Intraday Cross-Sectional Distributions of Systematic Risk[*]

Torben G. Andersen[†]   Raul Riva[‡]   Martin Thyrsgaard[§]   Viktor Todorov[¶]

November 28, 2021

Abstract

We develop a test for the detection of intraday changes in the cross-sectional distribution of assets' exposure to observable factors. The test is constructed for a panel of high-frequency asset returns, with the size of the cross-section and the sampling frequency increasing simultaneously. It is based on a comparison of the empirical characteristic functions of estimates of the assets' factor loadings at different parts of the trading day, formed from local blocks of asset returns and the corresponding factor realizations. The limiting behavior of the test statistic is governed by unobservable latent factors in the asset prices. The critical values of the test are constructed on the basis of a novel simulation-based procedure. Empirical implementation of the test to stocks in the S&P 500 index and the five Fama-French factors, as well as the momentum factor, reveals different intraday behavior of the factor loadings: assets' exposure to size, market and value risks vary systematically over the trading day while the three remaining factors do not exhibit statistically significant intraday variation. Finally, we document that time-varying correlations between the observable risk factors drive a wedge between the time-of-day pattern of market betas, estimated with and without control for the other observable risk factors.

Keywords: asset pricing, dynamic factor models, empirical characteristic function, high-frequency data, nonparametric inference, stable convergence.

JEL classification: C51, C52, G12.

1 Introduction

Estimating factor loadings, or betas, is an integral part of the empirical evaluation of asset pricing models. A standard approach, see, e.g., Cochrane (2009), is to assume that betas remain constant for a limited time window, typically several quarters or years, and to estimate them through a time series regression of excess stock returns on the set of observable factors inside the given window. This approach, however, ignores potential time variation in the factor loadings across the estimation window. One way to circumvent this problem is to model the time-variation of the assets' betas as functions of observable characteristics or macroeconomic variables, see, e.g., Shanken (1990), Jagannathan and Wang (1996), Connor et al. (2012), Fan et al. (2016), Gagliardini et al. (2016) and Kelly et al. (2019), among many others. An alternative is to take advantage of the availability of high-frequency intraday return data. As is well known, unlike the estimation of the drift term of a (multivariate) continuous process, the estimation of the second order moments is improved by the use of a higher sampling frequency. As a result, assets' betas can be estimated reliably from high-frequency data, even over very short intervals such as a day or a week, see, e.g., Barndorff-Nielsen and Shephard (2004).

Estimation of factor loadings from high-frequency data relies explicitly either on these quantities being constant during a trading day or on ignoring any such time variation. As a result, daily estimates of betas from high-frequency data may be computed either as ratios of daily integrated covariances and variances (as in Barndorff-Nielsen and Shephard (2004)) or as sums of local intraday beta estimates (as in Mykland and Zhang (2009)). However, Andersen et al. (2021) show that market betas vary in a systematic way during a trading day. More specifically, they document that the cross-sectional dispersion of market betas is high at market open and decays monotonically during the trading day, reaching its minimum at market close. This finding appears at odds with the standard approach of modeling time-variation in betas, based on quantities that evolve at a lower frequency. The analysis in Andersen et al. (2021) is for univariate market betas, i.e., the only observable factor in their analysis is the market portfolio. If there are additional systematic risk factors, which are partially correlated with the market portfolio, then one can potentially rationalize the finding in Andersen et al. (2021) through a shift in the volatility and covariance of the systematic factors over the course of the trading day. This is what we aim to investigate in this paper. More specifically, we derive a test for changes in the cross-sectional distribution of the assets' factor loadings in a setting with multiple observable factors in addition to unobservable or latent systematic risk factors in the asset prices.

To motivate our exploration, we provide an indication of the potential systematic intraday inter-

action between the market factor and five popular non-market factors, namely the four additional Fama and French (2015) factors and the momentum factor. Specifically, Figure 1 depicts estimates of the correlation between the market and each of these non-market factors across the trading day. The source of the underlying high-frequency data and the composition of the factors are explained in detail in Section 5.

Figure 1: **Correlation between the market and non-market risk factors**. We use rolling one-hour windows to compute the correlation of each factor with the market, pooling data across the full sample. Factor returns are at the 5-minute frequency, from 2010 until 2017. All factor series are obtained from Aït-Sahalia et al. (2020).
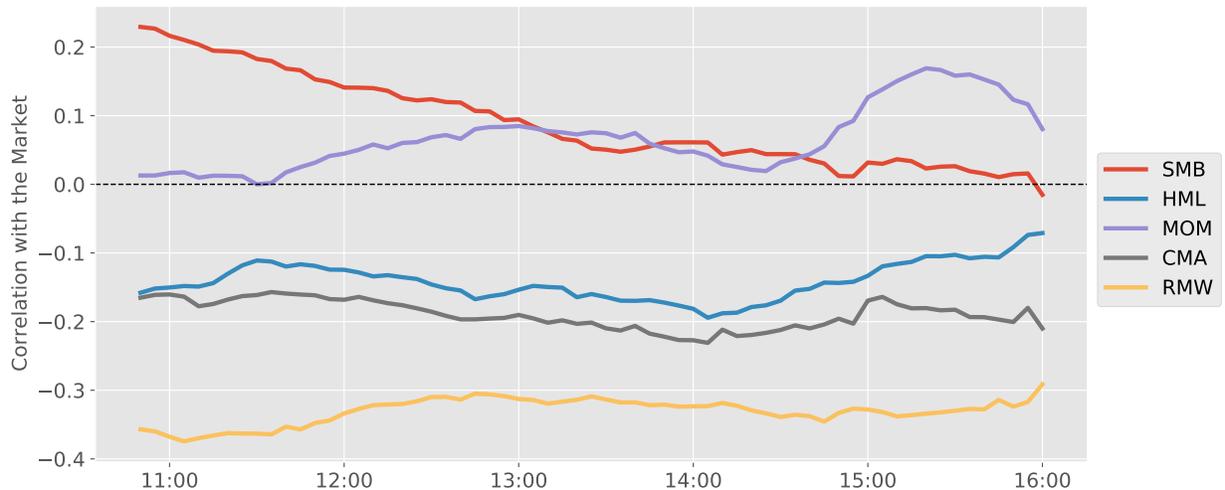
As Figure 1 shows, a number of the non-market factors display a modest and time-varying degree of correlation with the market across the trading day. The signs and relative magnitudes of these correlations are fairly well aligned with the numbers reported in Table 4 of Fama and French (2015), who relies on monthly data covering a much longer sample, 1963-2013. Likewise, we deem the mild positive correlation of the market with momentum to be consistent with prior evidence.[1]

Interestingly, the reversion towards zero of the market correlation for the SMB and HML factors across the trading day is reminiscent of the contractionary pattern in the market beta dispersion reported in Andersen et al. (2021), suggesting a potential relationship. In contrast, the (unconditional) market correlations of the CMA and RMW factors appear fairly stable across the trading day, while the momentum factor displays a higher correlation with the market towards the close of trading.

---

[1]However, this is not a straightforward assessment. The momentum factor is known to display extreme fluctuations, with the market correlation turning sharply negative during the so-called momentum crashes, see, e.g., Daniel and Moskowitz (2016). Thus, the moderately negative unconditional correlation reported in the literature is generated through averaging of a few starkly negative correlation episodes and the more modest correlation in regular scenarios. Since our data is void of momentum crashes, we deem a mild positive correlation to be consistent with prior findings.

The above evidence suggests that time-varying correlation between the market factor and other observable risk factors can play a role in explaining the observed pattern in univariate market betas in Andersen et al. (2021). Our theoretical developments in this paper provide a formal basis for exploring this. The analysis is cast within a joint asymptotic setting for which the sampling frequency increases along with the size of the cross-section of the return panel. Using standard linear regressions of asset returns on observable factors in local windows of time, we begin by estimating the factor loadings during the trading day. Using these noisy beta proxies, we then construct an estimate of the characteristic function of the cross-sectional distribution of the factor loadings for different points in time during a trading day. We derive a functional Central Limit Theorem (CLT) for this quantity in a complex-valued weighted $\mathcal{L}^2$ Hilbert space. For each cross-sectional estimate during the trading day, the limiting distribution is a mixed-Gaussian process and is determined by latent (unobservable) systematic risk factors in the asset prices. Moreover, the limiting processes of these cross-sectional estimates are $\mathcal{F}$-conditionally independent across different points in time, for $\mathcal{F}$ being the σ-algebra of the probability space.

Using this limiting result, we propose a test for the null hypothesis of no difference in the cross-sectional distribution in the factor loadings at distinct points during a trading day. The test is based on the distance in the Hilbert space of the empirical characteristic functions of the factor loadings during two different points in time. The limiting distribution of the test statistic is non-standard and we develop a novel simulation-based approach for computing the critical values of the test. In particular, we generate independent random variables, with mean zero and variance of one, for each increment in the local window used in the factor loading estimation. These random variables are aimed at mimicking the shocks in the unobservable latent systematic risks in the asset prices, the dimension of which is not known to the econometrician. These simulated random variables are then weighted appropriately using suitable transforms of the observed return data to match the $\mathcal{F}$-conditional variance of our empirical characteristic functions of the factor loadings.

We evaluate the finite sample performance of our test on simulated data from an asset pricing model calibrated to match key features of the data in our empirical study. We implement the test on 5-minute data on stocks in the S&P 500 index during the period 2010 to 2017. As observable factors, we use the five Fama-French risk factors plus momentum, constructed as high-frequency time-series by Aït-Sahalia et al. (2020). We find that the intraday variation in the loadings on the market factor is more pronounced than for other observable factors. We further analyze if large macroeconomic announcements during the trading day, such as FOMC announcements, change the distribution of factor loadings from Open to Close and find this to be the case for some of the risk factors.

The current paper relates to several strands of existing work. First, there has been a lot of work on estimating betas/factor loadings from high-frequency data, see e.g., Barndorff-Nielsen and Shephard (2004), Andersen et al. (2005a,b), Mykland and Zhang (2009), Li et al. (2017) and Aït-Sahalia et al. (2020). Second, Reiß et al. (2015), Kalnina (2020) and Zhang et al. (2020) develop tests for detecting time-variation in betas using high-frequency data. Unlike these strands of work, our interest here is cross-sectional distributions of betas. As a result, unlike the above-mentioned papers, the size of the cross-sectional dimension of the return panel in our paper is also growing asymptotically (together with the increase in the sampling frequency). Another difference, due to our focus on cross-sectional beta distributions, is that we need to consider functional convergence in our analysis. Third, the current paper is related to Andersen et al. (2021) which study cross-sectional dispersion of market betas. The analysis in Andersen et al. (2021) is for a univariate observable factor while in our case we allow for a multivariate systematic risk factor. This has nontrivial consequences empirically as discussed above. In addition, our focus here is on the cross-sectional distribution of betas while the main focus of Andersen et al. (2021) is the cross-sectional second moment (dispersion) of market betas. This difference also leads to the development here of a novel and easy-to-implement method for conducting feasible functional inference that avoids knowledge of the dimension of the latent systematic risk factor.

The rest of the paper is organized as follows. The setup and notation are presented in Section 2. Section 3 develops our estimator for the characteristic function of the cross-sectional distribution of factor exposures at a point in time and derives functional feasible limit theory for it. This result is used to develop a test for intraday changes in the cross-sectional distribution of factor loadings. Section 4 contains a Monte Carlo study and Section 5 our empirical application. Section 6 concludes. Assumptions and proofs are given in the Appendix.

2 Setup and Notation

We first introduce the setup. We consider a set of stocks, indexed by $j = 1, ..., N$, whose log-prices, $X^{(j)}$, are defined on some filtered probability space $\left(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P}\right)$. The stock return dynamics is impacted by q observable risk factors, collected into the vector, F, as well as r (unknown) latent factors. The value of these processes at time t is generated by the following system,

$$F_t \ = \ F_0 + \int_0^t \alpha_s \, ds \ + \int_0^t \sigma_s \, dW_s \ + \sum_{s \leq t} \Delta F_s, \tag{1}$$

$$
\begin{aligned}
X_t^{(j)} \;=\; & X_0^{(j)} + \int_0^t \alpha_s^{(j)}\, ds + \int_0^t \beta_s^{(j)\top}\sigma_s\, dW_s + \int_0^t \gamma_s^{(j)}\, dB_s \\
& + \int_0^t \widetilde{\sigma}_s^{(j)}\, d\widetilde{W}_s^{(j)} + \sum_{s \le t} \Delta X_s^{(j)}, \qquad j = 1,..,N,
\end{aligned}
\tag{2}
$$

where W_t, $(\widetilde{W}_t^{(j)})_{j=1,\ldots,N}$ and B_t are independent Brownian motions, the dimension of W_t and B_t is $q \times 1$ and $r \times 1$, respectively, while the remaining Brownian motions are univariate; α_t, $\{\alpha_t^{(j)}\}_{j=1,\ldots,N}$, σ_t, $\{\beta_t^{(j)}\}_{j=1,\ldots,N}$, $\{\widetilde{\sigma}_t^{(j)}\}_{j=1,\ldots,N}$ and $\{\gamma_t^{(j)}\}_{j=1,\ldots,N}$ are processes with càdlàg paths, with α_t and $\beta_t^{(j)}$ being of dimension $q \times 1$, σ_t is $q \times q$, $\gamma_t^{(j)}$ is $1 \times r$ dimensional, representing the loadings on the latent factors, and the remaining processes are scalar-valued functions of time. Finally, for any process Y, $\Delta Y_t = Y_t - Y_{t-}$ denotes the size of a jump at time t.

The above setup is quite general and accommodates most asset pricing models used in prior work. We allow for the presence of both observable and unobservable (latent) systematic factors, with the latter driven by the Brownian motion B_t, and we note that the dimension of that vector is arbitrary, but fixed. In our setup, we do not impose any structure on the cross-sectional dependence between the assets' jumps and, consequently, we avoid using any information regarding price jumps extracted from the data in our inference procedures.

Our focus in this paper is on the cross-sectional distribution of the factor loading vector, $\beta_t^{(j)}$, along the trading day. We let the unit of time be one trading day, and denote a fraction of the trading day by $\kappa \in [0,1]$. Thus, for $t \in \mathbb{N}_+$ and $\kappa \in [0,1]$, we will use the shorthand notation $\beta_{t,\kappa}^{(j)} = \beta_{t+\kappa}^{(j)}$, and we obtain the loading of asset j on factor k at time κ during day t as $\beta_{t-1,\kappa}^{(j,k)} = \iota_k^\top \beta_{t-1,\kappa}^{(j)}$, where ι_k is the k'th unit vector of the standard basis of $\mathbb{R}^q$, i.e., a vector of zeros except for one in the k'th entry. In most existing work, either explicitly or implicitly, these factor loadings are assumed to be constant over short intervals of time, such as weeks or months, or to be functions of firm-specific characteristics or macroeconomic variables that only exhibit low-frequency dynamics.

3 Inference for Cross-Sectional Distributions of Factor Exposures

In this section, we propose an estimator for the characteristic function of the cross-sectional distribution of the factor exposures at fixed points in time, establish a feasible limit theory for the estimator, and develop a test for differences in the cross-sectional distribution of factor exposures at different time points.

3.1 Cross-Sectional Empirical Characteristic Functions of Factor Exposures

Our inference is based on discrete observations of $\{X^{(j)}\}_{j=1,...,N}$ and $\{F^{(k)}\}_{k=1,...,q}$ at equidistant times $0, \frac{1}{n}, \frac{2}{n}, ..., T$, where T refers to the time span of our data, which is fixed throughout, and the integer n is the number of times we sample within each trading day. We denote the length of the sampling interval by $\Delta_n = 1/n$. The high-frequency increment of the stocks and observable factors are denoted, for $t \in \mathbb{N}_+$,

$$\Delta_{t,i}^n X^{(j)} = X_{(t-1)+i/n}^{(j)} - X_{(t-1)+(i-1)/n}^{(j)} \quad \text{and} \quad \Delta_{t,i}^n F^{(k)} = F_{(t-1)+i/n}^{(k)} - F_{(t-1)+(i-1)/n}^{(k)},$$

where $i = 1, ..., n$, $j = 1, ..., N$ and $k = 1, ..., q$. Our estimation procedure for the cross-sectional distribution of the factor exposures starts with estimating the latter from the high-frequency returns over local blocks of size k_n, where,

$$k_n \asymp n^\varrho, \quad \varrho \in (0, 1). \tag{3}$$

This implies that $k_n \to \infty$ with $k_n/n \to 0$, as $n \to \infty$. The index of returns used to compute the factor exposures at a given point $\kappa \in [0, 1]$ on a given trading day is then given by,

$$\mathcal{I}_\kappa^n = \{\lfloor \kappa n \rfloor - k_n + 1, ..., \lfloor \kappa n \rfloor\}, \quad \kappa \in [0, 1].$$

To filter out jumps along the sample paths, we adopt the truncation approach of Mancini (2001). In particular, we define a truncation parameter for each asset and the vector of factors by,

$$\nu_{t,n}^{(j)} \asymp \Delta_n^\varpi, \quad \nu_{t,n} \asymp \Delta_n^\varpi, \quad \varpi \in (0, 1/2), \quad t \geq 1, \ j = 1, ..., N.$$

To simplify the notation in what follows, we introduce the following sets,

$$\mathcal{A}_{t,i} = \{\|\Delta_{t,i}^n F\| \leq \nu_{t,n}\}, \quad \mathcal{B}_{t,i}^{(j)} = \{|\Delta_{t,i}^n X^{(j)}| \leq \nu_{t,n}^{(j)} \cap \|\Delta_{t,i}^n F\| \leq \nu_{t,n}\}, \tag{4}$$

for $j = 1, ..., N$, where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^q$. To estimate the individual factor exposures for a stock, i.e., the $\beta^{(j)}$s, we need estimates of the variance-covariance matrix of the factors as well as the covariance between the factors and the asset return. These are given by,

$$\widehat{V}_{t,\kappa} = \frac{n}{k_n} \sum_{i \in \mathcal{I}_\kappa^n} \Delta_{t,i}^n F \Delta_{t,i}^n F^\top 1_{\{\mathcal{A}_{t,i}\}}, \quad \widehat{C}_{t,\kappa}^{(j)} = \frac{n}{k_n} \sum_{i \in \mathcal{I}_\kappa^n} \Delta_{t,i}^n X^{(j)} \Delta_{t,i}^n F 1_{\{\mathcal{B}_{t,i}^{(j)}\}}, \tag{5}$$

and we then define,

$$\widehat{\beta}_{t,\kappa}^{(j,k)} = \iota_k^\top \, \widehat{V}_{t,\kappa}^{-1} \, \widehat{C}_{t,\kappa}^{(j)}. \tag{6}$$

For a fixed j, using standard results in high-frequency asymptotics, see e.g., Theorem 9.3.2 in Jacod and Protter (2012), we have $\widehat{\beta}_{t,\kappa}^{(j,k)} \xrightarrow{\mathbb{P}} \iota_k^\top \beta_{t,\kappa}^{(j)}$. Our interest here, however, is the cross-sectional distribution of the factor exposures.

Estimated factor loadings based on intraday data covering a single trading day can be rather noisy. For this reason, we also consider estimates that pool observations across trading days. Specifically, we denote a finite set of trading days by $\mathcal{T}$, and define,

$$\widehat{C}_{\mathcal{T},\kappa}^{(j)} = \frac{1}{T} \sum_{t \in \mathcal{T}} \widehat{C}_{t,\kappa}^{(j)}, \quad \widehat{V}_{\mathcal{T},\kappa} = \frac{1}{T} \sum_{t \in \mathcal{T}} \widehat{V}_{t,\kappa}, \quad \widehat{\beta}_{\mathcal{T},\kappa}^{(j,k)} = \iota_k^\top \, \widehat{V}_{\mathcal{T},\kappa}^{-1} \, \widehat{C}_{\mathcal{T},\kappa}^{(j)}. \tag{7}$$

Note that, in this aggregation, we utilize intraday returns from different trading days, but they are all sampled at the same time of day. This is motivated by our focus on the evolution of the cross-sectional distribution of systematic risk exposures viewed as a function of time-of-day. In analogy to the single-day case, Theorem 9.3.2 in Jacod and Protter (2012) implies,

$$\widehat{\beta}_{\mathcal{T},\kappa}^{(j,k)} \xrightarrow{\mathbb{P}} \beta_{\mathcal{T},\kappa}^{(j,k)} \equiv \iota_k^\top \left(\sum_{t \in \mathcal{T}} \sigma_{t+\kappa} \, \sigma_{t+\kappa}^\top \right)^{-1} \sum_{t \in \mathcal{T}} \sigma_{t+\kappa} \, \sigma_{t+\kappa}^\top \, \beta_{t+\kappa}^{(j)}.$$

If $\mathcal{T}$ is a set of consecutive trading days covering a relatively short time window (e.g., one month), it is natural to expect $\beta_{t+\kappa}^{(j)}$ to be constant for $t \in \mathcal{T}$. If this is not the case, however, then $\beta_{\mathcal{T},\kappa}^{(j,k)}$ is a weighted average of the betas over the trading days in $\mathcal{T}$. Hence, in that scenario, the inference procedures developed below are valid for these weighted averages of the factor loadings.

To quantify the cross-sectional distribution of $\beta_{\mathcal{T},\kappa}^{(j,k)}$ at a specific point in time, we rely on characteristic functions, as they are a distribution-determining class. In particular, we denote,

$$\mathcal{L}_{\mathcal{T},\kappa,k}^N(u) = \frac{1}{N} \sum_{j=1}^N \exp\left(iu \, \beta_{\mathcal{T},\kappa}^{(j,k)} \right), \quad \kappa \in [0,1], \quad k = 1,..,q, \quad u \in \mathbb{R}. \tag{8}$$

When $N \to \infty$, under standard conditions, $\mathcal{L}_{\mathcal{T},\kappa,k}^N(u)$ converges to the characteristic function of the cross-sectional distribution of $\beta_{\mathcal{T},\kappa}^{(j,k)}$. Of course, $\mathcal{L}_{\mathcal{T},\kappa,k}^N(u)$ is infeasible, but we can estimate it as follows,

$$\widehat{\mathcal{L}}_{\mathcal{T},\kappa,k}^N(u) = \frac{1}{N} \sum_{j=1}^N \exp\left(iu \, \iota_k^\top \, \widehat{V}_{\mathcal{T},\kappa}^{-1} \, \widehat{C}_{\mathcal{T},\kappa}^{(j)} \right), \quad t \geq 1, \quad \kappa \in [0,1], \quad k = 1,...,q, \quad u \in \mathbb{R}. \tag{9}$$

3.2 Feasible Limit theory

One can show that the difference $\widehat{\mathcal{L}}^N_{\mathcal{T},\kappa,k}(u) - \mathcal{L}^N_{\mathcal{T},\kappa,k}(u)$ is asymptotically negligible in a functional sense, as both $n \to \infty$ and $N \to \infty$. Our goal here is to derive an associated feasible Central Limit Theorem to quantify the size of this gap. The limit theory for our estimator of the empirical characteristic function is derived in the complex-valued Hilbert space $\mathcal{L}^2(w)$, defined as,

$$\mathcal{L}^2(w) = \left\{ f : \mathbb{R} \to \mathbb{C} \int_{\mathbb{R}} |f(u)|^2 w(u) du < \infty \right\}, \tag{10}$$

where w is a continuous weight function with exponential tail decay. We denote the norm associated with this Hilbert space by $\| \cdot \|_w$. We have the following limiting result in which $\xrightarrow{\mathcal{L}-s}$ denotes stable convergence in law, implying that the convergence holds jointly with any $\mathcal{F}$-measurable random variable, see, e.g., Jacod and Shiryaev (2003).

Theorem 1. *Let $\kappa \in [0,1]$ be given and let $\mathcal{T} \subset \mathbb{N}$ be a fixed set, such that $|\mathcal{T}| < \infty$. Suppose Assumptions A and B in the appendix hold. As $N \to \infty$ and $n \to \infty$, with $\omega \in (1/4, 1/2)$ and $\varrho \in (1 - 2\omega, 1/2)$, we have,*

$$\sqrt{k_n} \left(\widehat{\mathcal{L}}^N_{\mathcal{T},\kappa,k} - \mathcal{L}^N_{\mathcal{T},\kappa,k} \right) \xrightarrow{\mathcal{L}-s} \mathfrak{Z}_{\mathcal{T},\kappa,k}, \tag{11}$$

where $\mathfrak{Z}_{\mathcal{T},\kappa,k}$ is a $\mathcal{F}$-conditional mean-zero complex normal process in $\mathcal{L}^2(w)$, with covariance and relations operators given by,

$$
\begin{aligned}
\Psi^{(1)}_{\mathcal{T},\kappa,k} h(z) &= \int_{\mathbb{R}} \psi^{(1)}_{\mathcal{T},\kappa,k}(u, z) \, h(u) \, w(u) \, du \,, \\
\Psi^{(2)}_{\mathcal{T},\kappa,k} h(z) &= \int_{\mathbb{R}} \psi^{(2)}_{\mathcal{T},\kappa,k}(u, z) \, h(u) \, w(u) \, du \,,
\end{aligned}
$$

for all $h \in \mathcal{L}^2(w)$, with the kernels $\psi^{(1)}_{\mathcal{T},\kappa,k}$ and $\psi^{(2)}_{\mathcal{T},\kappa,k}$ provided in the appendix. The limiting processes $\mathfrak{Z}_{\mathcal{T},\kappa,k}$, for different values of κ, are $\mathcal{F}$-conditionally independent.

This limiting result warrants several comments. First, the convergence is functional in the characteristic exponent u, but it is finite-dimensional in the time-of-day parameter κ. The rate of convergence is determined by the size of the local block, k_n. The latter needs to be of smaller order than $\sqrt{n}$, so that biases in estimation stemming from the time-variation of the stochastic variance and factor exposures are of higher asymptotic order. Finally, the limit process is determined by the presence of latent systematic factors, that are orthogonal (in a martingale sense) to the observable factors captured by F.

We now turn to the issue of feasible inference. For this, we make use of a simple simulation-based procedure. It involves a sequence of i.i.d. random variables with mean zero and variance one, denoted $(e_{i,t,\kappa})_{i\geq 1}$. For each draw of $(e_{i,t,\kappa})_{i\geq 1}$, we compute,

$$\widehat{W}_{\mathcal{T},\kappa,k}^{N*}(u) \;=\; \frac{iu}{N} \sum_{j=1}^{N} \exp\left(iu\, \iota_k^\top \widehat{\beta}_{\mathcal{T},\kappa}^{(j)}\right) \widetilde{Z}_{t,\kappa,k}^{(j)}, \tag{12}$$

$$\widetilde{Z}_{t,\kappa,k}^{(j)} \;=\; \frac{n}{k_n|\mathcal{T}|} \iota_k^\top \widehat{V}_{\mathcal{T},\kappa}^{-1} \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{I}_\kappa^n} \left(\Delta_{t,s}^n X^{(j)} - \widehat{\beta}_{\mathcal{T},\kappa}^{(j)\top} \Delta_{t,s}^n F\right) \Delta_{t,s}^n F \; 1_{\{\mathcal{B}_{t,s}^{(j)}\}} \, e_{s,t,\kappa}. \tag{13}$$

The next theorem shows that $\widehat{W}_{\mathcal{T},\kappa,k}^{N*}(u)$ has the same $\mathcal{F}$-conditional limiting distribution as that of $\mathfrak{Z}_{\mathcal{T},\kappa,k}$ in Theorem 1.

Theorem 2. *Under the conditions of Theorem 1, we have,*

$$\sqrt{k_n}\, \widehat{W}_{\mathcal{T},\kappa,k}^{N*} \;\xrightarrow{\mathcal{L}|\mathcal{F}}\; \mathfrak{Z}_{\mathcal{T},\kappa,k}, \tag{14}$$

where $\mathfrak{Z}_{\mathcal{T},\kappa,k}$ is the $\mathcal{F}$-conditional Gaussian process defined in Theorem 1.

The idea behind the approximation of the $\mathcal{F}$-conditional distribution of $\mathfrak{Z}_{\mathcal{T},\kappa,k}$ by that of $\sqrt{k_n}\, \widehat{W}_{\mathcal{T},\kappa,k}^{N*}$ is the following. The random variables $e_{s,t,\kappa}$ are independent of j and are aimed at mimicking the systematic diffusive shocks that determine $\mathfrak{Z}_{\mathcal{T},\kappa,k}$. The weights assigned to each $e_{s,t,\kappa}$ in the summation in equation (12) are $\mathcal{F}$-adapted and designed to ensure that the $\mathcal{F}$-conditional variance of $\sqrt{k_n}\, \widehat{W}_{\mathcal{T},\kappa,k}^{N*}$ converges to that of $\mathfrak{Z}_{\mathcal{T},\kappa,k}$.

The result of Theorem 2 renders inference feasible by simply generating many copies of $\widehat{W}_{\mathcal{T},\kappa,k}^{N*}(u)$, and then exploiting the resulting empirical distribution of $\widehat{W}_{\mathcal{T},\kappa,k}^{N*}$. Importantly, from a practical perspective, we do not need to know the dimension r of the latent systematic risk captured by the Brownian motion B to conduct feasible inference on the basis of $\widehat{W}_{\mathcal{T},\kappa,k}^{N*}$.

3.3 Testing for Changes in the Cross-Sectional Distribution of Factor Loadings

Using the feasible limit theory, we can test for equal cross-sectional distribution of the factor loadings at two different points κ and κ' of the trading day. The natural null hypothesis is that these two distributions are identical. Letting $\mathcal{L}_{\mathcal{T},\kappa,k}(u) = \plim_{N\to\infty} \mathcal{L}_{\mathcal{T},\kappa,k}^N(u)$, we define the null and alternative hypotheses through the following partition of the sample probability space,

$$\Omega^{(0)} \;=\; \left\{\omega : \|\mathcal{L}_{\mathcal{T},\kappa,k} - \mathcal{L}_{\mathcal{T},\kappa',k}\|_w = 0\right\} \quad \text{and} \quad \Omega^{(a)} \;=\; \left\{\omega : \|\mathcal{L}_{\mathcal{T},\kappa,k} - \mathcal{L}_{\mathcal{T},\kappa',k}\|_w \neq 0\right\}, \tag{15}$$

where $\|\cdot\|_w$ is the norm associated with $\mathcal{L}^2(w)$. Our goal is to design a test that allows us to decide to which of the two sets the sample path belongs. In developing the test, we will assume that the following is true,

$$\| \mathcal{L}^N_{\mathcal{T},\kappa,k} - \mathcal{L}^N_{\mathcal{T},\kappa',k} \|_w = o_p(1/\sqrt{k_n}), \quad \text{under the restriction to } \Omega^{(0)}. \tag{16}$$

The above restriction amounts to assuming that the error in measuring $\mathcal{L}_{\mathcal{T},\kappa,k} - \mathcal{L}_{\mathcal{T},\kappa',k}$ from the infeasible $\mathcal{L}^N_{\mathcal{T},\kappa,k} - \mathcal{L}^N_{\mathcal{T},\kappa',k}$, stemming from the cross-sectional variation in the factor loadings, is of higher asymptotic order than the error in $\widehat{\mathcal{L}}^N_{\mathcal{T},\kappa,k} - \widehat{\mathcal{L}}^N_{\mathcal{T},\kappa',k}$ due to the estimation of the factor loadings. Under standard conditions, e.g., when drawing at random from the cross-sectional distribution of stocks, $\|\mathcal{L}^N_{\mathcal{T},\kappa,k} - \mathcal{L}^N_{\mathcal{T},\kappa',k}\|_w = O_p(1/\sqrt{N})$, and equation (16) will hold, as long as $k_n/N \to 0$, which is a fairly weak condition on the relative size of the two dimensions of the return panel. Furthermore, in the typical situation (under the null), that the factor loadings remain constant across the trading day, $\| \mathcal{L}^N_{\mathcal{T},\kappa,k} - \mathcal{L}^N_{\mathcal{T},\kappa',k} \|_w$ will be exactly zero.

Turning next to the test, we rely on the following statistic,

$$TS_{\mathcal{T},\kappa,\kappa',k} = k_n \| \widehat{\mathcal{L}}^N_{\mathcal{T},\kappa,k} - \widehat{\mathcal{L}}^N_{\mathcal{T},\kappa',k} \|_w. \tag{17}$$

Stable convergence of the above statistic is a direct corollary of Theorem 1. The limiting distribution will, conditional of $\mathcal{F}$, be an infinite weighted sum of $\chi^2(1)$ random variables. Rather than seeking to estimate the weights and use those estimates to approximate the limiting distribution, we can make direct use of the simulation-based approach coupled with Theorem 2. In particular, for a draw of $(e_{i,t,\kappa})_{i\geq1}$ and $(e_{i,t,\kappa'})_{i\geq1}$, we compute the statistic,

$$TS^*_{\mathcal{T},\kappa,\kappa',k} = k_n \| \widehat{W}^{N*}_{\mathcal{T},\kappa,k} - \widehat{W}^{N*}_{\mathcal{T},\kappa',k} \|_w. \tag{18}$$

This procedure is repeated $\overline{W}$ times. For $\alpha \in (0,1)$, the α-level critical value of our test statistic, $TS_{\mathcal{T},\kappa,\kappa',k}$, can then be approximated by the $1-\alpha$ quantile of the sequence of $TS^*_{\mathcal{T},\kappa,\kappa',k}$, which we denote $\widehat{Q}^*_{1-\alpha}$. Using Theorems 1 and 2, if condition (16) holds, we have,

$$\mathbb{P}\left(TS_{\mathcal{T},\kappa,\kappa',k} > \widehat{Q}^*_{1-\alpha}\big|\Omega^{(0)}\right) \to \alpha \quad \text{and} \quad \mathbb{P}\left(TS_{\mathcal{T},\kappa,\kappa',k} > \widehat{Q}^*_{1-\alpha}\big|\Omega^{(a)}\right) \to 1. \tag{19}$$

4 Monte Carlo Study

This section explores the finite-sample properties of our new inference procedures. To this end, we conduct a Monte Carlo study in which we vary the sampling frequency, window length, number of assets in the cross section, and number of observed systematic factors.

4.1 Setup

The Monte Carlo analysis is based on simulating the following $q + 1$ factor affine jump-diffusion model for each of the N assets,

$$
\begin{aligned}
dF_t &= \sqrt{V_t}\, dW_t + diag(Z_t)\, dN_t, \qquad V_t = V_t^{(1)} + V_t^{(2)}, \\
dV_t^{(i)} &= \kappa_i \left(\theta - V_t^{(i)} \right) dt + \xi_i \sqrt{V_t^{(i)}}\, dB_t^{(i)}, \qquad i = 1, 2 \\
dX_t^{(j)} &= \beta_t^{(j)\top} dF_t + \widetilde{\beta}_t^{(j)} \sqrt{V_t}\, d\widetilde{B}_t + \sqrt{V_t}\, d\widetilde{W}_t^{(j)}, \qquad j = 1, ..., N,
\end{aligned}
\tag{20}
$$

where $B^{(1)}, B^{(2)}, \widetilde{B}, B, W, \widetilde{W}^{(1)}, ..., \widetilde{W}^{(N)}$ are independent standard Brownian motions, in which W has dimension $q \times 1$ and the remaining processes are of dimension 1, N_t is a q-dimensional Poisson process with intensity $\lambda_J^{(1)} = ... = \lambda_J^{(K)} = \lambda_J$ capturing the arrival of jumps in the factors with size given by an i.i.d. sequence $(Z_s)_{s \geq 1}$, with $Z_s \sim N\left(0, \sigma_J^2 I_q\right)$, and I_q denoting the q-dimensional identity matrix. The parameters take the following values,

$$
\left(\kappa_1, \kappa_2, \theta, \xi_1, \xi_2, \lambda_J, \sigma_J^2 \right) = \left(0.0128, 0.6930, 0.4068, 0.0954, 0.7023, 0.2, 0.932 \right).
$$

This calibration follows Bollerslev and Todorov (2011). Consequently, the spot volatility consists of a persistent and a quickly mean-reverting component, following insights from Chernov et al. (2003).

We consider two separate sampling frequencies, $n = 78$ and $n = 130$, corresponding to sampling prices every 5 and 3 minutes, respectively. We consider two different setups for the number of factors, $q = 3$ and $q = 6$. The first case resembles the common empirical setup of a 3-factor model, while the second concerns a 6-factor model, matching the specification in our empirical analysis. Throughout, we set $|\mathcal{T}| = 22$, representing estimation across approximately one month. Results based on using approximately one year of data ($|\mathcal{T}| = 252$) are available upon request. To assess the importance of the size of the cross-section, we set N to $100, 300$, and 500. We fix the length of the estimation window, $k_n \Delta_n$, to equal a 2-hour interval.

The truncation threshold for the individual assets are set at $\nu_{t,n}^{(j)} = 4\sqrt{BV_{t,n}^{(j)}}\Delta_n^{0.49}$ for $j = 1, ..., N$, where $BV_{t,n}^{(j)}$ is the bipower variation of asset j, i.e.,

$$BV_{t,n}^{(j)} = \frac{\pi}{2}\sum_{i=2}^{n}|\Delta_{t,i}^n X^{(j)}||\Delta_{t,i-1}^n X^{(j)}|. \tag{21}$$

Similarly, we let $\nu_{t,n} = 4\sqrt{\iota^\top BV_{t,n}\iota}\Delta_n^{0.49}$, where ι is a $q \times 1$ dimensional vector of ones. Here, $BV_{t,n}$ is an estimate of the integrated variance of the factor vector. The (k,k') entry, $k, k' = 1, ..., q$, of this matrix may be estimated by,

$$BV_{t,n}^{(k,k')} = \frac{\pi}{8}\sum_{i=2}^{n}\left(|\Delta_{t,i}^n F^{(k)} + \Delta_{t,i}^n F^{(k')}||\Delta_{t,i-1}^n F^{(k)} + \Delta_{t,i-1}^n F^{(k')}|\right.$$
$$\left. - |\Delta_{t,i}^n F^{(k)} - \Delta_{t,i}^n F^{(k')}||\Delta_{t,i-1}^n F^{(k)} - \Delta_{t,i-1}^n F^{(k')}|\right). \tag{22}$$

To facilitate the evaluation of the integral involved in the computation of TS and TS^*, we truncate the integral at $\bar{u}$, which is set so that,

$$\bar{u} = \min\left(u : \left|\widehat{\mathcal{L}}_{\mathcal{T},\kappa,k}(u) + \widehat{\mathcal{L}}_{\mathcal{T},\kappa',k}(u)\right|/2 < 0.1, \ u > 0\right). \tag{23}$$

Finally, we set the weight function w to be the density function of a normal distribution with mean zero and variance such that $\int_{-\bar{u}}^{\bar{u}} w(u)\,du = 0.995$. This truncation, in practice, allows us to mitigate the noise from estimation in the tails of the characteristic functions.

4.2   Test Size

Given our objective, we carefully calibrate the cross-sectional dispersion of factor loadings. Under our null hypothesis, there is no change in the cross-sectional dispersion of factor loadings along the trading day. This is the case, for example, when the betas of individual assets do not change over the entire time window which is, in fact, a common assumption in the empirical asset pricing literature. To assess how our test behaves under this scenario, we sample a set of preliminary factor loadings and then adjust the first and second moments, for all t,

$$\begin{aligned}
\breve{\beta}^{(j,k)} &\sim \text{Uniform}[-0.5, 0.5], & j &= 1, ..., N; & k &= 1, ..., q \\
\beta_t^{(j,1)} &= 1 + \psi\breve{\beta}^{(j,1)}, & j &= 1, ..., N \\
\beta_t^{(j,k)} &= \psi\breve{\beta}^{(j,k)}, & j &= 1, ..., N; & k &= 2, ..., q,
\end{aligned} \tag{24}$$

where $\psi = 1.2964$. This parameter controls the dispersion of the factor loadings in the cross-section. The calibration is identical to the one in Andersen et al. (2021), which was chosen to match the average daily cross-sectional dispersion of the market beta in the data. Our results are robust to the choice of ψ. Factor 1 is centered at one to replicate the fact that the market beta has a cross-sectional mean around unity in our data. The other factors are long-short portfolios and, hence, the associated loadings have a zero cross-sectional mean.

We focus on testing whether the cross-sectional distribution of the factor loadings is the same during the first and last two hours of trading. Henceforth, we refer to these intervals as the *Open* and *Close*. Our nominal test level is $\alpha = 0.05$. In all simulations, we use $\overline{W} = 5{,}000$ to compute standard errors through our simulation-based method. For each specification, we repeat the process 1,000 times to obtain the empirical rejection probabilities.[2]

Table 1 displays our simulation results concerning test size. We note that the test statistic and associated critical values are computed for each factor independently. In the implementation of the test, we set $|\mathcal{T}| = 22$, with results for $|\mathcal{T}| = 252$ available upon request. Across the board, the empirical rejection probabilities are fairly close to the nominal value of 0.05. The highest rejection probability across all factors and specifications is 0.064 while the lowest one is 0.021. Increasing the number of parameters to be estimated, from $q = 3$ to $q = 6$, for instance, has only a minimal impact on the test performance under the null. Similarly, the rejection probabilities do not systematically change, as we increase the sampling frequency or the size of the cross-section. If anything, our test is slightly conservative in finite samples, with only a slight tendency to generate more undersized tests for scenarios with a large number of assets and only three factors.

4.3 Test Power

We now turn to the power properties of our test. For conciseness, we retain the null hypothesis for all factors apart from factor 1. This implies that, in all the simulations in this section, only factor 1 will display a time-varying cross-sectional dispersion, while the other factors are calibrated as before, and the parameters for the affine jump-diffusion model are identical to those in the last section. Furthermore, we continue to use $\overline{W} = 5{,}000$ to compute critical values and we approximate the finite sample rejection probabilities using 1,000 Monte Carlo trials.

---

[2]For implementation, the user may find it beneficial to exploit two separate margins at which the computations can be parallelized. The first obvious one is related to the independent trials in the Monte Carlo simulation. The second margin concerns our simulation-based method. For each of the $\overline{W}$ runs, the results are independent, so they can also be parallelized in straightforward fashion. This can reduce the computational cost of calculating the critical values for each Monte Carlo run substantially. On the other hand, computing the test statistic is basically instantaneous.

Table 1: **Monte Carlo results under the null hypothesis**. The table reports empirical rejection rates of the test statistic in equation (17) for nominal size 0.05 obtained using $1,000$ simulations with $|\mathcal{T}| = 22$. We use $\overline{W} = 5000$ to compute the critical values for each run. The time windows for the test are the first and last 2 hours of the trading day.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6
			$n = 78,\ N = 100$			
3-factor model	0.043	0.064	0.044			
6-factor model	0.040	0.060	0.046	0.048	0.050	0.045
			$n = 78,\ N = 300$			
3-factor model	0.045	0.032	0.043			
6-factor model	0.039	0.030	0.048	0.030	0.032	0.039
			$n = 78,\ N = 500$			
3-factor model	0.021	0.040	0.031			
6-factor model	0.036	0.038	0.042	0.035	0.047	0.027
			$n = 130,\ N = 100$			
3-factor model	0.029	0.032	0.046			
6-factor model	0.034	0.037	0.047	0.041	0.051	0.048
			$n = 130,\ N = 300$			
3-factor model	0.028	0.032	0.041			
6-factor model	0.024	0.026	0.036	0.042	0.047	0.034
			$n = 130,\ N = 500$			
3-factor model	0.029	0.028	0.034			
6-factor model	0.035	0.042	0.045	0.026	0.035	0.039

As in Section 4.2, we test whether the cross-sectional distribution of factor loadings is the same at *Open* and at *Close*. To assess test performance under different degrees of violation of the null hypothesis, we fix the *ratio* between the cross-sectional variance of loadings on Factor 1 at the *Open* and *Close*. It is implemented through the multiplicative factor ψ in equation (24). Specifically, we set $\psi_{\text{Close}} = \psi$, and we then fix the corresponding value at *Open* to a higher value. That is, we have the following relation,

$$\psi_{\text{Open}} = \sqrt{\text{ratio}} \cdot \psi_{\text{Close}}.$$

If the *ratio* is close to unity, the deviation from the null hypothesis is minor while, conversely, a high *ratio* implies that the cross-sectional distribution of loadings differs substantially.

We emphasize that, for each Monte Carlo trial, we sample the loadings for all factors only once across the trading days. This entails the assumption that, for each asset j, the loadings may change from *Open* to *Close*, but they are invariant across trading days for any given time-of-day, denoted by κ.

Table 2: **Empirical ratio between the cross-sectional variance of factor loadings at the Open and Close**.
We compute the factor loadings in the first and last two hours of trading and compute the cross-sectional variance
for each factor. We rely on all trading days in the sample, and follow the definitions provided in equation (7).

| Factor	Variance of loadings at Open	Variance of loadings at Close	Empirical Ratio
Market	0.081	0.043	1.878
SMB	0.022	0.019	1.165
HML	0.170	0.193	0.877
MOM	0.046	0.042	1.100
CMA	0.098	0.089	1.097
RMW	0.137	0.117	1.162

Figure 2 displays empirical rejection probabilities for different specifications and ratios between the cross-sectional variance of loadings at *Open* and *Close*. All panels use $|\mathcal{T}| = 22$, or about one month of data. The top row displays results for $q = 3$, while the second row reports results for $q = 6$. We increase the sampling frequency n from left to the right. The panels in the first column rely on a 5-minute sampling frequency, matching our data in the empirical section below, while the middle and right columns sample prices at the 3- and 1-minute frequencies, respectively. Each panel displays curves referring to simulations involving a different numbers of stocks in the cross-section. For reference, Table 2 reports the empirical ratios between the cross-sectional variances of factor loadings at the *Open* and *Close*, computed using all trading days in our sample.

The first immediate conclusion from Figure 2 is that, across specifications, the rejection probability increases with the sampling frequency n, as expected given our theoretical results. The rationale is that more frequent sampling enables us to estimate second return moments with higher precision, thus improving the inference for the covariance and variance terms reflected in the beta estimates, and thus also for the cross-sectional beta distribution.

A second consistent pattern is that the rejection probabilities decline, as we increase the number of observable factors from $q = 3$ (first row) to $q = 6$ (second row). For the same amount of data, our test procedure relies on an increasingly larger number of parameters, as the number of factors rises. The factor spot covariance matrix, for example, has 21 free parameters with $q = 6$ and only 6 when $q = 3$, naturally translating into lower power in the second row of Figure 2. Nonetheless, our procedure continues to display good power properties, even for $q = 6$. Specifically, for $n = 78$, we mostly reject the null when the variance ratio is close to 2, which corresponds to our reference value from Table 2.

Finally, from the reported results, it is evident that a larger cross-section enhances the power of our test, particularly when the deviation from the null hypothesis is small. Larger values of

Figure 2: **Monte Carlo results under different alternatives**. The horizontal axis displays the ratio between the cross-sectional variance of the loadings on factor 1 at *Open* and at *Close*. The loadings for all other factors remain constant throughout and across days (given by their specification under the null). The vertical axis displays empirical rejection probability of the test statistic in equation (17), obtained using $1,000$ simulations. We use $\overline{W} = 5,000$ to compute the critical values for each run.
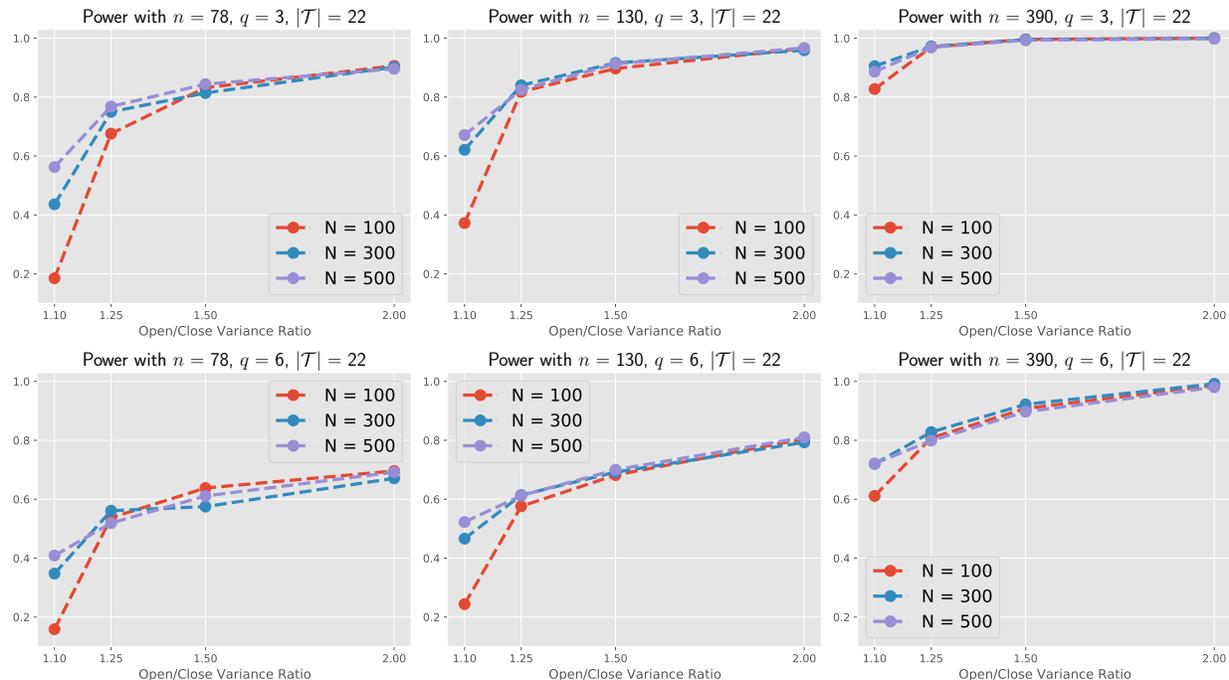
N help stabilize our estimates of the characteristic function, and this effect is critical at lower sampling frequencies, where the individual second order moments (and hence their risk loadings) are less precise. However, we also note that the discrepancy in power as a function of the number of assets vanishes, as the sampling frequency increases to one minute. The drawback of high-frequency sampling in practice is, of course, the potential for biases induced by market microstructure effects.[3]

5 Empirical Analysis

We apply our newly develop test methodology for stocks in the S&P 500 index. We rely on data from two sources. The first is stock prices extracted at the 5-minute frequency ($n = 78$) from the TAQ database, covering 2010 until 2017. We use a balanced panel, consisting of the stocks that belonged to the S&P 500 index throughout the sample period, which implies $N = 335$ and $T = 1,969$. The second dataset contains factor returns at the 5-minute frequency, obtained from

---

[3]As with the results under the null, an analogous analysis for $|\mathcal{T}| = 252$ is available upon request. With more data, the estimates are smoother and we get more power across the different specifications.

Aït-Sahalia et al. (2020). This source provides data for six factors: the Market, Size (SMB), Value (HML), Investment (CMA), Profitability (RMW) and Momentum (MOM) over the 2010-2017 period. Thus, we have high-frequency observations on the five Fama and French (2015) factors plus the Momentum factor of Carhart (1997) up until the end of 2017.

5.1 Market Beta Dispersion in a Multifactor Setting

A common procedure in empirical asset pricing is to regress stock returns on the time-series of factors in order to estimate factor loadings. The underlying assumption is that these loadings move slowly and can be treated as constant across the given quarter or year. Andersen et al. (2021) document that this assumption is grossly violated for a panel of stocks in the SP500 index, when one concentrates on a single observable factor model (and orthogonal latent risk factors), with only the market factor included on the right-hand side of the time-series regression. Instead, they find that the distribution of individual stock loadings on the market is highly dispersed at the *Open*, and then tends to decline steadily across the trading day, generating a more narrow cross-sectional market beta distribution at the *Close*.

One potential explanation for the above finding is that the market factor is correlated with other factors, whose intensity varies systematically across the trading day. The estimated loadings for the market factor would then be subject to an omitted variable bias, which is changing over time. Indeed, if estimation of beta is based only on the first element of F_t, then the univariate counterpart of $\widehat{\beta}_t^{(j,1)}$ will be a consistent estimate for $(\iota_1^\top \sigma_t \sigma_t^\top \iota_1)^{-1} \iota_1^\top \sigma_t \sigma_t^\top \beta_t^{(j)}$. This equals the first element of $\beta_t^{(j)}$ only if σ_t can be partitioned as block diagonal with diagonal blocks of dimension 1×1 and $(q-1) \times (q-1)$. Figure 1 suggests that this type of explanation for the pattern of the univariate market betas may, in fact, be plausible. By extending our theory to cover multiple observed factors, we are able to assess the validity of this conjecture. For this reason, we include the most prominent observable factors from the empirical asset pricing literature in our study. If these factors fully account for the shift in the market factor loadings across the trading day, then standard assumptions would imply the absence of any intraday changes in the cross-sectional dispersion of market betas, up to estimation error.

Using a forward-moving one-hour window, we run time-series regressions of returns for all stocks in our sample on two different sets of explanatory variables, pooling the data at the same time-of-day across the full sample. First, we regress individual stock returns on the market factor. For each window, we obtain N loading estimates. For each of the windows, we then compute the cross-sectional variance across these N market beta loadings. The black dashed line on Figure 3

Figure 3: **Cross-sectional dispersion of loadings on the market factor across the trading day**. We use rolling one-hour windows to estimate betas, with data pooled across all trading days. We regress individual stock returns on the market return (single-factor model), and on all six observable factors (multi-factor model). For each window, the cross-sectional variance of the loadings is plotted as a function of time-of-day.
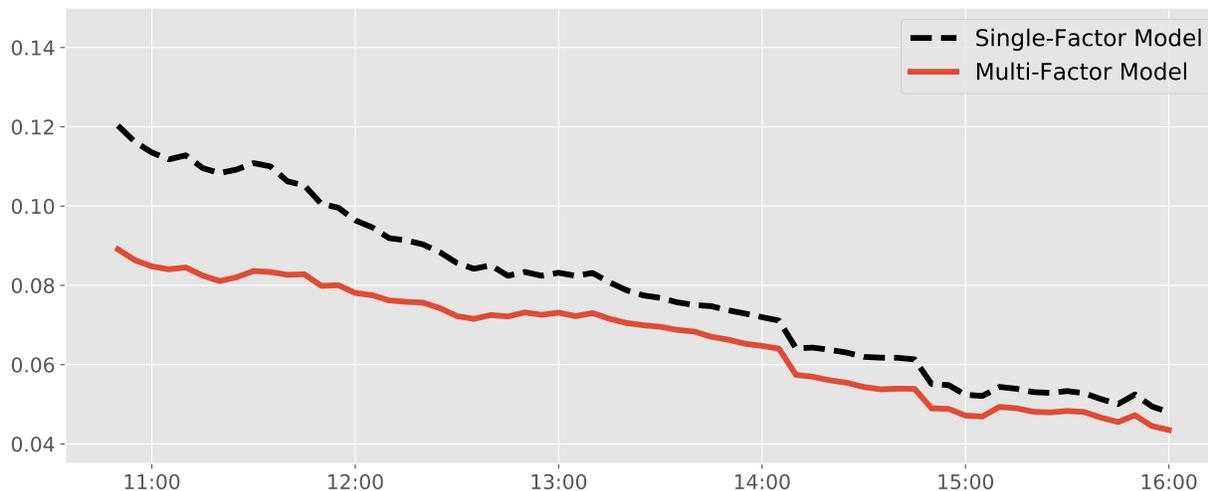

display the resulting series, starting from 10:45am (based on data from 9:45-10:45), so the curves avoid potential irregularities stemming from market microstructure issues immediately following the market *Open*. This essentially replicates the exercise in Andersen et al. (2021). Next, we repeat this procedure, but now regress the individual stock returns on all six observed factors. Again, for each window and factor, we obtain N estimates for the loadings. The red solid line on Figure 3 displays the cross-sectional variance across the N market factor loadings over the course of the trading day. Both curves in Figure 3 display a declining pattern and, at the end of trading, they are close. However, at the *Open*, controlling for the observable risk factors alters the market loading estimates substantially, with the cross-sectional dispersion dropping by roughly 25%.

We note that the inclusion of additional risk factors has reduced the cross-sectional dispersion of the market betas. This interpretation is consistent with the idea that, at *Open*, firms typically are exposed to a diverse set of innovations to fundamental risk factors, whose intensity declines as the trading day evolves. In contrast, the active trading towards the *Close* may be due in large part to trading by institutions, who need to reallocate asset position in light of the daily price changes to satisfy specific target allocations or hedging positions. There are several reasons for the accumulation of information at the *Open*. Major news from around the globe arrive during the U.S. overnight hours, and firm-specific news are predominantly released outside trading hours to avoid temporary trading suspensions due to the impending arrival of pertinent new information.

To the extent additional risk factors capture some dimensions of such risks and they are correlated with the market returns, a multi-factor model should alleviate this bias in the univariate market beta estimates.

More generally, Figure 3 suggests that controlling for additional sources of risk is important in assessing the time-variation in asset factor exposures. It provides a first step in addressing the questions posed by the striking empirical findings in Andersen et al. (2021). Nonetheless, despite the inclusion of the Fama-French and momentum factors, the cross-sectional market beta dispersion drops monotonically over the trading day. This points towards the possibility of additional missing factors correlated with the market, along with other complementary reasons behind the systematic variation in assets' intraday market exposures.[4] We provide formal tests for the discrepancy in the cross-sectional distribution of factor exposures at the *Open* and *Close* in Section 5.3.

Figure 4: **Cross-sectional distribution of loadings on the market factor**. We use the first and last two hours of trading data to run time-series regressions pooling all trading days. The black dashed line refers to regressions that use the market return as the only factor. The red solid line is constructed based on a six-factor model. The densities are computed by direct Fourier inversion of the estimated characteristic functions.
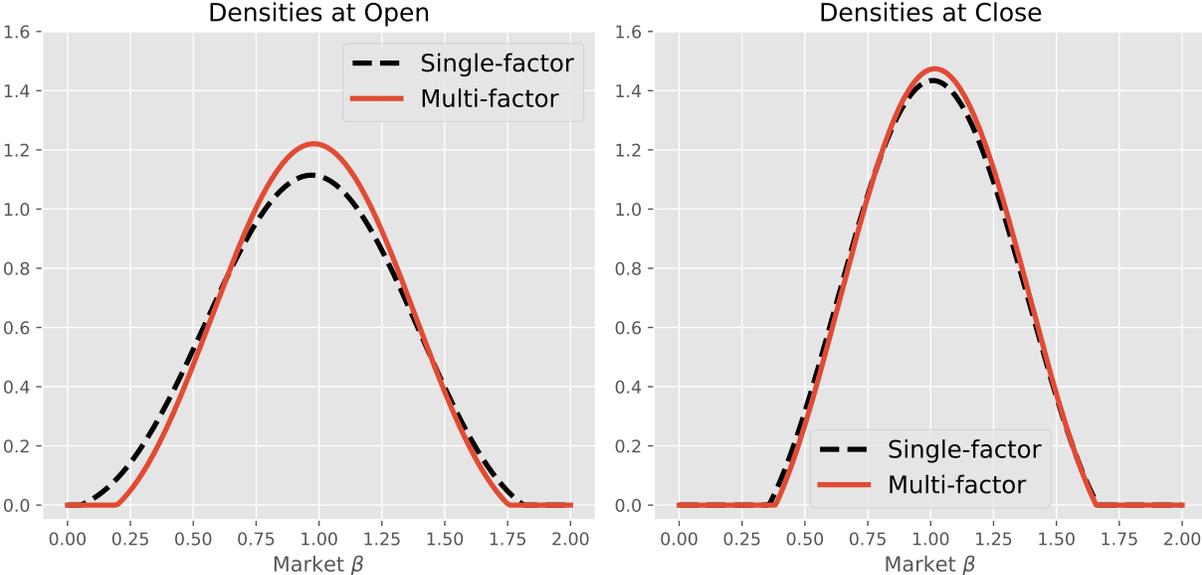
Another way of visualizing the above phenomenon is to inspect the entire distribution of loadings. Figure 4 displays the estimated density for the cross-sectional market beta distributions at *Open* and at *Close*. The black dashed line displays the density, when the loadings are estimated from a single-factor model, while the red solid curve is constructed from the six-factor model con-

---

[4]Note that we allow for unobservable factors that are orthogonal (in a martingale sense) to the observable ones in our analysis.

Figure 5: **Examples of intraday market beta dynamics**. The figure plots the intraday evolution of the market betas for two different stocks. The black dashed lines are constructed with the market serving as the sole factor, while the red solid lines are constructed from a six-factor model. We use a rolling one-hour window to estimate the loadings, pooling the data across all trading days.
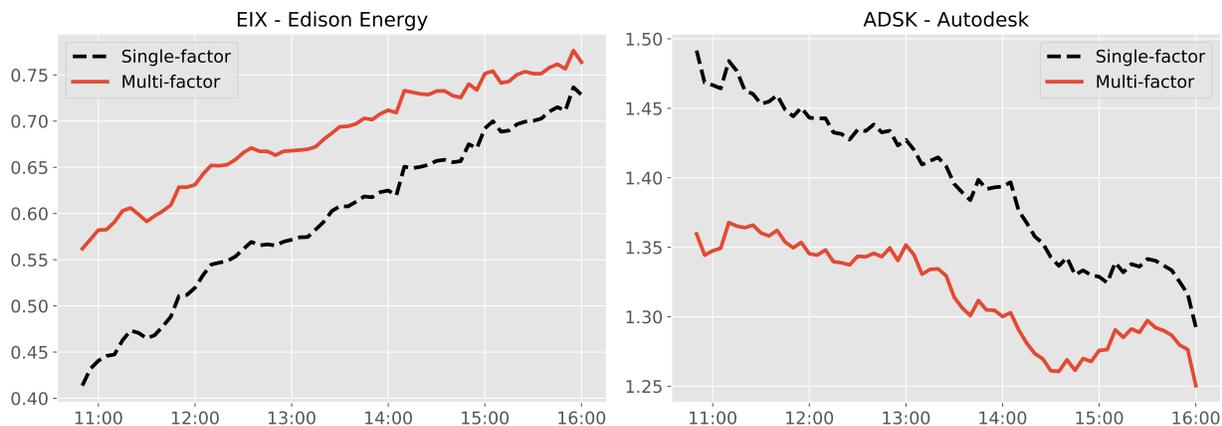
sisting of the observed factors from Aït-Sahalia et al. (2020). At the *Close*, the estimated densities are nearly identical, suggesting that the added factors do not change the market betas in a significant way. At *Open*, however, the market beta distribution obtained from the multi-factor setting is distinctly more concentrated around the mean, in line with the dispersion profile in Figure 3. Finally, we note the stark discrepancy between the densities in the left and right panels, again confirming the decreasing market beta dispersion over the course of the trading day.

To exemplify how the introduction of additional risk factors impacts the estimated market exposure for individual stocks, Figure 5 displays the intraday evolution of the estimated market beta for Edison Energy (`EIX`) on the left and Autodesk (`ADSK`) on the right panel. Again, controlling for additional sources of risk does not render the market betas flat across the trading day, but reduces the slope substantially in both instances. For the low-beta stock EIX, the one-factor model delivers an *Open* beta of around 0.35 (0.50) without (with) controls, while it ends up near 0.73 (0.75) at *Close*. Likewise, for the high-beta stock ADSK, we observe a pronounced flattening of the declining intraday pattern in the market beta estimates for the six-factor compared to one-factor model. Again, accounting for multiple sources of risk explains a non-trivial part of the intraday variation in the univariate market factor loadings, but it is still far from eliminating the pattern altogether.

The evidence of pronounced and systematic intraday shifts in factor risk exposures provides a challenge for standard asset pricing procedures such as factor risk premium measurement through

the so-called two-step approach.[5] Moreover, time-varying intraday betas also add a new wrinkle to studies that explore cross-sectional variation in expected returns across the trading day. Specifically, relying on constant loadings throughout the day, as in, e.g., the interesting recent work by Bogousslavsky (2021), might induce spurious findings.

5.2 Densities of Cross-Sectional Factor Exposures

Our approach readily generates estimates for the densities of the factor loadings, as we can invert the characteristic functions computed as an integral part of our procedure for calculating our test statistic. In Figure 6, we plot the estimated densities for the cross-sectional distribution of the factor loadings, obtained by pooling all trading days in our sample. The solid lines display our density estimates at *Open*, while the dashed lines represent those at *Close*. First, we note that the mean of the loading on the market factor is near unity, while it is close to zero for the other factors. This is reassuring, as all factors besides the market are generated from long-short portfolios. It is consistent with our beta estimates being free of systematic biases, which may arise if market microstructure effects are substantial at the chosen sampling frequencies. Inspecting the estimated densities for the non-market factor exposures, we identify indications of shifts in the cross-sectional distribution in some cases, although the effect is less pronounced than for the market factor. Specifically, as for the market exposure, most of the other factors display densities that become more concentrated at the *Close*, with the value (HML) and, possibly, the momentum factors (MOM) representing exceptions. Given the more modest shifts in the estimated densities for the non-market factor exposures, the formal inference procedures applied in Section 5.3 will shed additional light on the strength of the evidence in this regard.

The tentative interpretation of the intraday pattern in the cross-sectional market beta dispersion offered by Andersen et al. (2021) rests on the fact that firm-specific and global news arrive disproportionally prior to and during the *Open* period. However, some prescheduled macroeconomic news releases, expected to convey important new information, occur within the trading day. The primary example is the monetary policy announcements following the Federal Open Market Committee (FOMC) meetings, which take place about every six weeks. These news releases are currently scheduled for 2:00pm ET, but may be followed by a press conference. There is a broad empirical literature documenting that such "FOMC days" feature unusual return dynamics.[6]

---

[5]This technique estimates risk loadings by time-series regressions, in the spirit of the current paper, and then uses these loadings as explanatory variables to account for *cross-sectional* variation of expected returns. See Cochrane (2009) for an overview on factor risk premium measurement.

[6]See, e.g., Savor and Wilson (2014), Lucca and Moench (2015) and Cieslak et al. (2019) for empirical accounts of the FOMC announcement effects and their institutional underpinnings. Ai and Bansal (2018) provide a model to

Figure 6: **Cross-sectional distribution of the factor loadings**. The picture displays the estimated densities for the factor loadings in the first and last two hours of trading using data from the full sample. The densities are estimated through a direct Fourier inversion of the estimated characteristic functions.
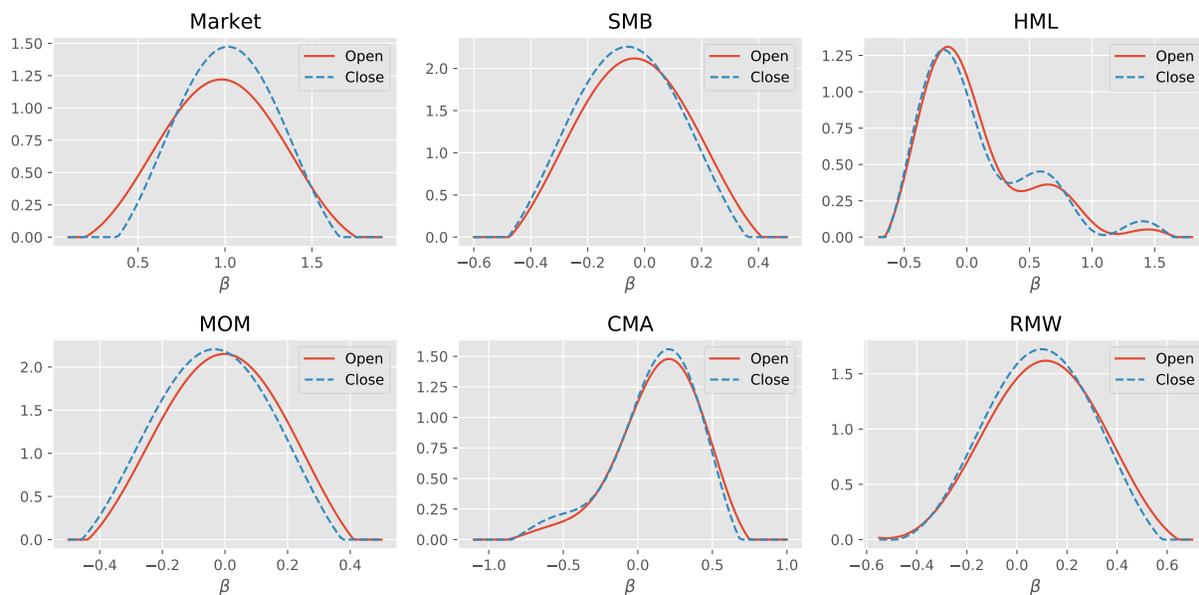

Because FOMC announcements are released after the *Open*, but before the *Close*, the anticipation about and – especially – the subsequent reaction to the monetary policy news is likely to render the factors, that are sensitive to interest or discount rate innovations, particularly volatile and influential on such days. Bernanke and Kuttner (2005) and Nakamura and Steinsson (2018), among others, exploit this fact to identify the real effects of monetary policy shocks. We complement this literature by exploring whether the FOMC days display any unusual traits concerning the cross-sectional distributions of factor betas. In Section 5.3, we test explicitly for changes in the loadings across all the observed factors between the *Open* and *Close* on the subsample of days with FOMC announcements. In total, we have $|\mathcal{T}| = 64$ such FOMC meetings in our data.[7]

As an initial illustration, Figure 7 replicates Figure 6, but only for the subset of FOMC days. The compression in the distribution of market betas towards the *Close* is even stronger for FOMC Days. Similarly, the distributions of loadings on Value (HML) and Investment (CMA) seem to change more between the *Open* and *Close*, while those on Size (SMB), Momentum and Profitability (RMW) do not appear to shift significantly on FOMC days relative to regular trading days.

---

account for some of the stylized facts documented in this literature.

[7]The dates were obtained from the FOMC website https://www.federalreserve.gov/monetarypolicy/fomc_historical_year.htm.

Figure 7: **Cross-sectional distribution of the factor loadings on FOMC Days**. The picture displays the estimated densities for the factor loadings in the first and last two hours of trading, but only for days with pre-scheduled FOMC announcements. The densities are estimated through direct Fourier inversion of the estimated characteristic functions.
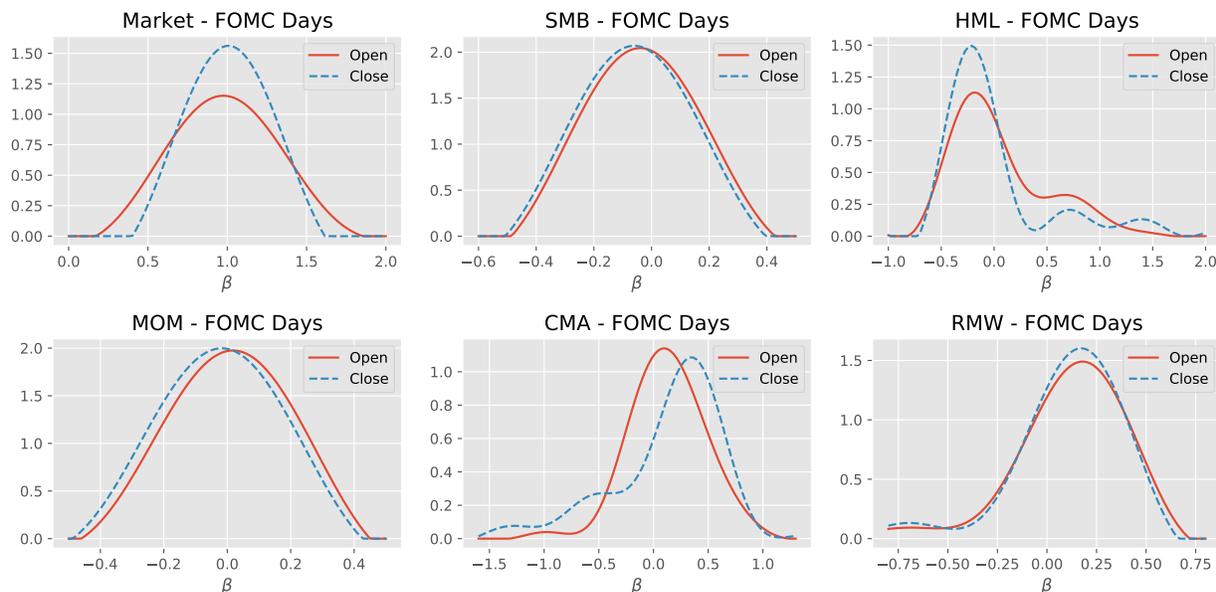
5.3 Test of the Null Hypothesis of No Intraday Beta Variation

The illustrations in Sections 5.1 and 5.2 help us develop a sense of the intraday variability of the factor loadings, but they do not tell us whether the apparent discrepancies are significant or a result of noise stemming from random sample variation. Our new test developed in Section 3.3 enables us to address this issue formally. The null hypothesis is that the cross-sectional distribution of factor loadings is unchanged from *Open* to *Close*. We test the hypothesis separately for each factor, using the test statistic TS from equation (17) and the corresponding inference procedure outlined in equations (18) and (19). As discussed above, we also conduct separate tests for the subsample of FOMC days, the days that contain prescheduled FOMC announcements.

Table 3 displays the p-values for each factor and subsample. As in the Monte Carlo study, we use $\overline{W} = 5{,}000$. The first column refers to p-values obtained exploiting data across the full sample. The results confirm the impression from our prior visual displays in Figures 3, 4 and 6. Assessed at the 1% level, the market, value, and profitability factors all feature significant shifts in the cross-sectional distribution of factor exposure between the *Open* and *Close*, while the size factor displays significance at the 5% level. Only for the momentum and investment factors, we

cannot reject the hypothesis of an invariant intraday factor exposure.

Table 3: **Test results for no intraday variation in factor loadings.** The first column uses all days in our sample. The second and third columns run the test pooling all days with and without pre-scheduled FOMC announcements, respectively. We use $\overline{W} = 5000$ to compute the critical values. Only the associated p-values are reported in the table.

Sample/ Factor	Full Sample	FOMC Days	Non FOMC Days
Market	0.000	0.000	0.000
SMB	0.027	0.451	0.049
HML	0.000	0.013	0.000
MOM	0.180	0.335	0.212
CMA	0.594	0.023	0.502
RMW	0.009	0.869	0.017

The second column of Table 3 focuses solely on the FOMC days. The sharp reduction in the sample size lowers test power dramatically. Nonetheless, the null hypothesis continues to be strongly rejected for the market factor, and we still have fairly strong evidence against invariance of the exposure to the value factor, generating a p-value of 1.3%. In contrast, there is now no evidence of a shift for the cross-sectional distribution in loadings on the size and profitability factors. Finally, and most strikingly, the investment factor is now significant at the 5% level. Overall, these findings are well aligned with the visual impression conveyed by Figure 7.

The result concerning the investment factor warrants additional reflection. The CMA factor is a long-short portfolio with positive exposure to firms considered conservative in terms of investments and negative exposure to firms that are considered aggressive on the same margin. If the CMA exposure reflects a mixture of underlying exposures to a set of latent economic factors, with discount rate innovations being an important one among those and carrying a large (unobserved) exposure coefficient, then an enhanced signal from monetary policy announcements maps into a stronger loading on the CMA factor. Hence, the variability of asset exposures to the investment factor on FOMC days is consistent with the intensity of discount rate news being elevated and fluctuating sharply across the day. Inspecting the shift between the estimated distribution for the CMA loadings at the *Open* and *Close*, we observe a pronounced increase in the left tail, balanced with a slight shift in the mode of the distribution into the positive region. This suggests a substantial degree of heterogeneity across the firms in their sensitivity to monetary policy shocks, with a subset of firms seeing their exposure to CMA drop markedly following the FOMC announcement.

In order to more closely scrutinize the hypothesis that the FOMC announcements induce a shift in the loadings on the CMA and the HML factors, Figure 8 displays the cross-sectional variance

Figure 8: **Cross-sectional dispersion of loadings on observed factors across the trading day**. We use rolling one-hour windows to estimate betas. The panel on the left excludes days with an FOMC announcement, while the right panel relies solely on "FOMC days" (64 days in our sample). The windows are backward looking. For example, data plotted at 15:00 is associated with beta estimates obtained from the 14:00–15:00 time window.
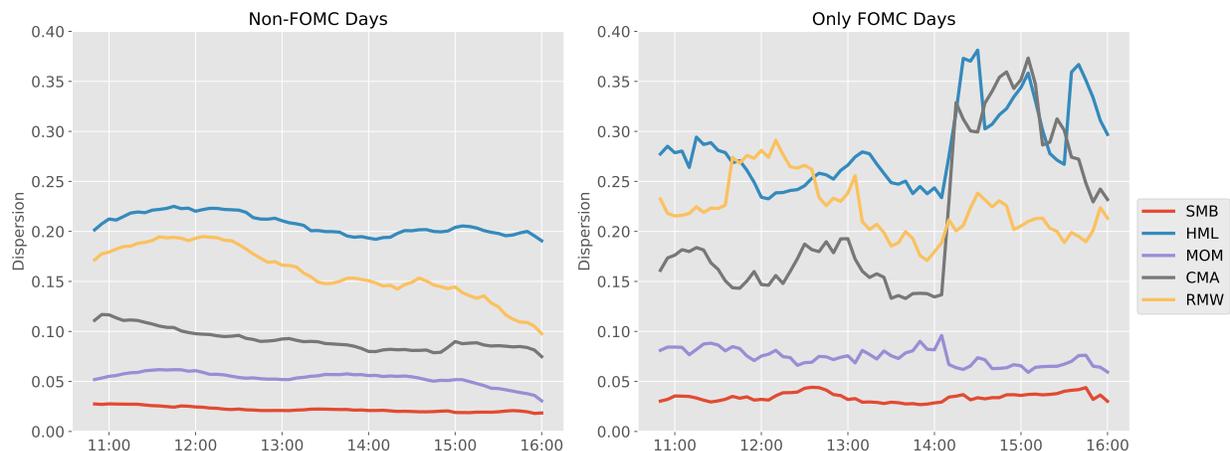
of risk loadings across the trading day for all our observable non-market factors, replicating the technique used to compute the dispersion measure in Figure 3. In the left panel, we provide the estimated dispersion measures across the full sample, excluding only the FOMC days, while the right panel depicts the same measures computed exclusively on FOMC days. The contrast between the panels is telling. On FOMC days, the estimated dispersion measures are naturally somewhat noisy due to the sharply reduced sample size, yet there is no indication of any striking shift in exposures for the SMB, RMW, and MOM factors. In contrast, there is a seemingly dramatic change in the distribution of the CMA and HML factor exposures *exactly* at 2:00pm, when the FOMC statement is released. Moreover, the effect persists, consistent with a response to the press conference conducted by the FOMC chair, starting at 2:30pm ET.[8]

We previously discussed the logic behind the shift in the CMA factor exposures around FOMC announcements. That we observe a similar, albeit slightly muted, effect on the HML factor is also intuitive. The value factor is long value and short growth stocks, implying a systematic cash-flow duration mismatch between the long and short positions. If the intensity of monetary policy news increases, we should expect a more dominant role for this factor in driving stock returns. In summary, firms that differ in investment strategy or cash-flow duration are naturally systematically different in their exposure to discount rate shocks. The timing of the shocks to the estimated dispersion measures for the factor loadings in Figure 8 are consistent with this reasoning.

---

[8]See Nakamura and Steinsson (2018) for a detailed discussion concerning the type of information embedded within these statements and their empirical implications for the effect of monetary policy shocks.

We further note that we cover a period in which the Federal Reserve deployed so-called "forward guidance," intended to better inform market participants about the future path of monetary policy. It is of interest, but outside the scope of this paper, to analyze whether the CMA and HML factor exposures display similar cross-sectional intraday variation on FOMC days prior to the Global Financial Crisis, for example, when central bank communication was structured differently.[9]

Finally, the last column of Table 3 confirms that the general conclusions obtained across the full sample are robust to the exclusion of the, potentially influential, FOMC days. Given the relatively low number of FOMC days, this result is not surprising.

5.4   Heterogeneity in Factor Exposures across Time

So far, our primary analysis has pooled data across the full sample to maximize the power of the test. We now explore whether there is heterogeneity across the years in terms of our ability to reject the null hypothesis. Consequently, we apply our procedure for each factor and each year within the sample. The associated p-values are reported in Figure 9. The dashed horizontal line represents a 5% reference level, and we continue to rely on $\overline{W} = 5,000$ to compute the p-values.

Figure 9: **Testing for year-by-year intraday variation in factor loadings**. We test the null hypothesis of a constant distribution of factor loadings between the *Open* and *Close* for each year and factor. The p-values are reported by red bars. The black dashed line references the 5% level. We use $\overline{W} = 5000$ to compute critical values.

---
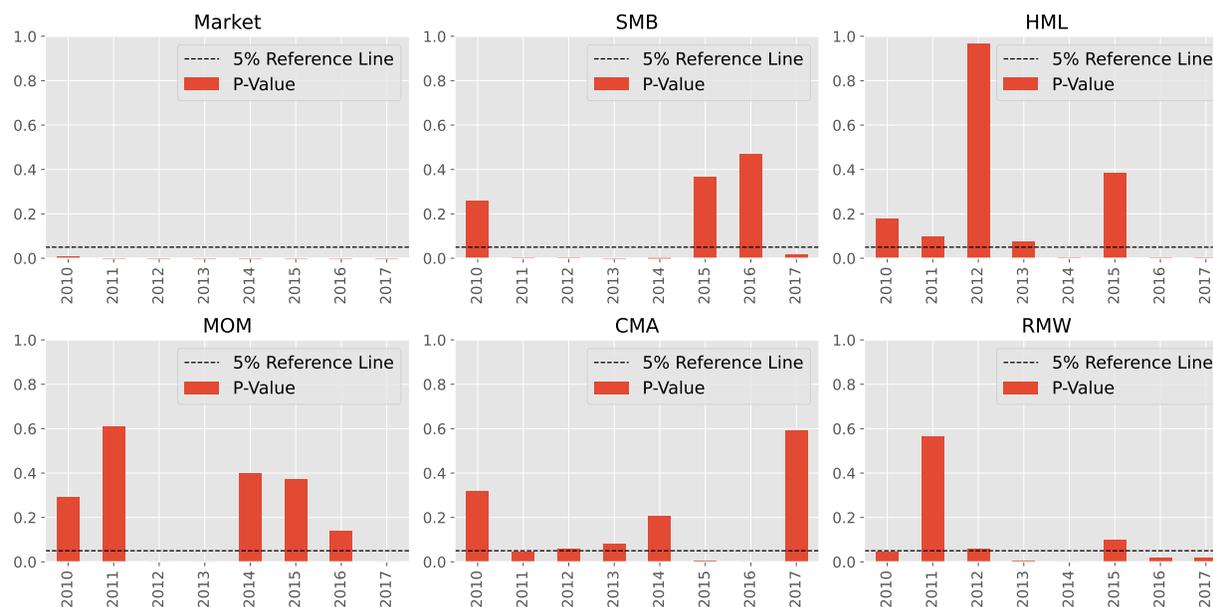
[9]See Garcia-Schmidt and Woodford (2019) for an updated view on forward guidance policies.

The first notable fact is that the p-values for the market factor, effectively, are zero uniformly across all years in the sample. Thus, we confidently reject the hypothesis of a constant distribution for market betas between *Open* and *Close* throughout the sample. There is considerably more variation in the results for the other factors, consistent with the smaller observed variation depicted in Figure 4 and a less powerful test due to the shorter annual samples. For the three factor exposures associated with a significant distributional shift over the full sample, namely SMB, HML, and RMW, we find them all rejecting the null hypothesis in 2014 and 2017, and at the 10% level in 2013. Moreover, for each of these factors, we reject the null for about half of the eight years in our sample. Interestingly, the same observation is valid for the investment factor, CMA, while the p-values associated with the momentum factor fluctuate rather randomly from effectively zero in some years to large values in others, indicating the existence of periods of limited intraday cross-sectional variation in the exposure to this factor.

Given the exceedingly strong evidence for intraday cross-sectional variation in the exposure to the market factor, we go one step further and perform our test for this factor on a month-by-month basis. The result is provided in Figure 10. Indeed, we do find that the effect is sufficiently strong, that it remains almost uniformly significant, even when assessed over short monthly horizons.[10]

In summary, the evidence for highly significant intraday variation in the distribution of the cross-sectional asset exposure to the market factor is extremely robust. Moreover, there is evidence of similar intraday fluctuations in the exposures to various popular observed asset pricing factors, but the latter are less sizable and certainly less universal, potentially reflecting a dependence of these factor exposures on the broader economic and market environment. Overall, these findings run counter to commonly invoked – yet hardly ever tested – assumptions in the empirical asset pricing literature stipulating stability in factor exposures for asset returns across longer time windows.

6 Conclusion

This paper develops a novel test for stability of the intraday cross-sectional distribution of asset exposures to observable factors. The test procedure generates very robust evidence for this type of intraday variation for the market betas across firms in the S&P 500 index. Similar, albeit weaker, evidence is obtained in favor of systematic intraday variation in the asset exposures towards a number of other popular non-market pricing factors.

---

[10]Given the more limited degree of variation for the non-market factor exposures and the lower power associated with shorter samples, the corresponding month-by-month tests for those factors are expected to display uninformative p-values, that are close to uniformly distributed over the unit interval. We have confirmed this conjecture. The figures are available upon request.

Figure 10: **Testing for month-by-month intraday variation in the market factor**. We test the null of constant market factor exposure in the 6-factor model for every month in our sample, and report the associated p-values as a time-series. We use $\overline{W} = 5000$ to compute the critical values.
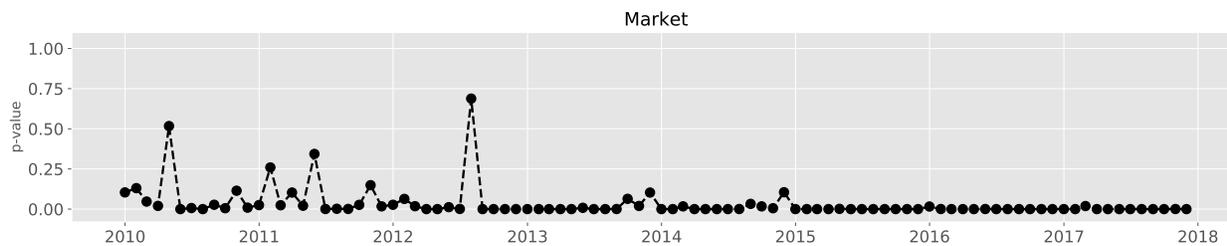
These findings extend the empirical evidence in Andersen et al. (2021) to cover the entire cross-sectional distribution of intraday factor exposures, rather than simply the second moments. In addition, the results are shown to carry through in a multi-factor setting with observable and unobservable factors. We find that controlling for observed factors that are correlated with the market, the distributional heterogeneity is mitigated, yet remains highly significant. Furthermore, our findings suggest an interesting variation in the factor exposures around major macroeconomic announcements, as exemplified by the prescheduled statements following regular FOMC meetings.

From a theoretical perspective, our inference procedure exploits a newly developed estimator for the characteristic function of the cross-sectional distribution for the factor loadings at a given point within the trading day. The corresponding feasible functional central limit theory is derived in a complex-valued weighted $\mathcal{L}^2$ Hilbert space. In turn, we use these results to characterize the limiting distribution for the distance between the empirical characteristic functions at two distinct times in the trading day under the null hypothesis. The test statistic has a non-standard limiting distribution, but we develop a simple approach to compute the associated critical values via a novel simulation-based procedure.

Our empirical results point towards interesting new areas of inquiry. In particular, the determination of the underlying sources of variation in the asset factor loadings should be investigated in further detail. Progress along those lines should facilitate a cleaner identification of factor loadings and risk premiums in general. Furthermore, it should help address recent empirical findings in the literature concerning systematic intraday variation in expected returns.

Appendix: Assumptions and Proofs

A   Assumptions

Assumption A. *For the processes $(X^{(j)})_{j \geq 0}$ we have:*

(a) For a sequence of stopping times, $(T_m)_{m \geq 1}$, increasing to infinity, the processes $(\alpha^{(j)})_{j \geq 0}$, $(\beta^{(j)})_{j \geq 1}$, $(\gamma^{(j)})_{j \geq 1}$, and $(\widetilde{\sigma}^{(j)})_{j \geq 1}$, are all uniformly bounded on $[0, T \wedge T_m]$.

(b) The process $\lambda_{\min}(V_t)$ take positive values on $[0, T]$.

(c) For a sequence of stopping times, $(T_m)_{m \geq 1}$, increasing to infinity and a sequence of constants, $(K_m)_{m \geq 1}$, we have uniformly in $j \geq 1$ and $k_1, k_2 \in \{1, ..., q\}$:

$$\mathbb{E}\Bigg[\sup_{s,t \in [0, T \wedge T_m]} |\sigma_t^{(k_1, k_2)} - \sigma_s^{(k_1, k_2)}|^2 + \sup_{s,t \in [0, T \wedge T_m]} |\beta_t^{(j)} - \beta_s^{(j)}|^\top |\beta_t^{(j)} - \beta_s^{(j)}| + \sup_{s,t \in [0, T \wedge T_m]} |\widetilde{\sigma}_t^{(j)} - \widetilde{\sigma}_s^{(j)}|^2$$
$$+ \sup_{s,t \in [0, T \wedge T_m]} |\gamma_t^{(j)} - \gamma_s^{(j)}||\gamma_t^{(j)} - \gamma_s^{(j)}|^\top\Bigg] \leq K_m |t - s|, \tag{25}$$

$$|\mathbb{E}(\beta_{t \wedge T_m}^{(j)} - \beta_{s \wedge T_m}^{(j)})| + |\mathbb{E}(\sigma_{t \wedge T_m}^{(k_1, k_2)} - \sigma_{s \wedge T_m}^{(k_1, k_2)})| \leq K_m |t - s|, \tag{26}$$

$$|\mathbb{E}(\chi_{s \wedge T_m, t \wedge T_m}^{(1)} \chi_{s \wedge T_m, t \wedge T_m}^{(2)})| \leq K_m |t - s|^2, \tag{27}$$

for $\chi_{s,t}^{(1)}$ equal to $\beta_t^{(j)} - \beta_s^{(j)}$ or $\sigma_t^{(k_1, k_2)} - \sigma_s^{(k_1, k_2)}$, and $\chi_{s,t}^{(2)}$ equal to one of $(W_t^{(k)} - W_s^{(k)})^2 - (t - s)$, $(W_t^{(k)} - W_s^{(k)})(B_t - B_s)$ and $(W_t^{(k)} - W_s^{(k)})(\widetilde{W}_t^{(j)} - \widetilde{W}_s^{(j)})$, $k = 1, ..., q$.

(d) We have $\sum_{s \leq t} \Delta X_s = \int_0^t \int_E \delta_X(s, u) \mu(ds, du)$ and $\sum_{s \leq t} \Delta F_s = \int_0^t \int_E \delta_F(s, u) \mu(ds, du)$, where μ is a Poisson random measure on $\mathbb{R}_+ \times E$ with compensator $ds \otimes \nu(du)$, for some σ-finite measure ν on a Polish space E. Furthermore, the jump size functions δ_X and δ_F are mappings $\Omega \times \mathbb{R}_+ \times E \to \mathbb{R}^N$ and $\Omega \times \mathbb{R}_+ \times E \to \mathbb{R}^q$, respectively, which are locally predictable. We have $\int_0^T \int_E 1(\delta_X^{(j)}(s, u) \neq 0)ds\nu(du) < \infty$, for $j \geq 1$ and $\int_0^T \int_E 1(\delta_F^{(k)}(s, u) \neq 0)ds\nu(du) < \infty$, for $k \geq 1$. For a sequence of stopping times, $(T_m)_{m \geq 1}$, increasing to infinity and a sequence of nonnegative functions $(\Gamma_m(u))_{m \geq 1}$ satisfying $\int_E (1 \vee \Gamma_m^2(u))\nu(du) < \infty$, we have $\sup_{j \geq 1} |\delta_X^{(j)}(s, u)| + \sup_{k \geq 1} |\delta_F^{(k)}(s, u)| \leq \Gamma_m(u)$, for $s \in [0, T \wedge T_m]$.

To state the next assumption, we define,

$$
\psi_{t,\kappa,k,N}^{(1)}(z,u) = \frac{-uz}{N}\sum_{j=1}^{N}\exp\left(i(u+z)\beta_{t-1+\kappa}^{(j,k)}\right)\iota_k^\top V_{t-1+\kappa}^{-1}\sigma_{t-1+\kappa}\gamma_{t-1+\kappa}^{(j)\top}\gamma_{t-1+\kappa}^{(j)}\sigma_{t-1+\kappa}V_{t-1+\kappa}^{-1}\iota_k, \quad (28)
$$

$$
\psi_{t,\kappa,k,N}^{(2)}(z,u) = \frac{uz}{N}\sum_{j=1}^{N}\exp\left(i(u-z)\beta_{t-1+\kappa}^{(j,k)}\right)\iota_k^\top V_{t-1+\kappa}^{-1}\sigma_{t-1+\kappa}\gamma_{t-1+\kappa}^{(j)\top}\gamma_{t-1+\kappa}^{(j)}\sigma_{t-1+\kappa}V_{t-1+\kappa}^{-1}\iota_k. \quad (29)
$$

Assumption B. *There exists random functions $\psi_{t,\kappa,k}^{(j)}(z,u)$, $j=1,2$, such that the following convergence in probability as $N \to \infty$ holds, with $t \in \mathbb{N}_+ \cap [0,T]$, and $\kappa \in [0,1]$:*

$$
\int_{\mathbb{R}}\int_{\mathbb{R}}|\psi_{t,\kappa,k,N}^{(j)}(z,u) - \psi_{t,\kappa,k}^{(j)}(z,u)|^2 w(u)w(z)dudz \xrightarrow{\mathbb{P}} 0, \ \ j=1,2, \quad (30)
$$

where w is the weight function for the $\mathcal{L}^2(w)$ space defined in (10).

B   Notation

Throughout the proofs, we will take $\mathcal{T}$ to have a single element t. This is without loss of generality because the limiting variables of Theorems 1 and 2 over different elements in $\mathcal{T}$ in the general case, are $\mathcal{F}$-conditionally independent. In the proofs, we denote with K a positive constant that does not depend on n and N, and can change from one line to another.

We now introduce some notation, which will be used throughout the proofs. We denote by $\|\cdot\|$ and $\|\cdot\|_F$ the Euclidean norm on $\mathbb{R}^q$ and the Frobenius norm on $\mathbb{R}\otimes\mathbb{R}$, respectively. For a generic process Z_t, we set,

$$
Z_{t,\kappa,n} = Z_{t-1+\frac{\lfloor \kappa n\rfloor - k_n}{n}}, \quad t \in \mathbb{N}_+, \ \kappa \in [0,1] \quad (31)
$$

and define with Z^c the continuous part of Z. Similarly, we use the shorthand notation $\mathbb{E}_{t,s}^n(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_{t-1+(s-1)/n})$. The spot variance-covariance matrix of the factors is denoted by,

$$
V_t = \sigma_t\sigma_t^\top. \quad (32)
$$

As for the estimators of β, we also introduce the shorthand notation $\beta_t^{(j,k)} = \iota_k^\top\beta_t^{(j)}$.

We set $C_t^{(j)} = V_t\beta_t^{(j)}$, $j = 1,...,N$, and define,

$$
\overline{C}_{t,\kappa}^{(j)} = \frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\Delta_{t,s}^n X^{(j)c}\Delta_{t,s}^n F^c, \ \ \overline{V}_{t,\kappa} = \frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\Delta_{t,s}^n F^c\Delta_{t,s}^n F^{c\top}, \ j=1,...,N, \quad (33)
$$

$$\widetilde{C}_{t,\kappa}^{(j)} = \frac{n}{k_n} \sum_{s \in \mathcal{I}_\kappa^n} \int_{t-1+(s-1)\Delta_n}^{t-1+s\Delta_n} C_u^{(j)} du, \quad \widetilde{V}_{t,\kappa} = \frac{n}{k_n} \sum_{s \in \mathcal{I}_\kappa^n} \int_{t-1+(s-1)\Delta_n}^{t-1+s\Delta_n} V_u du, \tag{34}$$

$$\mathcal{C}_{t,\kappa}^{(j)} = \frac{1}{k_n} \sum_{s \in \mathcal{I}_\kappa^n} \left(\sigma_{t,\kappa,n} \left(n\Delta_{t,s}^n W \Delta_{t,s}^n W^\top - I_q \right) \sigma_{t,\kappa,n}^\top \beta_{t,\kappa,n}^{(j)} + n\sigma_{t,\kappa,n}\Delta_{t,s}^n W \Delta_{t,s}^n B^\top \gamma_{t,\kappa,n}^{(j)\top} \right.$$

$$\left. + n\sigma_{t,\kappa,n}\Delta_{t,s}^n W \Delta_{t,s}^n \widetilde{W}^\top \widetilde{\sigma}_{t,\kappa,n}^{(j)\top} \right), \tag{35}$$

and

$$\mathcal{V}_{t,\kappa} = \frac{1}{k_n} \sum_{s \in \mathcal{I}_\kappa^n} \sigma_{t,\kappa,n} \left(n\Delta_{t,s}^n W \Delta_{t,s}^n W^\top - I_q \right) \sigma_{t,\kappa,n}^\top. \tag{36}$$

C   Localization

Assumption SA. *We have Assumption A with $T_1 = \infty$. Furthermore, the processes $(\alpha^{(j)})_{j\geq 0}$, $(\beta^{(j)})_{j\geq 1}$, $(\gamma^{(j)})_{j\geq 1}$, and $(\widetilde{\sigma}^{(j)})_{j\geq 1}$ are uniformly bounded on $[0,T]$, and $\lambda_{\min}(V_t)$ is bounded from below by a positive constant on $[0,T]$.*

We will prove the results under the stronger Assumption SA. A standard localization argument then can be used to show that they continue to hold under the weaker Assumption A.

D   Preliminary results

In this section, we provide a number of preliminary results, which we subsequently will use to prove the theorems presented in the paper.

Lemma 1. *Under Assumption SA, we have for each $j = 1, ..., N$, fixed $q \in \mathbb{N}$, $t = 1, ..., T$, $\kappa \in [0,1]$, and $p \geq 1$ that,*

$$\mathbb{E}_{t,\lfloor \kappa n \rfloor - k_n + 1}^n \left[\|\widehat{C}_{t,\kappa}^{(j)} - \overline{C}_{t,\kappa}^{(j)}\|^p + \|\widehat{V}_{t,\kappa} - \overline{V}_{t,\kappa}\|_F^p \right] \leq K_p \left(k_n^{1-p}\Delta_n^{1+p(2\omega-1)} \vee \Delta_n^{p+p(2\omega-1)} \right), \tag{37}$$

where $K_p > 0$ is a constant which depends on p, but not on j.

Proof. The result follows from an application of Minkowski's inequality combined with Lemma 1 in Andersen et al. (2021). □

Lemma 2. *Assume that Assumption SA holds, then we have for each $j = 1, ..., N$, fixed $q \in \mathbb{N}$, $t = 1, ..., T$, $\kappa \in [0, 1]$, and $p \geq 2$ that,*

$$\mathbb{E}^n_{t, \lfloor \kappa n \rfloor - k_n + 1} \left[\| \overline{C}^{(j)}_{t,\kappa} - \widetilde{C}^{(j)}_{t,\kappa} \|^p + \| \overline{V}_{t,\kappa} - \widetilde{V}_{t,\kappa} \|^p_F + \| \mathcal{C}^{(j)}_{t,\kappa} \|^p + \| \mathcal{V}_{t,\kappa} \|^p_F \right] \leq K_p k_n^{-p/2}, \qquad (38)$$

where $K_p > 0$ is a fixed constant which does not depend on j.

Proof. The result follows from an application of Minkowski's inequality and Lemma 2 in Andersen et al. (2021). $\qquad\square$

Lemma 3. *Assume that Assumption SA holds, then we have the following decompositions,*

$$\overline{C}^{(j)}_{t,\kappa} = \widetilde{C}^{(j)}_{t,\kappa} + \mathcal{C}^{(j)}_{t,\kappa} + R^{(j)}_{t,\kappa,n}, \quad \overline{V}_{t,\kappa} = \widetilde{V}_{t,\kappa} + \mathcal{V}^{(j)}_{t,\kappa} + R_{t,\kappa,n}, \qquad (39)$$

for each $j = 1, ..., N$, fixed $q \in \mathbb{N}$, $t = 1, ..., T$, $\kappa \in [0, 1]$, where the remainder terms $R^{(j)}_{t,\kappa,n}$ satisfies,

$$\mathbb{E}^n_{t, \lfloor \kappa n \rfloor - k_n + 1} \left[\| R^{(j)}_{t,\kappa,n} \|^p + \| R_{t,\kappa,n} \|^p_F \right] \leq K_p \frac{k_n^{2-p}}{n}, \ \ p \geq 2, \qquad (40)$$

Furthermore, we have that,

$$\mathbb{E}^n_{t, \lfloor \kappa n \rfloor - k_n + 1} \left[\| \widetilde{C}^{(j)}_{t,\kappa} - C^{(j)}_{t,\kappa} \|^p \right] + \mathbb{E}^n_{t, \lfloor \kappa n \rfloor - k_n + 1} \left[\| \widetilde{V}_{t,\kappa} - V_{t,\kappa} \|^p_F \right] \leq K_p \frac{k_n}{n}, \ \ p \geq 2, \qquad (41)$$

for some positive constant K_p, which does not depend on j.

Proof. The result follows from an application of Minkowski's inequality combined with Lemma 2 and 3 in Andersen et al. (2021). $\qquad\square$

Lemma 4. *Assume that Assumption SA holds. If $\Delta_n k_n \to 0$ and $k_n \Delta_n^{1-2\omega} \to \infty$, then it holds for each $t = 1, ..., T$ that,*

$$\mathbb{P}^n_{t, \lfloor \kappa n \rfloor - k_n + 1} \left[\lambda_{\min} \left(\widehat{V}_{t,\kappa} \right) \leq \alpha_n \right] \leq K_p \left(k_n^{p/2} \vee k_n^{1-p} \Delta_n^{1+p(2\omega-1)} \vee k_n/n \right), \ \ p \geq 2, \qquad (42)$$

where $\alpha_n \to 0$ is a deterministic sequence, and K_p is a constant that depends on p, but not n.

Proof. We can apply Weyl's inequality which states that for two symmetric $K \times K$ matrices A and B, $\lambda_1(A + B) \geq \lambda_1(A) + \lambda_K(B)$, where $\lambda_1(A) \leq ... \leq \lambda_K(A)$ are the ordered eigenvalues of A.

From here, we have for n sufficiently large,

$$1_{\{\lambda_{\min}(\widehat{V}_{t,\kappa})\leq\alpha_n\}} \leq 1_{\{\lambda_{\min}(V_{t,\kappa,n})\leq\underline{\lambda}\}} + 1_{\{|\lambda_{\min}(\widehat{V}_{t,\kappa}-\widetilde{V}_{t,\kappa})|\geq\underline{\lambda}/4\}} + 1_{\{|\lambda_{\min}(\widetilde{V}_{t,\kappa}-V_{t,\kappa,n})|\geq\underline{\lambda}/4\}}, \tag{43}$$

where $\underline{\lambda} > 0$ is some lower bound on the eigenvalues of the matrix process V by assumption SA. We now have,

$$\mathbb{P}^n_{t,\lfloor\kappa n\rfloor-k_n+1}\left(\left|\lambda_{\min}\left(\widehat{V}_{t,\kappa}-\widetilde{V}_{t,\kappa}\right)\right|\geq\underline{\lambda}/4\right) \leq K_p\mathbb{E}^n_{t,\lfloor\kappa n\rfloor-k_n+1}\|\widehat{V}_{t,\kappa}-\widetilde{V}_{t,\kappa}\|_F^p, \tag{44}$$

for each $p \geq 2$ and,

$$\mathbb{P}^n_{t,\lfloor\kappa n\rfloor-k_n+1}\left(\left|\lambda_{\min}\left(\widetilde{V}_{t,\kappa}-V_{t,\kappa,n}\right)\right|\geq\underline{\lambda}/4\right) \leq K\mathbb{E}^n_{t,\lfloor\kappa n\rfloor-k_n+1}\|\widetilde{V}_{t,\kappa}-V_{t,\kappa,n}\|_F^2, \tag{45}$$

The result then follows by an application of Lemmas 1-2. $\qquad\square$

Before stating the next lemma, we introduce the following notation:

$$\overline{Z}^{(j,k)}_{t,\kappa,n}(u) = iu\exp\left(iu\beta^{(j,k)}_{t,\kappa,n}\right)\frac{n}{k_n}\iota_k^\top V_{t,\kappa,n}^{-1}\sum_{s\in\mathcal{I}_\kappa^n}\left[\sigma_{t,\kappa,n}\Delta^n_{t,s}W\Delta^n_{t,s}B^\top\gamma^{(j)\top}_{t,\kappa,n} + \sigma_{t,\kappa,n}\Delta^n_{t,s}W\Delta^n_{t,s}\widetilde{W}^{(j)}\widetilde{\sigma}^{(j)}_{t,\kappa,n}\right], \tag{46}$$

and further denote,

$$\chi^{(a)}_{s,t,\kappa,k,n}(u) = \frac{iu}{N}\sum_{j=1}^N\exp\left(iu\beta^{(j,k)}_{t,\kappa,n}\right)\frac{n}{\sqrt{k_n}}\iota_k^\top V_{t,\kappa,n}^{-1}\sigma_{t,\kappa,n}\Delta^n_{t,s}W\Delta^n_{t,s}B^\top\gamma^{(j)\top}_{t,\kappa,n},$$

$$\chi^{(b)}_{s,t,\kappa,k,n}(u) = \frac{iu}{N}\sum_{j=1}^N\exp\left(iu\beta^{(j,k)}_{t,\kappa,n}\right)\frac{n}{\sqrt{k_n}}\iota_k^\top V_{t,\kappa,n}^{-1}\sigma_{t,\kappa,n}\Delta^n_{t,s}W\Delta^n_{t,s}\widetilde{W}^{(j)}\widetilde{\sigma}^{(j)}_{t,\kappa,n}.$$

We have,

$$\frac{\sqrt{k_n}}{N}\sum_{j=1}^N\overline{Z}^{(j,k)}_{t,\kappa,n}(u) = \sum_{s\in\mathcal{I}_\kappa^n}\left(\chi^{(a)}_{s,t,\kappa,k,n}(u) + \chi^{(b)}_{s,t,\kappa,k,n}(u)\right). \tag{47}$$

Using the uniform boundedness of $\{\widetilde{\sigma}^{(j)}\}_{j\geq1}$ on $[0,T]$ by Assumption SA and the independence of the Brownian motions $\{\widetilde{W}^{(j)}\}_{j\geq1}$ and W, we obtain,

$$\left\|\sum_{s\in\mathcal{I}_\kappa^n}\chi^{(b)}_{s,t,\kappa,k,n}\right\|_w = O_p\left(\frac{1}{\sqrt{N}}\right). \tag{48}$$

Lemma 5. *Assume that Assumption SA and B hold. If $\Delta_nk_n \to 0$, $k_n\Delta_n^{1-2\omega} \to \infty$, and $N \to \infty$,*

34

then it holds that,

$$\frac{\sqrt{k_n}}{N}\sum_{j=1}^{N}\overline{Z}_{t,\kappa,n}^{(j,k)} \xrightarrow{\mathcal{L}-s} \mathfrak{Z}_{t,\kappa,k}, \tag{49}$$

where $\mathfrak{Z}_{t,\kappa,k}$ is a complex-valued Gaussian process in $\mathcal{L}^2(w)$ with variance and relation operators given by $\psi_{t,\kappa,k}^{(1)}$ and $\psi_{t,\kappa,k}^{(2)}$.

Proof. First, we recall that for arbitrary functions f and g in $\mathcal{L}^2(w)$, the inner product is given by $\langle f, g \rangle_w = \int_{\mathbb{R}} f(u)\overline{g(u)}w(u)du$. Then, we will first show that the convergence holds finite-dimensionally. That is, that $\langle \frac{\sqrt{k_n}}{N}\sum_{j=1}^{N}\overline{Z}_{t,\kappa,n}^{(j)}, h \rangle_w \xrightarrow{\mathcal{L}-s} \langle \mathfrak{Z}_{t,\kappa}, h \rangle_w$, for arbitrary h in $\mathcal{L}^2(w)$, which can be shown by an application of Theorem IX.7.3 of Jacod and Shiryaev (2003). Given the result in (48), we need to establish that,

$$\sum_{s\in\mathcal{I}_\kappa^n}\mathbb{E}_{t,s}^n\left(\langle\chi_{s,t,\kappa,k,n}^{(a)},h\rangle_w\right)^2 \xrightarrow{\mathbb{P}} \int_{\mathbb{R}}\int_{\mathbb{R}}\psi_{t,\kappa,k}^{(1)}(u,z)\overline{h(u)}\,\overline{h(z)}w(u)w(z)dudz, \tag{50}$$

$$\sum_{s\in\mathcal{I}_\kappa^n}\mathbb{E}_{t,s}^n\left[\langle\chi_{s,t,\kappa,k,n}^{(a)},h\rangle_w\overline{\langle\chi_{s,t,\kappa,k,n}^{(a)},h\rangle_w}\right] \xrightarrow{\mathbb{P}} \int_{\mathbb{R}}\int_{\mathbb{R}}\psi_{t,\kappa,k}^{(2)}(u,z)\overline{h(u)}h(z)w(u)w(z)dudz, \tag{51}$$

$$\sum_{s\in\mathcal{I}_\kappa^n}\mathbb{E}_{t,s}^n\left[\langle\chi_{s,t,\kappa,k,n}^{(a)},h\rangle_w\Delta_{t,s}^n M\right] \xrightarrow{\mathbb{P}} 0, \tag{52}$$

where, for an arbitrary variable X, $\overline{X}$ denotes its complex conjugate, and M is a component of W, a component of B, or a bounded martingale which is orthogonal to those two processes. The convergence results in (50)-(51) follow upon noting that the volatility, beta, and gamma processes all are assumed to have càdlàg paths (Assumption A) and because of Assumption B. We are thus left with establishing (52). This result follows immediately if M is a component of W or B, due to the symmetry of the standard normal distribution. This leaves the case where M is a bounded martingale which is orthogonal to W and B. In this case, the conditional expectation in (52 is zero and the result in (52) trivially follows.

We are thus left with establishing tightness of the sequence. For this, we will make use of Theorem 1.8 in van der Vaart and Wellner (1996). Denote with $\{e_k\}_{k\geq 1}$ an orthonormal basis in $\mathcal{L}^2(w)$, then we need to show that,

$$\limsup_{n\to 0}\mathbb{P}\left(\sum_{k>J}\left|\langle\frac{\sqrt{k_n}}{N}\sum_{j=1}^{N}\overline{Z}_{t,\kappa,n}^{(j,k)},e_k\rangle_w\right|^2 > \delta\right) \to 0, \quad \text{as } J\to\infty, \tag{53}$$

for any $\delta > 0$ and its counterpart where $\left|\langle\frac{\sqrt{k_n}}{N}\sum_{j=1}^{N}\overline{Z}_{t,\kappa,n}^{(j,k)},e_k\rangle_w\overline{\langle\frac{\sqrt{k_n}}{N}\sum_{j=1}^{N}\overline{Z}_{t,\kappa,n}^{(j,k)},e_k\rangle_w}\right|$ replaces

$\left| \langle \frac{\sqrt{k_n}}{N} \sum_{j=1}^{N} \overline{Z}_{t,\kappa,n}^{(j,k)}, e_k \rangle_w \right|^2$. To show this, we note that,

$$
\mathbb{E} \left| \langle \frac{\sqrt{k_n}}{N} \sum_{j=1}^{N} \overline{Z}_{t,\kappa,n}^{(j,k)}, e_k \rangle_w \right|^2 + \mathbb{E} \left| \langle \frac{\sqrt{k_n}}{N} \sum_{j=1}^{N} \overline{Z}_{t,\kappa,n}^{(j,k)}, e_k \rangle_w \overline{\langle \frac{\sqrt{k_n}}{N} \sum_{j=1}^{N} \overline{Z}_{t,\kappa,n}^{(j,k)}, e_k \rangle_w} \right|
\tag{54}
$$

$$
\leq K \left(\int_{\mathbb{R}} |u| e_k(u) w(u) du \right)^2,
$$

because of the boundedness assumptions in SA. From here, the result to be shown follows by application of Bessel's inequality. $\qquad\square$

E   Proof of Theorem 1

A second order Taylor expansion, for $j = 1, ..., N$ and $k = 1, ..., q$, yields,

$$
\exp \left(i u \iota_k^\top \widehat{V}_{t,\kappa}^{-1} \widehat{C}_{t,\kappa}^{(j)} \right) - \exp \left(i u \iota_k^\top \widetilde{V}_{t,\kappa}^{-1} \widetilde{C}_{t,\kappa}^{(j)} \right) = Z_{t,\kappa,n}^{(j,k)}(u) + \overline{R}_{t,\kappa,n}^{(j,k)}(u),
$$

where α_n is a deterministic sequence, with $\alpha_n \asymp 1/log(n)$,

$$
Z_{t,\kappa,n}^{(j,k)}(u) = iu \exp \left(iu \widetilde{\beta}_{t,\kappa}^{(j,k)} \right) \left(\iota_k^\top \widetilde{V}_{t,\kappa}^{-1} \left(\widehat{C}_{t,\kappa}^{(j)} - \widetilde{C}_{t,\kappa}^{(j)} \right) - vec \left(\widetilde{\beta}_{t,\kappa}^{(j)} \iota_k^\top \widetilde{V}_{t,\kappa}^{-1} \right)^\top vec \left(\widehat{V}_{t,\kappa} - \widetilde{V}_{t,\kappa} \right) \right),
$$

and the residual term $\overline{R}_{t,\kappa,n}^{(j,k)}$ satisfies,

$$
|\overline{R}_{t,\kappa,n}^{(j,k)}(u)| \leq K \left[1_{\{\lambda_{min}(\widehat{V}_{t,\kappa}) \leq \alpha_n\}} + \alpha_n^{-2} \left(1 \vee u^2 \right) \left(\|\widehat{V}_{t,\kappa} - \widetilde{V}_{t,\kappa}\|_F^2 + \|\widehat{C}_{t,\kappa}^{(j)} - \widetilde{C}_{t,\kappa}^{(j)}\|^2 \right) \right],
\tag{55}
$$

for a constant K that does not depend on j or u. By an application of Lemmas 1-2 and using the exponential tail decay of w, we have,

$$
\mathbb{E} \left\| \frac{1}{N} \sum_{j=1}^{N} \overline{R}_{t,\kappa,n}^{(j,k)} \right\|_w \leq \frac{K}{\alpha_n^2} \left(\frac{1}{k_n} \vee \frac{k_n}{n} \vee \frac{\Delta_n^{4\varpi-1}}{k_n} \right).
\tag{56}
$$

We proceed with the difference $\frac{1}{N}\sum_{j=1}^{N}\left(Z_{t,\kappa,n}^{(j,k)} - \overline{Z}_{t,\kappa,n}^{(j,k)}\right)$. We start with making the following decomposition,

$$\widehat{C}_{t,\kappa}^{(j)} - \widetilde{C}_{t,\kappa}^{(j)} = \widehat{C}_{t,\kappa}^{(j)} - \overline{C}_{t,\kappa}^{(j)} + \mathcal{C}_{t,\kappa}^{(j)} + R_{t,\kappa,n}^{(j)},$$
$$\widehat{V}_{t,\kappa} - \widetilde{V}_{t,\kappa} = \widehat{V}_{t,\kappa} - \overline{V}_{t,\kappa} + \mathcal{V}_{t,\kappa} + R_{t,\kappa,n},$$

and further note that,

$$\mathcal{C}_{t,\kappa}^{(j)} = \mathcal{V}_{t,\kappa}\beta_{t,\kappa,n}^{(j)} + \frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\left[\sigma_{t,\kappa,n}\Delta_{t,s}^n W\Delta_{t,s}^n B^\top\gamma_{t,\kappa,n}^{(j)\top} + \sigma_{t,\kappa,n}\Delta_{t,s}^n W\Delta_{t,s}^n\widetilde{W}^{(j)}\widetilde{\sigma}_{t,\kappa,n}^{(j)}\right]. \tag{57}$$

Therefore, we have the decomposition,

$$Z_{t,\kappa,n}^{(j,k)} - \overline{Z}_{t,\kappa,n}^{(j,k)} = \left(iu\exp\left(iu\widetilde{\beta}_{t,\kappa}^{(j,k)}\right)\iota_k^\top\widetilde{V}_{t,\kappa}^{-1} - iu\exp\left(iu\beta_{t,\kappa,n}^{(j,k)}\right)\iota_k^\top V_{t,\kappa,n}^{-1}\right)$$

$$\times\left[\frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\left[\sigma_{t,\kappa,n}\Delta_{t,s}^n W\Delta_{t,s}^n B^\top\gamma_{t,\kappa,n}^{(j)\top} + \sigma_{t,\kappa,n}\Delta_{t,s}^n W\Delta_{t,s}^n\widetilde{W}^{(j)}\widetilde{\sigma}_{t,\kappa,n}^{(j)}\right]\right]$$

$$+ iu\exp\left(iu\widetilde{\beta}_{t,\kappa}^{(j,k)}\right)\left[\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\mathcal{V}_{t,\kappa}\beta_{t,\kappa,n}^{(j)} - vec\left(\widetilde{\beta}_{t,\kappa}^{(j)}\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\right)^\top vec\left(\mathcal{V}_{t,\kappa}\right)\right] \tag{58}$$

$$+ iu\exp\left(iu\widetilde{\beta}_{t,\kappa}^{(j,k)}\right)\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\left[\widehat{C}_{t,\kappa}^{(j)} - \overline{C}_{t,\kappa}^{(j)} + R_{t,\kappa,n}^{(j)}\right]$$

$$- iu\exp\left(iu\widetilde{\beta}_{t,\kappa}^{(j,k)}\right)vec\left(\widetilde{\beta}_{t,\kappa}^{(j)}\iota^\top\widetilde{V}_{t,\kappa}^{-1}\right)^\top vec\left(\widehat{V}_{t,\kappa} - \overline{V}_{t,\kappa} + R_{t,\kappa,n}\right).$$

Next, we note that,

$$vec\left(\widetilde{\beta}_{t,\kappa}^{(j)}\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\right)^\top vec\left(\frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\sigma_{t,\kappa,n}\left(n\Delta_{t,s}^n W\Delta_{t,s}^n W^\top - I_q\right)\sigma_{t,\kappa,n}^\top\right)$$
$$= \iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\frac{n}{k_n}\sum_{s\in\mathcal{I}_\kappa^n}\sigma_{t,\kappa,n}\left(n\Delta_{t,s}^n W\Delta_{t,s}^n W^\top - I_q\right)\sigma_{t,\kappa,n}^\top\widetilde{\beta}_{t,\kappa}^{(j)}, \tag{59}$$

from which it follows that,

$$\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\mathcal{V}_{t,\kappa}\beta_{t,\kappa,n}^{(j)} - vec\left(\widetilde{\beta}_{t,\kappa}^{(j)}\iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\right)^\top vec\left(\mathcal{V}_{t,\kappa}\right) = \iota_k^\top\widetilde{V}_{t,\kappa}^{-1}\mathcal{V}_{t,\kappa}\left[\beta_{t,\kappa,n}^{(j)} - \widetilde{\beta}_{t,\kappa}^{(j)}\right]. \tag{60}$$

From here, by application of Lemmas 1-3, we obtain,

$$\mathbb{E}\left\|\frac{1}{N}\sum_{j=1}^{N}\left(Z_{t,\kappa,n}^{(j,k)} - \overline{Z}_{t,\kappa,n}^{(j,k)}\right)\right\|_w \leq K\left(k_n^{-1} \vee \Delta_n^{2\omega} \vee \sqrt{\frac{k_n}{n}}\right). \tag{61}$$

Combining the above bounds, together with Lemma 3, we have the result of the theorem.

F Proof of Theorem 2

We notice that given the properties of $\{e_{s,t,\kappa}\}_{s\geq 1}$, we have $\mathbb{E}\left(Z_{t,\kappa,k}^{(j)}\big|\mathcal{F}\right)=0$ and,

$$
\begin{aligned}
\mathbb{E}\left(\left(\widetilde{Z}_{t,\kappa,k}^{(j)}\right)^2\bigg|\mathcal{F}\right)=\frac{n^2}{k_n^2}\iota_k^\top\widehat{V}_{\mathcal{T},\kappa}^{-1}\sum_{s\in\mathcal{I}_\kappa^n}&\left(\Delta_{t,s}^n X^{(j)}-\widehat{\beta}_{t,\kappa}^{(j)\top}\Delta_{t,s}^n F\right)\Delta_{t,s}^n F \\
&\times\Delta_{t,s}^n F^\top\left(\Delta_{t,s}^n X^{(j)}-\widehat{\beta}_{t,\kappa}^{(j)\top}\Delta_{t,s}^n F\right)^\top 1_{\{\mathcal{B}_{t,s}^{(j)}\}}\widehat{V}_{\mathcal{T},\kappa}^{-1}\iota_k.
\end{aligned}
\tag{62}
$$

Using this and Lemmas 1-4, it is easy to show that $\widehat{\beta}_{t,\kappa}^{(j)}\xrightarrow{\mathbb{P}}\beta_{t-1+\kappa}^{(j)}$, and further that,

$$
k_n\mathbb{E}\left(\left(\widetilde{Z}_{t,\kappa,k}^{(j)}\right)^2\bigg|\mathcal{F}\right)\xrightarrow{\mathbb{P}}\iota_k^\top V_{t-1+\kappa}^{-1}\sigma_{t-1+\kappa}\gamma_{t-1+\kappa}^{(j)\top}\gamma_{t-1+\kappa}^{(j)}\sigma_{t-1+\kappa}V_{t-1+\kappa}^{-1}\iota_k.
\tag{63}
$$

From here, the proof of the theorem follows exactly the same steps as those in the proof of Lemma 5.

References

Ai, H. and R. Bansal (2018). Risk preferences and the macroeconomic announcement premium. *Econometrica 86*(4), 1383–1430.

Aït-Sahalia, Y., I. Kalnina, and D. Xiu (2020). High-frequency factor models and regressions. *Journal of Econometrics 216*(1), 86–105.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and J. G. Wu (2005a). A framework for exploring the macroeconomic determinants of systematic risk. *American Economic Review 95*, 398–404.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and J. G. Wu (2005b). Realized beta: persistence and predictability. In T. Fomby (Ed.), *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series*, Volume B, pp. 1–40.

Andersen, T. G., M. Thyrsgaard, and V. Todorov (2021). Recalcitrant betas: Intraday variation in the cross-sectional dispersion of systematic risk. *Quantitative Economics 12*(2), 647–682.

Barndorff-Nielsen, O. E. and N. Shephard (2004). Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica 72*(3), 885–925.

Bernanke, B. S. and K. Kuttner (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of Finance 60*(3), 1221–1257.

Bogousslavsky, V. (2021). The cross-section of intraday and overnight returns. *Journal of Financial Economics 141*(1), 172–194.

Bollerslev, T. and V. Todorov (2011). Estimation of jump tails. *Econometrica 79*(6), 1727–1783.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative models for stock price dynamics. *Journal of Econometrics 116*(1), 225–257.

Cieslak, A., A. Morse, and A. Vissing-Jorgensen (2019). Stock returns over the fomc cycle. *The Journal of Finance 74*(5), 2201–2248.

Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton university press.

Connor, G., M. Hagmann, and O. Linton (2012). Efficient semiparametric estimation of the fama–french model and extensions. *Econometrica 80*(2), 713–754.

Daniel, K. and T. J. Moskowitz (2016). Momentum crashes. *Journal of Financial Economics 122*, 221–247.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics 116*(1), 1–22.

Fan, J., Y. Liao, and W. Wang (2016). Projected principal component analysis in factor models. *Annals of statistics 44*(1), 219.

Gagliardini, P., E. Ossola, and O. Scaillet (2016). Time-varying risk premium in large cross-sectional equity data sets. *Econometrica 84*(3), 985–1046.

Garcia-Schmidt, M. and M. Woodford (2019, January). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review 109*(1), 86–120.

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. *Journal of Finance 51*(1), 3–53.

Kalnina, I. (2020). Inference for nonparametric high-frequency estimators with an application to time variation in betas.

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.

Li, J., V. Todorov, and G. Tauchen (2017). Adaptive estimation of continuous-time regression models using high-frequency data. *Journal of Econometrics 200*(1), 36–47.

Lucca, D. O. and E. Moench (2015). The pre-fomc announcement drift. *The Journal of Finance 70*(1), 329–371.

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.

Mykland, P. and L. Zhang (2009). Inference for continuous semimartingales observed at high frequency. *Econometrica 77*, 1403–1445.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics 133*(3), 1283–1330.

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.

Savor, P. and M. Wilson (2014). Asset pricing: A tale of two days. *Journal of Financial Economics 113*(2), 171–201.

Shanken, J. (1990). Intertemporal asset pricing: An empirical investigation. *Journal of Econometrics 45*(1-2), 99–120.

van der Vaart, A. W. and J. A. Wellner (1996). *Weak convergence and empirical processes*. New York: Springer.

Zhang, C., J. Li, V. Todorov, and G. Tauchen (2020). Variation and efficiency of high-frequency betas. *Journal of Econometrics*.