

Jump factor models in large cross-sections

JIA LI

Department of Economics, Duke University

VIKTOR TODOROV

Department of Finance, Kellogg School of Management, Northwestern University

GEORGE TAUCHEN

Department of Economics, Duke University

We develop tests for deciding whether a large cross-section of asset prices obey an exact factor structure at the times of factor jumps. Such jump dependence is implied by standard linear factor models. Our inference is based on a panel of asset returns with asymptotically increasing cross-sectional dimension and sampling frequency, and essentially no restriction on the relative magnitude of these two dimensions of the panel. The test is formed from the high-frequency returns at the times when the risk factors are detected to have a jump. The test statistic is a cross-sectional average of a measure of discrepancy in the estimated jump factor loadings of the assets at consecutive jump times. Under the null hypothesis, the discrepancy in the factor loadings is due to a measurement error, which shrinks with the increase of the sampling frequency, while under an alternative of a noisy jump factor model this discrepancy contains also nonvanishing firm-specific shocks. The limit behavior of the test under the null hypothesis is non-standard and reflects the strong-dependence in the cross-section of returns as well as their heteroskedasticity which is left unspecified. We further develop estimators for assessing the magnitude of firm-specific risk in asset prices at the factor jump events. Empirical application to S&P 100 stocks provides evidence for exact one-factor structure at times of big market-wide jump events.

KEYWORDS. Factor model, panel, high-frequency data, jumps, semimartingale, specification test, stochastic volatility.

JEL CLASSIFICATION. C51, C52, G12.

1. INTRODUCTION

Asset prices often “jump” simultaneously in response to important market-wide events such as macroeconomic announcements, political news, and natural disasters. Linear

Jia Li: j1410@duke.edu

Viktor Todorov: v-todorov@northwestern.edu

George Tauchen: george.tauchen@duke.edu

Li's work was partially funded by NSF grant SES-1326819. Todorov's work was partially funded by NSF grant SES-1530748. We would like to thank three anonymous referees as well as seminar participants and discussants at various conferences for their thoughtful comments, which helped greatly improve the paper.

factor models which are widely used in asset pricing have strong implications regarding this co-jump structure of assets in the cross-section. To fix ideas, consider a continuous-time factor model given by

$$dY_t = \alpha_t dt + \mathbf{B} dF_t + d\tilde{Y}_t, \quad (1.1)$$

where Y_t is a $N \times 1$ vector of asset prices, F_t is a $r \times 1$ vector of factors capturing systematic risk, α_t is a drift term due to the compensation for risk and time demanded by investors for holding the stocks, \mathbf{B} is a $N \times r$ matrix of factor loadings, and \tilde{Y}_t is a $N \times 1$ vector of idiosyncratic risks. Idiosyncratic risk is formally defined as being orthogonal to the systematic one, that is, having zero quadratic covariation with it. Regarding the jumps, this orthogonality condition means $\Delta F_t \Delta \tilde{Y}_t = 0$ for every t , where for a generic process Z we denote with $\Delta Z_t = Z_t - Z_{t-}$ its jump at time t .¹ Therefore, we have

$$\Delta Y_\tau = \mathbf{B} \Delta F_\tau, \quad \text{for } \tau \text{ such that } \Delta F_\tau \neq 0. \quad (1.2)$$

That is, at the times of the factor jumps, the whole cross-section of asset prices follows an exact factor model with no idiosyncratic risk, that is, \tilde{Y} is not present in (1.2).²

Moreover, if the jumps in the components of the vector F are normalized such that they are orthogonal to each other, that is, have zero quadratic covariation, then we also have

$$\Delta Y_\tau = b \Delta Z_\tau, \quad \text{for } \tau \in \mathcal{T} \equiv \{t \in [0, T] : \Delta Z_t \neq 0\}, \quad (1.3)$$

where Z_t is one of the elements of the vector F_t and b is the corresponding vector of factor loadings in the matrix \mathbf{B} .³ That is, we have an exact one-factor structure of the jumps in Y at the jump times of the univariate factor Z_t . This one-factor structure of jumps in Y holds true for the jump times of each of the factors in F , but with different factor loadings in general.

Our goal in this paper is to test whether a hypothesis of exact factor structure for the systematic jumps in (1.3) is true.⁴ The test is based on a large cross-section of assets sampled at high-frequency on a time interval of fixed length. In its most basic form,

¹In many applications, orthogonality of systematic and idiosyncratic risk is defined in a much stronger sense by imposing independence between F and \tilde{Y} .

²Put differently, if $\tau_1, \tau_2, \dots, \tau_k$ are the jump times of the systematic factors in the observation interval $[0, T]$ (these times are in general random) and if without loss of generality we assume $k > r$ and $N > r$, then the matrix $[\Delta Y_{\tau_1}, \Delta Y_{\tau_2}, \dots, \Delta Y_{\tau_k}]$ is of reduced-rank r and the same holds true for the quadratic variation of Y at the factor jump times $\sum_{i=1}^k \Delta Y_{\tau_i} \Delta Y_{\tau_i}^\top$. Only when we aggregate the quadratic variation of *all* jump risk in Y over the interval $[0, T]$, given by $\sum_{t \in [0, T]} \Delta Y_t \Delta Y_t^\top$, we restore the full rank of N . This is because $\sum_{t \in [0, T]} \Delta Y_t \Delta Y_t^\top$ includes also the quadratic variation due to the jumps in the idiosyncratic component \tilde{Y} (which happen outside the set \mathcal{T} of factor jump times) which is a diagonal matrix.

³Note that b is a $N \times 1$ vector that contains the slope coefficients (i.e., betas) for the N assets. There is no restriction on the relationship among individual assets' betas.

⁴When the factors in the vector F do not have orthogonal jump components, then for $\tau \in \mathcal{T}$, we will have in general an exact multifactor model. The analysis of the paper can be extended to cover this more general setup but at the cost of significantly more involved derivations and notation. Given the empirical evidence presented in Section 6, however, such an extension seems to be of little practical relevance.

which will be generalized in our formal analysis, our test discriminates (1.3) from its “noisy” counterpart given by

$$\Delta Y_\tau = b\Delta Z_\tau + \tilde{\chi}_\tau, \quad \text{for } \tau \in \mathcal{T}, \quad (1.4)$$

where $\tilde{\chi}_\tau$ is a random $N \times 1$ vector that captures deviations from the exact factor model at the systematic jump events; we refer to (1.4) as the noisy linear factor model henceforth. This model can be equivalently written in the form of a random jump beta model as

$$\beta_\tau \equiv \frac{\Delta Y_\tau}{\Delta Z_\tau} = b + \tilde{\chi}_\tau, \quad (1.5)$$

where $\tilde{\chi}_\tau \equiv \tilde{\chi}_\tau/\Delta Z_\tau$.⁵ There are many reasons for the presence of $\tilde{\chi}_\tau$ in the jump factor loadings (i.e., betas). One of them is predictable time-variations in these loadings. Indeed, factor loadings in asset pricing are often modeled as functions of assets’ characteristics and/or macro state variables; see, for example, Connor, Hagmann, and Linton (2012) and Gagliardini, Ossola, and Scaillet (2016b), and the theoretical analysis of Hansen and Richard (1987). Typically, such time-variation in the factor loadings happens at low frequencies and by performing our tests on intervals of short time span we will minimize the possibility for such a violation of the null hypothesis in (1.3). Another reason for $\tilde{\chi}_\tau$ in (1.4) is the presence of (locally unpredictable) firm-specific shocks in Y even at the times of systematic jumps. In this case, we have $\mathbb{E}_{\tau-}[\tilde{\chi}_\tau\Delta Z_\tau] = 0$, that is, $\tilde{\chi}_\tau$ is uncorrelated with the jump ΔZ_τ conditional on the information prior to the jump time τ , and (1.3) gets replaced by the weaker linear projection condition

$$\mathbb{E}_{\tau-}[\Delta Y_\tau\Delta Z_\tau] = b\mathbb{E}_{\tau-}[\Delta Z_\tau^2]. \quad (1.6)$$

Obviously, (1.3) implies (1.6) but the reverse is generally not true. The test we propose here will be able to discriminate (1.3) from alternatives under which we have firm-specific shocks at the systematic jump events and only (1.6) holds.

From a practical point of view, separating the null from such an alternative hypothesis is important as this has strong implications regarding the inference for the matrix of jump factor loadings \mathbf{B} . Indeed, if the null hypothesis holds, we can use a fixed (small) number of systematic jump times to identify the (constant) jump factor loadings. On the other hand, if only (1.6) holds, then we need to use a long span of data to identify the permanent component of \mathbf{B} , which would be further complicated by the possible time-variation in the factor loadings at low frequencies. The separation of the null from the alternative hypothesis is also important for practical risk management decisions and more generally can help shed light on the sources of priced jump risk in the cross-section of asset prices.

To motivate empirically our theoretical analysis, we show in Figure 1 the cross-sectional relationship between stock returns versus their market jump betas for S&P

⁵The random jump betas may be viewed as “random coefficients” like in classical panel data analysis. However, we note that the jump betas only concern the jump returns, which form a small subset of all high-frequency returns.




FIGURE 1. Stock jump returns versus jump beta. *Note:* The vertical axis is the time series average 2-minute stock return at times when a positive jump in the S&P 500 index futures is detected; the horizontal axis is the average jump beta measured from 2-minute asset returns at times when a negative jump in the index futures is detected. The straight line is a linear fit implied by the exact linear jump factor model in (1.3) and the recovered jump betas. The panel comprises S&P 100 stocks, 2007–2015, and the market jump returns are selected according the thresholding procedure described in Section 5.

100 stocks: the 2-minute average stock returns are computed at the times when a positive market jump is detected (i.e., $\Delta Z_\tau > 0$), and the jump betas are calculated from the 2-minute asset returns at the times when a negative market jump is detected (i.e., $\Delta Z_\tau < 0$). We use negative jump betas because the estimation error is relatively lower for negative jumps, while we measure average stock returns at times of positive market jumps to maintain complete separation of the data sets for computing the variables on the x and y axes of the figure. This cross-sectional relationship appears quite tight, even though the market betas were computed from only forty-four 2-minute returns when negative market jumps were detected in our sample—hence the question: Is this fit consistent with the exact linear jump factor model (1.3)?

To address this question, we develop a formal test for deciding whether such hypothesis is true and we further propose a measure for the magnitude of the firm-specific risk at factor jump events. Our asymptotic theory employs a joint asymptotic setting in which both the cross-sectional dimension of the panel and the sampling frequency increase simultaneously, while keeping the time span of the sample fixed. To compute the test statistic, we first use high-frequency returns to nonparametrically estimate the individual assets' betas at the jump times of Z ; the test statistic is then formed as the cross-

sectional average of the temporal variations in these beta estimates. Under the null hypothesis, changes in the beta estimates across jump times are due to high-frequency estimation error that shrinks to zero asymptotically, but is nondegenerate under the alternative. We characterize theoretically the asymptotic behavior of the test statistic and propose an easy-to-implement algorithm for computing its critical values.

The asymptotic theory underlying our test is nonstandard for several reasons. First, we allow individual assets' returns to be *strongly* dependent in the cross-section, through their loadings on the common diffusive factors. Importantly, these systematic shocks are not “averaged out” under the cross-sectional aggregation (only the loadings on them are), but they remain to have nondegenerate distributions in the limit. As a result, unlike conventional econometric settings, the limiting behavior of the aggregated systematic shocks is not obtained from a central limit theory for weakly dependent data; instead, it is implied by the local Gaussianity of diffusive processes. Second, the limit distribution of our test statistic depends on the spot covariance process of the *latent* diffusive factors and individual assets' time-varying loadings on them within local windows around jump times. Therefore, our feasible inference on the panel of jumps actually involves estimating characteristics of the diffusive factor component as a by-product. We provide novel theoretical results in this direction, which are further used to construct critical values of our test. Third, the limit behavior of our test statistic is of “mixing” type. More precisely, the limit distribution can be realized on an extension of the original probability space and changes (in general) depending on the realizations of various sources of randomness on the original space, for example, it depends on the realized value of the systematic risk factors. Finally, we impose essentially no assumption on the relative growth rates of the cross-section dimension with respect to the sampling frequency. This accommodates situations in which either the cross-section or the sampling frequency is much higher than the other dimension of our panel and is very convenient empirically.

Going one step further, we develop an estimator for assessing the magnitude of the firm-specific risk at the times of systematic jump events. In particular, we measure the cross-sectional mean of the squared difference $(\tilde{\chi}_{j,\tau} - \tilde{\chi}_{j,\eta})^2$ for two jump times τ and η , where $\tilde{\chi}_{j,\tau}$ denotes the j th element of $\tilde{\chi}_\tau$. This measure provides an estimate of the temporal variation in jump betas for a “representative” asset.

We examine the performance of our inference procedures on simulated data from a model that captures salient features of typical financial data sets. In an empirical application to high-frequency data on stocks in the S&P 100 index covering the period 2007–2015, we test the exact jump factor model in (1.3) with the sole jump factor being the market jump. Our results provide strong empirical support for the model with no statistically significant role for firm-specific shocks in the assets at the market-wide extreme events.⁶ These findings are consistent with the empirical results in Savor and Wilson (2014) regarding the validity of CAPM around macroeconomic announcement days as well as earlier empirical evidence; see, for example, Longin and Solnik (2001) and Ang and Chen (2002), for increased asset correlation during extreme market events.

⁶We note, however, that our evidence is for the “big” jumps only as separating the “small” jumps from the diffusive component of returns is statistically much harder.

The theoretical results of the current paper are related to several strands of literature. First, the nonparametric separation of jumps from diffusive volatility using high-frequency data was initiated in work by Barndorff-Nielsen and Shephard (2004, 2006) and Mancini (2001, 2009).

Second, in Li, Todorov, and Tauchen (2017), we developed univariate methods for estimating and testing the validity of the exact jump factor model for a single asset, while being silent about the co-jump behavior of the cross-section of assets. We further extended these results to a cross-section of *fixed* size in Li, Todorov, Tauchen, and Lin (2017b). The asymptotics in these papers is of in-fill type for high-frequency data. By sharp contrast, here we consider a *joint* asymptotic setup in which the number of assets also grows asymptotically so as to accommodate the large cross-section. Relative to prior work, a key theoretical challenge in the current paper is the strong dependence among a large cross-section of asset returns resulting from the latent factor structure of the diffusive price components. We address this problem by developing factor-analytical tools in the spirit of Stock and Watson (2002) and Bai (2003, 2009) (see also the recent work of Gagliardini, Ossola, and Scaillet (2016a)), but in the nonstandard (infill, nonergodic, and nonstationary) setting for a large panel of high-frequency data. This theoretical innovation is absent from Li, Todorov, and Tauchen (2017) and Li et al. (2017b). The novel econometric setup of the current paper thus leads to a very different inference procedure and asymptotic theory than prior work, reflected in particular in the distinct roles of the systematic and idiosyncratic diffusive shocks in the limiting behavior of the test. In addition, unlike our prior work, the current setup allows us to study how the whole cross-section of assets reacts to economy-wide events, and further make inference for the firm-specific risk at the systematic factor jump times in the case when the null hypothesis of an exact jump factor model is not satisfied.

Third, Pelger (2015, 2019) developed methods for determining the number of systematic jump factors and further proposes inference procedures for recovering the latent systematic jump factors from a large cross-section of high-frequency return data. The inference in Pelger (2015, 2019) is based on the quadratic variation of the jumps of the assets over the observation interval. The latter includes the contribution of the idiosyncratic jump risks, and hence, it cannot be used to separate between the null in (1.3) and alternatives that satisfy (1.6).

Fourth, our results are also related to the panel data literature but with very distinctive features. Unlike conventional panel regressions, the identification strategy underlying our inference for jump betas can be viewed as “identification-by-discontinuity.”⁷ An important consequence is that we can use high-frequency data at a *fixed* number of jump events to estimate firm-specific jump betas nonparametrically. The “error terms” generated from this nonparametric measurement consist of diffusive returns in the vicinity of jumps, which shrink with the sampling interval, and the asymptotics of our statistics is mainly driven by their local Gaussianity (which is a generic property

⁷Identification-by-discontinuity is also widely used in microeconomic applications with regression discontinuity designs (RDD); see, for example, Lee and Lemieux (2010) for a review. However, we use jumps to identify the equilibrium relationship between asset prices, instead of a causal effect as is typically done in the RDD.

of Itô processes). Our in-fill asymptotic theory works under very weak conditions allowing for essentially unrestricted nonstationarity and heteroskedasticity for the underlying stochastic processes in a nonergodic setting. The theory in the current paper is thus very different from conventional panel data settings where weakly dependent disturbance terms in the regression equation are “averaged out” with a Gaussian limit distribution, which in turn drives the asymptotics underlying the econometric inference.

Finally, since we allow the diffusive “error terms” to have a latent factor structure, our results are closely related to the growing literature in panel data analysis with interactive fixed effects; see, for example, [Pesaran \(2006\)](#) and [Bai \(2009\)](#). Besides the very different settings mentioned above, we stress an important unique feature in our asymptotic theory relative to this literature. Unlike [Bai \(2009\)](#) where the asymptotic distribution of the estimator is only driven by the aggregated idiosyncratic shocks,⁸ here in our setting systematic shocks in the error term also have a non-degenerate contribution to the limit distribution. Intuitively, the systematic shocks are clearly not “averaged out” in the cross-section and, because we consider a fixed number of jump events (as they are rare), they also “survive” the time-series aggregation. Nevertheless, we show that feasible inference is still possible by relying on the local Gaussianity of the underlying Itô process, formally represented as a stable convergence in law under the in-fill asymptotics. In addition, we allow the latent diffusive factors to have general stochastic volatility and the stocks’ loadings on them to be (nonparametrically) time-varying; indeed, we allow these processes to jump, which leads to an additional source of “nondiversifiable” sampling variability due to the indeterminacy of exact systematic jump times within a discrete observation interval. Another useful feature of our theory is that it does not require knowing exactly the number of latent diffusive factors; this feature is clearly desirable in practice and has been recently studied by [Moon and Weidner \(2015\)](#) in standard linear panel regressions with interactive fixed effects using different techniques.

The rest of this paper is organized as follows. In Section 2, we introduce the econometric setting. Our main theoretical results are given in Section 3 where we propose a pairwise test for the exact jump factor model and derive feasible limit theory for it. In this section, we also develop estimators for assessing the magnitude of the firm-specific risk in assets at the times of the systematic jump events. Section 4 proposes various extensions to our main theory. Section 5 contains results from a Monte Carlo study and Section 6 presents our empirical application. Section 7 concludes. All proofs are given in the Online Supplementary Material ([Li, Todorov, and Tauchen \(2019\)](#)).

2. THE ECONOMETRIC SETUP AND ASSUMPTIONS

2.1 *The panel of jump betas*

We start with introducing formally the asset price dynamics on the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. The dynamics for the vector of asset prices Y_t , to which our analysis applies, generalizes the linear factor model in (1.1) and is given by

$$dY_{j,t} = \alpha_{j,t} dt + \lambda_{j,t}^\top df_t + dJ_{Y,j,t} + d\epsilon_{j,t}, \quad 1 \leq j \leq N, \quad (2.1)$$

⁸See Proposition A.3 of [Bai \(2009\)](#), which is used in the proofs of the main theorems in that paper.

where as in the [Introduction](#) $\alpha_{j,t}$ is a drift term in the asset price. Here, we denote with f_t a r -dimensional diffusive factor process and with $\lambda_{j,t}$ the time-varying loading on it, and $J_{Y,j,t}$ denotes the jump component of $Y_{j,t}$. Finally, $\epsilon_{j,t}$ is a one-dimensional diffusive idiosyncratic component that is orthogonal to the systematic factor f_t (in the sense that it has zero quadratic covariation with f_t ; see [Assumption 3](#) below).

The main object of interest of the current paper is the factor structure of the jump component $J_{Y,j,t}$. The diffusive component $\int_0^t \lambda_{j,s}^\top df_s + \epsilon_{j,t}$, on the other hand, plays the role of a “disturbance” as in classical regression settings.⁹ We stress that these “disturbance” terms in the asset prices are allowed to be strongly dependent in the cross-section through their loadings on the common factors. This has important implications for the asymptotic theory and is a nontrivial departure from classical panel data applications in econometrics.

Turning to the jump component $J_{Y,j,t}$, we will be interested in its behavior only at the jump times of an one-dimensional reference asset (systematic factor) which we denote with Z . In our application, Z will be the market portfolio, and hence our attention will be on the co-jump behavior of the cross-section of assets during market-wide events, located at times of “big” market jumps. Economically speaking, such jumps are mostly due to important public news arrival, such as macroeconomic announcements, major political events, and natural disasters. In addition to co-jumping with Z , we will also allow Y to jump at different times. As the model in the [Introduction](#), this can allow for asset prices to have idiosyncratic jumps as in [Merton \(1976\)](#) (in the sense that these jumps are asset-specific and arrive at different times) as well as to have exposure to systematic jump factors which jump at different times than Z .

The dynamics of the reference asset Z is given by

$$dZ_t = \lambda_{Z,t}^\top df_t + dJ_{Z,t}, \quad (2.2)$$

with $J_{Z,t}$ denoting its jump component. When Z is the market portfolio, it will typically be one of the factors itself as in the market factor model and its extensions. Hence, in this case, one can set the first element of f to be the continuous part of Z and $\lambda_{Z,t}$ will be a vector with the first entry being 1 and the rest being zero.

We collect the jump times of Z in the random set $\mathcal{T} \equiv \{t \in [0, T] : |\Delta Z_t| \neq 0\}$. The relationship between the asset price jumps and those of the one-dimensional Z is succinctly summarized by the jump betas of individual assets with respect to Z . At each jump time $\tau \in \mathcal{T}$, the *spot jump beta* of asset j is simply defined by

$$\beta_{j,\tau} \equiv \frac{\Delta Y_{j,\tau}}{\Delta Z_\tau}.$$

We note that these betas are defined in a nonparametric fashion and in general they are random quantities. However, when a factor model as noted in the [Introduction](#) is in

⁹The diffusive component of returns can have different factor structure and factor loadings from that of the jump component, and this can have asset pricing implications; see, for example, [Bollerslev, Li, and Todorov \(2016\)](#) for the case when the factor is the return on the market portfolio. We leave for future work the development of tests for deciding whether diffusive and jump factor loadings are the same and efficient inference techniques under such a scenario.

force, the above spot jump betas will be constant. That is, we will have

$$\beta_{j,\tau} = \beta_j \iff \Delta Y_{j,\tau} = \beta_j \Delta Z_\tau \quad \text{for all } \tau \in \mathcal{T}.$$

We note that for the above to be true, we do not require that Z is the only jump factor that affects the asset prices in the vector Y . This is because the above exact one-factor relationship *only* concerns the set of jump times \mathcal{T} . Jumps of Y at times outside the set \mathcal{T} do not need to obey this exact factor structure.

In this paper, we are interested in testing whether the above constant beta restriction holds *jointly* for a large cross-section of assets (N will be asymptotically increasing). This is clearly a very strong hypothesis, especially when \mathcal{T} involves many jump times. Hence, we start with the shortest possible event horizon by focusing on pairs of consecutive jump times of the systematic factor Z . That is, we examine, for the successful jump times $\eta, \tau \in \mathcal{T}$ with $\eta < \tau$,

$$\beta_{j,s} = \beta_j, \quad s \in \{\eta, \tau\} \text{ all } 1 \leq j \leq N. \quad (2.3)$$

The extension to the more general case with multiple jump times is discussed in Section 4.1.

The hypothesis of the constant jump beta in (2.3) can be nested within a more general alternative of the random jump beta model given by

$$\beta_{j,s} = \beta_j + \tilde{\chi}_{j,s}, \quad 1 \leq j \leq N. \quad (2.4)$$

The constant beta model (2.3) is a special case of (2.4) with $\tilde{\chi}_{j,s} = 0$ identically. The random term $\tilde{\chi}_{j,s}$ captures the possible violation of an exact linear one-factor model for the jumps $\{\Delta Y_{j,s}\}_{s \in \{\eta, \tau\}}$. In particular, if there are firm-specific shocks at the times when Z jumps, then this risk will be reflected in $\tilde{\chi}_{j,s}$. Our test will be able to discriminate this alternative from the exact jump factor model in (2.3).

We stress the importance of separating the null from the alternative hypothesis on a practical level. If the null hypothesis is true, then one can in principle estimate consistently the jump beta from a single jump event or, more generally from a fixed number of jump events. Intuitively, by zooming into the short window around the big jump of Z , we obtain high signal-to-noise measurement of the shocks to the underlying efficient prices. This gives rise to a type of “identification-by-discontinuity.” On the other hand, if $\tilde{\chi}_{j,s}$ in (2.4) is nondegenerate, then we cannot estimate consistently the permanent component β_j from a fixed number of jump events. This is because $\tilde{\chi}_{j,s}$ is $O_p(1)$, and hence, one would need to appeal to long-span asymptotics in this case in order to “average out” $\tilde{\chi}_{j,s}$ (this is exactly as in the first-step of the classical Fama–MacBeth regression that is widely used in cross-sectional asset pricing).

By focusing on two jump times, the structure of (2.4) resembles that of the “large N small T ” setting in microeconomic panel data analysis. In particular, the time-invariant coefficient β_j plays the role of a fixed effect (or the permanent component) and $\tilde{\chi}_{j,s}$ plays the role of a random shock (or the transitory component) in the spot jump betas.

The challenge in our analysis is that the jump betas $(\beta_{j,\eta}, \beta_{j,\tau})$ are not directly observed from discrete return data. We will use data sampled at high frequencies to recover them nonparametrically. The econometric problem at hand may thus be classified as one with large cross-sectional dimension, short, and fixed time span, and with nonparametrically generated dependent variables (i.e., the high-frequency beta estimates). In addition, due to the factor structure in the diffusive returns, the estimation errors for these betas are *strongly* cross-sectionally dependent with nonstandard asymptotic distribution, which further sets our study apart from prior work in microeconomic panels.

We now describe our test statistic. If the jump betas were observed, we could use the following (infeasible) statistic to detect deviations from the constant beta restriction:

$$V_N^* \equiv \frac{1}{N} \sum_{j=1}^N L(|\beta_{j,\tau} - \beta_{j,\eta}|),$$

where $L(\cdot)$ is a loss function that we use to gauge the temporal variation in the jump betas. In particular, if the innovation terms $\tilde{\chi}_{j,\tau} - \tilde{\chi}_{j,\eta}$ are nondegenerate, we expect V_N^* to converge to the cross-sectional mean of the loss $L(|\beta_{j,\tau} - \beta_{j,\eta}|)$, which summarizes the temporal variation in beta for any “average” stock.

In this paper, we consider loss functions that are “quadratic near zero.” More precisely, we maintain the following assumption for the loss function.

ASSUMPTION 1. *The function $L : \mathbb{R} \mapsto [0, \infty)$ satisfies the following: (i) for some fixed $\bar{x} \in [0, \infty]$, $L(x) = x^2$ for all $x \in [0, \bar{x}]$; (ii) $L(x)$ is nondecreasing in $|x|$; (iii) $L(\cdot)$ is Lipschitz on bounded sets.*

Assumption 1 ensures that $L(\cdot)$ behaves for our testing purposes like the quadratic loss function under the null hypothesis. On the other hand, the finite-sample power of the test naturally depends on the global shape of the loss function. In this way, Assumption 1 provides some flexibility in “directing” the power of the test. For example, one may consider a “Huber-like” loss function $L(x) = \min\{|x|, x^2\}$ that is less sensitive to outliers than the quadratic one (Huber (2004)).

Our feasible test statistic is constructed as a sample analogue of V_N^* by replacing the spot betas with their high-frequency estimates. We suppose that the processes Y_j and Z are observed at discrete times $\{i\Delta_n : 0 \leq i \leq [T/\Delta_n]\}$, where Δ_n is the sampling interval. For the limit theory that we develop below, we consider an asymptotic setting in which the sample span T is fixed, but both $\Delta_n \rightarrow 0$ and $N \rightarrow \infty$ jointly. Since the number of assets N also depends on the asymptotic stage n , we write N_n to emphasize this dependency. Below, we index all estimators by n and all limits are for $n \rightarrow \infty$.

For each $\tau \in \mathcal{T}$, we denote by $i(n, \tau)$ the unique integer i such that $\tau \in ((i-1)\Delta_n, i\Delta_n]$. We then collect these indices using $\mathcal{I}_n = \{i(n, \tau) : \tau \in \mathcal{T}\}$, which is finite almost surely. This set can be consistently recovered by

$$\hat{\mathcal{I}}_n \equiv \{i : |\Delta_i^n Z| > u_n\},$$

where the truncation threshold satisfies $u_n \asymp \Delta_n^\varpi$ for some $\varpi \in (0, 1/2)$.¹⁰ In light of this jump detection result, we can assume that $i(n, \tau)$ is observed for each $\tau \in \mathcal{T}$ without loss of generality.

The spot jump beta estimator associated with each jump time τ is then given by

$$\hat{\beta}_{n,j,\tau} = \frac{\Delta_n^{i(n,\tau)} Y_j}{\Delta_n^{i(n,\tau)} Z}, \quad 1 \leq j \leq N_n.$$

The corresponding feasible test statistic can be constructed naturally as

$$\hat{V}_n^* \equiv \frac{1}{N_n} \sum_{j=1}^{N_n} L(\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}). \tag{2.5}$$

Under the null hypothesis of (2.3) and in the absence of idiosyncratic jumps in Y_j , the leading term of $\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}$ is due to the diffusive component in $Y_j - \beta_j Z$ in the high-frequency increments containing the two jumps. Therefore, this difference is of order $O_p(\sqrt{\Delta_n})$ under the null, and thus \hat{V}_n^* should be $O_p(\Delta_n)$. On the other hand, under the alternative hypothesis, $\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}$ will contain also the difference $\tilde{\chi}_{j,\tau} - \tilde{\chi}_{j,\eta}$ which is $O_p(1)$, and hence so is \hat{V}_n^* . This explains on an intuitive level how \hat{V}_n^* can discriminate between the exact and noisy linear jump factor model. Below, we develop formal statistical tests based on this intuition.

The situation becomes more complicated when individual assets also contain idiosyncratic jumps. Although the idiosyncratic jumps of each asset do not occur exactly at systematic jump times, there is a small probability that these jumps occur in the same high-frequency sampling interval containing the jump in Z . This effect would be asymptotically negligible if the number of assets were fixed, but this will generally not be the case in the current setting in which the number of assets goes to infinity. This issue can be addressed via winsorization. Formally, we consider a sequence $q_n^w \rightarrow 0$ which specifies the proportion of data to be winsorized. We denote by $\bar{B}_{n,\eta,\tau}$ the $1 - q_n^w$ quantile of the variables $(|\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}|)_{1 \leq j \leq N_n}$. The winsorized test statistic is then

$$\hat{V}_n = \frac{1}{N_n} \sum_{j=1}^{N_n} L(|\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}| \wedge \bar{B}_{n,\eta,\tau}). \tag{2.6}$$

Below, we focus on the asymptotic properties of the above winsorized statistic, while noting that the unwinsorized statistic can be used equivalently if the jump arrivals in Y outside of the set \mathcal{T} are driven by a finite counting measure.¹¹ We maintain the following condition for the winsorization quantile q_n^w .

ASSUMPTION 2. $q_n^w \asymp \Delta_n^\kappa$ for some constant $\kappa \in (0, 1)$.

¹⁰It can be shown that $\hat{\mathcal{I}}_n = \mathcal{I}_n$ with probability approaching one; see Proposition 1 in Li, Todorov, and Tauchen (2017).

¹¹That is, when the total number of idiosyncratic jump times in all stocks is finite almost surely.

2.2 Regularity conditions

We proceed with some regularity conditions that will be used throughout. We start with some standard conditions regarding the pathwise regularities of the underlying processes.

ASSUMPTION 3. *The processes $(Y_j)_{1 \leq j \leq N_n}$ and Z are given by (2.1) and (2.2) such that the following conditions hold:*

(i) *The jump processes J_Z and $(J_{Y,j})_{1 \leq j \leq N_n}$ have the form*

$$J_{Z,t} = \int_0^t \int_E \delta_Z(s, u) \mu(ds, du), \tag{2.7}$$

$$J_{Y,j,t} = \int_0^t \int_E \delta_{Y,j}(s, u) \mu(ds, du) + \int_0^t \int_E \tilde{\delta}_{Y,j}(s, u) \tilde{\mu}_j(ds, du), \tag{2.8}$$

where μ is a Poisson random measure on $\mathbb{R}_+ \times E$ with compensator $\nu(ds, du) = ds \otimes \nu'(du)$ for some finite measure ν' on a Polish space E ; $(\tilde{\mu}_j)_{1 \leq j \leq N_n}$ are Poisson random measures that satisfy the same conditions as μ ; and the jump size functions $\delta_Z, (\delta_{Y,j})_{j \geq 1}$ and $(\tilde{\delta}_{Y,j})_{j \geq 1}$, which are mappings $\Omega \times \mathbb{R}_+ \times E \mapsto \mathbb{R}$, are predictable. Moreover, the jump processes $\tilde{J}_{Y,j,s} \equiv \int_0^t \int_E \tilde{\delta}_{Y,j}(s, u) \tilde{\mu}_j(ds, du)$ are uniformly locally bounded and satisfy $\Delta \tilde{J}_{Y,j,s} \Delta J_{Z,s} = 0$ almost surely for $s \in [0, T]$.

(ii) *The diffusive factor process f is a r -dimensional continuous Itô semimartingale of the form:*

$$f_t = \int_0^t b_{f,s} ds + \int_0^t \sigma_{f,s} dW_s,$$

where the processes b_f and σ_f are locally bounded and W is a r -dimensional Brownian motion. Moreover, the spot covariance process $\Sigma_f = \sigma_f \sigma_f^\top$ is nonsingular almost surely.

(iii) *For each j , ϵ_j is a one-dimensional continuous local martingale given by*

$$\epsilon_{j,t} = \int_0^t \tilde{\sigma}_{j,s} d\tilde{W}_{j,s},$$

where $(\tilde{W}_j)_{1 \leq j \leq N_n}$ are one-dimensional Brownian motions that are orthogonal to W and the processes $(\tilde{\sigma}_j)_{1 \leq j \leq N_n}$ are locally uniformly bounded.¹²

(iv) *The processes α_j, λ_j , and λ_Z are locally bounded and the spot jump betas $(\beta_{j,\tau})_{\tau \in \mathcal{T}}$ are bounded, uniformly for $1 \leq j \leq N_n$.*

(v) *For an increasing sequence $(T_m)_{m \geq 1}$ of stopping times that goes to infinity and a sequence $(K_m)_{m \geq 1}$ of constants,*

$$\mathbb{E} \left[\sup_{s,t \in [0, T \wedge T_m]} (|\tilde{\sigma}_{j,s}^2 - \tilde{\sigma}_{j,t}^2|^2 + |\lambda_{j,s} - \lambda_{j,t}|^2 + |\lambda_{Z,s} - \lambda_{Z,t}|^2) \right] \leq K_m |s - t|,$$

uniformly in $1 \leq j \leq N_n$.

¹²That is, for a sequence $(T_m)_{m \geq 1}$ of stopping times that increases to infinity, the processes $\tilde{\sigma}_j, 1 \leq j \leq N_n$, are uniformly bounded on $[0, T \wedge T_m]$.

Assumption 3 imposes a set of regularity conditions that are commonly used in the analysis of high-frequency data. Conditions (i)–(iii) mainly require that the price processes are Itô semimartingales. We allow for leverage effect, price-volatility co-jumps, as well as co-jumps between the factor volatility and idiosyncratic volatility. We further allow for idiosyncratic jump risk in Y and multiple systematic jump factors (in addition to Z) in the jump component $\tilde{J}_{Y,j,t}$. Condition (iv) requires (local) boundedness for the factor loadings and the jump betas, which is reasonable from an empirical point of view. Condition (v) requires that the idiosyncratic volatility and the factor loading processes are $(1/2)$ -Hölder continuous under the L_2 -norm and locally in time. Note that these processes are allowed to have jumps in their sample paths with arbitrary activity. This condition holds for many stochastic processes such as Itô semimartingales and long-memory processes driven by the fractional Brownian motion.

We note that Assumption 3 imposes a finite activity restriction on the jump measures μ and $\tilde{\mu}_j$. Therefore, our analysis here applies for the “big” jumps in asset price. Allowing for jumps of infinite activity is nontrivial and we leave such an extension for future work.

In order to obtain well-defined asymptotic limits, we shall assume that the factor loadings and the idiosyncratic variances are “moderately heterogeneous” in the sense that they are well-behaved under cross-sectional aggregation; see Assumption 4 below. In the analysis that follows, it is convenient to introduce factor loadings for the residual process that are defined as

$$\tilde{\lambda}_{j,\tau\pm} = \lambda_{j,\tau\pm} - \beta_{j,\tau} \lambda_{Z,\tau\pm}, \quad \tau \in \mathcal{T}. \tag{2.9}$$

In particular, $\tilde{\lambda}_{j,\tau-}$ and $\tilde{\lambda}_{j,\tau+}$ are the factor loadings on f for the residual process $Y_j - \beta_{j,\tau} Z$ before and after the jump time τ , respectively. For notational simplicity, below we abuse our notation slightly by writing $\tilde{\lambda}_{j,q}$ in place of $\tilde{\lambda}_{j,\tau-}$ (resp., $\tilde{\lambda}_{j,\tau+}$) with $q = \tau-$ (resp., $q = \tau+$) indicating the pre-jump (resp., post-jump) window; the same convention also applies to other variables.

ASSUMPTION 4. For $p, q \in \{\tau-, \tau+, \eta-, \eta+\}$, the following hold:

- (i) $N_n^{-1} \sum_{j=1}^{N_n} \tilde{\lambda}_{j,p} \tilde{\lambda}_{j,q}^\top \xrightarrow{\mathbb{P}} M_\Lambda(p, q)$ for some \mathcal{F} -measurable $r \times r$ random matrix $M_\Lambda(p, q)$;
- (ii) $N_n^{-1} \sum_{j=1}^{N_n} \tilde{\sigma}_{j,q}^2 \xrightarrow{\mathbb{P}} M_\epsilon(q)$ for some \mathcal{F} -measurable random variable $M_\epsilon(q)$.
- (iii) The factor loadings $(\tilde{\lambda}_{j,q})_{1 \leq j \leq N_n}$ are independent of the idiosyncratic diffusive components $(\epsilon_j)_{1 \leq j \leq N_n}$ conditional on the jump times in \mathcal{T} .

Conditions (i), (ii) in Assumption 4 concern the convergence of cross-sectional sample averages. These conditions can be verified with an appeal to the law of large numbers provided that the loadings $\tilde{\lambda}_j$ and the idiosyncratic variances $\tilde{\sigma}_j^2$ are conditionally weakly dependent in the cross-section. We note that these variables can still be unconditionally strongly dependent, in which case the limiting variables $M_\Lambda(p, q)$ and $M_\epsilon(q)$ are generally nondegenerate random variables. We allow this complication in our econometric inference, which requires characterizing asymptotic distributions in terms of stable

convergence in law. Condition (iii) appears to be mild; indeed, this assumption holds automatically if the factor loadings are nonrandom, which is typically assumed in the factor analysis literature.

Assumption 4 allows us to characterize precisely the aggregate behavior of the residual diffusive disturbances. Indeed, the cross-sectional average of the spot variances of the residual returns (i.e., $dY_{j,t} - \beta_{j,t} dZ_t$) before/after each jump time τ can be written as

$$\frac{1}{N_n} \sum_{j=1}^{N_n} (\tilde{\lambda}_{j,\tau\pm}^\top \Sigma_{f,\tau\pm} \tilde{\lambda}_{j,\tau\pm} + \tilde{\sigma}_{j,\tau\pm}^2) = \text{Trace} \left[\Sigma_{f,\tau\pm} \left(\frac{1}{N_n} \sum_{j=1}^{N_n} \tilde{\lambda}_{j,\tau\pm} \tilde{\lambda}_{j,\tau\pm}^\top \right) \right] + \frac{1}{N_n} \sum_{j=1}^{N_n} \tilde{\sigma}_{j,\tau\pm}^2,$$

which converges in probability (under Assumption 4) to

$$M_{\text{total}}(\tau\pm) \equiv \text{Trace}[\Sigma_{f,\tau\pm} M_A(\tau\pm, \tau\pm)] + M_\epsilon(\tau\pm). \quad (2.10)$$

The two components on the right-hand side of (2.10) are due to the diffusive factor component in the residual process and the diffusive idiosyncratic components in Y , respectively.

Finally, we maintain a very mild condition for the relative growth rates of Δ_n and N_n .

ASSUMPTION 5. N_n grows to infinity at most polynomially in Δ_n^{-1} , that is, $N_n = O(\Delta_n^{-k})$ for some (arbitrary but fixed) constant $k > 0$.

Assumption 5 is very weak because it does not require any specific rate at which N_n grows to infinity relative to Δ_n . This assumption is used to show that the cross-section of diffusive increments around the jump time fall uniformly in any fixed (small) neighborhood around zero, so that the loss function $L(\cdot)$ behaves like the quadratic function when acting on these returns. This assumption is not needed if the loss function is (globally) quadratic.

3. TESTING THE EXACT LINEAR JUMP FACTOR MODEL

This section contains the core of our econometric analysis. We derive a feasible test on the basis of the statistic \hat{V}_n to test the null hypothesis of the exact linear jump factor model at two jump times of the systematic factor Z . We further derive a measure for the magnitude of the firm-specific shocks in the jump betas when the null might be violated.

3.1 Asymptotic properties of the test statistic

We start with characterizing the asymptotic behavior of our test statistic. We first show the convergence in probability of the test statistic under the general random jump beta model (2.4), for which we need the following condition.

ASSUMPTION 6. *The innovations in jump betas* $\{\chi_{j,\eta,\tau} \equiv \tilde{\chi}_{j,\tau} - \tilde{\chi}_{j,\eta}; 1 \leq j \leq N_n\}$ *are* $\mathcal{F}_{\eta-}$ -*conditionally independent on the cross-section.*

Assumption 6 mainly requires that the innovations in the jump beta (i.e., $\chi_{j,\eta,\tau}$) are cross-sectionally independent conditional on the information set before the jumps occur. We can further accommodate in the analysis situations in which $\chi_{j,\eta,\tau}$ are cross-sectionally dependent. Such a situation can happen, if for example, there are extra factors at the jump times of Z (in addition to the jump size of Z). In this case, under some weak regularity conditions for the factor loadings on the omitted factor(s), the asymptotic behavior of our test statistic $L(\chi_{j,\eta,\tau})$ will be similar to that under Assumption 6 in terms of asymptotic order of magnitude. For ease of exposition, we do not consider this situation here.

Proposition 1 below describes the first-order asymptotic behavior of the test statistic.

PROPOSITION 1. *Under Assumptions 1, 2, 3, and 6, $\hat{V}_n = N_n^{-1} \sum_{j=1}^{N_n} \mathbb{E}[L(\chi_{j,\eta,\tau})|\mathcal{F}_{\eta-}] + o_p(1)$. In particular, under the constant beta restriction (2.3), $\hat{V}_n = o_p(1)$.*

Proposition 1 shows that \hat{V}_n converges to the cross-sectional average of the conditional expected loss that is due to the temporal variation in jump betas. This convergence result readily accommodates cross-sectional heterogeneity in the data. Of course, if the loss $L(\chi_{j,\eta,\tau})$ has identical conditional mean over the cross-section, the limit of \hat{V}_n is simply $\mathbb{E}[L(\chi_{j,\eta,\tau})|\mathcal{F}_{\eta-}]$. This proposition suggests that the test statistic \hat{V}_n is able to detect on-average deviations from the exact jump factor model.

In view of Proposition 1, we can state the hypotheses to be tested precisely as follows. The testing problem is to decide in which of the two events the observed sample path falls:

$$\Omega_0 \equiv \{\beta_{j,\eta} = \beta_{j,\tau} = \beta_j \text{ for all } j\},$$

$$\Omega_a \equiv \left\{ \liminf_{n \rightarrow \infty} \frac{1}{N_n} \sum_{j=1}^{N_n} \mathbb{E}[L(\chi_{j,\eta,\tau})|\mathcal{F}_{\eta-}] > 0 \right\},$$

where Ω_0 and Ω_a play the role of the null and alternative hypotheses, respectively. Stating the hypotheses in terms of random events is typical in the infill asymptotic setting because the “population” quantities are (random) sample paths; see [Jacod and Protter \(2012\)](#) and [Aït-Sahalia and Jacod \(2014\)](#) for many similar examples. In addition, formulating the alternative hypothesis in terms of the limit inferior is also common in testing problems involving heterogeneous data (see, e.g., [Giacomini and White \(2006\)](#)).

In order to construct a proper test, we need to characterize the asymptotic distribution of the test statistic under the null hypothesis. Theorem 1, below, presents the result. We need some additional notation to represent the limit distribution. To this end, we consider random variables $(\kappa_\tau, \zeta_{\tau-}, \zeta_{\tau+})_{\tau \in \mathcal{T}}$ that, conditional on \mathcal{F} , are mutually independent with the following marginal distributions: κ_τ is uniformly distributed on $[0, 1]$ and $\zeta_{\tau\pm}$ are r -dimensional standard normal. With each $s \in \mathcal{T}$, we associate the following weights:

$$w_{s-} \equiv \frac{\sqrt{\kappa_s}}{\Delta Z_s}, \quad w_{s+} \equiv \frac{\sqrt{1 - \kappa_s}}{\Delta Z_s}. \tag{3.1}$$

Following our notational convention in Assumption 4, we write the weights in (3.1) as w_q indexed by $q \in \{s-, s+\}$; similarly, we denote for $p, q \in \{\eta-, \eta+, \tau-, \tau+\}$,

$$M_C(p, q) = \Sigma_{f,p}^{1/2} M_\Lambda(p, q) \Sigma_{f,q}^{1/2}. \tag{3.2}$$

Finally, we set

$$\begin{cases} \mathcal{A}(s) \equiv \sum_{p,q \in \{s-, s+\}} w_p w_q \zeta_p^\top M_C(p, q) \zeta_q + \sum_{q \in \{s-, s+\}} w_q^2 M_\epsilon(q), & s \in \{\eta, \tau\}, \\ \mathcal{B}(\eta, \tau) \equiv \sum_{p \in \{\tau-, \tau+\}} \sum_{q \in \{\eta-, \eta+\}} w_p w_q \zeta_p^\top M_C(p, q) \zeta_q. \end{cases} \tag{3.3}$$

THEOREM 1. *Suppose that Assumptions 1–5 hold. In restriction to Ω_0 , the sequence $\Delta_n^{-1} \hat{V}_n$ of variables converges \mathcal{F} -stably in law toward $\mathcal{L}(\eta, \tau)$ given by*

$$\mathcal{L}(\eta, \tau) \equiv \mathcal{A}(\eta) + \mathcal{A}(\tau) - 2\mathcal{B}(\eta, \tau). \tag{3.4}$$

COMMENTS. (i) The limiting variable $\mathcal{L}(\eta, \tau)$ given in (3.4) captures two types of sampling variability. The first type is generated by the diffusive factor process and it takes the form of quadratic functions of the Gaussian variables $\zeta_{\eta\pm}$ and $\zeta_{\tau\pm}$. The \mathcal{F} -conditional distributions of these terms are nondegenerate. Intuitively, the variable $\zeta_{\tau-}$ (resp., $\zeta_{\tau+}$) represents the asymptotic distribution of the normalized diffusive factor returns before (resp., after) the jump time τ , that is,

$$\frac{\Sigma_{f,\tau-}^{-1/2} (f_\tau - f_{(i(n,\tau)-1)\Delta_n})}{\sqrt{\tau - (i(n, \tau) - 1)\Delta_n}} \quad \left(\text{resp. } \frac{\Sigma_{f,\tau}^{-1/2} (f_{i(n,\tau)\Delta_n} - f_\tau)}{\sqrt{i(n, \tau)\Delta_n - \tau}} \right).$$

Importantly, the factor diffusive returns represent systematic risk, and hence, they have a nontrivial conditional distribution even after the cross-sectional aggregation. This is in sharp contrast to conventional econometric settings, because the (mixed) normality of the limiting variable here is not obtained from the aggregation of weakly dependent variables via a central limit theorem, but is implied from the local Gaussianity of the continuous Itô process f .

(ii) The second type of sampling variability reflected in $\mathcal{L}(\eta, \tau)$ (more specifically the second term $\sum_{q \in \{s-, s+\}} w_q^2 M_\epsilon(q)$ in the definition of $\mathcal{A}(s)$) is attributed to the idiosyncratic diffusive components of Y . Although the idiosyncratic Brownian shocks $d\tilde{W}_j$ are “averaged out” in the cross-sectional averaging, this term can still have a nondegenerate \mathcal{F} -conditional distribution because of the variables $(\kappa_\eta, \kappa_\tau)$ that appear in the weights w_q for $q = \{\eta-, \eta+, \tau-, \tau+\}$. Intuitively, κ_τ represents the relative location of the exact jump time τ within its observation interval. This source of sampling variability is, again, systematic (as it is due to Z), so it “survives” the cross-sectional aggregation. In the case when average idiosyncratic spot variance does not co-jump with the price (i.e., $M_\epsilon(s-) = M_\epsilon(s+)$, for $s = \eta, \tau$), the \mathcal{F} -conditional distribution of this term becomes degenerate.

(iii) Jumps in the vector Y that happen outside the set \mathcal{T} of jump times of Z play no role in the limit result for our test statistic in Theorem 1. This is because jump arrivals are “sparsely” scattered in time. Hence, with probability approaching one, each high-frequency interval containing one of the jump times in the set \mathcal{T} will not contain jump times outside of \mathcal{T} for “most” components in Y ; the number of exceptions is small and their effect is “regularized” by the winsorization.

We note that the test statistic \hat{V}_n is nonnegative, and hence, the limiting variable $\mathcal{L}(\eta, \tau)$ concentrates on the positive line. We can gauge the relative contribution of the components in (3.4) by computing their \mathcal{F} -conditional means. It is easy to see that $\mathbb{E}[\mathcal{B}(\eta, \tau)|\mathcal{F}] = 0$ and, for $s \in \{\eta, \tau\}$, we have

$$\mathbb{E}[\mathcal{A}(s)|\mathcal{F}] = \frac{M_{\text{total}}(s-) + M_{\text{total}}(s+)}{2\Delta Z_s^2},$$

where $M_{\text{total}}(s\pm)$ is the (large-sample) cross-sectional average of individual stocks’ residual spot variances after/before the jump time s ; recall (2.10). Consequently, the limit distribution $\mathcal{L}(\eta, \tau)$ has a strictly positive conditional mean depending on the relative magnitude of the residual spot variances with respect to the jump size. In the current setting, it is therefore natural to define the signal-to-noise ratio for the post- and pre-jump window as

$$SNR_{s\pm} \equiv \frac{\Delta Z_s^2}{M_{\text{total}}(s\pm)}. \tag{3.5}$$

The \mathcal{F} -conditional mean of the limiting variable can then be succinctly written as

$$\mathbb{E}[\mathcal{L}(\eta, \tau)|\mathcal{F}] = \sum_{s \in \{\eta, \tau\}} \left(\frac{1}{2 \cdot SNR_{s-}} + \frac{1}{2 \cdot SNR_{s+}} \right). \tag{3.6}$$

From this expression, we see that under the null hypothesis, the test statistic tends to be centered “more positively” when the signal-to-noise ratios around the jump times are lower. These are exactly the times when it will be more difficult to separate the null hypothesis from the alternative.

In order to implement the test, we need to consistently estimate the critical values defined as the \mathcal{F} -conditional quantiles of the limiting variable $\mathcal{L}(\eta, \tau)$. This is a nontrivial task because the \mathcal{F} -conditional distribution of $\mathcal{L}(\eta, \tau)$ is highly nonstandard. Most importantly, it depends on the unknown random quantities $M_C(\cdot, \cdot)$ and $M_\epsilon(\cdot)$, which in turn involve the spot covariance of the *latent* diffusive factor process f , the factor loadings of a large cross-section of assets and the average idiosyncratic spot variances, before and after each jump time. The next subsection is devoted to addressing this problem.

3.2 Critical values

For a matrix A , we denote $\|A\| = \text{Trace}[A^\top A]^{1/2}$. To proceed, we suppose that some preliminary estimator $\tilde{\beta}_{n,j,\tau}$ of the jump beta is available and it satisfies the following assumption.

ASSUMPTION 7. *The following conditions hold for some variables $(\beta_{j,\tau}^*)_{1 \leq j \leq N_n, \tau \in \mathcal{T}}$:*

- (i) $N_n^{-1} \sum_{j=1}^{N_n} |\tilde{\beta}_{n,j,\tau} - \beta_{j,\tau}^*|^2 = o_p(1)$;
- (ii) for all $\tau \in \mathcal{T}$, $\tilde{\beta}_{n,j,\tau}$ and $\beta_{j,\tau}^*$ are \mathcal{F} -measurable and bounded with probability approaching one;
- (iii) $\beta_{j,\tau}^* = \beta_j$ in restriction to Ω_0 .

One possible choice for the preliminary beta estimators is to set $\tilde{\beta}_{n,j,\tau}$ to be the spot beta estimator $\hat{\beta}_{n,j,\tau}$. In this case, $\beta_{j,\tau}^* = \beta_{j,\tau}$ over the entire sample space. Condition (i) requires that these preliminary estimates are consistent in a cross-sectional average sense, which is easy to verify; see Lemma SA3(b) for the formal result. More generally, the variables $\beta_{j,\tau}^*$ in Assumption 7 are interpreted as pseudo-true parameters in the sense that they coincide with the true parameter under the null hypothesis but are allowed to differ from the latter under the alternative. For example, we can take $\tilde{\beta}_{n,j,\eta} = \tilde{\beta}_{n,j,\tau} = (\hat{\beta}_{n,j,\eta} + \hat{\beta}_{n,j,\tau})/2$ which corresponds to $\beta_{j,\eta}^* = \beta_{j,\tau}^* = (\beta_{j,\eta} + \beta_{j,\tau})/2$. This type of averaging typically produces less noisy preliminary estimates under the null, which tends to have better size control in finite samples.¹³

We now describe how to calculate the critical values. The first step is to extract information concerning the spot covariance of the factors (i.e., Σ_f) and the factor loadings (i.e., λ_j and λ_Z) from the diffusive returns in local windows before and after the jump times. We consider an integer sequence k_n of local window sizes such that

$$k_n \rightarrow \infty \quad \text{and} \quad k_n \Delta_n \rightarrow 0.$$

For each detected jump time τ , we denote $i(n, \tau-) \equiv i(n, \tau) - k_n - 1$ and $i(n, \tau+) \equiv i(n, \tau)$. We use $X_n^*(\tau \pm)$ to denote the $N_n \times k_n$ matrix whose (j, l) element is given by

$$X_n^*(\tau \pm)_{j,l} \equiv \frac{\Delta_{i(n,\tau \pm)+l}^n Y_j - \beta_{j,\tau}^* \Delta_{i(n,\tau \pm)+l}^n Z}{\sqrt{\Delta_n}}. \tag{3.7}$$

These matrices collect the (scaled) residual diffusive returns of the assets in the local windows.

We note that $X_n^*(\tau \pm)$ have an approximate factor structure

$$X_n^*(\tau \pm)_{j,l} \approx (\lambda_{j,\tau \pm} - \beta_{j,\tau}^* \lambda_{Z,\tau \pm})^\top \frac{\Delta_{i(n,\tau \pm)+l}^n f}{\sqrt{\Delta_n}} + \frac{\Delta_{i(n,\tau \pm)+l}^n \epsilon_j}{\sqrt{\Delta_n}}. \tag{3.8}$$

This is only an approximation because the (time-varying) factor loadings (λ_j, λ_Z) before and after each jump time τ are only approximately $(\lambda_{j,\tau-}, \lambda_{Z,\tau-})$ and $(\lambda_{j,\tau}, \lambda_{Z,\tau})$, respectively.¹⁴ It is convenient to rewrite (3.8) in matrix form as

$$X_n^*(\tau \pm) \approx \Lambda_n^*(\tau \pm) F_n(\tau \pm)^\top + \mathcal{E}_n(\tau \pm),$$

¹³We can similarly use an average of jump betas over a larger number of systematic jump events, provided we are willing to assume the null hypothesis of the exact jump factor model holds true for this larger number of events.

¹⁴The approximation error is also contributed by the drift and the idiosyncratic jump terms in Y_j . All of these approximation errors are accounted for in our analysis.

where we denote

$$\begin{cases} F_n(\tau\pm) \equiv \Delta_n^{-1/2}(\Delta_{i(n,\tau\pm)+1}^n f, \dots, \Delta_{i(n,\tau\pm)+k_n}^n f)^\top, \\ \Lambda_n^*(\tau\pm) \equiv (\lambda_{1,\tau\pm} - \beta_{1,\tau}^* \lambda_{Z,\tau\pm}, \dots, \lambda_{N_n,\tau\pm} - \beta_{N_n,\tau}^* \lambda_{Z,\tau\pm})^\top, \\ \mathcal{E}_n(\tau\pm) \equiv \Delta_n^{-1/2}[\Delta_{i(n,\tau\pm)+l}^n \epsilon_j]_{1 \leq j \leq N_n, 1 \leq l \leq k_n}. \end{cases} \quad (3.9)$$

Assumption 8 extends Assumption 4 in order to accommodate the more general situation with pseudo-true parameters.

ASSUMPTION 8. (i) For $p, q \in \{\eta-, \eta+, \tau-, \tau+\}$, $N_n^{-1} \Lambda_n^*(p)^\top \Lambda_n^*(q) \xrightarrow{\mathbb{P}} M_\Lambda^*(p, q)$ for some \mathcal{F} -measurable $r \times r$ random matrix $M_\Lambda^*(p, q)$.

(ii) For each $q \in \{\eta-, \eta+, \tau-, \tau+\}$, the eigenvalues of $M_\Lambda^*(q, q)$ are distinct almost surely.

(iii) The factor loadings in $\Lambda_n^*(\tau\pm)$ are independent of the idiosyncratic diffusive components $(\epsilon_j)_{1 \leq j \leq N_n}$.

Since $\beta_{j,\tau}^* = \beta_j$ for each $\tau \in \mathcal{T}$ under the null hypothesis, the matrix $M_\Lambda^*(p, q)$ coincides with $M_\Lambda(p, q)$ as well (recall Assumption 4). With this notation, we further complement the definition in (3.2) by setting

$$M_C^*(p, q) \equiv \Sigma_{f,p}^{1/2} M_\Lambda^*(p, q) \Sigma_{f,q}^{1/2}. \quad (3.10)$$

Then the key to conducting feasible inference is to consistently estimate $M_C^*(p, q)$. To do this, we assume that the dimension r of the diffusive factors is bounded by a known constant \bar{r} . This assumption is much weaker than knowing the exact value of r . Since the number of the diffusive factors is not the main object of interest in our analysis of jumps, we aim to be agnostic on its exact value. Instead, we only assume that a bound \bar{r} is known, for which the vast empirical literature provides reliable guidance. We shall show theoretically that knowing the upper bound is enough for constructing valid critical values for our test statistic.¹⁵

We now introduce our estimator for $M_C^*(p, q)$. We denote the sample analogue of $X_n^*(q)$ by $\hat{X}_n(q)$, which is constructed as

$$\hat{X}_n(\tau\pm)_{j,l} \equiv \left(\frac{\Delta_{i(n,\tau\pm)+l}^n Y_j \wedge u_n \vee (-u_n) - \tilde{\beta}_{n,j,\tau} \Delta_{i(n,\tau\pm)+l}^n Z}{\sqrt{\Delta_n}} \right), \quad (3.11)$$

where we recall that u_n is the sequence of threshold used for winsorizing jumps. We set $\hat{F}_n(q)$ as the $k_n \times \bar{r}$ matrix that consists of the eigenvectors associated with the \bar{r} largest

¹⁵The number of diffusive factors may be consistently estimated by adapting the general strategy of Bai and Ng (2002). Like our setting, such methods also assume the availability of an upper bound for the number of factors (see the recent work of Gagliardini, Ossola, and Scaillet (2016a) which proposes a diagnostic criterion that does not need such a bound), subject to which a criterion function is minimized for estimating the number of factors. The resulting estimate may differ when using different methods (often associated with different penalty functions for model complexity). We avoid this potential ambiguity by only relying on the upper bound.

eigenvalues of $\hat{X}_n(q)^\top \hat{X}_n(q)$ under the normalization $\hat{F}_n(q)^\top \hat{F}_n(q)/k_n = I_{\bar{r}}$, where $I_{\bar{r}}$ denotes the \bar{r} -dimensional identity matrix. We then set

$$\hat{\Lambda}_n(q) \equiv \hat{X}_n(q)\hat{F}_n(q)/k_n. \tag{3.12}$$

Due to the normalization of the eigenvector matrix $\hat{F}_n(q)$, it no longer “carries” information concerning the spot covariance $\Sigma_{f,q}$ of the factor process f . Instead, this information is now embedded in the estimator $\hat{\Lambda}_n(q)$, which also contains information about the loading matrix $\Lambda_n^*(q)$. With this intuition in mind, we shall show that $M_C^*(p, q)$ can be estimated (in a proper sense) using

$$\widehat{M}_{C,n}(p, q) \equiv \frac{1}{N_n} \hat{\Lambda}_n(p)^\top \hat{\Lambda}_n(q), \quad p, q \in \{\eta-, \eta+, \tau-, \tau+\}. \tag{3.13}$$

It is interesting to note that the above estimator is the cross-variation of the extracted diffusive factor loadings across (possibly) different local windows, on which our feasible inference is based. Subsequently, the average idiosyncratic variance $M_\epsilon(q)$ can be estimated by

$$\widehat{M}_{\epsilon,n}(q) \equiv \frac{1}{N_n k_n} \|\hat{X}_n(q)\|^2 - \text{Trace}[\widehat{M}_{C,n}(q, q)]. \tag{3.14}$$

In Theorem 2 below, we describe the asymptotic properties of the estimators $\hat{\Lambda}_n$, $\widehat{M}_{C,n}$, and $\widehat{M}_{\epsilon,n}$. We present explicitly the intermediate result for $\hat{\Lambda}_n$ in order to streamline the intuition underlying our constructions. Below, we partition $\hat{\Lambda}_n(q) = [\hat{\Lambda}_n^*(q); \hat{\Lambda}_n^0(q)]$ where the two blocks contain r and $\bar{r} - r$ columns, respectively. Similarly, we partition $\hat{F}_n(q) = [\hat{F}_n^*(q); \hat{F}_n^0(q)]$. We denote the sign function by $\text{sign}(x) = 1_{\{x \geq 0\}} - 1_{\{x < 0\}}$ and apply it component-wise on matrices.

THEOREM 2. *Suppose that Assumptions 3, 4, 7, and 8 hold. For $p, q \in \{\eta-, \eta+, \tau-, \tau+\}$, we have the following:*

(a) $N_n^{-1} \|\hat{\Lambda}_n^*(q) - \Lambda_n^*(q) \Sigma_{f,q}^{1/2} H_q S_n^*(q)\|^2 = o_p(1)$ where

$$S_n^*(q) = \text{diag}(\text{sign}(\hat{F}_n^*(q)^\top F_n(q)(F_n(q)^\top F_n(q)/k_n)^{-1/2} H_q)),$$

and H_q is the ordered orthogonal eigenvector matrix of $M_C^*(q, q)$,¹⁶

(b) $N_n^{-1} \|\hat{\Lambda}_n^0(q)\|^2 = o_p(1)$;

(c) the $\bar{r} \times \bar{r}$ matrix $\widehat{M}_{C,n}(p, q)$ satisfies

$$\widehat{M}_{C,n}(p, q) = \begin{pmatrix} S_n^*(p) H_p^\top M_C^*(p, q) H_q S_n^*(q) & 0 \\ 0 & 0 \end{pmatrix} + o_p(1),$$

where we note that the upper-left block $S_n^*(p) H_p^\top M_C^*(p, q) H_q S_n^*(q)$ is $r \times r$ dimensional;

(d) $\widehat{M}_{\epsilon,n}(q) \xrightarrow{\mathbb{P}} M_\epsilon(q)$.

¹⁶That is, $M_C^*(q, q) = H_q D H_q^\top$, where D is the diagonal matrix that collects the eigenvalues of $M_C^*(q, q)$ in descending order.

COMMENTS. (i) Part (a) shows that, in a cross-sectional average sense, $\hat{\Lambda}_n^*(q)$ approximates $\Lambda_n(q)\Sigma_{f,q}^{1/2}$, up to the transformation $H_q S_n^*(q)$. Despite the presence of the latter “unobserved nuisance,” this result is sufficient for our purpose of estimating the critical values. This is because the law of an \mathcal{F} -conditional standard normal vector is invariant to the transformation associated with the orthogonal matrix $H_q S_n^*(q)$.¹⁷

(ii) Part (b) shows that the “null” columns collected by $\hat{\Lambda}_n^0(q)$ are “asymptotically small.” Hence, estimating the factor structure conservatively by (potentially over-) extracting factors leads to no effect asymptotically.

(iii) Parts (c), (d) follow from parts (a), (b) and establish a type of consistent estimation result for $\hat{M}_{C,n}$ and $\hat{M}_{\epsilon,n}$. Again, we note that the result of part (c) is sufficient for our feasible inference described below in spite of the appearance of the transformation $H_q S_n^*(q)$. We also note that when $p \neq q$, $\hat{M}_{C,n}(p, q)$ estimates the cross-variation of factor loadings between two distinct blocks of data. For example, $\hat{M}_{C,n}(\eta-, \tau+)$ concerns data blocks before the η -jump and after the τ -jump. This type of estimation is needed because we allow the factor loadings to change across jump times and actually we also allow them to co-jump with Z . This is why we need to characterize precisely the behavior of $\hat{\Lambda}_n^*(q)$ in the form of part (a). Such analysis is not common in conventional factor analysis (cf. Stock and Watson (2002), Bai (2003)).¹⁸

We remind the reader that our analysis is not about the diffusive factors per se. Rather, Theorem 2 is an intermediate (but important) step for making inference regarding the cross-sectional behavior of jump betas, which is the focus of the current paper. We are now ready to introduce the critical value $cv_{n,\alpha}$ for our test statistic \hat{V}_n at some nominal level $\alpha \in (0, 1)$. Algorithm 1 below provides the details and is quite intuitive as it follows closely from Theorem 1. The asymptotic validity of this algorithm is provided in Theorem 3 that follows.

ALGORITHM 1. Step 1. Simulate $(\tilde{\kappa}_s, \tilde{\zeta}_{s-}, \tilde{\zeta}_{s+})_{s \in \{\eta, \tau\}}$ independently such that $\tilde{\kappa}_s \sim \text{Uniform}[0, 1]$ and $\tilde{\zeta}_{s\pm} \sim \mathcal{N}(0, I_{\tilde{r}})$.

Step 2. Set $\tilde{w}_{n,s-} = \sqrt{\tilde{\kappa}_s} / \Delta_{i(n,s)}^n Z$ and $\tilde{w}_{n,s+} = \sqrt{1 - \tilde{\kappa}_s} / \Delta_{i(n,s)}^n Z$ for $s \in \{\eta, \tau\}$. Then compute

$$\begin{cases} \tilde{A}_n(s) = \sum_{p,q \in \{s-, s+\}} \tilde{w}_{n,p} \tilde{w}_{n,q} \tilde{\zeta}_p^\top \hat{M}_{C,n}(p, q) \tilde{\zeta}_q + \sum_{q \in \{s-, s+\}} \tilde{w}_{n,q}^2 \hat{M}_{\epsilon,n}(q), & s \in \{\eta, \tau\}, \\ \tilde{B}_n(\eta, \tau) = \sum_{p \in \{\tau-, \tau+\}} \sum_{q \in \{\eta-, \eta+\}} \tilde{w}_{n,p} \tilde{w}_{n,q} \tilde{\zeta}_p^\top \hat{M}_{C,n}(p, q) \tilde{\zeta}_q, \\ \tilde{L}_n(\eta, \tau) = \tilde{A}_n(\eta) + \tilde{A}_n(\tau) - 2\tilde{B}_n(\eta, \tau). \end{cases}$$

Step 3. Repeat step 1 and step 2 for a large number of times and report the critical value $cv_{n,\alpha}$ as the $1 - \alpha$ quantile of $\tilde{L}_n(\eta, \tau)$ in the Monte Carlo sample.

¹⁷In Bai (2003) and Pelger (2015), the transformation matrix is only known to be invertible. Such a result would be sufficient for the consistency of the column space spanned the estimator, but is not enough for the inference problem considered here.

¹⁸See also Pelger (2015, 2019) and Ait-Sahalia and Xiu (2018) for recent development in the high-frequency setting.

THEOREM 3. *The following statements hold under Assumptions 1–8:*

- (a) *The sequence $cv_{n,\alpha}$ of critical values is $O_p(1)$. Moreover, in restriction to Ω_0 , $cv_{n,\alpha} \xrightarrow{\mathbb{P}} cv_\alpha$ where cv_α is the \mathcal{F} -conditional $1 - \alpha$ quantile of $\mathcal{L}(\eta, \tau)$.*
- (b) *The test associated with the critical region $\{\Delta_n^{-1}\hat{V}_n > cv_{n,\alpha}\}$ has asymptotic level α under the null hypothesis and has asymptotic power one under the alternative hypothesis, that is,*

$$\mathbb{P}(\Delta_n^{-1}\hat{V}_n > cv_{n,\alpha} | \Omega_0) \rightarrow \alpha, \quad \mathbb{P}(\Delta_n^{-1}\hat{V}_n > cv_{n,\alpha} | \Omega_a) \rightarrow 1.$$

COMMENT. Part (a) of Theorem 3 shows that the critical value $cv_{n,\alpha}$ consistently estimates the corresponding \mathcal{F} -conditional quantile of the limit distribution $\mathcal{L}(\eta, \tau)$ of the test statistic, and it remains stochastically bounded (i.e., tight) under the alternative. As a direct consequence, part (b) shows that the proposed test has correct size control under the null hypothesis, and is consistent under the alternative hypothesis.

3.3 Measuring temporal variation in jump betas

We have so far developed tests for the exact jump factor model in the cross-section. The test detects positive cross-sectional average losses resulting from innovations in the jump beta, that is, $\chi_{j,\eta,\tau} = \tilde{\chi}_{j,\tau} - \tilde{\chi}_{j,\eta}$. From an estimation point of view, this loss also serves naturally as a measure for the variation of jump beta for a “representative” asset. Since this measure is of economic interest in its own, we further develop in this subsection an estimator for it. We focus on the quadratic loss function, in which case the test statistic is given by

$$\hat{V}_n = \frac{1}{N_n} \sum_{j=1}^{N_n} (|\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}| \wedge \bar{B}_{n,\eta,\tau})^2.$$

As shown in equation (3.6), in the boundary case with $\chi_{j,\eta,\tau} = 0$ (i.e., when jump betas are time-invariant), the “raw” estimator \hat{V}_n is not asymptotically centered, but has a positive bias that is high in situations with low signal-to-noise ratio. While this bias term is accounted for in the critical values of the test described in Theorem 3, it is of course more conventional from an estimation point of view to use a centered estimator in practice. To this end, we consider a bias-corrected estimator given by

$$\hat{v}_n \equiv \frac{1}{N_n} \sum_{j=1}^{N_n} (|\hat{\beta}_{n,j,\tau} - \hat{\beta}_{n,j,\eta}| \wedge \bar{B}_{n,\eta,\tau})^2 - \Delta_n \hat{b}_n(\eta, \tau), \tag{3.15}$$

where the correction term $\hat{b}_n(\eta, \tau)$ is defined as

$$\hat{b}_n(\eta, \tau) \equiv \frac{1}{N_n k_n} \sum_{s \in \{\eta, \tau\}} \left(\frac{\|\hat{X}_n(s-)\|^2 + \|\hat{X}_n(s+)\|^2}{2(\Delta_{i(n,s)}^n Z)^2} \right). \tag{3.16}$$

We note that $\hat{b}_n(\eta, \tau)$ is constructed as the sample analogue of the asymptotic bias described in (3.6).¹⁹ In fact, as an intermediate step in the proof of Theorem 2, we have shown that

$$\frac{1}{N_n k_n} \|\hat{X}_n(s\pm)\|^2 \xrightarrow{\mathbb{P}} M_{\text{total}}^*(s\pm) \equiv \text{Trace}[\Sigma_{f,s\pm} M_A^*(s\pm, s\pm)] + M_\epsilon(s\pm).$$

We note that $M_{\text{total}}^* = M_{\text{total}}$ holds under the null hypothesis, and hence,

$$\hat{b}_n(\eta, \tau) \xrightarrow{\mathbb{P}} \mathbb{E}[\mathcal{L}(\eta, \tau) | \mathcal{F}]. \tag{3.17}$$

The estimator \hat{v}_n is consistent for the cross-sectional average (conditional) second moment of $\chi_{j,\eta,\tau}$.

4. EXTENSIONS

In this section, we briefly discuss two extensions of the main theoretical results developed in Section 3 above. Section 4.1 proposes a test for the joint hypothesis of constant jump beta involving multiple jump times. Section 4.2 describes how to use a mixed-scale approach to address the issue of gradual response to systematic jump events of individual stock prices, which is relevant for empirical studies about jumps of less liquid assets.

4.1 Joint tests at multiple jump times

The test described in Theorem 3 involves only a pair of jump times. The corresponding pairwise test can be easily extended to a joint analysis involving multiple jump times.²⁰ Among many possibilities, we describe here a specific implementation in which we simply compute the consecutive pairwise test statistics and then take their average. The theoretical justification for this procedure can be adapted straightforwardly from the main theory in Section 3; we thus omit the formal statements and instead provide only the algorithm for conducting the test.

To fix ideas, we consider $P + 1$ jump times $\{\tau_1, \dots, \tau_{P+1}\}$ of Z . The pairwise test statistic for the p th pair of consecutive jump times $\{\tau_p, \tau_{p+1}\}$, $1 \leq p \leq P$, is denoted by

$$\hat{V}_n(\tau_p, \tau_{p+1}) \equiv \frac{1}{N_n} \sum_{j=1}^{N_n} L(|\hat{\beta}_{n,j,\tau_{p+1}} - \hat{\beta}_{n,j,\tau_p}| \wedge \bar{B}_{n,\tau_p,\tau_{p+1}}).$$

¹⁹The correction term $\Delta_n \hat{b}_n(\eta, \tau)$ in (3.15) is of order $O_p(\Delta_n)$. This rate is consistent with that in Theorem 1, because this bias term is asymptotically relevant only when the jump betas are constant. That said, this correction should also be effective in finite-samples when innovations in jump betas are “local to zero,” though a complete analysis under drifting sequences of data generating processes is beyond the scope of the current paper.

²⁰We consider a Kolmogorov-type test. Another approach is to use multiple-testing techniques for controlling the false discovery rate; see, for example, Barras, Scaillet, and Wermers (2010) and Bajgrowicz, Scaillet, and Treccani (2016) for this type of applications on jumps.

The null asymptotic distribution of $\Delta_n^{-1}\hat{V}_n(\tau_p, \tau_{p+1})$ is $\mathcal{L}(\tau_p, \tau_{p+1})$ as shown in Theorem 1. By a simple extension of that theorem, we can show that these convergences hold jointly, that is,

$$(\Delta_n^{-1}\hat{V}_n(\tau_p, \tau_{p+1}))_{1 \leq p \leq P} \xrightarrow{\mathcal{L}\text{-}s} (\mathcal{L}(\tau_p, \tau_{p+1}))_{1 \leq p \leq P}, \tag{4.1}$$

where $\xrightarrow{\mathcal{L}\text{-}s}$ denotes stable convergence in law. Importantly, the limiting variables are represented using the same variables $(\kappa_{\tau_p}, \zeta_{\tau_p-}, \zeta_{\tau_{p+1}})_{1 \leq p \leq P+1}$ that are described in Section 3.1; in particular, $\mathcal{L}(\tau_{p-1}, \tau_p)$ and $\mathcal{L}(\tau_p, \tau_{p+1})$ are \mathcal{F} -conditionally dependent because they both involve $(\kappa_{\tau_p}, \zeta_{\tau_p-}, \zeta_{\tau_{p+1}})$.

One possibility to implement a joint test is to “combine” the pairwise test statistics by taking their average, $P^{-1} \sum_{p=1}^P \hat{V}_n(\tau_p, \tau_{p+1})$. One drawback of this (joint) test statistic, however, is that it is not scale-invariant. In particular, it will tend to put a lot of weight to the pairwise tests which happen in very volatile periods and much less weight to those during calm periods within the testing time window.

For this reason, we consider instead a scaled statistic. In view of (3.6), we use the conditional mean of the limit distribution $\mathcal{L}(\tau_p, \tau_{p+1})$ as the scaling factor of the pairwise test, which can be estimated by $\hat{b}_n(\tau_p, \tau_{p+1})$ as shown in (3.17). The resulting joint test statistic becomes

$$\widehat{JT}_n \equiv \frac{1}{P} \sum_{p=1}^P \frac{\hat{V}_n(\tau_p, \tau_{p+1})}{\hat{b}_n(\tau_p, \tau_{p+1})}.$$

By the joint convergence (4.1), we see that under the joint null hypothesis, $\Delta_n^{-1}\widehat{JT}_n$ converges stably in law to

$$\frac{1}{P} \sum_{p=1}^P \frac{\mathcal{L}(\tau_p, \tau_{p+1})}{\mathbb{E}[\mathcal{L}(\tau_p, \tau_{p+1})|\mathcal{F}]}$$

Under the alternative, $\Delta_n^{-1}\widehat{JT}_n$ diverges to $+\infty$ in probability. We reject the joint null hypothesis at significance level $\alpha \in (0, 1)$ if $\Delta_n^{-1}\widehat{JT}_n > cv_{n,\alpha}^{JT}$, where the critical value $cv_{n,\alpha}^{JT}$ is computed using the following algorithm.

ALGORITHM 2. Step 1. Simulate independent $(\tilde{\kappa}_s, \tilde{\zeta}_{s\pm})_{s \in \{\tau_1, \dots, \tau_{P+1}\}}$ such that $\tilde{\kappa}_s \sim \text{Uniform}[0, 1]$ and $\tilde{\zeta}_{s\pm} \sim \mathcal{N}(0, I_{\tilde{r}})$.

Step 2. Compute $\tilde{L}_n(\tau_p, \tau_{p+1})$ as in Algorithm 1 and set $\tilde{L}_n^{JT} = P^{-1} \sum_{p=1}^P \tilde{L}_n(\tau_p, \tau_{p+1})/\hat{b}_n(\tau_p, \tau_{p+1})$.

Step 3. Repeat step 1 and step 2 for a large number of times and report the critical value $cv_{n,\alpha}^{JT}$ as the $1 - \alpha$ quantile of \tilde{L}_n^{JT} in the Monte Carlo sample. \square

We note that for the joint test we can take $\tilde{\beta}_{n,j,\tau_p} = P^{-1} \sum_{p=1}^P \hat{\beta}_{n,j,\tau_p}$ in the construction of $\hat{X}_n(\tau_{\pm})$ (see (3.11)) which in turn is used for computing $\hat{M}_{C,n}(p, q)$. This is because the hypothesis now is that the jump beta remains constant across the $P + 1$ jump events.

4.2 Gradual jumps and the mixed-scale approach

In practice, the analysis of the co-jump behavior among assets is often complicated by the so-called “gradual jump” phenomenon (Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009)). That is, jumps in the efficient price (e.g., due to a surprising macroeconomic announcement) may take a longer time to be fully incorporated in the observed prices of less liquid stocks than is the case for the highly liquid market index. Clearly, the resulting asynchronicity between jumps can lead to (downward) biases in the jump beta estimates, and thus distort our testing results. In the jump regression setting, Li, Todorov, Tauchen, and Chen (2017a) proposed a mixed-scale approach to address this issue. The idea is to precisely detect jump times in the market portfolio at the fine time scale Δ_n , and then to analyze the relationship between individual assets’ price jumps with market jumps at a coarse time scale $k\Delta_n$ so as to mitigate the aforementioned asynchronicity. Li et al. (2017a) documented the empirical effectiveness of this simple approach. In this subsection, we introduce the mixed-scale strategy into our current problem of testing and estimating the noisy linear factor model. The theory underlying the mixed-scale extension is straightforward, and hence, we only focus on the implementation but we leave out the technical details for brevity.

The mixed-scale version of our econometric procedure goes as follows. As in Section 2.1, we detect the jumps in the systematic factor Z at the (fine) scale Δ_n . However, we estimate the jump betas at a (possibly) coarse scale $k\Delta_n$ for some $k \geq 1$. A larger k makes the procedure more robust against the presence of gradual jumps, but at the cost of reducing the signal-to-noise ratio for making inference about the jumps.²¹ We denote the coarsely sampled returns of Y_j by $\Delta_{i,k}^n Y_j = Y_{j,(i+k-1)\Delta_n} - Y_{j,(i-1)\Delta_n}$ and define $\Delta_{i,k}^n Z$ similarly. The corresponding jump beta estimates are given by

$$\hat{\beta}_{n,j,\tau}^{(k)} \equiv \frac{\Delta_{i(n,\tau),k}^n Y_j}{\Delta_{i(n,\tau),k}^n Z},$$

where we emphasize the degree k of mixed-scale in our notation. Analogous to (2.5), we can define the mixed-scale test statistic $\hat{V}_n^{(k)}$ using these mixed-scale beta estimates.

The null asymptotic distribution of $\hat{V}_n^{(k)}$ is similar to that of \hat{V}_n , except that it uses different weights in the before/after jump windows as we now describe. We first generalize the definitions in (3.1) and (3.3) by setting

$$\left\{ \begin{aligned} w_{s-}^{(k)} &\equiv \frac{\sqrt{\kappa_s}}{\Delta Z_s}, & w_{s+}^{(k)} &\equiv \frac{\sqrt{k - \kappa_s}}{\Delta Z_s}, & s \in \mathcal{T}, \\ \mathcal{A}^{(k)}(s) &\equiv \sum_{p,q \in \{s-,s+\}} w_p^{(k)} w_q^{(k)} \zeta_p^\top M_C(p,q) \zeta_q + \sum_{q \in \{s-,s+\}} (w_q^{(k)})^2 M_\epsilon(q), \\ \mathcal{B}^{(k)}(\eta, \tau) &\equiv \sum_{p \in \{\tau-, \tau+\}} \sum_{q \in \{\eta-, \eta+\}} w_p^{(k)} w_q^{(k)} \zeta_p^\top M_C(p,q) \zeta_q. \end{aligned} \right. \tag{4.2}$$

²¹Empirically, we find that taking the coarse scale to be 3 minutes is often sufficiently conservative. For robustness, it is also advisable to report results for different levels of mixed-scales which we do in our empirical section.

By a straightforward adaptation of Theorem 1, we can show that under the null hypothesis, the normalized mixed-scale test statistic $\Delta_n^{-1} \hat{V}_n^{(k)}$ converges \mathcal{F} -stably in law toward

$$\mathcal{L}^{(k)}(\eta, \tau) \equiv \mathcal{A}^{(k)}(\eta) + \mathcal{A}^{(k)}(\tau) - 2\mathcal{B}^{(k)}(\eta, \tau).$$

In the general case with mixed-scale k , the \mathcal{F} -conditional mean of this limiting variable is now

$$\mathbb{E}[\mathcal{L}^{(k)}(\eta, \tau) | \mathcal{F}] = \sum_{s \in \{\eta, \tau\}} \left(\frac{1}{2} \frac{1}{SNR_{s-}} + \left(k - \frac{1}{2} \right) \frac{1}{SNR_{s+}} \right), \tag{4.3}$$

with the signal-to-noise ratio SNR defined as in (3.5). From (4.3), we see that the conditional mean of mixed-scaled test statistic depends more on the signal-to-noise ratio in the post-jump than the pre-jump window.

The critical values for the mixed-scaled test can be calculated by modifying the procedure in Section 3.2 as follows. First, we extend the definition of residual return matrices $\hat{X}_n(\tau \pm)$ to the k -mixed-scale version $\hat{X}_n^{(k)}(\tau \pm)$ by replacing $i(n, \tau +)$ with $i(n, \tau +) + k - 1$ in equation (3.11). Note that this actually only affects the definition of $\hat{X}_n^{(k)}(\tau +)$ while leaving $\hat{X}_n^{(k)}(\tau -) = \hat{X}_n(\tau -)$. The estimators $\hat{M}_{C,n}$ and $\hat{M}_{\epsilon,n}$ are then modified correspondingly. Second, we change the weights $\tilde{w}_{n,s \pm}$ in Algorithm 1 as follows:

$$\tilde{w}_{n,s-}^{(k)} = \frac{\sqrt{\tilde{\kappa}_s}}{\Delta_{i(n,s),k}^n \mathbf{Z}}, \quad \tilde{w}_{n,s+}^{(k)} = \frac{\sqrt{k - \tilde{\kappa}_s}}{\Delta_{i(n,s),k}^n \mathbf{Z}}.$$

With these modifications, we can follow the same steps as in Algorithm 1 to compute the critical values for the mixed-scale test statistic $\hat{V}_n^{(k)}$.

Next, we describe how to adapt the estimator \hat{v}_n for the variance of beta innovations (recall (3.15)) to the mixed-scale setting. Besides using the mixed-scale jump beta estimates, we also need to adjust the bias-correction term in view of (4.3). The mixed-scale generalization of \hat{v}_n is given by

$$\left\{ \begin{aligned} \hat{v}_n^{(k)} &\equiv \frac{1}{N_n} \sum_{j=1}^{N_n} (\hat{\beta}_{n,j,\tau}^{(k)} - \hat{\beta}_{n,j,\eta}^{(k)})^2 - \Delta_n \hat{b}_n^{(k)}, \quad \text{where} \\ \hat{b}_n^{(k)} &\equiv \frac{1}{N_n k_n} \sum_{s \in \{\eta, \tau\}} \left(\frac{1}{2} \frac{\|\hat{X}_n^{(k)}(s-)\|^2}{(\Delta_{i(n,s),k}^n \mathbf{Z})^2} + \left(k - \frac{1}{2} \right) \frac{\|\hat{X}_n^{(k)}(s+)\|^2}{(\Delta_{i(n,s),k}^n \mathbf{Z})^2} \right). \end{aligned} \right. \tag{4.4}$$

Similar to \hat{v}_n , the estimator $\hat{v}_n^{(k)}$ is consistent for the cross-sectional average (conditional) second-moment of $\chi_{j,\eta,\tau}$.

5. MONTE CARLO STUDY

We now examine the finite-sample performance of the proposed test in simulation settings that are calibrated to match some key features of the data that is used in our empirical analysis.

5.1 Simulation settings

We start with describing the Monte Carlo setup. We consider $N = 100$ and $\Delta_n = 1/400$ which are very close to the corresponding numbers for the data set used in the empirical application. The unit of time is one day, so this sampling frequency corresponds roughly to sampling every minute. The log returns of the assets are generated from the following model:

$$\begin{cases} dY_{j,t} = \lambda_j^\top df_t + \tilde{\sigma}_{j,t} d\tilde{W}_{j,t} + \beta_{j,t} \varphi_t dN_t, \\ dZ_{j,t} = \lambda_Z^\top df_t + \varphi_t dN_t, \end{cases} \quad (5.1)$$

where the process N_t is a counting process for the jumps. The diffusive systematic factor process f is a three-factor model given by

$$df_t = \Sigma_{f,t}^{1/2} dW_t, \quad \Sigma_{f,t} = \text{diag}(\sigma_{f,1,t}^2, \sigma_{f,2,t}^2, \sigma_{f,3,t}^2),$$

and $W_t = (W_{1,t}, W_{2,t}, W_{3,t})^\top$ is a three-dimensional standard Brownian motion. The log-variance process of factor $k \in \{1, 2, 3\}$ is generated from an Ornstein–Uhlenbeck process of the form

$$\begin{aligned} d \log(\sigma_{f,k,t}^2) &= 0.3(\log(\bar{\sigma}_{f,k}^2) - \log(\sigma_{f,k,t}^2)) dt \\ &\quad + 0.3(\rho(dW_{1,t} + dW_{2,t} + dW_{3,t}) + \sqrt{1 - \rho^2} dB_{k,t}) + \tilde{\varphi}_{k,t} dN_t, \end{aligned}$$

where $\rho = -0.6$ captures the so-called leverage effect between the volatility and the price process and $(B_{k,t})_{1 \leq k \leq 3}$ are standard Brownian motions that are independent of W . We set

$$(\bar{\sigma}_{f,1}^2, \bar{\sigma}_{f,2}^2, \bar{\sigma}_{f,3}^2) = (0.92, 0.21, 0.07),$$

and draw the log-volatility jump sizes $\tilde{\varphi}_{k,t}$ independently from a Uniform[0, 0.2] distribution. We draw the price jump size φ_t independently from

$$\varphi_t \sim \text{Uniform}[-\underline{\varphi} - 0.4, -\underline{\varphi}] \cup [\underline{\varphi}, \underline{\varphi} + 0.4].$$

We set the lower bound $\underline{\varphi} = 0.4$, which is about 8 times of the local diffusive standard deviation of Z for $\Delta_n = 1/400$. This is close to what we observe in the empirical application. The frequency of jump arrivals is set to one jump per month.

In each simulation, we draw the diffusive factor loadings λ_j independently from

$$\lambda_{j,1} \sim \text{Uniform}[0.3, 1.7], \quad \lambda_{j,2}, \lambda_{j,3} \sim \text{Uniform}[-1, 1],$$

and the factor loading of Z is fixed at $\lambda_Z = (1, 0, 0)^\top$.

Turning next to the idiosyncratic diffusive component of Y , we set

$$\tilde{\sigma}_{j,t}^2 = \tilde{\gamma}_j (\sigma_{f,1,t}^2 + \sigma_{f,2,t}^2 + \sigma_{f,3,t}^2). \quad (5.2)$$

In this setting, the idiosyncratic variances co-move with the factor variances, which is empirically realistic. The “variance betas” $\tilde{\gamma}_j$ are drawn independently from a

Uniform[0.8, 1.2] distribution. We note that the idiosyncratic diffusive variance in our simulation is quite sizable; indeed, it is equal to the sum of all diffusive factor variances on average, so it is greater, in particular, than the variance of the leading factor Z . This simulation setting thus corresponds to a low signal-to-noise scenario for estimating the spot jump betas and, hence, presents a nontrivial challenge to the proposed test.

The above specification of Z and the diffusive part of Y is designed to mimic key features of the data set that we use in the empirical application. In particular, the volatility dynamics of Z , the frequencies of jumps and their magnitude in the Monte Carlo match approximately those of the market portfolio which we use in the next section as the reference process Z . Further, the diffusive factor structure and the magnitude of the idiosyncratic variance in Y match those of the empirical data set (recall that these quantities determine the limiting distribution of our test statistic).²²

Finally, our specification of the jump betas is as follows. At the first jump time τ_1 , the jump betas are drawn independently from

$$\beta_{j,\tau_1} \sim \text{Uniform}[0.7, 1.3].$$

Under the null hypothesis, we set $\beta_{j,\tau_p} = \beta_{j,\tau_1}$ for $p > 1$, and under the alternative we set:

$$\begin{aligned} \beta_{j,\tau_p} &= \beta_{j,\tau_{p-1}} + \chi_{j,\tau_{p-1},\tau_p}, & \chi_{j,\tau_{p-1},\tau_p} &\sim \mathcal{N}(0, v_{\tau_{p-1},\tau_p}), \\ v_{\tau_{p-1},\tau_p} &= \begin{cases} 0.1 & \text{Alternative 1,} \\ 0.2 & \text{Alternative 2.} \end{cases} \end{aligned}$$

Note that Alternative 2 implies larger deviation from the null hypothesis than Alternative 1.

Below, we examine the test specified in Theorem 3 and refer to it as the baseline test. We also consider the mixed-scale versions implemented at scales $k = 2$ and $k = 3$ as described in Section 4.2. We remind the reader that the baseline test is a special case of the mixed-scale test with $k = 1$. In the empirical analysis, the mixed-scale approach is more robust to the presence of gradual jumps, but at the cost of lowering the signal-to-noise ratio; we aim to examine this trade-off in the simulation exercise. Finally, we also assess the performance of the joint test of Section 4.1. We use nine consecutive jumps for the construction of this test which is similar to the average number of jumps used in our empirical implementation. The joint test is expected to have more power relative to the pairwise test and we will study this in our Monte Carlo experiment.

The simulation results are based on 1000 sample paths in total and the tuning parameters of the test are as follows. We set the loss function $L(x) = \min\{|x|, x^2\}$, which will also be used in our empirical analysis below. With the true number of diffusive factors being 3, we consider $\bar{r} \in \{1, 3, 5\}$ in our Monte Carlo experiments. The benchmark case is $\bar{r} = 3$, which coincides with the true number of factors. The case with $\bar{r} = 1$ is generally *not* justified by our asymptotic theory, but it allows us to understand the consequence

²²More specifically, the eigenvalues of the matrix $M_C(\tau_{\pm}, \tau_{\pm})$ decay on average as those in the real data and the ratio $M_e(\tau_{\pm})/\text{Trace}[M_C(\tau_{\pm}, \tau_{\pm})]$ is similar to that in the data as well.

TABLE 1. Rejection rates of pairwise constant beta tests.

Level	Null			Alternative 1			Alternative 2		
	10%	5%	1%	10%	5%	1%	10%	5%	1%
<i>Panel A: Case $\bar{r} = 3$ (correct number of diffusive factors)</i>									
$k = 1$	9.7	3.8	0.5	99.9	99.9	98.5	100	100	99.9
$k = 2$	9.4	4.2	0.6	98.9	97.0	84.4	99.7	99.1	96.2
$k = 3$	7.9	2.6	0.7	93.8	87.6	64.5	99.7	99.1	96.2
<i>Panel B: Case $\bar{r} = 1$ (underestimated number of diffusive factors)</i>									
$k = 1$	15.2	7.1	1.5	100	100	99.6	100	99.9	99.8
$k = 2$	10.8	5.4	0.7	98.6	96.9	87.7	99.7	99.2	96.5
$k = 3$	8.6	3.5	0.6	93.7	87.6	65.0	98.2	96.6	86.0
<i>Panel C: Case $\bar{r} = 5$ (overestimated number of diffusive factors)</i>									
$k = 1$	10.5	5.5	0.9	100	100	99.4	100	100	99.8
$k = 2$	7.1	3.0	0.3	99.0	97.0	84.7	99.5	99.1	96.0
$k = 3$	6.6	2.9	0.2	93.3	86.6	60.2	97.8	95.1	83.5

Note: The table presents rejection rates of pairwise tests for constant beta in the setting with $k_n = 30$. Critical values are based on Algorithm 1 with 1000 simulations.

of “undershooting” the number of diffusive factors. The case $\bar{r} = 5$ is justified by our theory and shows the effect of conservatively including “non-factors” in finite-samples. The local window parameter k_n is taken from $\{25, 30\}$. The algorithm for computing critical values is implemented using 1000 simulations.

Finally, we set the truncation u_n in an adaptive way. In particular, for determining the jumps in Z , u_n is equal to $7 \times \Delta_n^{0.49} \times \sqrt{RV \wedge BV}$, where RV and BV denote the daily realized variance and so-called bipower variation (Barndorff-Nielsen and Shephard (2004)), respectively.²³ The former is a measure of total daily quadratic variation and the latter is a measure of the continuous part of the quadratic variation. The above high threshold minimizes the probability of erroneously classifying a diffusive increment as one containing a jump. For the truncation of the elements in Y in the calculation of $\widehat{X}_n(\tau \pm)$ in (3.11), we use twice the threshold for Z in order to account for the additional idiosyncratic risk in Y .

5.2 Simulation results for the pairwise test

We start with the case with $k_n = 30$. Table 1 presents the Monte Carlo rejection rates for testing $\beta_{j,\tau} = \beta_{j,\eta}$, $1 \leq j \leq N$. We report rejection rates at nominal levels 10%, 5%, and 1%. Panel A shows results for the case when \bar{r} coincides with the number of diffusive factors (i.e., $\bar{r} = 3$). Under the null hypothesis, the finite-sample rejection rates for our baseline test are very close to the nominal levels. We also see similarly good performance at mixed-scale $k = 2$. These results are in line with our asymptotic theory. The test at mixed-scale $k = 3$ appears to be somewhat undersized under null hypothesis, reflecting the cost of using a larger mixed-scale.

²³We use a relatively stringent jump detection rule so as to avoid misclassifying diffusive returns as jumps. This is particularly important in our current testing context, because the asymptotic theory depends crucially on using only the jump returns.

Turning to the power analysis, we first observe that the baseline test rejects essentially with probability one under Alternative 1, which is consistent with the asymptotic theory. The mixed-scale tests are more conservative but still exhibit good power properties. The somewhat lower rejection rate at $k = 3$ is expected, because the jump signal gets weaker relative to the disturbance from the diffusive part, rendering sharp inference concerning the jumps more difficult. As the data generating process further deviates away from the null hypothesis (i.e., Alternative 2), the rejection rates of all tests become very close to one, even for the more conservative $k = 3$ mixed-scale test.

We now examine the effect of misspecifying the number of diffusive factors. Panel B of Table 1 shows the rejection rates when we “undershoot” the number of diffusive factors by setting $\bar{r} = 1$. We remind the reader that our asymptotic theory does *not* justify the validity of this implementation. Under the null hypothesis, the baseline test slightly overrejects at the 5% and 1% levels, but the overrejection appears to be more severe at the 10% level. For the mixed-scale tests with $k = 2$ and 3, the overrejection is offset by the underrejection seen in Panel A; such “offsetting” of course needs to be taken with a grain of salt. By contrast, Panel C of Table 1 shows that “overshooting” the number of diffusive factors does not lead to overrejection under the null. Indeed, the baseline test with $k = 1$ has rejection rates that are very close to the nominal levels, and the mixed-scale tests underreject only slightly in comparison with the same tests implemented using $\bar{r} = 3$. Under the alternative, results in Panel C are quite similar to those in Panel A, suggesting that the test has desirable power properties even if we set \bar{r} conservatively.

Next, we investigate the sensitivity of the rejection rates with respect to the local window parameter k_n . To do so, we repeat the exercise in Table 1 except that we replace $k_n = 30$ with $k_n = 25$. The rejection rates are presented in Table 2, which are generally very similar to those in Table 1. These results suggest that reducing moderately the window size does not have much of an effect on the test. We hence focus on the case $k_n = 30$ in our numerical work below.

In summary, we find that the proposed test controls size well under the null hypothesis and has high rejection rates under the alternatives. The performance of the baseline test is closely in line with our asymptotic theory. The mixed-scale approach with $k = 2$ performs very similarly as the baseline test. When we increase the mixed-scale to $k = 3$, the test tends to be slightly undersized but it still has nontrivial power under the alternative. Finally, we find that the test is robust with respect to the choice of \bar{r} and the local window k_n . Overall, we have seen that the finite-sample performance of the proposed test is satisfactory in empirically realistic settings.

5.3 Simulation results for the joint test

We finish this section with presenting simulation results for the performance of the joint test in Table 3. For brevity, we only consider the case with $k_n = 30$. Compared to the pairwise test, the under-rejection under the null hypothesis now becomes more significant for the mixed-scale $k = 3$. Intuitively, the small finite-sample biases in the pairwise test play a more prominent role in the joint test as pooling the pairwise tests across the different jump times reduces the sampling variability without affecting these biases. On the

TABLE 2. Rejection rates of pairwise constant beta tests.

Level	Null			Alternative 1			Alternative 2		
	10%	5%	1%	10%	5%	1%	10%	5%	1%
<i>Panel A: Case $\bar{r} = 3$ (correct number of diffusive factors)</i>									
$k = 1$	11.1	5.0	0.9	100	100	99.0	100	100	99.8
$k = 2$	7.3	3.3	0.4	98.3	96.4	84.8	99.8	99.3	96.3
$k = 3$	7.5	3.3	0.5	90.5	83.3	56.1	98.3	96.1	84.0
<i>Panel B: Case $\bar{r} = 1$ (underestimated number of diffusive factors)</i>									
$k = 1$	14.6	6.8	1.4	100	100	99.4	100	100	99.9
$k = 2$	12.3	4.5	0.9	98.4	97.4	85.9	99.7	99.4	96.7
$k = 3$	9.9	4.6	0.5	94.4	86.7	63.9	98.2	96.1	87.0
<i>Panel C: Case $\bar{r} = 5$ (overestimated number of diffusive factors)</i>									
$k = 1$	9.4	4.5	0.5	99.9	99.9	98.3	99.9	99.9	99.9
$k = 2$	9.2	3.4	0.1	98.7	95.8	84.0	99.7	99.3	95.3
$k = 3$	6.8	2.9	0.4	93.4	84.3	59.0	97.3	94.7	82.4

Note: The table presents rejection rates of pairwise tests for constant beta in the setting with $k_n = 25$. Critical values are based on Algorithm 1 with 1000 simulations.

other hand, and as expected, we now have uniformly significantly more power to reject the two alternatives.

6. EMPIRICAL APPLICATION

We now apply the econometric techniques developed above to a data set from the U.S. equity market. Our interest in this analysis will be the factor structure and presence of firm-specific risk in the cross-section of asset returns during market-wide jump events.

TABLE 3. Rejection rates of joint constant beta tests.

Level	Null			Alternative 1			Alternative 2		
	10%	5%	1%	10%	5%	1%	10%	5%	1%
<i>Panel A: Case $\bar{r} = 3$ (correct number of diffusive factors)</i>									
$k = 1$	10.0	4.1	0.8	100	100	100	100	100	100
$k = 2$	8.3	4.3	0.5	100	100	99.5	100	100	100
$k = 3$	4.3	1.0	0.0	99.3	97.7	85.6	99.9	99.6	97.5
<i>Panel B: Case $\bar{r} = 1$ (underestimated number of diffusive factors)</i>									
$k = 1$	12.6	6.1	1.3	100	100	100	100	100	100
$k = 2$	9.2	3.2	0.3	100	100	99.7	100	100	99.8
$k = 3$	6.2	3.0	0.4	99.4	97.6	84.0	99.8	99.3	94.6
<i>Panel C: Case $\bar{r} = 5$ (overestimated number of diffusive factors)</i>									
$k = 1$	11.0	5.0	1.2	100	100	100	100	100	100
$k = 2$	6.5	2.8	0.3	100	100	98.9	100	100	100
$k = 3$	4.5	1.5	0.1	99.4	97.1	85.0	99.7	99.1	94.3

Note: The table presents rejection rates of joint tests for constant beta in the setting with $k_n = 30$. The number of jump times included in the test is 9. Critical values are based on Algorithm 2 with 1000 simulations.

6.1 Data

We start with describing the data. Our sample consists of 1-minute returns for a panel of 93 stocks that were in the S&P 100 index for the entire period 2007–2015.²⁴ The data is from the Trade and Quote (TAQ) database. Standard cleaning procedures were used in assembling the data. Our proxy for the market (which is our systematic factor Z) is the S&P 500 E-Mini futures Index on the nearest contract with appropriate rollovers. The futures data were obtained from Tick Data. Of course, nothing in our theory requires Z to be the market factor, and empirical extensions could consider alternative sets of interesting event times such as macro and/or monetary announcement times, volatility jumps, as well as jumps in the various S&P 500 industry subcomponents, among others.

The detection of jumps in Z as well as the truncation needed for the construction of $\widehat{X}_n(\tau \pm)$ in (3.11) is done exactly as in the Monte Carlo, with the only difference being the scaling of the threshold by a time-of-day function in order to account for the well-known diurnal intraday pattern in volatility. We find 91 putative jump moves in Z in our sample. We further filter these moves based on the following criteria: remove detections that occur within 5 minutes of each other and remove detections of size less than 0.10 (i.e., 10 basis points); the latter occurred only in the very late part of the sample with exceptionally low volatility. These steps yield 83 (1-minute) jump returns over the 2007–2015 period.²⁵ Evidently, the 83 jump returns are very large moves in the futures index relative to its local volatility. Our analysis below will be based on these 83 market jump events in the sample.

6.2 Test outcomes

We first perform tests for the constant jump beta hypothesis to each of the 82 consecutive pairs of market jump events in our sample. As in the simulations, we use the Huber-type loss function $L(x) = \min\{|x|, x^2\}$ and take the local window parameter to be $k_n = 30$. This exercise consists of testing 82 hypotheses independently, and we summarize the results by reporting the proportion of tests that reject the null hypothesis at the 5% and 1% nominal levels when averaged over the whole sample. These separate pairwise tests are for the null hypothesis of no variation in adjacent jump betas. To further increase power, we also implement the joint test of Section 4.1 for the jumps in each of the calendar years in our sample and report the average of the p-values across the 9 years in the data set. As shown in our simulation experiment, the joint test is expected to have more power to detect violations of the null hypothesis as those considered in the Monte Carlo. We choose the time span of one calendar year for the joint test in order to avoid violations of the null hypothesis of the exact jump factor model that are due to predictable

²⁴There are 107 ticker symbols that were in the S&P 100 Index for at least 4 years over the same period. Our inference techniques do not require a balanced panel, but we retained the 93 stocks mainly for transparency and ease of exposition. In initial work, we did the pairwise tests for all 107 stocks with essentially no difference in outcomes relative to those reported here.

²⁵Similar to Bajgrowicz, Scaillet, and Treccani (2016), we do not find much evidence for clustering of these very big jumps although the statistical power for detection of such clustering is probably not very high given the rare nature of these jumps.

TABLE 4. Test results for U.S. equity data.

Mixed Scale	$\bar{r} = 3$			$\bar{r} = 5$		
	Rejection Rate		Joint Test	Rejection Rate		Joint Test
	5% Level	1% Level	Mean P-value	5% Level	1% Level	Mean P-value
<i>Panel A: Full Sample ($q_n^w = 0$)</i>						
$k = 1$	0.94	0.76	0.00	0.91	0.72	0.00
$k = 2$	0.38	0.23	0.04	0.35	0.21	0.05
$k = 3$	0.18	0.00	0.37	0.17	0.00	0.37
<i>Panel B: Winsorized Sample ($q_n^w = 0.05$)</i>						
$k = 1$	0.89	0.66	0.00	0.89	0.63	0.00
$k = 2$	0.32	0.21	0.08	0.32	0.21	0.08
$k = 3$	0.13	0.00	0.44	0.10	0.00	0.45

Note: Note: P-values are computed using Algorithm 1 based on 1000 simulations. The joint test is done by grouping jumps in periods of calendar years.

changes in the factor loadings. Indeed, the latter are typically associated with changes in the firms' characteristics and they are not expected to change much over periods of one year.²⁶ Therefore, violations of the null hypothesis in our setting will be more likely due to the presence of unpredictable firm-specific shocks at the market-wide jump events.

For implementing the tests, we choose $\bar{r} = 3$ as motivated by the popular Fama–French model (Fama and French (1993)) and consider $\bar{r} = 5$ as a more conservative robustness check. Since the results are not sensitive to the choice of \bar{r} , below we focus the discussion on the case $\bar{r} = 3$. Finally, we present results with and without winsorization; the winsorization is implemented at the 0.05 level (i.e., $q_n^w = 0.05$). Recall from Section 2 that the winsorization is needed in order to guard against the possible occurrence of idiosyncratic jumps in the high-frequency increment at which a systematic jump event has been detected. In finite samples, the winsorization also provides robustness against possible presence of outliers in the data.

We turn now to a discussion of the test results. From Panel A of Table 4, the baseline test at the 1-minute frequency (i.e., $k = 1$) rejects the constant beta hypothesis for a majority of consecutive jump pairs: the rejection rate is 0.94 at the 5% level and 0.76 at the 1% level. The average p-value of the joint test is very low, suggesting a strong rejection of the joint hypothesis. However, as recognized in Li et al. (2017a), the constant beta hypothesis at a fine scale will be rejected if there is a gradual response in individual assets to a market jump event due to less liquidity in the trading process of some of the stocks in the cross-section at such frequency. The mixed-scale version of our test described in Section 4.1 mitigates the effect of such microstructure issues, and we consider it therefore as the more reliable method. By “zooming out” and measuring returns at the 2-minute sampling interval (i.e., $k = 2$), we find that the rejection rate drops substantially: for example, only 23% of the tests reject at the 1% level, which is far below the rejection rate of the baseline test with $k = 1$. In addition, the averaged p-value of

²⁶Our pairwise test is most “conservative” in this regard as the time between the market jumps is typically around 1 month.

the joint test is slightly below the conventional significance level of 5%, suggesting only borderline rejection for the joint hypothesis. When we further zoom out by increasing the mixed-scale parameter to $k = 3$, the majority of pairwise tests no longer reject the constant beta hypothesis. Furthermore, the p-value for the joint test is far above the conventional significance level.

As seen from Panel B of Table 4, after performing winsorization, the rejection rates of the tests drop uniformly across all considered settings. However, qualitatively the results are the same as in the case of no winsorization.

Overall, the significant drop in the rejection rate of both the pairwise and joint tests, when going from $k = 1$ to $k = 3$ observed in our data, is inconsistent with the behavior of the tests under a typical alternative hypothesis of a noisy jump factor model as illustrated in our simulation experiments. Instead, such behavior is consistent with a null hypothesis of exact jump factor model which is “clouded” at the fine time scale of 1-minute by liquidity-related microstructure issues.²⁷

We can further contrast this result for the jumps with the factor structure of the diffusive returns before and after the systematic jump events. The cross-sectional average contribution of idiosyncratic diffusive risk in the total return variation of the residual component $X_n^*(\tau \pm)_{j,l}$ in (3.7) (computed as the time series of the ratio $\widehat{M}_{\epsilon,n}/(\text{Trace}[\widehat{M}_{C,n}(q, q)] + \widehat{M}_{\epsilon,n})$ for $\bar{r} = 3$) is a nontrivial 67%. This decomposition of systematic and idiosyncratic variance in $X_n^*(\tau \pm)_{j,l}$ reveals also that the market cannot solely account for the cross-sectional dependence in asset returns before and after the systematic jump events which is in sharp contrast to our empirical evidence in support of the exact jump factor model at market jump times.

6.3 The magnitude of firm-specific shocks at market-wide jump events

We next estimate the variance $v(\eta, \tau)$ of the innovations in jump betas using the estimator \hat{v}_n (see (3.15)), as well as the mixed-scaled version $\hat{v}_n^{(k)}$ (see (4.4)). Although the estimand $v(\eta, \tau)$ is nonnegative, the bias-corrected estimator $\hat{v}_n^{(k)}$ occasionally takes negative values in finite-samples because of the subtraction of the positive bias-correction term. Therefore, we report results for a sign-regularized version of $\hat{v}_n^{(k)}$ given by

$$\hat{v}_n^{(k)+} = \max\{\hat{v}_n^{(k)}, 0\}. \tag{6.1}$$

To evaluate economically the relative magnitude of these estimates, we also report the following measure:

$$\text{Normalized } \hat{v}_n^{(k)+} \equiv \hat{v}_n^{(k)+} \left(\frac{(\Delta_{i(n,\tau),k}^n Z)^2 + (\Delta_{i(n,\eta),k}^n Z)^2}{2} \right) / D_n, \tag{6.2}$$

where

$$D_n \equiv \frac{1}{N_n} \sum_{j=1}^{N_n} (\Delta_{i(n,\tau)}^n Y_j)^2 + \frac{1}{N_n} \sum_{j=1}^{N_n} (\Delta_{i(n,\eta)}^n Y_j)^2 + \frac{1}{2} \sum_{s \in \{\tau, \eta\}} \frac{1}{k_n N_n} \sum_{j=1}^{N_n} \sum_{l=1}^{k_n} (\Delta_{i(n,s-l)}^n Y_j)^2$$

²⁷Of course, such a behavior of the test can be also consistent with very short-lived volatility bursts which are confused for jumps; see, for example, Bajgrowicz, Scaillet, and Treccani (2016).

$$+ \left(k - \frac{1}{2}\right) \sum_{s \in \{\tau, \eta\}} \frac{1}{k_n N_n} \sum_{j=1}^{N_n} \sum_{l=1}^{k_n} (\Delta_{i(n,s^+)+k-1+l}^n Y_j)^2.$$

The ratio in (6.2) measures an average stock's price variation due to time-varying jump beta (i.e., the numerator) relative to the total price variation during market jump times (i.e., D_n).

Panel A of Table 5 displays the empirical quantiles (across jumps and stocks) of the estimates $\hat{v}_n^{(k)+}$ and the normalized $\hat{v}_n^{(k)+}$ for all 82 pairs of consecutive jumps; the quantiles are invariant to the monotonic transformation used in the regularization (6.1). Panel A1 suggests that the median value of $\hat{v}_n^{(k)+}$ is about 0.17 and this estimate varies little across various levels of mixed-scales. To gauge their relative magnitudes, we observe in Panel A2 that the median value of the normalized measure at mixed-scale $k = 1, 2,$ and 3 are 7.70%, 6.80%, and 5.38%, respectively. This means, for example, for jump returns measured at the 2-minute frequency, return variance due to the temporal variation of the jump betas explains only about 6.80% of the total jump return variation. This rather small percent is consistent with our testing result that the majority of pairwise tests do not reject the null hypotheses at the conventional significance levels (for $k = 2$).

In Panel B of Table 5, we report the same statistics but after performing winsorization in the cross-section. This winsorization is to guard against distortions resulting from the idiosyncratic jump risk. Indeed, the estimators of variance can be very sensitive to outliers and the estimates in Panel A may be overly influenced by a few stocks, which makes them of less interest as measures of jump beta variations for the majority of stocks. Clearly, the estimates after winsorization are much smaller than those based on the full sample, suggesting that the estimates in Panel A have likely exaggerated the variations in jump betas for a typical stock because of a few influential observations.

The estimates reported in Panel B1 of Table 5 indicate that between consecutive large rare jumps the typical betas differ randomly with a variance of about 0.02, based on the estimate for $k = 3$. For perspective on this estimate, the corresponding normalized measure (6.2) indicates that such variation of beta explains only 0.99% of the total return variation of the median stock in our sample at the systematic jump event. By contrast, the relative contribution of idiosyncratic risk in the total diffusive return variation in the local windows before and after the systematic jump events (based on $\bar{r} = 3$) is a rather nontrivial 48%.

Taken together, the results in Table 5 document a rather small role for firm-specific shocks in asset returns at the times of market-wide jump events. This evidence is in line with the support for the exact jump factor model reported earlier.

7. CONCLUSION

In this paper, we develop a formal test for deciding whether an exact factor model holds for asset returns in a large cross-section at the jump times of discretely-observed systematic risk factors. The inference is based on a panel of high-frequency asset returns in which the cross-sectional dimension and the sampling frequency increase simultaneously but with no restriction on their relative asymptotic order. The test is based on

TABLE 5. Idiosyncratic risk at market-wide jump events.

Mixed Scale	Percentiles					Percentiles				
	10%	25%	50%	75%	90%	10%	25%	50%	75%	90%
<i>Panel A: Full Sample ($q_n^w = 0$)</i>										
	<i>(A1) $\hat{v}_n^{(k)+}$</i>					<i>(A2) Normalized $\hat{v}_n^{(k)+}$ (%)</i>				
$k = 1$	0.09	0.12	0.17	0.23	0.30	3.92	5.11	7.70	11.26	14.91
$k = 2$	0.04	0.09	0.19	0.34	0.70	1.31	3.64	6.80	11.74	22.43
$k = 3$	0.00	0.03	0.16	0.30	0.85	0.00	0.91	5.38	11.58	24.04
<i>Panel B: Winsorized Sample ($q_n^w = 0.05$)</i>										
	<i>(B1) $\hat{v}_n^{(k)+}$</i>					<i>(B2) Normalized $\hat{v}_n^{(k)+}$ (%)</i>				
$k = 1$	0.06	0.09	0.13	0.19	0.22	2.51	3.61	5.84	8.37	10.86
$k = 2$	0.00	0.04	0.11	0.18	0.37	0.00	1.42	4.33	7.79	11.97
$k = 3$	0.00	0.00	0.02	0.15	0.35	0.00	0.00	0.99	6.11	10.64

comparing temporal variation of estimates for the factor loadings from detected consecutive jump times of the risk factor. This difference in factor loading estimates shrinks asymptotically when the null hypothesis is true and is $O_p(1)$ otherwise. The limit distribution of the test is nonstandard and depends on systematic and idiosyncratic diffusive risks around the jump times of the factors in distinct ways. We further develop an estimator that allows the econometrician to formally assess the magnitude of firm-specific shocks in assets at the times of systematic factor jump events. Empirical application to stocks in the S&P 100 index provides support for an exact market jump model over the period 2007–2015. This stands in sharp contrast to our evidence for the presence of multiple systematic risk factors in the asset returns prior to and following the systematic jump events.

REFERENCES

- Aït-Sahalia, Y. and J. Jacod (2014), *High-Frequency Financial Econometrics*. Princeton University Press. [433]
- Aït-Sahalia, Y. and D. Xiu (2018), “Principal component analysis of high frequency data.” *Journal of the American Statistical Association*, forthcoming. [439]
- Ang, A. and J. Chen (2002), “Asymmetric correlations of equity portfolios.” *Journal of Financial Economics*, 63, 443–494. [423]
- Bai, J. (2003), “Inferential theory for factor models of large dimensions.” *Econometrica*, 71 (1), 135–171. [424, 439]
- Bai, J. (2009), “Panel data models with interactive fixed effects.” *Econometrica*, 77 (4), 1229–1279. [424, 425]
- Bai, J. and S. Ng (2002), “Determining the number of factors in approximate factor models.” *Econometrica*, 70 (1), 191–221. [437]

- Bajgrowicz, P., O. Scaillet, and A. Treccani (2016), “Jumps in high-frequency data: Spurious detections, dynamics and news.” *Management Science*, 62, 2198–2217. [441, 450, 452]
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009), “Realized kernels in practice: Trades and quotes.” *The Econometrics Journal*, 12 (3), C1–C32. [443]
- Barndorff-Nielsen, O. E. and N. Shephard (2004), “Power and bipower variation with stochastic volatility and jumps.” *Journal of Financial Econometrics*, 2, 1–37. [424, 447]
- Barndorff-Nielsen, O. E. and N. Shephard (2006), “Econometrics of testing for jumps in financial economics using bipower variation.” *Journal of Financial Econometrics*, 4, 1–30. [424]
- Barras, L., O. Scaillet, and R. Wermers (2010), “False discoveries in mutual fund performance: Measuring luck in estimated alphas.” *Journal of Finance*, 65, 179–216. [441]
- Bollerslev, T., S. Li, and V. Todorov (2016), “Roughing up beta: Continuous versus discontinuous betas, and the cross-section of expected stock returns.” *Journal of Financial Economics*, 120, 464–490. [426]
- Connor, G., M. Hagmann, and O. Linton (2012), “Efficient semiparametric estimation of the Fama-French model and extensions.” *Econometrica*, 80 (2), 713–754. [421]
- Fama, E. and K. French (1993), “Common risk factors in the returns on stocks and bonds.” *Journal of Financial Economics*, 33. [451]
- Gagliardini, P., E. Ossola, and O. Scaillet (2016a), “A diagnostic criterion for approximate factor structure.” Discussion paper, Swiss Finance Institute. [424, 437]
- Gagliardini, P., E. Ossola, and O. Scaillet (2016b), “Time-varying risk premium in large cross-sectional equity data sets.” *Econometrica*, 84 (3), 985–1046. [421]
- Giacomini, R. and H. White (2006), “Tests of conditional predictive ability.” *Econometrica*, 74 (6), 1545–1578. [433]
- Hansen, L. P. and S. F. Richard (1987), “The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models.” *Econometrica*, 55 (3), 587–613. [421]
- Huber, P. (2004), *Robust Statistics*. John Wiley & Sons. [428]
- Jacod, J. and P. Protter (2012), *Discretization of Processes*. Springer. [433]
- Lee, D. S. and T. Lemieux (2010), “Regression discontinuity designs in economics.” *Journal of Economic Literature*, 48 (2), 281–355. [424]
- Li, J., V. Todorov, and G. Tauchen (2017), “Jump regressions.” *Econometrica*, 85, 173–195. [424, 429]
- Li, J., V. Todorov, and G. Tauchen (2019), “Supplement to ‘Jump factor models in large cross-sections.’” *Quantitative Economics Supplemental Material*, 10, <https://doi.org/10.3982/QE1060>. [425]

- Li, J., V. Todorov, G. Tauchen, and R. Chen (2017a), “Mixed-scale jump regressions with bootstrap inference.” *Journal of Econometrics*, 201, 417–432. [443, 451]
- Li, J., V. Todorov, G. Tauchen, and H. Lin (2017b), “Rank tests at jump events.” *Journal of Business and Economic Statistics*, forthcoming. [424]
- Longin, F. and B. Solnik (2001), “Extreme correlation of international equity markets.” *Journal of Finance*, 56, 649–676. [423]
- Mancini, C. (2001), “Disentangling the jumps of the diffusion in a geometric jumping Brownian motion.” *Giornale dell’Istituto Italiano degli Attuari*, LXIV, 19–47. [424]
- Mancini, C. (2009), “Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps.” *Scandinavian Journal of Statistics*, 36, 270–296. [424]
- Merton, R. C. (1976), “Option pricing when underlying stock returns are discontinuous.” *Journal of Financial Economics*, 3, 125–144. [426]
- Moon, H. R. and M. Weidner (2015), “Linear regression for panel with unknown number of factors as interactive fixed effects.” *Econometrica*, 83 (4), 1543–1579. [425]
- Pelger, M. (2015), “Understanding systematic risk: A high-frequency approach.” Discussion paper, Stanford University. [424, 439]
- Pelger, M. (2019), “Large-dimensional factor modeling based on high-frequency observations.” *Journal of Econometrics*, 208 (1), 23–42. [424, 439]
- Pesaran, M. H. (2006), “Estimation and inference in large heterogeneous panels with a multifactor error structure.” *Econometrica*, 74 (4), 967–1012. [425]
- Savor, P. and M. Wilson (2014), “Asset pricing: A tale of two days.” *Journal of Financial Economics*, 113, 171–201. [423]
- Stock, J. H. and M. W. Watson (2002), “Forecasting using principal components from a large number of predictors.” *Journal of the American Statistical Association*, 97 (460), 1167–1179. [424, 439]

Co-editor Frank Schorfheide handled this manuscript.

Manuscript received 28 January, 2018; final version accepted 1 November, 2018; available online 16 November, 2018.