ observations in each of the two time windows while under the ten-minute frequency, there are $k_n = 12$ observations per window.

In addition, in order to increase precision of estimating the number of factors, we consider pooling data from several consecutive trading days in the estimation. This is possible if the factor loadings across the pooled periods remain the same. More specifically, we pool data over a period of D consecutive days, which results in time-series dimension of our return panels of $D \times k_n$, for $D \in \{1, 2, \dots, 30\}$. This ranges from 12 to $30 \times 12 = 360$ observations under the ten-minute frequency, and from 24 to $30 \times 24 = 720$ observations under the five-minute frequency.

Finally, the tuning parameters of the test are set as follows. First, the truncation parameter is set in the following data-driven way:

$$\varpi = 0.49 \quad \text{and} \quad \gamma_j = 4 \times \sqrt{RV_{d,j} \wedge BV_{d,j}}, \quad j = 1, \dots, p, \quad (47)$$

where $RV_{d,j}$ and $BV_{d,j}$ are realized variance and bipower variation of asset j on the (trading) day d the increment belongs to, given by

$$RV_{d,j} = \sum_{i=\lfloor (d-1)/(252\Delta_n) \rfloor + 1}^{\lfloor d/(252\Delta_n) \rfloor} (\Delta_i^n X_j)^2, \quad BV_{d,j} = \frac{\pi}{2} \sum_{i=\lfloor (d-1)/(252\Delta_n) \rfloor + 2}^{\lfloor d/(252\Delta_n) \rfloor} |\Delta_i^n X_j \Delta_{i-1}^n X_j|. \quad (48)$$

The bipower variation is a jump-robust and tuning-free measure of volatility proposed by Barndorff-Nielsen and Shephard (2004b). This choice of truncation is commonly adopted in applied work using truncated variation. Next, as in Bai and Ng (2002), we set $g_{n,p} = \log\left(\frac{k_n p}{k_n + p}\right)$ in the penalty term for determining the number of factors. For determining the number of factors, we standardize the asset returns by estimates of their volatility in order to minimize the impact of idiosyncratic volatility.

5.2 Estimating the number of factors

We first analyze the accuracy of estimating the number of factors under the null and alternative hypotheses. Figure 1 plots the average of the estimated number of factors over 100 Monte Carlo replications. Recall that K_a , K_b and K_{mix} respectively denote the true number of factors for \bar{Y}_a , \bar{Y}_b and \bar{Y}_{mix} return matrices. The top three panels of the figure plot results under the null hypothesis, where the true number of factors is 3. The middle three panels plot results under alternative \mathbb{A}_1 where only the first columns of β_a and β_b are different, and $\|\beta_a - \beta_b\|_F^2 / \|\beta_b\|_F^2 \approx 0.05$. Therefore, $K_{mix} = 4$ in this case. The bottom three panels plot results under alternative \mathbb{A}_2 where the third column of β_a is zero, which leads to $K_a = 2$.

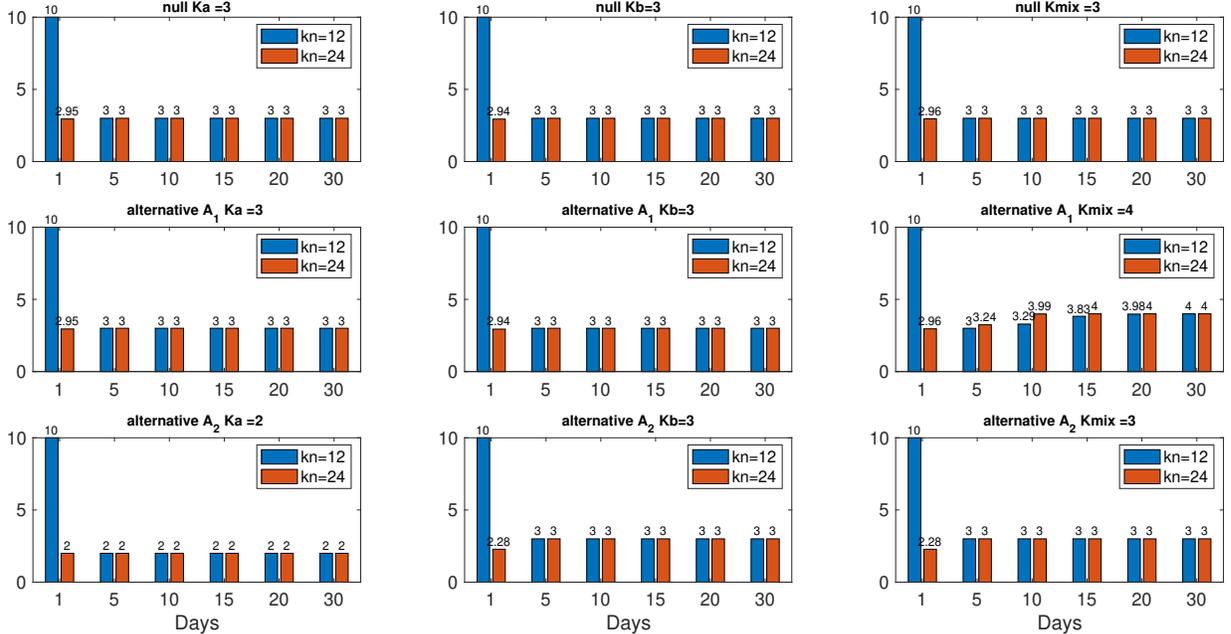


Figure 1: Estimated number of factors in the Monte Carlo averaged over 100 replications.

From Figure 1 we see that regardless of the scenario, when we use data for just one day ($D = 1$) and $k_n = 12$, the estimated number of factors is always equal to 10, the default upper

Table 1: Rejection Probabilities of the Test \mathcal{S} in Monte Carlo.

Hypothesis	Number of pooled days D					
	1	5	10	15	20	30
	$k_n = 12$: ten-minute frequency					
\mathbb{H}_0	0.140	0.038	0.038	0.043	0.041	0.035
A1, $\tau = 0.01$	0.148	0.340	0.828	0.982	1	1
A1, $\tau = 0.05$	0.220	0.992	1	1	1	1
A2, $\beta_a^{(2)} = 0$	0.872	1	1	1	1	1
A2, $\beta_a^{(3)} = 0$	0.652	0.982	1	1	1	1
	$k_n = 24$: five-minute frequency					
\mathbb{H}_0	0.077	0.035	0.033	0.042	0.044	0.037
A1, $\tau = 0.01$	0.154	0.808	0.996	1	1	1
A1, $\tau = 0.05$	0.658	1	1	1	1	1
A2, $\beta_a^{(2)} = 0$	0.854	1	1	1	1	1
A2, $\beta_a^{(3)} = 0$	0.594	1	1	1	1	1

The results are based on 2,000 Monte Carlo replications under the null, and on 1,000 replications under the alternative. We set the number of bootstrap replications to $\mathcal{B} = 1,000$. The nominal size of the test is 5%.

bound on the estimate. That is, in this case the number of factors is severely overestimated. Meanwhile, in all other configurations of k_n and D , the number of estimated factors is very close to its true value.

5.3 Results: size and power of the proposed test

We proceed with examining the rejection probabilities of the proposed test under the various scenarios. In each Monte Carlo simulation, like in the empirical application, the number of factors is estimated. The Monte Carlo results are reported in Table 1.

The finite sample size properties of the test appear good across most of the considered configurations in terms of sampling frequencies and levels of time aggregation, except for $D = 1$. In the case of $D = 1$, noticeable over-rejection occurs (0.141 for $k_n = 12$ and 0.083 for $k_n = 24$). This is mainly due to the fact that the estimated number of factors is severely overestimated in this case, recall Figure 1. Nevertheless, the size distortions caused by overestimating K are not very big. In addition, the results reported in Table 1 show robustness of the test in this regard with the test performing similarly under \mathbb{H}_0 for low and high levels of time aggregation.

Table 2: Rejection Probabilities of the Test $k_n\sqrt{p}\widehat{\mathcal{A}}$ in Monte Carlo.

Hypothesis	Number of pooled days D					
	1	5	10	15	20	30
	$k_n = 12$: ten-minute frequency					
\mathbb{H}_0	1	0.004	0.009	0.009	0.009	0.012
$\mathbb{A}1, \tau = 0.01$	1	0.244	0.932	0.998	1	1
$\mathbb{A}1, \tau = 0.05$	1	1	1	1	1	1
$\mathbb{A}2, \beta_a^{(2)} = 0$	1	1	1	1	1	1
$\mathbb{A}2, \beta_a^{(3)} = 0$	1	0.980	1	1	1	1
	$k_n = 24$: five-minute frequency					
\mathbb{H}_0	0.044	0.005	0.007	0.012	0.02	0.016
$\mathbb{A}1, \tau = 0.01$	0.068	0.946	1	1	1	1
$\mathbb{A}1, \tau = 0.05$	0.646	1	1	1	1	1
$\mathbb{A}2, \beta_a^{(2)} = 0$	0.922	1	1	1	1	1
$\mathbb{A}2, \beta_a^{(3)} = 0$	0.616	1	1	1	1	1

The results are based on 2,000 Monte Carlo replications under the null, and on 500 replications under the alternative. We set the number of bootstrap replications to $\mathcal{B} = 1,000$. The nominal size of the test is 5%.

Turning next to the behavior of the test under the two considered alternatives, we can note from the reported results that the test has overall good power. Not surprisingly, the power increases as we increase the sampling frequency and when we consider higher levels of time aggregation.

5.4 The role of $\widehat{\mathcal{A}}_{mix}$

To study the role of including the bias mimicking statistic in our test, Table 2 summarizes the rejection probabilities when the test statistic does not contain $\widehat{\mathcal{A}}_{mix}$ but is just $k_n\sqrt{p}\widehat{\mathcal{A}}$. Like the test based on \mathcal{S} , we perform cross-sectional bootstrap, and the rejection is based on the bootstrap critical value. The asymptotic null distribution is still normal. However, from the results in Table 2, we can now notice nontrivial size distortions.

The size distortions are most extreme in the case $D = 1$ and $k_n = 12$. Recall from Figure 1 that for this configuration the number of factors is severely overestimated. The test, with and without using $\widehat{\mathcal{A}}_{mix}$, will over-reject the null hypothesis. However, this effect is much more dramatic without using $\widehat{\mathcal{A}}_{mix}$. As D increases, the number of factors are estimated relatively well throughout the simulation replications. In these scenarios, the test

without $\hat{\mathcal{A}}_{mix}$ now becomes severely undersized, with type I errors below 0.011. This is mainly because of the presence of higher-order biases in $k_n\sqrt{\hat{p}}\hat{\mathcal{A}}$.

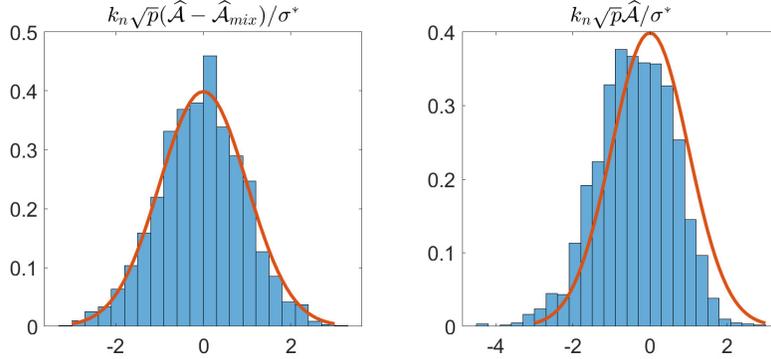


Figure 2: Histogram of t-statistics from 2,000 replications under the null, with (left panel) and without (right panel) the bias-mimicking statistics, superposed by the standard normal density. The t-statistics are standardized by the bootstrap interquartile range. The true number of factors are used in the construction of the test statistics, and we fix $k_n = 12$ and $D = 5$ days.

To illustrate this, we conduct another Monte Carlo experiment under the null, fixing $k_n = 12$ and $D = 5$ days and using the true number of factors in the construction of the test. Figure 2 plots the histograms of the standardized test statistics with (left panel) or without (right panel) the centering by the bias-mimicking statistic. The test statistics are standardized by the bootstrap interquartile range, defined as

$$\sigma^* = \frac{q_{0.25}^* - q_{0.75}^*}{z_{0.25} - z_{0.75}},$$

where q_τ^* is the τ -th upper quantile of the bootstrap statistic and z_τ is the τ -th upper quantile of the standard normal distribution. We see from Figure 2 that, without centering using the bias-mimicking statistic, $k_n\sqrt{\hat{p}}\hat{\mathcal{A}}$ is downward biased. This leads to the under-rejections documented in Table 2.

Finally, the test without $\hat{\mathcal{A}}_{mix}$ has somewhat higher power than the one with $\hat{\mathcal{A}}_{mix}$ in the various configurations of D and k_n and for the various alternative scenarios. Nevertheless, given the reported nontrivial size distortions of a test without using $\hat{\mathcal{A}}_{mix}$ above, it is important for applications to use $\hat{\mathcal{A}}_{mix}$ in the test.

5.5 Robustness to weak cross-sectional dependence in idiosyncratic risk

Our test is based on cross-sectional bootstrap which relies on independence of the Brownian motions driving the idiosyncratic risk in asset prices. In this section we study how sensitive is the performance of the test to presence of mild cross-sectional dependence in the idiosyncratic risk. We use the following modified model for this:

$$dY_t = \sigma_t \left(\beta \Sigma_f^{1/2} dW_t + \Sigma_\epsilon^{1/2} d\widetilde{W}_t \right), \quad (49)$$

where \widetilde{W}_t is a $p \times 1$ standard Brownian motion independent from W_t , $\beta = (\beta_j : j = 1, \dots, p)$ and Σ_f are exactly as in our original Monte Carlo setup, and Σ_ϵ is a $p \times p$ matrix with (i, j) entry equal to $\sigma_\epsilon^2 \rho^{|i-j|}$ for some parameters $\rho \in (0, 1)$ and σ_ϵ^2 . The parameter σ_ϵ^2 is chosen such that the contribution of idiosyncratic noise to total asset variance is around 40%. The parameter ρ governs the strength of the cross-sectional dependence in the idiosyncratic risk in asset prices. We experiment with two values for it: $\rho = 0.3$ and $\rho = 0.1$.

For brevity, we consider only the configuration $k_n = 12$ and $D = 5$. When $\rho = 0.1$, the results are quantitatively similar to those reported in Tables 1 and 2. When $\rho = 0.3$, the size is slightly distorted: the rejection probability is 0.063 if the test uses $\widehat{\mathcal{A}}_{mix}$, and it is 0.015 if it does not. The power is about 0.5 for both tests under $\mathbb{A}1$ and about 1 under all other alternatives. Overall, these results show that our test procedure, with centering using $\widehat{\mathcal{A}}_{mix}$, is robust to weak form of cross-sectional dependence in the idiosyncratic risk.

6 Empirical application

We use the developed test to study intraday variation in the linear span of systematic risk exposures of assets. The sample in our study covers the period from January 1, 2015 till December 31, 2021. On each day, we sample the asset prices every five minutes and we exclude the first five minutes of each trading day. The cross-section of assets changes over the calendar years in the sample and at each point in time it consists of the 500 largest stocks by market capitalization as of the end of the previous calendar year. The data is extracted from the TAQ database.

Is the span of systematic risk exposures at market open different from that at market close? To answer this question, we implement our test using return data over two local windows. One window is the first two hours after market open and the other one is the last two trading hours prior to market close. To gain power, as in the Monte Carlo, we aggregate

data over 20 consecutive trading days when implementing the test. The jump truncation parameters and the determination of the number of factors for each return panel is done exactly as in the Monte Carlo.

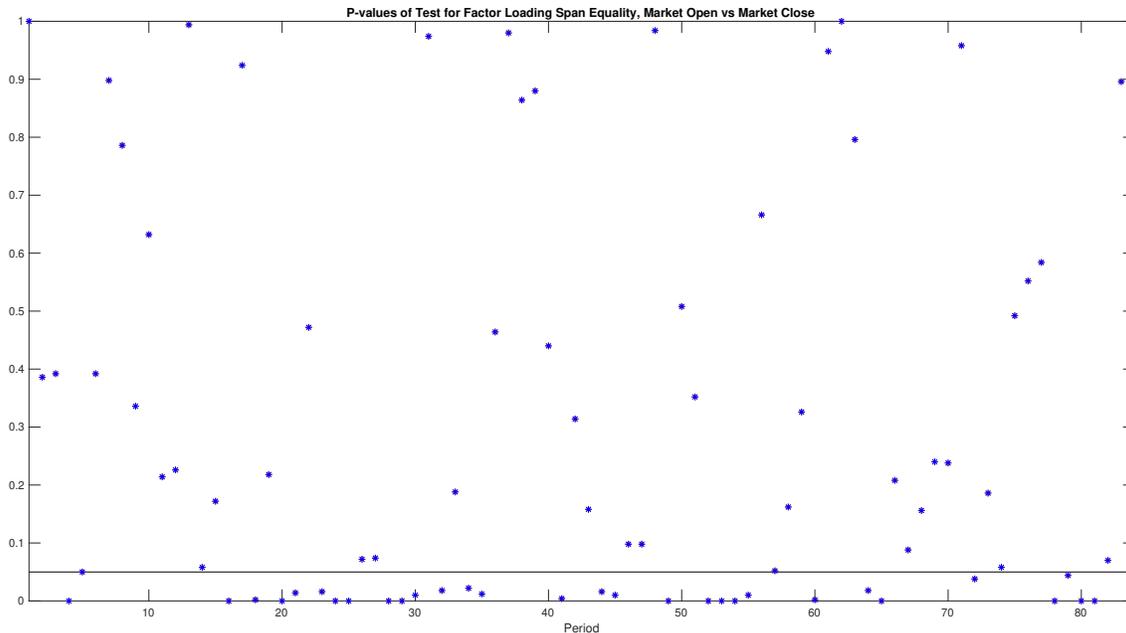


Figure 3: P-values of test for equal linear spans of factor exposures at market open and market close. Each period consists of 20 consecutive trading days. The solid horizontal line is at 0.05.

We report the p-values of the test in Figure 3. As seen from the figure, there is a nontrivial number of periods with very low p-values. More specifically, in 30 out of a total of 84 periods in our sample, our test rejects the null hypothesis at the conventional 5% level. That said, there is a nontrivial number of periods for which there is no evidence for the linear span of the systematic risk exposures of assets changing from market open to market close. We note also that, even though there is some evidence for clustering in time of the low p-values, we can observe low p-values throughout the sample.

We next study the difference in the linear spans of factor exposures at market open and market close on the days of Federal Open Market Committee (FOMC) announcement. We consider only the days for which the announcement happens at 2pm Eastern Time.⁸ We drop the first ten minutes for the market close window and we do the same for the time window at market open. In order to gain power, we group the announcement days by year and we conduct the test on the aggregated by year data. Our test rejects strongly the null

⁸This includes the majority of the announcements in our sample with the exception of the unscheduled ones in 2020.

hypothesis. In particular, at 5% significance level we reject in 3 out of the 7 years in our sample and this rejection increases to 6 out of 7 years for 10% significance level.

We can contrast the above test results for FOMC days with ones when the test is performed on neighboring days. Figure 4 plots the means and standard deviations of the p-values of the test across the seven years in our sample, on the days of FOMC announcement $\pm t$ days, for $t \in \{0, 1, 2\}$. The figure reveals markedly different behavior of the test on and around FOMC days, with p-values of the test on FOMC days significantly lower than the ones for days around the FOMC days. In fact, the mean p-values of the test for the neighbouring days of the FOMC announcement are very close to 0.5, which is the expected value of the test p-value under the null hypothesis.

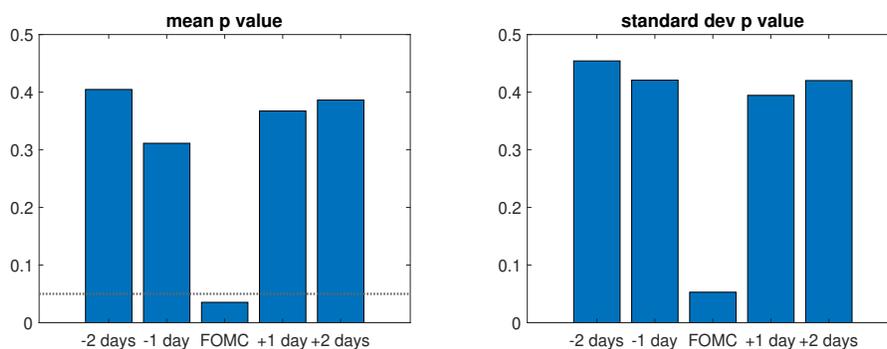


Figure 4: P-values of test for equal linear spans of factor exposures at market open and market close on days around FOMC announcement. In the horizontal axis, “ t days” means the test is conducted for data on the FOMC+ t day. The means and standard deviations are calculated over seven years from 2015 through 2021. The dashed horizontal line is at 0.05.

An interesting question is if the above-documented rejections of the test are due to different factors being present in the different periods compared in the analysis or if they are due to changing exposures to common factors over the two periods. Both scenarios seem economically plausible. Indeed, FOMC announcements trigger a monetary policy shock while at the same time abrupt changes in aggregate macro variables can lead to shifts in betas.⁹ If there is a new factor present only in the hours after the policy announcement, then one would expect a one day change of factor exposures. That is, if we test equality of factor exposures in the afternoons of $t = 0$ against $t = -1$ or $t = 1$, then our test should reject in both cases. On the other hand, if only exposures to factors present throughout the period around the FOMC announcement change and this change persists at least locally, then we would expect that a test for equal factor exposures at $t = 0$ vs. $t = 1$ will not reject while

⁹As mentioned in the introduction, prior work has used macro variables to model dynamics of assets’ exposures to systematic risk.

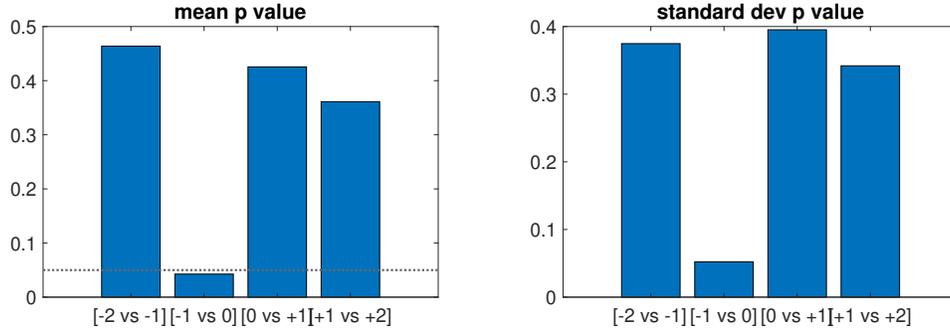


Figure 5: P-values of test for equal linear spans of factor exposures at market close on days around FOMC announcement. In the horizontal axis, “[t_1 vs t_2]” means the test compares factor exposures at market close on the FOMC+ t_1 day and the FOMC+ t_2 day. The means and standard deviations are calculated over seven years from 2015 through 2021. The dashed horizontal line is at 0.05.

that for equal factor exposures at $t = 0$ vs. $t = -1$ will do.

Figure 5 plots the mean p-values over the sample of tests for equal span of factor exposures in the afternoons of days around FOMC announcements. The mean p-value of a test for equal linear span of factor exposures in the afternoons of $t = 0$ and $t = -1$ is 0.04 while the corresponding number for $t = 0$ vs $t = 1$ is 0.43. Based on the reasoning above, these empirical results are more aligned with a shift in the span of factor exposures that persists (at least over one day) than with the appearance of a new factor on the day of the FOMC announcement.

7 Conclusion

We propose a nonparametric test for deciding whether the linear spans of factor exposures of a large cross-section of assets towards latent systematic risk factors at two distinct points in time differ. The test is derived under a joint in-fill and large cross-sectional asymptotic setting, implying that both dimensions of the return panels of high-frequency return observations within local windows of the two points in time are growing. We allow for the two dimensions of the return panels to grow at different rates and impose weak conditions for the dynamics of asset prices. The test is based on the projection discrepancy between the factor loadings estimated separately from the two return panels, which converges asymptotically to zero under the null hypothesis and diverges otherwise. Suitable centering of the statistic is performed to eliminate higher-order asymptotic biases and cross-sectional bootstrap method is developed for feasible implementation. An empirical application of the test reveals that

the linear spans of factor exposures at market open and market close can differ, and the evidence for this is particularly strong on days with FOMC announcements.

References

- Ahn, S. and A. Horenstein (2013). Eigenvalue ratio test for the number of factors. *Econometrica* 81, 1203–1227.
- Aït-Sahalia, Y., J. Jacod, and D. Xiu (2023). Continuous-time fama-macbeth regressions.
- Aït-Sahalia, Y. and D. Xiu (2017). Using principal component analysis to estimate a high dimensional factor model with high-frequency data. *Journal of Econometrics* 201(2), 384–399.
- Andersen, T. G., R. Riva, M. Thyrgaard, and V. Todorov (2023). Intraday cross-sectional distributions of systematic risk. *Journal of Econometrics* 235(2), 1394–1418.
- Andersen, T. G., T. Su, V. Todorov, and Z. Zhang (2023). Intraday periodic volatility curves. *Journal of the American Statistical Association*, 1–11.
- Andersen, T. G., M. Thyrgaard, and V. Todorov (2021). Recalcitrant betas: Intraday variation in the cross-sectional dispersion of systematic risk. *Quantitative Economics* 12(2), 647–682.
- Andrews, D. (2005). Cross-sectional regression with common shocks. *Econometrica* 73, 1551–1585.
- Ang, A. and D. Kristensen (2012). Testing conditional factor models. *Journal of Financial Economics* 106(1), 132–156.
- Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica* 71, 135–171.
- Bai, J., J. Duan, and X. Han (2022). Likelihood ratio test for structural changes in factor models. *arXiv preprint arXiv:2206.08052*.
- Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica* 70(1), 191–221.
- Baltagi, B. H., C. Kao, and F. Wang (2017). Identification and estimation of a large factor model with structural instability. *Journal of Econometrics* 197(1), 87–100.

- Barndorff-Nielsen, O. E. and N. Shephard (2004a). Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica* 72(3), 885–925.
- Barndorff-Nielsen, O. E. and N. Shephard (2004b). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2(1), 1–37.
- Bollerslev, T., S. Z. Li, and V. Todorov (2016). Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns. *Journal of Financial Economics* 120(3), 464–490.
- Breitung, J. and S. Eickmeier (2011). Testing for structural breaks in dynamic factor models. *Journal of Econometrics* 163(1), 71–84.
- Campbell, J. Y. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review* 94(5), 1249–1275.
- Chen, L., J. J. Dolado, and J. Gonzalo (2014). Detecting big structural breaks in large factor models. *Journal of Econometrics* 180(1), 30–48.
- Cheng, X., Z. Liao, and F. Schorfheide (2016). Shrinkage estimation of high-dimensional factor models with structural instabilities. *The Review of Economic Studies* 83(4), 1511–1543.
- Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton university press.
- Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of finance* 66(4), 1047–1108.
- Connor, G., M. Hagmann, and O. Linton (2012). Efficient semiparametric estimation of the fama–french model and extensions. *Econometrica* 80(2), 713–754.
- Connor, G. and R. A. Korajczyk (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics* 15(3), 373–394.
- Corradi, V. and N. R. Swanson (2014). Testing for structural stability of factor augmented forecasting models. *Journal of Econometrics* 182(1), 100–118.
- Eisfeldt, A. L., E. Kim, and D. Papanikolaou (2022). Intangible value. *Critical Finance Review* 11, 299–332.

- Embrechts, P., C. Klüppelberg, and T. Mikosch (2013). *Modelling extremal events: for insurance and finance*, Volume 33. Springer Science & Business Media.
- Fama, E. F. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy* 81(3), 607–636.
- Fan, J., Y. Liao, and W. Wang (2016). Projected principal component analysis in factor models. *Annals of Statistics* 44(1), 219–254.
- Feng, G., S. Giglio, and D. Xiu (2020). Taming the factor zoo: A test of new factors. *The Journal of Finance* 75(3), 1327–1370.
- Ferson, W. E. and C. R. Harvey (1999). Conditioning variables and the cross section of stock returns. *The Journal of Finance* 54(4), 1325–1360.
- Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica* 84(3), 985–1046.
- Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy* 129(7), 1947–1990.
- Hallin, M. and R. Liška (2007). Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association* 102, 603–617.
- Hansen, L. P. and S. F. Richard (1987). The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* 55(3), 587–613.
- Harvey, C. R., Y. Liu, and H. Zhu (2016). ... and the cross-section of expected returns. *The Review of Financial Studies* 29(1), 5–68.
- Jacod, J., H. Lin, and V. Todorov (2022). Systematic jump risk.
- Jacod, J. and P. Protter (2012). *Discretization of Processes*. Berlin: Springer-Verlag.
- Jacod, J. and A. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.
- Jagannathan, R. and Z. Wang (1996). The conditional capm and the cross-section of expected returns. *The Journal of finance* 51(1), 3–53.
- Kalnina, I. (2023). Inference for nonparametric high-frequency estimators with an application to time variation in betas. *Journal of Business & Economic Statistics* 41, 538–549.

- Kapetanios, G. (2010). A testing procedure for determining the number of factors in approximate factor models with large datasets. *Journal of Business & Economic Statistics* 28(3), 397–409.
- Kelly, B. T., S. Pruitt, and Y. Su (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134(3), 501–524.
- Lintner, J. (1965a). Security prices, risk, and maximal gains from diversification. *The journal of finance* 20(4), 587–615.
- Lintner, J. (1965b). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics* 47, 13–37.
- Mancini, C. (2001). Disentangling the jumps of the diffusion in a geometric jumping brownian motion. *Giornale dell'Istituto Italiano degli Attuari* 64(19-47), 44.
- Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society* 5, 867–887.
- Mykland, P. A. and L. Zhang (2006). Anova for diffusions and ito processes. *The Annals of Statistics* 34(4), 1931–1963.
- Mykland, P. A. and L. Zhang (2009). Inference for continuous semimartingales observed at high frequency. *Econometrica* 77, 1403–1445.
- Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics* 92(4), 1004–1016.
- Onatski, A. (2012). Asymptotics of the principal components estimator of large factor models with weakly influential factors. *Journal of Econometrics* 168(2), 244–258.
- Pelger, M. (2019). Large-dimensional factor modeling based on high-frequency observations. *Journal of Econometrics* 208(1), 23–42.
- Pelger, M. (2020). Understanding systematic risk: a high-frequency approach. *The Journal of Finance* 75(4), 2179–2220.
- Reiß, M., V. Todorov, and G. Tauchen (2015). Nonparametric test for a constant beta between Itô semi-martingales based on high-frequency data. *Stochastic Processes and their Applications* 125(8), 2955–2988.
- Shanken, J. (1990). Intertemporal asset pricing: An empirical investigation. *Journal of Econometrics* 45(1-2), 99–120.

- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19(3), 425–442.
- Stock, J. H. and M. W. Watson (2002). Forecasting using principal components from a large number of predictors. *Journal of the American statistical association* 97(460), 1167–1179.
- Su, L. and X. Wang (2020). Testing for structural changes in factor models via a nonparametric regression. *Econometric Theory* 36(6), 1127–1158.
- TAQ (2022). NYSE Trade and Quote (TAQ) database.
- Yamamoto, Y. and S. Tanaka (2015). Testing for factor loading structural change under common breaks. *Journal of Econometrics* 189(1), 187–206.