Intraday Variation in Systematic Risks and Information Flows

Andrew J. Patton[a,b] and Yasin Simsek[b]

[a] *School of Banking & Finance, UNSW Sydney*
[b] *Department of Economics, Duke University*

This version: October 31, 2024

Abstract

This paper analyzes variation in the factor structure of asset returns within a trade day by combining non-parametric kernel methods with principal component analysis. We estimate the model on a collection of over 400 high frequency U.S. equity returns over the period 1996-2020 and show that the proposed model has superior explanatory power relative to a collection of well-known observable factor models and standard PCA. We present a stylized model of asset prices and information flows and show that the factor structure of asset returns varies with the arrival of news. Using data on individual firm earnings announcements, FOMC announcements, and other macroeconomic announcements, we provide evidence consistent with our stylized model, that the superior performance of the proposed model is due to time variation in the factor structure of asset returns around times of information flows.

Keywords: News, principal component analysis, high-frequency data, announcements.

J.E.L. Codes: C38, C58, G12, G14.

*We thank Torben Andersen, Tim Bollerslev, Anna Cieslak, Roger Edelen, Peter Hansen, Jihyun Kim (discussant), Jia Li, Bradley Paye, Viktor Todorov, Guofu Zhou and seminar participants at Duke University, Singapore Management University, Toulouse School of Economics, Virginia Tech, Washington University in St. Louis, the 2024 IAAE annual meeting and the 2024 Society for Financial Econometrics annual meeting for helpful comments.
Email: yasin.simsek@duke.edu.

1 Introduction

Factor models are fundamental tools for understanding systematic risks. The widely used Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965) remains a workhorse for analyzing financial data, and extensions, including Fama and French (1992), Carhart (1997), Fama and French (2015) and many more, have provided additional insights into the risk and returns of financial assets. Nevertheless, these models are known to be imperfect, failing to capture some forms of comovements in stock returns, see Fama and French (2004) and Giglio and Xiu (2021). This has given rise to a proliferation of factor models, famously called the "factor zoo" by Cochrane (2011).

An alternative to building a model with a pre-specified set of observable factors is to extract the factors directly from the data, resulting in the class of unobservable, or latent, factor models. The estimation of the latent factors is most commonly carried out by principal component analysis (PCA), see Hastie et al. (2009) for a textbook treatment. This approach was first used by Connor and Korajczyk (1986) in asset pricing and can be motivated by the arbitrage pricing theory of Ross (1976). Recent noteworthy extensions of PCA for financial applications include Kelly et al. (2019), Pelger (2020) and Lettau and Pelger (2020b), amongst others, which we discuss further below.

Standard implementations of factor models, whether based on observable or latent factors, assume that the factor loadings are constant, yet the battery of evidence of time-varying asset return volatility (Bollerslev et al. (1994)), correlations (Engle (2009)), and risk premia (Shanken (1990)), questions this assumption, and motivates allowing for a time-varying factor structure, see for example Jagannathan and Wang (1996); Connor et al. (2012); Kelly et al. (2019); Gu et al. (2021), among others. These papers consider factor structures that evolve at lower frequencies, such as monthly or daily intervals. Recent work by Andersen et al. (2021, 2023) and Liao and Todorov (2024) goes further and suggests that the factor structure in asset returns also varies at higher frequencies, *within* the trading day.

Motivated by this empirical evidence, we propose a new framework to allow for intraday variation in latent factor models. We draw on recent econometric advances, see Su and Wang (2017) and Fan et al. (2022), that combine non-parametric kernel methods with principal component analysis, generalizing the latent factor model of Connor and Korajczyk (1986) by treating factor exposures as an unknown function of the time of day. We estimate the new factor model on 15-minute data for more than 400 individual U.S equity returns over the period from January 1996 to December 2020. We show that our proposed "Intraday PCA" has superior explanatory power, both in-sample and out-of-sample, economically as well as statistically, relative to a collection of well-known observable factor models and standard PCA. Indeed, the improvement in fit from using Intraday PCA compared with the famous Fama-French (1992) three-factor model (FF3) is larger than the improvement from using the FF3 compared with the workhorse CAPM.

Our empirical analysis unveils a striking intraday pattern in the improvement of performance metrics, with gains concentrated predominantly around first periods of the trading day, with smaller and sometimes insignificant gains at the market close. We hypothesize that the overnight flow of information, and its incorporation into asset prices at the market open, leads to variation in the factor structure of asset returns that standard factor models cannot capture. We formalize this hypothesis via a stylized model of high frequency returns and information flows, and demonstrate that information flows can lead to changes in the factor structure of asset returns. The model further predicts that these changes are larger when the surprise component of the news is greater.

We test our "information flow" hypothesis in three distinct empirical studies. First, we analyze our model's behavior following overnight periods with more surprising earnings announcements (relative to equity analysts' expectations). We find that new model achieves significantly greater improvements on days following larger earning surprises, supporting our hypothesis.

Second, we analyze the relative performance of the proposed Intraday PCA model on Federal Open Market Committee (FOMC) announcement days. This study is particularly revealing because, unlike earnings announcements which mostly occur in the overnight period, FOMC announcements occur *within* the trade day, typically at 2pm. Consistent with our "information flow" hypothesis, we find the largest improvements from our model in the periods immediately after 2pm. Moreover, we make use of the "monetary policy surprise" measure of Bauer and

Swanson (2023) to quantify how surprising a given announcement was to financial markets, and we find that our proposed model works particularly well when the surprise component FOMC announcement is large.

Finally, we consider a set of other macroeconomic announcements: non-farm payrolls, consumer price index (CPI), the Conference Board's consumer confidence index, and the ISM manufacturing report. We find the improvements in fit from using local PCA relative to standard PCA is larger for all four of these announcement types, relative to the full sample, with the improvements being between two and nearly four times as large as in the full sample.

1.1 Related literature

This work engages with several strands of the asset pricing literature. First and foremost, our estimation method links this paper to high-dimensional models in asset pricing. Starting with Ross (1976), this branch analyzes latent factors for returns, e.g. Connor and Korajczyk (1986), Aït-Sahalia and Xiu (2019), Pelger (2019, 2020), Lettau and Pelger (2020a,b), and Pelger and Xiong (2022), amongst others. For example, Lettau and Pelger (2020b) employs a penalized PCA method to extract latent factors using not only second moments but also first moments of returns. Pelger and Xiong (2022) proposes a semi-parametric framework, estimating factors and exposures as a function of states. Pelger (2019, 2020) apply standard PCA to stock returns and estimate overnight, intraday and jump factors separately. Our paper extends these two papers to consider whether, and why, the factor structure of asset returns varies within the trade day.

This paper is also related to recent work by Andersen et al. (2021, 2023) and Liao and Todorov (2024) exploring intraday variation in market beta and factor structures. Specifically, Andersen et al. (2021) show that dispersion in CAPM betas is high during the market open and gradually subsiding as market approaches the close. Andersen et al. (2023) shows that this heightened dispersion in market beta is attenuated when controlling for additional factors such as size, value and momentum. Liao and Todorov (2024) propose a test to determine whether the span of factor exposures is the same at the market open and the market close. Our paper complements these papers by proposing a method to capture within-day variation in a factor model and showing that doing so leads to economically and statistically significant forecast improvements. Furthermore, we conjecture an economic channel for the improved performance, namely our "information flow" hypothesis, and provide a battery of evidence consistent with this channel.

The idea of intraday variation is also related to work on capturing lower-frequency (monthly or daily) variation. For example, Connor et al. (2012) model factor loadings in observable factor models as a function of observed firm characteristics, estimating the function using semi-parametric methods. Gagliardini et al. (2016) model factor loadings as a function of macro variables and firm specific quantities. In a similar vein, Fan et al. (2016); Kelly et al. (2019); Gu et al. (2021) introduce time varying loadings in a latent factor model by linking them to firm characteristics. Like this paper, Pelger and Xiong (2022) combines local kernel methods with PCA and estimates the factor structure of (lower-frequency) asset returns as a function of economy-wide variables such the Chicago Board of Exchange's "volatility index" (VIX).

We also connect to the high-frequency econometrics literature, which utilizes the high frequency data to estimate factor loadings within in each day, see for example Barndorff-Nielsen and Shephard (2004); Andersen et al. (2005); Bollerslev et al. (2016), among many others. Our contribution to this body of research is to capture intraday variation in the factor structure, and explain its importance around times of information flows.

Finally, our paper is related to work that contrasts the behavior of asset returns in the overnight and intraday periods. For example, Lou et al. (2019) find that overnight and intraday returns show strong persistence along with an offsetting reversal effect. Bogousslavsky (2021) demonstrates that a mispricing portfolio earns a positive return for most of the trade day but yields a significant and negative alpha in the last 30 minutes of trading. We diverge from this branch of the literature as we focus solely on intraday periods, and study variation at a more granular level.

The rest of the paper is organized as follows. Section 2 presents our approach for capturing intraday variation in the latent factor structure of a panel of asset returns and presents initial evidence that such variation is statistically significant empirically. Section 3 presents a simple model of asset returns and information flows that shows how such flows can alter the factor structure of returns. Section 4 shows that the gains from allowing for a time-varying factor structure are particularly large around times of information flows. Section 5 presents some extensions and robustness checks and Section 6 concludes. An appendix contains additional details and derivations.

2 Intradaily variation in the factor structure of equity returns

This section introduces our model for the intraday factor structure in equity returns, describes the data that we use, and presents evidence that the factor structure changes predictably over the trade day.

Throughout, $r_{i,t-1+\tau}$ denotes the return on asset $i \in \{1, 2, ..., N\}$, observed at day $t \in \{1, 2, ..., T\}$ and intraday period $\tau \in \{\frac{1}{M}, \frac{2}{M}, ..., 1\}$. We use $\mathbf{r}_{t-1+\tau}$ to denote the $(N \times 1)$ vector of returns on all N assets at that time, and $\mathbf{R}$ to denote the $(TM \times N)$ matrix of returns on all assets across all time periods. We use K to denote the number of factors.

2.1 Data

This study employs high-frequency return data from the New Stock Exchange's Trade and Quote (TAQ) database, over the period January 1996 to December 2020 ($T = 6254$). We implement standard cleaning procedures following Barndorff-Nielsen et al. (2008) and merge with CRSP open-close prices. Days with incomplete trading hours are excluded from the analysis. High-frequency returns are sampled at 15-minute intervals, with overnight returns omitted.[1] As a result, there are $M = 26$ observations for each trading day. The analysis considers all stocks ever listed in the S&P 500 index in the period 1996-2020 and which traded throughout the entire sample period, resulting in a balanced panel comprising $N = 407$ stocks.[2] Summary statistics

[1]There is ample evidence in the literature, see e.g. Lou et al. (2019) and Bogousslavsky (2021), that overnight asset returns exhibit different properties to intradaily returns. The focus of this paper is whether intradaily returns, which might be thought to be similar to each other, exhibit important differences in their factor structure over the trade day, and if so, why.

[2]Omitting stocks that did not trade for the full sample period introduces survivorship bias, however our focus is on the second-moment structure of intradaily returns, not the cross-section of first moments, which are very close to zero at high frequencies, and so this is not a concern for our analysis. The methods we consider can be extended to accommodate an unbalanced panel of returns, for example estimating the principal components via alternating least squares, at the cost of an increased computational burden.

Table 1: **Summary statistics.** This table presents summary statistics on the $6,254 \times 26 \times 407$ panel of 15-minute returns we use in our analysis. We compute the mean, standard deviation, skewness and kurtosis for each of the 407 stocks, and report cross-sectional descriptive statistics (the cross-sectional mean and various percentiles) of these measures. The bottom row presents summary statistics for the $82,261$ pairwise correlations across the 407 stocks.

	Cross-sectional statistics					
	Mean	1%	5%	Median	95%	99%
Mean$\times 10^4$	0.135	-0.435	-0.173	0.143	0.419	0.559
Std Dev $\times 10^2$	0.412	0.066	0.195	0.394	0.669	0.802
Skewness	0.237	-3.067	-0.848	0.151	1.310	4.038
Kurtosis	73.127	12.622	15.454	31.291	192.388	669.695
Pairwise Corr	0.174	0.029	0.049	0.171	0.300	0.389

on these returns are presented in Table 1.

2.2 PCA and Local PCA for asset returns

Following Chamberlain and Rothschild (1983), Connor and Korajczyk (1986) and Bai (2003), a popular alternative to specifying a set of observable factors (which may be incomplete, see e.g. Giglio and Xiu, 2021), is to use the asset returns themselves to estimate the underlying factors, which can be done using principal components analysis (PCA). We assume that asset returns follow:

$$r_{i,t-1+\tau} = \boldsymbol{\beta}_i^\top \mathbf{f}_{t-1+\tau} + \varepsilon_{i,t-1+\tau} \tag{1}$$

where $\boldsymbol{\beta}_i$ is $K \times 1$ vector of exposures (or factor loadings), $\mathbf{f}_{t-1+\tau}$ is $K \times 1$ vector of factors, and $\varepsilon_{i,t-1+\tau}$ is the residual. The model can be written in matrix notation as:

$$\underbrace{\mathbf{r}_{t-1+\tau}}_{N\times 1} = \underbrace{\boldsymbol{\beta}}_{N\times K}\underbrace{\mathbf{f}_{t-1+\tau}}_{K\times 1} + \underbrace{\boldsymbol{\varepsilon}_{t-1+\tau}}_{N\times 1} \ \forall\ (t,\tau) \in \{1,...,T\} \times \{\frac{1}{M}, \frac{2}{M}, ..., 1\}. \tag{2}$$

A "strict" factor structure holds if the residuals are cross-sectionally uncorrelated. Following Bai (2003) we exploit the "large N" nature of our data set and assume only an "approximate" factor structure, which allows for some weak cross-sectional correlation in the residuals.[3] PCA allows us to find the factors and factor loadings that best explain the panel of asset returns, solving the following optimization problem:

$$\hat{\mathbf{F}}, \hat{\boldsymbol{\beta}} = \arg\min_{\mathbf{F},\beta} \frac{1}{NTM} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{\tau=1}^{M} (r_{i,t-1+\tau} - \boldsymbol{\beta}_i^\top \mathbf{f}_{t-1+\tau})^2 \tag{3}$$

The model in equation (1) is identified only up to a rotation of the factors, and to overcome this we impose an orthogonality constraint on exposures: $\boldsymbol{\beta}^\top \boldsymbol{\beta} = I_K$. The outputs of this estimation are the $TM \times K$ matrix of estimated risk factors, $\hat{\mathbf{F}}$, and the $N \times K$ matrix of estimated factor loadings, $\hat{\boldsymbol{\beta}}$. As described in Bai (2003), factor loadings can be recovered as the eigenvectors associated with the K largest eigenvalues of $\mathbf{R}^\top \mathbf{R}$, multiplied by $\sqrt{N}$. Given these, we can estimate factors by regressing the returns on the estimated loadings: $\hat{\mathbf{F}} = \mathbf{R}\hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\beta}}^\top \hat{\boldsymbol{\beta}})^{-1}$ or simply $\mathbf{R}\hat{\boldsymbol{\beta}}$, given the orthogonality constraint on the factor loadings.

Motivated by recent work on intraday variation in factor loadings in observable factor models (Andersen et al., 2021, 2023) and in the span of factors in latent factor models (Liao and Todorov, 2024), we next consider a factor model for high frequency asset returns that varies across the time of day:

$$\hat{\mathbf{F}}(\tau), \hat{\boldsymbol{\beta}}(\tau) = \arg\min_{\mathbf{F},\beta} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (r_{i,t-1+\tau} - \boldsymbol{\beta}_i^\top \mathbf{f}_{t-1+\tau})^2 \tag{4}$$

for each $\tau \in \{1/M, 2/M, ..., 1\}$. In our empirical work we use 15-minute returns, implying $M = 26$ for stocks traded on the New York Stock Exchange. The model in equation (4) can be estimated by using PCA for each the $T \times N$ matrix of returns over the same 15-minute window in our sample. That is, we use PCA on all of the 9:30-9:45am returns, then separately on the collection of 9:45-10am returns, etc. Estimating the factors and factor loadings separately for each of the M high-frequency windows is clearly more flexible than standard PCA: the factors and loadings can vary across each of the high frequency windows rather than being imposed to be constant.

Whilst flexible, the approach in equation (4) fails to incorporate the economic knowledge that adjacent high-frequency windows are *similar*. That is, while features of returns are known

[3] In detail, an approximate factor structure holds if $K < N$ eigenvalues of the covariance matrix of $\mathbf{R}$ diverge as as $N \to \infty$, while the remaining eigenvalues are bounded.

to vary across the trade day (see, e.g., Wood et al., 1985; Harris, 1986; Andersen and Bollerslev, 1998), this variation tends to be smooth, and we can draw on information from neighboring windows to improve estimation accuracy. To accomplish this, we propose using "local PCA" (see, e.g., Su and Wang, 2017, Fan et al., 2022 and Pelger and Xiong, 2022), which combines nonparametric kernel methods and principal component analysis. For a some kernel function, ϕ, and bandwidth, $h > 0$, local PCA solves:

$$\tilde{\mathbf{F}}(\tau), \tilde{\boldsymbol{\beta}}(\tau) = \arg\min_{\mathbf{F},\beta} \frac{1}{NTM} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{s=1}^{M} \phi\left(\frac{\tau - s}{h}\right) (r_{i,t-1+s} - \boldsymbol{\beta}_i^\top \mathbf{f}_{t-1+s})^2 \tag{5}$$

We dub the model obtained by applying local PCA to high frequency returns "Intraday PCA."

It is simple to show that the solution to the optimization problem in equation (5) can be obtained by applying standard PCA to the $TM \times N$ matrix of *weighted* returns, $\tilde{\mathbf{R}}(\tau)$, with representative row given by:

$$\tilde{\mathbf{r}}_{t-1+s}(\tau) = \phi\left(\frac{\tau - s}{h}\right)^{1/2} \mathbf{r}_{t-1+s} \tag{6}$$

These weighted returns lower the importance of returns from high-frequency windows that are far from the window of interest (that is, when $\tau - s$ is large).[4] As the bandwidth parameter in equation (5) shrinks ($h \to 0$), we recover the period-by-period PCA objective function in equation (4), and as it grows ($h \to \infty$), we recover the baseline PCA objective function in equation (3). In between these limiting cases, Intraday PCA allows for variation in the factor structure across the trade day, drawing on information from adjacent periods. In the next section we describe how we select the bandwidth and other hyperparameters.

2.3 Hyperparameters for PCA and Intraday PCA

A key hyperparameter in principal component analyses is the number of factors, K, to use in the model. We adopt the widely-used information criterion proposed in Bai and Ng (2002), extended

[4]This can be considered as standard PCA on the full panel of asset returns, as in equation (3), but downweighting estimation errors that occur in periods far from the period of interest and upweighting errors that occur at or near the period of interest, so that the fit is best near the period of interest.

Figure 1: **Number of factors across years.** This figure shows the number of factors selected for standard PCA using the information criterion proposed in Bai and Ng (2002) and Liao and Todorov (2024) for each year in our sample.

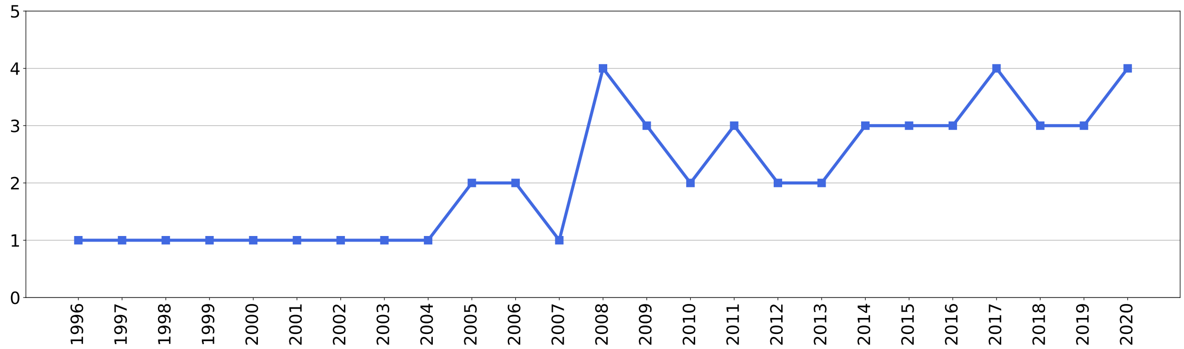

to the high frequency setting by Liao and Todorov (2024), for this purpose. In Section 5.2 we show that our results are very similar when using the method of Pelger (2019) to choose the number of factors.

First, consider the choice of K for the PCA model, where all intradaily periods are treated identically. Let $\lambda_1 \geq \lambda_2 \geq ... \geq \lambda_N$ denote the sorted eigenvalues of $\frac{1}{TN}\mathbf{R}^\top\mathbf{R}$, where $\mathbf{R}$ is the matrix of high-frequency returns. The number of factors is then estimated as:

$$\hat{K} = \text{argmin}_{1 \leq K \leq K_{max}} \log \sum_{k>K} \lambda_k + K\frac{T+N}{TN} \log\left(\frac{TN}{T+N}\right) \tag{7}$$

K_{max} is a pre-defined upper bound, which we set to 10.

We estimate $\hat{K}$ year-by-year and present the estimated values in Figure 1. The estimated number of factors varies over time, and increases almost monotonically over the sample period. Prior to 2008, we find a single factor, with the exception of the years 2005 and 2006 where our analysis identifies two factors as optimal. Subsequent to 2008, we find the presence of between two and four factors, with at least three factors found starting from 2014.

To choose the number of factors to use in the Intraday PCA model, we estimate number of factors following the same steps outlined above separately for each intradaily period and each year. Figure 2 presents the estimated number of factors across the 26 15-minute windows, for

Figure 2: **Number of factors across years and intraday periods.** This figure shows the number of factors selected for Intraday PCA using the information criterion proposed in Bai and Ng (2002) and Liao and Todorov (2024) for each 15-minute window of the New York Stock Exchange trade day. For clarity we present results for three subsamples, rather than each year separately, and report the median value in each subsample.

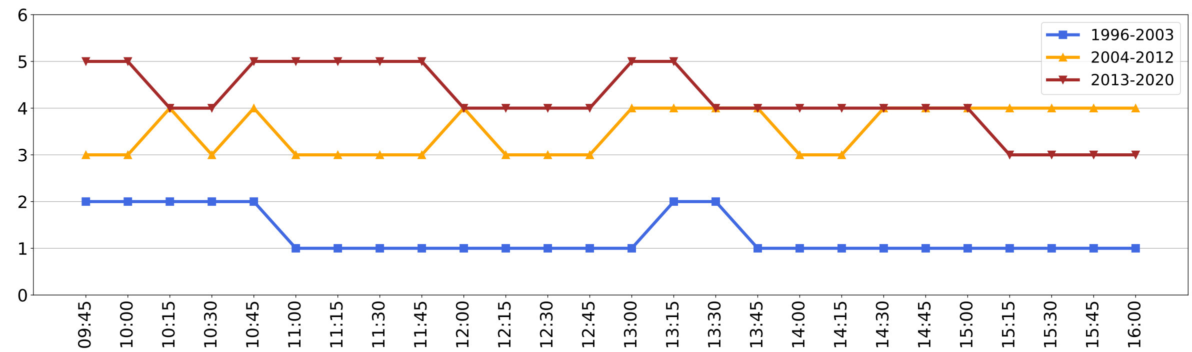

three subsamples of the data (1996-2003, 2004-2012, 2013-2020). We report the median number of factors selected for each of these three subsamples. Consistent with Figure 1, we observe that selected number of factors is notably lower for the first subsample compared with the latter two subsamples. Moreover, we find that optimal number of factors varies across the trade day. For example, in the latest subsample, the optimal number of factors at the market open is five, while it is only three at the market close. This intraday variation in the optimal number of factors is the first indication that the factor structure of these asset returns varies over the trade day. We investigate this more formally in the next subsection.

In Intraday PCA, we also need to choose a kernel to weight the intradaily periods. We adopt the Gaussian kernel:

$$\phi\left(\frac{\tau - s}{h}\right) = \exp\left\{-\frac{1}{2}\left(\frac{\tau - s}{h}\right)^2\right\} \tag{8}$$

and we impose that the kernel is one-sided at the open and close, so that the 9:30-9:45am window is not treated as adjacent to the 3:45-4pm window, by imposing zero weight before the market open and after the market close, and re-weighting the remaining windows so that the weights sum to one. Figure 3 illustrates the optimal kernel (with bandwidth estimated via cross-validation)

Figure 3: **Kernel weights across the trading day.** This figure shows the weights assigned to each intraday period for the first, middle and last 15-minute window of the trade day. The weights are computed using the Gaussian kernel described in Section 2.3 with bandwidth parameter $h = 0.17$ obtained via cross-validation.

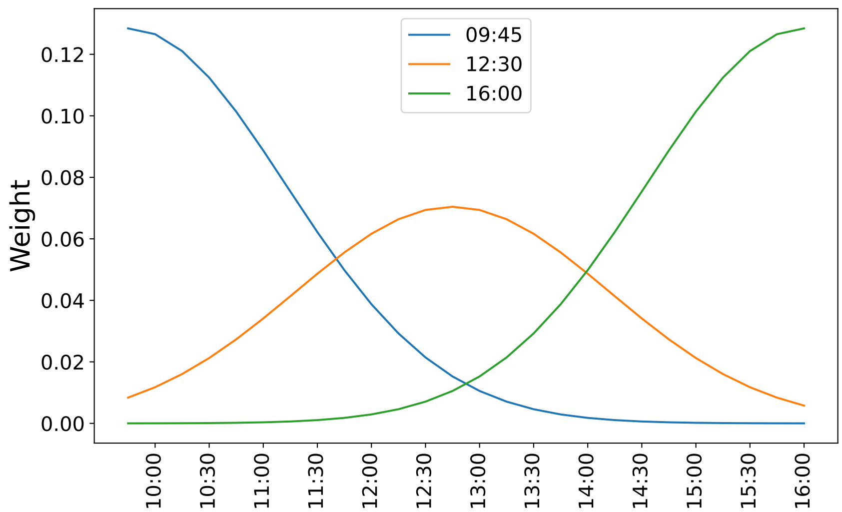

for three periods in the trade day.[5]

The final hyperparameter, the bandwidth h, controls the smoothness of the Intraday PCA model as a function of the time of day, i.e. how much information from adjactent intraday periods is used, with smaller values for h indicating a more localized model. We consider choices for the bandwidth ranging from 0.001 to 5, and we select the optimal bandwidth using a validation sample. For out-of-sample forecasts for year Y, we estimate the Intraday PCA model on data from year $Y - 2$ for each of the bandwidths, and construct forecasts for data in year $Y - 1$. The bandwidth that performs best in year $Y - 1$ is then used for the forecasts made in year Y. This year-by-year estimation procedure has the benefit that it allows for low-frequency changes in the factor structure.

[5]We also considered allowing the bandwidth parameter to vary over the trade day, but did not find much improvement in fit despite the greatly increased computational burden, and so we impose a common bandwidth parameter across all intraday periods.

2.4 Forecast performance evaluation

We compare performance of the factor models under consideration both in-sample and out-of-sample. In-sample, we compare the fit of the models using the R^2 associated with the grand mean-squared error that appears in the PCA objective function, and we test for significant differences in in-sample fit using the test of Rivers and Vuong (2002).

$$R^2_{\text{is}} = 1 - \frac{\sum_i \sum_{(t,\tau)\in\mathcal{T}} (r_{i,t-1+\tau} - \hat{\boldsymbol{\beta}}_i^\top \hat{\mathbf{f}}_{t-1+\tau})^2}{\sum_i \sum_{(t,\tau)\in\mathcal{T}} r^2_{i,t-1+\tau}} \tag{9}$$

where $\mathcal{T}$ denotes the collection of days and intradaily observations in the in-sample period. This measure captures the proportion of observed return variation explained by the model, aggregated across all assets and time observations. One can consider $\hat{\boldsymbol{\beta}}_i^\top \hat{\mathbf{f}}_{t-1+\tau} \equiv \hat{r}_{i,t-1+\tau}$ as the projection of contemporaneous returns onto the space spanned by estimated factor loadings. Therefore, comparing the R^2's of different models is equivalent to comparing the projection of factor loadings for these models. Put differently, differences in R^2 implies a discrepancy in factor loadings' projection matrices. This connects our in-sample comparisons to Liao and Todorov (2024), who use the distance between projections of factor loadings to test for differences in the factor structure between two periods of the trade day.

To compare models out-of-sample, we use an approximate "leave-one-out" R^2 metric that can be viewed as lying between the "predictive" and "total" R^2 metrics considered in Kelly et al. (2019). The "total R^2" metric uses estimation sample factor loadings and realized factor returns to construct the forecasts, and thus implicit in these predictions is information on the realized value of the target variable, meaning that they are not true forecasts. The "predictive R^2" measure, on the other hand, uses the in-sample factor loadings and replaces the realized factor return with the in-sample factor mean when constructing the forecasts, yielding proper forecasts. However, this metric is not informative in applications where the variables have means close to zero, as is the case for high-frequency asset returns, as in such applications all models, regardless of quality, are assigned an R^2 close to zero. To overcome this, we propose a simple approach where the forecast of variable i is based on factor(s) that use the realized values of *other* variables and the estimation-sample mean for variable i, leading to the "approximate leave-one-out" forecast[6]:

$$\tilde{r}_{i,t} = \hat{\boldsymbol{\beta}}_i^\top \hat{\mathbf{f}}_t - \hat{\boldsymbol{\beta}}_i^\top \hat{\boldsymbol{\beta}}_i (r_{i,t} - \hat{\mu}_i) \tag{10}$$

Appendix A presents details on the derivation of this forecast. This is a proper out-of-sample forecast of variable i, using contemporaneous information from other other variables in the panel. We then compute:

$$R^2_{\text{oos}} = 1 - \frac{\sum_i \sum_{(t,\tau)\in\mathcal{T}} (r_{i,t-1+\tau} - \tilde{r}_{i,t-1+\tau})^2}{\sum_i \sum_{(t,\tau)\in\mathcal{T}} r^2_{i,t-1+\tau}} \tag{11}$$

where $\mathcal{T}$ denotes the collection of days and intradaily observations in the out-of-sample period. We test for significant differences in out-of-sample forecast performance using the test of Diebold and Mariano (1995).

Table 2 shows in-sample fit and out-of-sample predictive accuracy of the competing models. Below each R^2 value we report the t-statistic comparing the fit with that of Intraday PCA, using Rivers and Vuong (2002) for the in-sample comparisons and Diebold and Mariano (1995) for the out-of-sample comparisons. A positive t-statistic indicates superiority of Intraday PCA in terms of R^2. This table shows that Intraday PCA exhibits superior performance compared to the observable factor models and the standard PCA specification. In-sample, Intraday PCA provides an R^2 gain of nearly 11% over CAPM and a 2.5% gain over standard PCA, both of which are significant at the 1% level. Out-of-sample R^2s are lower for all models than their in-sample counterparts, as expected, but the gains from using Intraday PCA remain significant and similarly large: 9.0% relative to the CAPM and 1.8% relative to standard PCA.

Our high frequency data enables us to delve into the details of the above performance gains and examine the *within-day* variation in predictive performance of Intraday PCA. We take the R^2 differences between Intraday PCA and standard PCA, the latter being the competing method with the best in-sample and out-of-sample fit, as shown above, and present in Figure 3 the R^2 differences across the trade day. This figure reveals that improved performance of Intraday PCA is not uniform across intraday periods, rather gains are larger and more significant near the

[6]The equation for the leave-one-out forecast exploits the normalization that $\boldsymbol{\beta}^\top\boldsymbol{\beta} = I_K$. If the other common normalization, that the factors are uncorrelated, is imposed, the functional form will differ.

Table 2: **In-sample and out-of-sample performance of factor models.** This table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) for three models based on observable factors and two models based on latent factors. The CAPM uses solely the market factor; FF3 additionally uses the size and value factors of Fama and French (1992); FF6 additionally uses the momentum factor of Carhart (1997) and the quality and profitability factors of Fama and French (2015). The column titled PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns. The last column uses the kernel-based local PCA introduced in Section 2.2. t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of Intraday PCA are reported in parentheses.

	Factor model				
	CAPM	FF3	FF6	PCA	Intraday PCA
In-sample R^2	19.56	22.60	24.40	24.44	26.12
(RV t-stat)	(18.39)	(14.11)	(9.57)	(19.41)	
OOS R^2	18.27	20.53	20.77	22.49	23.76
(DM t-stat)	(16.43)	(12.12)	(12.07)	(18.69)	

market open, and smaller and less significant at the market close. For example, in the first period Intraday PCA provides a 2% gain in OOS R^2, significant at the 0.01 level, while the improvement is between 0.5% and 1% in the last hour of the trade day.[7]

The significant improvement in fit of Intraday PCA compared with standard PCA represents evidence that the factor structure of high frequency asset returns varies within the trade day. We next propose an to answer the deeper question of *why*: what drives this variation in the factor structure over the trade day? Noting that the improvements achieved by Intraday PCA are greatest in the first hours of the trade day, and we conjecture a connection with the influx of new information in the overnight period, such as earnings announcements, Consumer Price Index releases, and non-farm payroll reports. We hypothesize that incorporation of new information into asset prices generates a change in the factor structure of asset returns in the opening hours

[7] We also implemented the test of Liao and Todorov (2024), testing that the span of latent factors is the same in the first two hours of the trade day as in the last two hours. Consistent with the conclusion from our forecast comparison results in Table 2, we reject the null at the 1% level.

Figure 4: **Comparing the fit of Intraday PCA and PCA.** This figure shows the difference in out-of-sample R^2 between Intraday PCA, introduced in Section 2.2, and standard PCA, for each of the 26 15-minute periods during the New York Stock Exchange trade day. Positive values indicate that Intraday PCA has a higher R^2 than standard PCA. The vertical lines at the end of each bar are 95% confidence intervals for the difference in R^2 based on a Diebold-Mariano (1995) test.

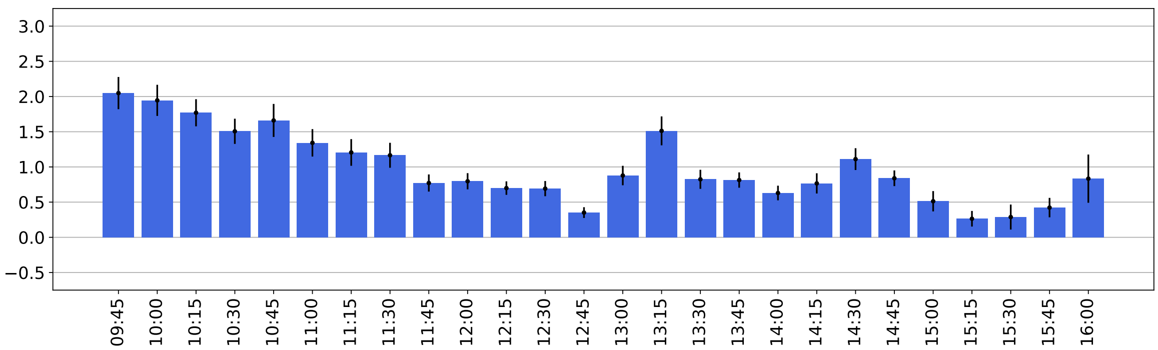

of the trade day relative to the middle and closing hours. Section 3 formalizes this conjecture and generates additional testable implications, which we take to the data in Section 4.

3 A model of high-frequency asset prices and information flows

This section presents a simple model of asset prices and information flows that reflects the three key features of our empirical environment. Firstly, multiple types of news impacts asset prices, e.g., earnings announcements, statements by the Federal Reserve, announcements by the government or an industry regulator, etc. Secondly, different types of news induce possibly different factor structures in returns, e.g., the factor structure in firm cashflows may differ from that generated by other types of news. Thirdly, different types of news arrive at different times in the trade day, e.g., earnings announcements typically occur in the overnight period, while announcements from the Federal Open Market Committee (FOMC) typically occur in the middle of the trade day.

Following the influential work of Vuolteenaho (2002) and Campbell and Vuolteenaho (2004), we model returns as driven by changes in expectations about cashflows and discount rates. We next describe how we model each of these types of news and how they impact asset returns.

We assume that log-earnings follow a random walk, consistent with Watts and Leftwich (1977), Kormendi and Lipe (1987) and Kothari (2001), who note that earnings are indeed very persistent. This implies that "cashflow news" can be modeled as updates to expectations about the current level of log-earnings for a given firm. Our model for the change in log-earnings is:

$$\Delta \log X_{i,t} = g_i + \gamma_i^\top \mathbf{Z}_t + u_{i,t} \tag{12}$$

where g_i is a firm-specific average growth rate of earnings, $\mathbf{Z}_t$ are the K factors driving earnings growth, γ_i is firm i's exposure to the common factors, and $u_{i,t}$ is the unpredictable component of earnings growth for firm i.

We assume that firms report their earnings every C days, always in the overnight (or 0^{th}) period of the day.[8] Since the level of earnings of a given firm are not observed every day, and since there are common factors in earnings growth, investors update their expectations about profitability using information contained in the earnings announcements of other firms. We adopt the model for cross-stock learning proposed in Patton and Verardo (2012), who use a Kalman filter designed to accommodate the intermittent nature of earnings announcements (Sinopoli et al., 2004). We describe the filtering problem in detail in Appendix B. The output of this filter is a sequence of cashflows news for each day and each firm:

$$N_{i,t}^{CF} = \mathbb{E}_t[\log X_{i,t}] - \mathbb{E}_{t-1}[\log X_{i,t-1}] \tag{13}$$

where the information set available at time t includes the earnings announcements of all firms up to and including day t, and $\mathbb{E}_t[\log X_{i,t}]$ denotes the "nowcast" of the earnings of firm i in period t based on the Kalman filter (see equation A.38 in the appendix). It varies even in the absence of an earnings announcement from firm i due to updates to expectations based on announcements by related firms.

[8] We discuss specific choices for the parameters of the model at the end of this subsection.

We model discount rate news as a "monetary policy surprise" (e.g., see Bauer and Swanson, 2023), which can affect each firm differently through its sensitivity to the surprise. We assume that the Fed makes an announcement every D days, and that this announcement always occurs in the d^{th} period of the announcement day, consistent with the 2pm timing of most announcements from the FOMC. On non-announcement days, and in non-announcement trading periods, the monetary policy suprise variable is equal to zero. The output of this part of the model is a sequence of discount rate news for each day and each firm:

$$N_{i,t-1+d/M}^{DR} = \beta_i^{MPS} S_t \tag{14}$$

Finally, to capture the slow diffusion of information (e.g., Brennan et al., 1993; Hong and Stein, 1999; Boguth et al., 2016), we assume that news is incorporated into prices with a possible delay: with probability $\pi_{CF} \in (0,1]$ the price of firm i reacts to cashflow news in the first trading period after the news is released, with probability $(1-\pi_{CF})\pi_{CF}$ it reacts in the second period, $(1-\pi_{CF})^2\pi_{CF}$ in the third, etc. We introduce an analogous parameter, $\pi_{DR} \in (0,1]$, for the incorporation of discount rate news into prices. This feature means that even though cashflow news is assumed to arrive in the overnight period, and we focus only on intradaily returns thus omitting the overnight return, it is still possible for intra-day returns to react to overnight cashflow news.

Combining these effects, our model for asset returns is:

$$r_{i,t-1+\tau} = N_{i,t-1+\tau}^{CF} + N_{i,t-1+\tau}^{DR} + \varepsilon_{i,t-1+\tau} \tag{15}$$

where $r_{i,t-1+\tau}$ is the return on firm i on day $t \in \{1, 2, ..., T\}$ and intraday period $\tau \in \{\frac{1}{M}, \frac{2}{M}, ..., 1\}$, $N_{i,t-1+\tau}^{CF}$ and $N_{i,t-1+\tau}^{DR}$ are the cashflow news (equation 13) and discount rate news (equation 14) incorporated into prices in this period, and $\varepsilon_{i,t}$ is the return attributable to noise traders, which we model as independent across time and firms.

This model implies that the periods that reflect cashflow news "inherit" the factor structure in earnings: for announcing firms, earnings growth realizations have the factor structure in equation

(12), and for non-announcing firms the Kalman filter updates to expectations also inherit this factor structure, see Appendix B. In the absence of news, returns are purely due to noise traders and no factor structure is present. The d^{th} period in the trading day *occasionally* has a one-factor structure, driven by the release of monetary policy news, but on non-announcement days returns in that period behave like those in the other intraday periods, and no factor structure is present.

Given the Kalman filtering step required to capture updates to investors' expectations in the presence of intermittent, correlated, earnings announcements, no closed-form expressions for the factor structure are available. Instead, as in Patton and Verardo (2012), we rely on 1000 simulations of this model. We assume $T = 3000$ trade days and $M = 20$ intradaily periods, both comparable to our data. Due to computational memory constraints, we assume only $N = 100$ firms and $C = 10$ days between earnings announcements.[9] To match the empirical feature that FOMC announcements occur about twice as often as earnings announcements (eight versus four times per year) we set $D = 5$, and we model them as arriving in intradaily period 13, approximately matching the 2pm announcements by the Fed. We set the diffusion of information parameters to $\pi_{CF} = \pi_{DR} = 0.25$. The choices for the remaining parameters of the model are presented in Appendix B.

Figure 5 presents the improvements in R^2 from using Intraday PCA compared with standard PCA, for each of the intradaily trading periods, averaged across the 1000 simulations.[10] For this first simulation, we shut down the discount rate news channel and focus on cashflow news only.[11] We see that the figure is qualitatively comparable to using real data. The gains are greatest in the periods immediately after flows of information, which in this model are those at the start of the day, after the overnight earnings announcements. For the rest of the trade day there is no difference in fit between the two models, as they both correctly select zero factors.

Next, we consider the implications of allowing for both cashflow and discount rate news. Panel A of Figure 6 and shows that the gains for allowing for a time-varying factor structure are

[9]The framework of Sinopoli et al. (2004), which we use, requires defining a state variable of size $N \times C$, and attempting to match our data on this dimension, with $N \approx 400$ and $C \approx 63$ makes the model too large to be stored in memory. The qualitative predictions with this choice of N and C are unaffected.

[10]As in our empirical application, we use Bai and Ng (2002) to select the number of factors to use in each of the models, and we use cross validation to choose the optimal bandwidth for Intraday PCA.

[11]In our data only approximately $8/252 \approx 3.2\%$ of days have an FOMC announcement, while our simulation design has a monetary policy surprise on 20% of days, roughly six times as often.

Figure 5: **Comparing the fit of Intraday PCA and PCA, simulated data.** This figure shows the difference in out-of-sample R^2 between Intraday PCA and standard PCA, using the simulation design described in Section 2. Positive values indicate that Intraday PCA has a higher R^2 than standard PCA.

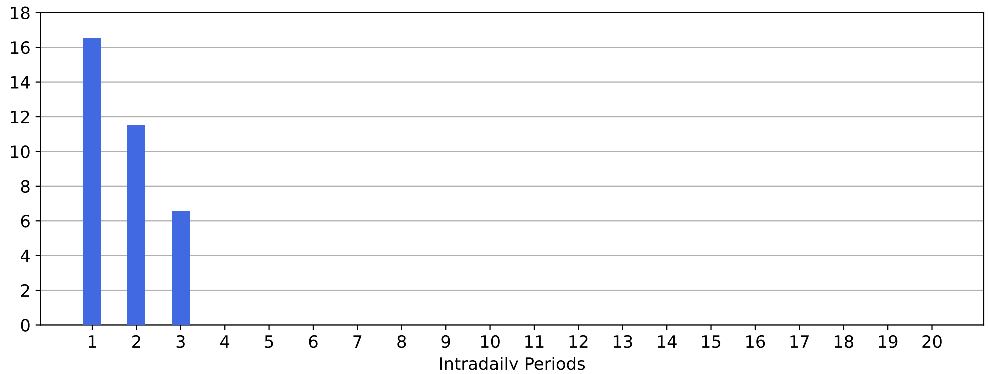

greatest in the periods immediately after flows of information, namely after the overnight earnings announcements at the start of the day, and those starting from period 13, when monetary policy surprises are realized. This suggests that the gains from using Intraday PCA will differ on FOMC announcement days versus other days, a prediction we test in Section 4.

We can also use the model to identify the types of news that lead to the greatest changes in the intraday factor structure. For example, in Panel B of Figure 6 we increase the variance of the unpredictable component of earnings, thereby making each earnings announcement "more surprising." We see an increase in the gains from using Intraday PCA when earnings announcements resolve more uncertainty. Similarly, we can increase the size of the monetary policy surprise, and in Panel C of Figure 6 we see that the gains from using Intraday PCA are also greater in that case. We test these two predictions from this model in the next section.

4 Information flows and the factor structure of equity returns

We now empirically consider some of the insights into the impact of information flows on asset return factor structures afforded by the model presented in the previous section. We first test whether days following a large number of firm earnings announcements are also ones where

Figure 6: **Comparing the fit of Intraday PCA and PCA, simulated data.** This figure shows the difference in out-of-sample R^2 between Intraday PCA and standard PCA, using the simulation designs described in Section 2. Positive values indicate that Intraday PCA has a higher R^2 than standard PCA. Panel (a) shows the results when both cashflow and discount rate news are present. Panel (b) shows the impact of larger cashflow news. Panel (c) shows the impact of larger discount rate news.

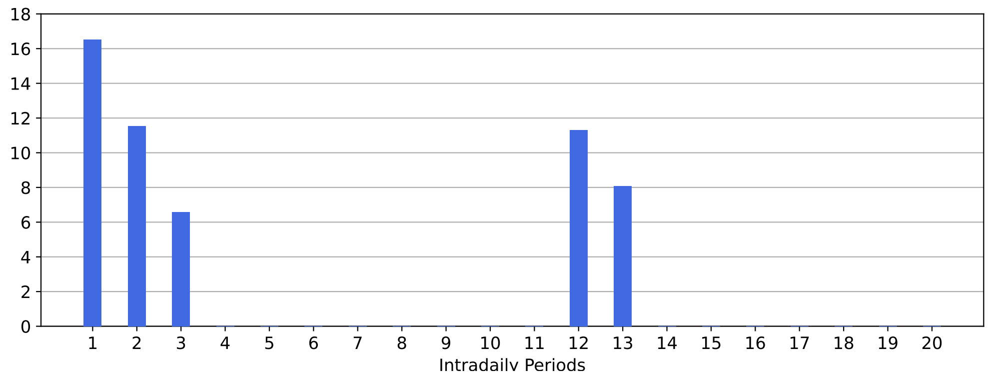

(a)

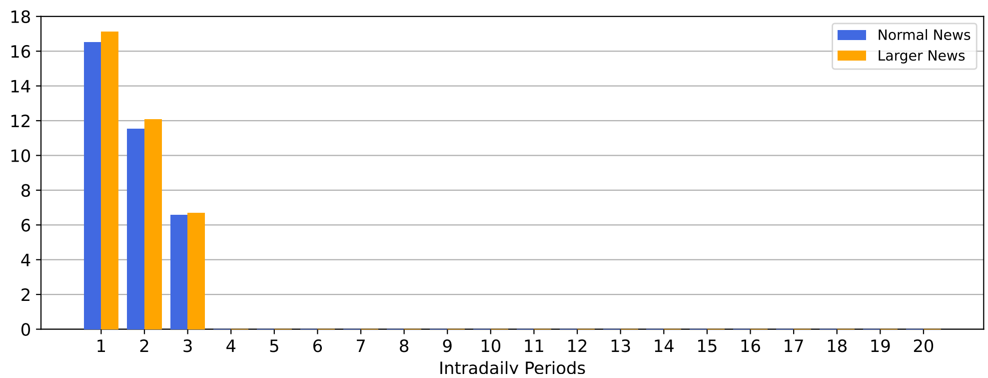

(b)

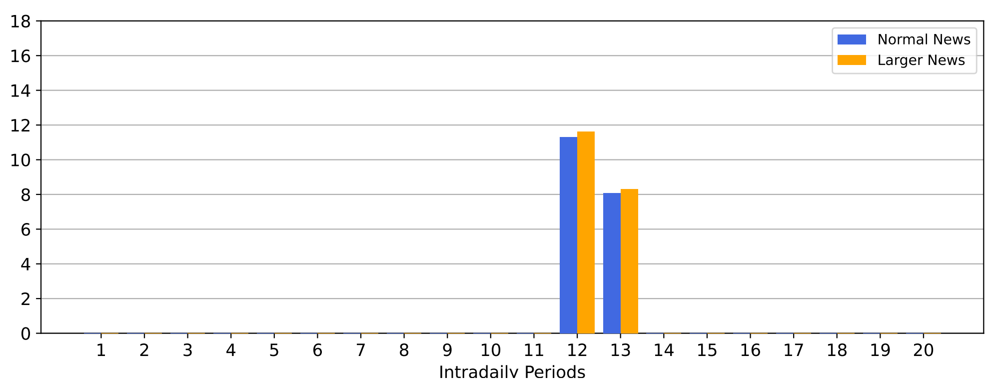

(c)

Intraday PCA outperformance is particularly large. Using equity analysts' forecasts of earnings, we are also able to sort earnings announcements by how surprising they were to analysts. Second, we consider FOMC announcements, which have the added benefit of occuring *during* the trade day, and so generate a clear testable implication for the information flow explanation for changes in the factor structure of asset returns. Finally, we consider four other important macroeconomic announcements, and compare the gains from using Intraday PCA on announcement days with the gains on average days.

4.1 Earnings anouncements

Firms' quarterly earnings announcements are well-documented disclosures of essential information regarding their operational and financial status, and prior research has found significant impacts of earnings announcements on equity prices, see for example Collins and Kothari (1989), Livnat and Mendenhall (2006) and Patton and Verardo (2012). We investigate whether earnings announcement surprises lead to changes in the *factor structure* of asset returns. The model in Section 3 implies that the factor structure of asset returns varies over the trade day, and that the variation is stronger when earnings announcements are "more surprising." Following Livnat and Mendenhall (2006), we measure the surprise component of a given earnings announcement as the difference between realized earnings-per-share and the median analyst forecast in the 90 days prior to the announcement, scaled by current quarter's price for cross-sectionaly comparability. We sort earnings days into quintiles using according to this measure of earnings surprise.

Table 3 shows R^2 improvement of Intraday PCA over PCA for each quintile of earnings surprise, and the last column shows the difference in R^2 improvements between the top and bottom quintiles. Table 3 reveals that the outperformance of Intraday PCA over standard PCA is significant for all types of earnings announcements (the t-statistics in the first five columns are all well above 1.96). Moreover, the outperformance is greatest for the most surprising earnings announcements and smallest for the least surprising announcements. The difference in the improvement in fit is 1.2% in-sample and 1.1% out-of-sample, and both differences are significant at the 1% level. The increased outperformance from allowing for a time-varying factor structure

Table 3: **Performance of Intraday PCA over standard PCA on earnings announcement days by quintile of earnings surprise.** This table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) on earnings announcement days, sorted into quintiles by absolute earnings surprise as described in Section 4.1. Each cell in the first five columns reports the difference in R^2 of Intraday PCA versus PCA. PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns, and Intraday PCA uses kernel-based local PCA introduced in Section 2.2. The last column shows the difference of top and bottom quintiles. t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of Intraday PCA are reported in parentheses.

	Quintile of absolute earnings surprise					
	Low	2	3	4	High	High-Low
In-sample R^2 gain	1.23	1.49	1.92	1.67	2.44	1.21
(RV t-stat)	(16.45)	(17.44)	(8.88)	(18.06)	(14.48)	(6.89)
OOS R^2 gain	0.89	1.26	1.46	1.30	2.01	1.12
(DM t-stat)	(15.78)	(16.16)	(14.04)	(16.02)	(13.71)	(8.03)

following more surprising news is consistent with the model in the previous section.

4.2 FOMC Announcements

We next investigate the factor structure of asset returns around perhaps the most important single piece of news for financial markets: the announcement by the Federal Open Market Committee (FOMC) of the target range for the federal funds rate.[12] FOMC announcements generally occur eight times per year, and in our sample we have a total of 196 announcement days. We obtain FOMC date and time information from the Bloomberg Economic Calendar.[13] In the analyses decribed above we use one-year rolling windows for estimation, however given the small number of FOMC announcement dates per year we use an expanding window estimation strategy for the

[12] An extensive body of research shows that monetary policy announcements have significant impacts on the stock market, see for example Bernanke and Kuttner (2005); Cieslak and Schrimpf (2019); Bollerslev et al. (2018).

[13] Almost 70% of the FOMC announcements in our sample occur at 2pm; another 20% occur at 2:15pm, and the remainder occuring at other times. Importantly for our analysis, all FOMC announcements in our sample period occur within the NYSE trade day.

Table 4: **In-sample and out-of-sample performance of factor models in FOMC announcements days.** This table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) using only FOMC announcement days, described in Section 4.2, for three models based on observable factors and two models based on latent factors. The CAPM uses solely the market factor; FF3 additionally uses the size and value factors of Fama and French (1992); FF6 additionally uses the momentum factor of Carhart (1997) and the quality and profitability factors of Fama and French (2015). The column titled PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns. The last column uses the kernel-based local PCA introduced in Section 2.2. t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of Intraday PCA are reported in parentheses.

	Factor model				
	CAPM	FF3	FF6	PCA	Intraday PCA
In-sample R^2	21.95	25.04	26.23	24.30	35.38
(RV t-stat)	(16.20)	(15.14)	(14.19)	(16.19)	
OOS R^2	26.42	29.02	29.36	28.91	33.69
(DM t-stat)	(13.47)	(11.35)	(9.68)	(12.90)	

study of these dates. Specifically, considering the days post-2010 as the out-of-sample period, the factor models are estimated for each date using all available data up to that date.

Table 4 shows the in-sample and out-of-sample performances of the competing factor models on FOMC days. We find that Intraday PCA exhibits the highest in-sample and OOS R^2, 35.4% and 33.7% respectively. Intraday PCA provides an R^2 gain of over 11% compared to standard PCA in-sample, and a gain of 4.8% out-of-sample, with both improvements being significant at the 1% level. Comparing the results on FOMC days with the full-sample results from Table 2, we see that the improvements from allowing for a time-varying factor structure are much greater on FOMC days: 11.1% versus 1.7% in-sample, and 4.8% versus 1.3% out-of-sample, consistent with the predictions of the model presented in Section 3.

Next, we exploit a feature of FOMC announcements that is particularly informative for our "information flows" hypothesis: unlike earnings announcements, which predominantly occur outside of trading hours, FOMC announcements typically happen at 2pm, *during* the trade

Figure 7: **Comparing the fit of Intraday PCA and standard PCA on FOMC announcement days.** This figure shows the difference in out-of-sample total R^2 between Intraday PCA, introduced in Section 2.2, and standard PCA on FOMC announcement days, described in Section 4.2, for each of the 26 15-minute periods during the New York Stock Exchange trade day. Positive values indicate that Intraday PCA has a higher R^2 than standard PCA. The vertical lines at the end of each bar are 95% confidence intervals for the difference in R^2 based on a Diebold-Mariano (1995) test.

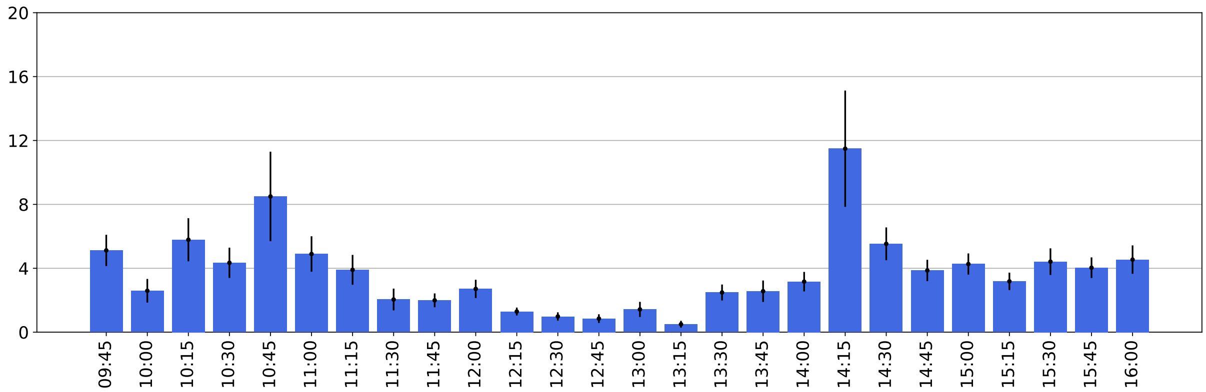

day. Thus, if it is the arrival of new information to the market and its subsequent processing by market participants, we would expect to see larger gains from using Intraday PCA in the windows immediately after 2pm. Figure 7 presents the OOS R^2 differences between Intraday PCA and PCA across the trade day, and reveals a pattern consistent with our conjecture: the gains from using Intraday PCA are particularly large, as high as 11%, in the first two 15-minute windows after the FOMC announcement.

The model presented in Section 3 predicts that the gains from allowing for changes in the factor structure of asset returns will be larger for when the news is "larger". To gauge the surprise component of a given FOMC announcement we use the "monetary policy surprise" series of Bauer and Swanson (2023), which is estimated as the first principal component of changes in the first four quarterly Eurodollar futures contracts around the announcement.[14]

Figure 8 presents the improvement in OOS R^2 for each of the 196 FOMC announcement days as a function of the monetary policy surprise (measured as the log-squared surprise from Bauer and Swanson, 2023). The figure reveals a broadly increasing pattern, especially for larger

[14]The monetary policy surprise data is available at `www.michaeldbauer.com`.

Figure 8: **Comparing the fit of Intraday PCA and PCA across monetary policy surprises.** This figure shows the difference in out-of-sample total R^2 between Intraday PCA, introduced in Section 2.2, and standard PCA for each of the 97 FOMC announcement days between 2010 and 2020 as a function of the monetary policy surprises (MPS) of Bauer and Swanson (2023). We use $log\ MPS^2$ as the measure of the magnitdue of the surprise. Positive values on the y-axis indicate that intraday PCA has a higher R^2 than standard PCA. The dashed line is a fitted local polynomial curve.

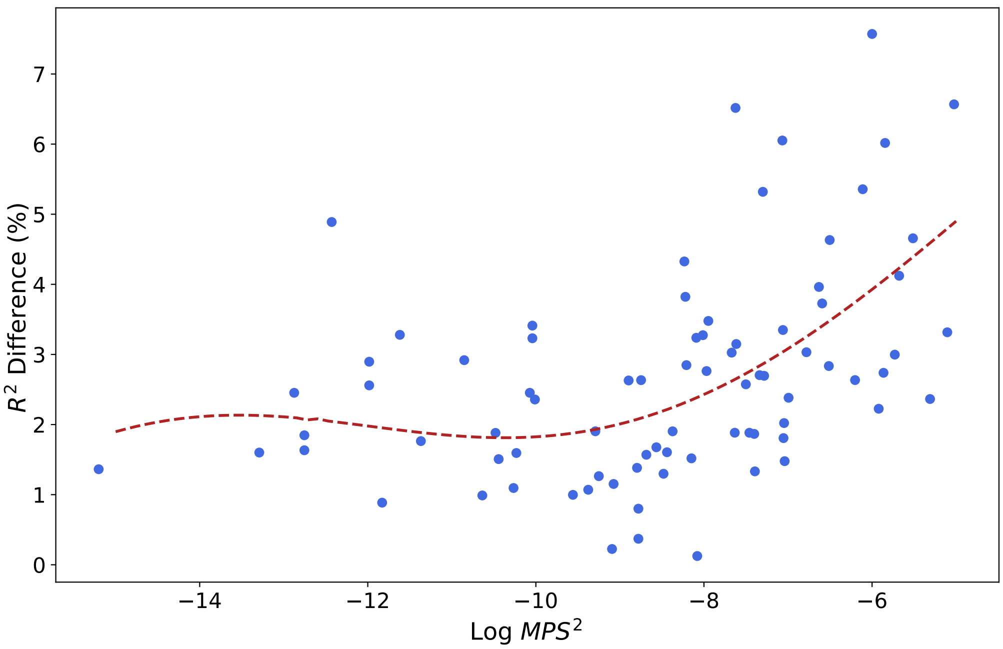

surprises (to the right-hand side of the plot). The red dashed line is a fitted local polynomial, indicating an approximately flat relationship for surprises below the median surprise and a clear positive relationship for surprises above the median. This represents further evidence that larger information flows cause changes in the factor structure of asset returns, which in turn lead to a greater improvement from using Intraday PCA relative to the standard PCA alternative.

4.3 Other macroeconomic announcements

While Federal Open Market Committee (FOMC) announcements are the most impactful announcements and are widely watched by investors and market analysts, other sources of macroeconomic information also impact asset prices. We next consider four other macroeconomic announcements: non-farm payrolls (typically released at 8:30am on the first Friday of each month); the consumer price index or CPI (8:30am in the second week of each month); the Conference Board's consumer confidence index (10am on the last Tuesday of each month); and the ISM manufacturing report (10am on the first business day of each month). The times of these announcements are noteworthy as two of them occur in the overnight period (similar to earnings announcements), and two are within the trade day (similar to FOMC announcements).

Table 5 reveals that for all four types of announcements Intraday PCA yields better forecasts than the competing methods, and for three out of four announcement types the improvement is statistically significant at the 5% level. Only for the consumer confidence announcement is the out-of-sample improvement not significantly different from zero. (The in-sample gains in R^2 are significant at the 1% level for all four types of announcements.) It is noteworthy that the OOS R^2 improvement offered by Intraday PCA is greater for all four types of announcements than it is in the full sample (see Table 2): in the full sample the improvement over standard PCA is 1.3%, while for these four announcements it ranges from 2.7% for nonfarm payroll announcements to 4.9% for CPI announcements. Thus, consistent with the model in Section 3, the gains from allowing for intradaily variation in the factor structure are greater following times of information flows. For these four macroeconomic announcements, the gains are between two and nearly four times as great on macroeconomic announcement days than on average.

5 Extensions and robustness checks

5.1 Subsample analysis

This section investigates whether our findings of a time-varying factor structure within equity returns varies over our sample period. We repeat our full sample analyses, reported in Table 2, separately for the first and second halves of our data (1998-2009, and 2010-2020) and present the results in Table 6. A first finding from Table 6 is that the factor structure in asset returns, however modeled, is stronger in the second period than the first: the OOS R^2s are higher for each of the five models, and the cross-model average R^2 rises from 18.4% in the first subsample to 27.6% in

Table 5: **In-sample and out-of-sample performance of factor models on other macroeconomic announcements days.** The table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) using macroeconomic announcement days including non-farm payrolls, CPI, consumer confidence and ISM Manufacturing releases, described in Section 4.2, for three models based on observable factors and two models based on latent factors. The CAPM uses solely the market factor; FF3 additionally uses the size and value factors of Fama and French (1992); FF6 additionally uses the momentum factor of Carhart (1997) and the quality and profitability factors of Fama and French (2015). The column titled PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns. The last column uses the kernel-based local PCA introduced in Section 2.2. t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of Intraday PCA are reported in parentheses.

	Factor model				
	CAPM	FF3	FF6	PCA	Intraday PCA
Panel A: Non-farm payrolls (8:30am)					
In-sample R^2	17.85	20.64	21.31	19.58	25.67
(RV t-stat)	(19.69)	(16.80)	(15.06)	(18.46)	
OOS R^2	24.72	27.93	27.90	26.81	29.46
(DM t-stat)	(3.95)	(2.63)	(2.07)	(3.48)	
Panel B: CPI (8:30am)					
In-sample R^2	17.98	21.29	22.27	22.19	29.31
(RV t-stat)	(16.01)	(14.25)	(13.18)	(15.45)	
OOS R^2	22.61	26.15	26.28	25.50	30.39
(DM t-stat)	(2.62)	(2.07)	(1.85)	(2.01)	
Panel C: Consumer confidence (10am)					
In-sample R^2	15.88	18.57	19.30	17.61	25.52
(RV t-stat)	(18.77)	(16.59)	(15.22)	(18.45)	
OOS R^2	20.09	22.92	22.78	22.12	24.96
(DM t-stat)	(1.08)	(1.06)	(1.06)	(1.07)	
Panel D: ISM Manufacturing (10am)					
In-sample R^2	15.67	18.34	19.10	17.36	23.90
(RV t-stat)	(21.26)	(18.78)	(16.92)	(20.51)	
OOS R^2	21.66	24.94	25.16	23.87	26.99
(DM t-stat)	(3.34)	(1.66)	(1.78)	(2.05)	

Table 6: **In-sample and out-of-sample performance of factor models.** The table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) for three models based on observable factors and two models based on latent factors. The CAPM uses solely the market factor; FF3 additionally uses the size and value factors of Fama and French (1992); FF6 additionally uses the momentum factor of Carhart (1997) and the quality and profitability factors of Fama and French (2015). The column titled PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns. The last column uses the kernel-based local PCA introduced in Section 2.2. t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of Intraday PCA are reported in parentheses. Panels A and B present results analogous to the full-sample results in Table 2, for the first and second halves of the sample period.

	Factor model				
	CAPM	FF3	FF6	PCA	Intraday PCA
Panel A: First subsample (1998-2009)					
In-sample R^2	17.67	19.85	21.35	20.68	22.11
(RV t-stat)	(16.79)	(11.69)	(5.23)	(18.17)	
OOS R^2	16.43	18.12	18.40	18.96	20.07
(DM t-stat)	(15.55)	(10.02)	(8.70)	(17.78)	
Panel B: Second subsample (2010-2020)					
In-sample R^2	23.96	28.97	31.49	33.14	35.43
(RV t-stat)	(16.03)	(12.94)	(10.83)	(13.13)	
OOS R^2	22.55	26.14	26.24	30.68	32.31
(DM t-stat)	(14.20)	(11.41)	(12.22)	(12.61)	

the second. In both subsamples we find strong evidence of a time-varying factor structure, with Intraday PCA significantly outperforming all competing models, including standard PCA.[15] The outperformance of Intraday PCA is slightly greater in the latter subsample (e.g., the improvement in OOS R^2 relative to standard PCA is 1.1% in the first subsample and 1.6% in the second.) The Diebold-Mariano t-statistic comparing Intraday PCA and standard PCA is greater than 18 in both subsamples, confirming our results from the full sample.

[15]The test of Liao and Todorov (2024), that the span of the latent factors is the same in the opening two hours of trade as in the last two hours, also rejects the null in both subsamples, at the 1% level.

Figure 9: **Comparing the fit of Intraday PCA and PCA, first and second subsamples.** This figure shows the difference in out-of-sample R^2 between Intraday PCA, introduced in Section 2.2, and standard PCA, for each of the 26 15-minute periods during the New York Stock Exchange trade day. Positive values indicate that Intraday PCA has a higher R^2 than standard PCA. The vertical lines at the end of each bar are 95% confidence intervals for the difference in R^2 based on a Diebold-Mariano (1995) test. Panels (a) and (b) present the results for the first (1998-2009) and second (2010-2020) subsamples.

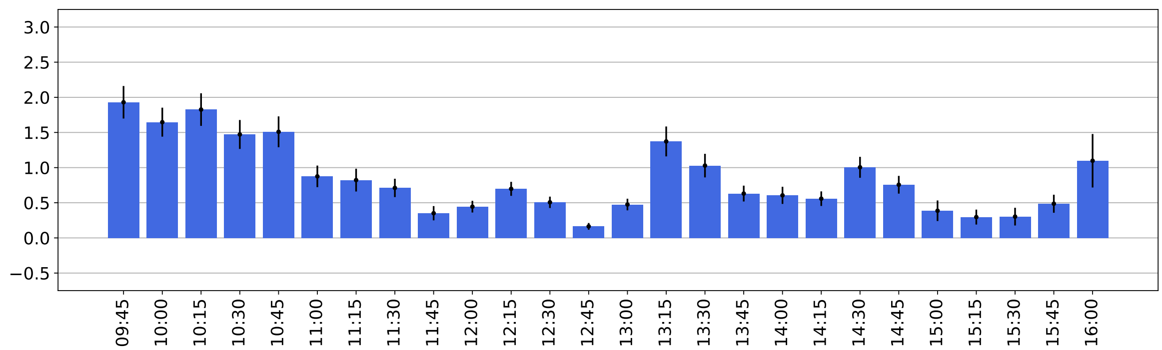

(a)

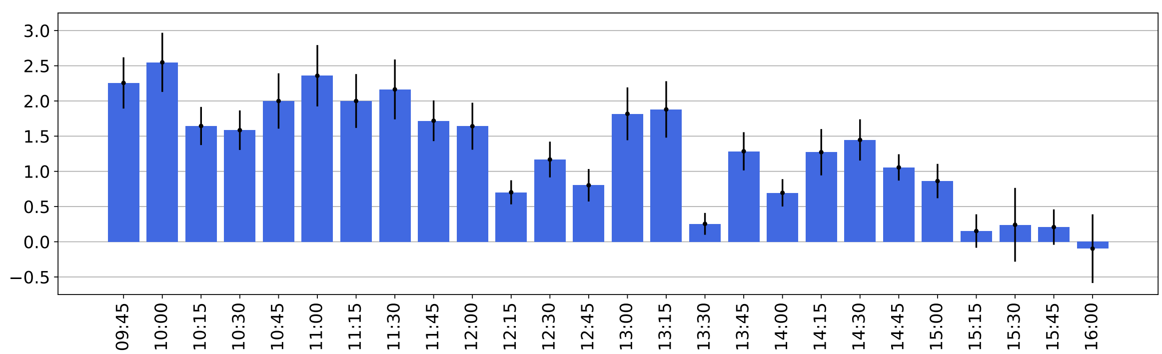

(b)

In Figure 9 we show how the improvement offered by Intraday PCA over standard PCA varies over the trade day for each of the two subsamples. The pattern is similar in both subsamples, and similar to the full-sample pattern presented in Figure 3: the gains are greatest in the first trading periods of the day, and diminish towards the end of the trade day. Consistent with the results in Table 6, we observe that the gains are slightly greater on average in the latter subsample. Comparing the top and bottom panels of Figure 9 we can observe a difference at the end of the trade day: in the first subsample Intraday PCA offered a significant improvement over standard PCA in the last hour of the trade day, while in the latter subsample the improvement is small and not significantly different from zero.

Overall, we conclude that our finding of significant variation in the factor structure of asset returns is not sensitive to the choice of sample period, with both halves of our sample yielding qualitatively similar results.

5.2 Choosing the number of factors

To select the number of factors in the standard PCA and Intraday PCA models in our main analysis we adopt the information criterion proposed in Bai and Ng (2002) and extended to high-frequency setting by Liao and Todorov (2024), see equation (7). In this section we consider the estimator for the number of factors in a high frequency factor setting proposed by Pelger (2019, Theorem 6). This estimator examines ratios of adjacent (ordered) eigenvalues, similar to the method of Ahn and Horenstein (2013). We follow the recommendation of Pelger (2019) and use median of eigenvalues as the "perturbation" and 0.2 as the "cutoff" hyperparameters for this estimator.

Table 7 presents results for the full sample and two subsamples, when the PCA-based methods choose the number of factors using Pelger (2019). The results in this table can be compared with those in Tables 2 and 6. The observable factor models (CAPM, FF3 and FF6) are unchanged in this analysis, so their in-sample and OOS R^2s are identical to the earlier tables, however their comparisons with Intraday PCA differ as the latter model is different in this analysis.

The performance of standard PCA and Intraday PCA when the number of factors are chosen

Table 7: **Performance of models with number of factors chosen using Pelger (2019).** The table reports measures of in-sample fit and out-of-sample accuracy (R^2 metrics defined in equations (9) and (11), measured in percent) for three models based on observable factors and two models based on latent factors. The CAPM uses solely the market factor; FF3 additionally uses the size and value factors of Fama and French (1992); FF6 additionally uses the momentum factor of Carhart (1997) and the quality and profitability factors of Fama and French (2015). The column titled PCA uses standard principal components analysis on the $TM \times N$ panel of high frequency stock returns. The last column uses the kernel-based local PCA introduced in Section 2.2. The latter two models set the number of factors using Pelger (2019). t-statistics from Rivers-Vuong (2002) and Diebold-Mariano (1995) tests comparing the fit of each model to that of intraday PCA are reported in parentheses. Panel A presents results analogous to the full-sample results in Table 2, and panels B and C present results for the first and second halves of the sample period, analogous to Table 6.

	Factor model				
	CAPM	FF3	FF6	PCA	Intraday PCA
Panel A: Full sample (1998-2020)					
In-sample R^2	19.56	22.60	24.40	24.18	25.55
(RV t-stat)	(18.91)	(14.15)	(8.11)	(18.25)	
OOS R^2	18.27	20.53	20.77	22.29	23.19
(DM t-stat)	(16.65)	(11.48)	(11.20)	(16.99)	
Panel B: First subsample (1998-2009)					
In-sample R^2	17.67	19.85	21.35	20.62	21.74
(RV t-stat)	(17.24)	(11.54)	(3.31)	(16.65)	
OOS R^2	16.43	18.12	18.40	18.94	19.74
(DM t-stat)	(15.93)	(9.45)	(7.82)	(14.20)	
Panel C: Second subsample (2010-2020)					
In-sample R^2	23.96	28.97	31.49	32.42	34.41
(RV t-stat)	(16.39)	(12.97)	(10.10)	(12.17)	
OOS R^2	22.55	26.14	26.24	30.06	31.17
(DM t-stat)	(13.96)	(10.61)	(11.33)	(10.72)	

using Pelger (2019) rather than Bai and Ng (2002) are very similar: in the full sample the in-sample and OOS R^2s fall by between 0.2% and 0.6%, but the outperformance of Intraday

PCA relative to the competing models remains significant at the 1% level. This is true also for both subsamples: the R^2 values decline slightly, but the improvement of Intraday PCA over the competing models remains significant at the 1% level.

6 Conclusion

Motivated by recent work on within-day variation in CAPM factor loadings (Andersen et al., 2021, 2023) we propose using a combination of kernel-based methods and principal components analysis (PCA), known as "local PCA" (Su and Wang, 2017; Fan et al., 2022; Pelger and Xiong, 2022) to allow the latent factor structure of asset returns to vary over the trade day. This approach nests both standard PCA, which imposes no variation in the factor structure across the trade day, and PCA applied to each high-frequency period separately, which ignores the fact that adjacent high-frequency periods are likely to be similar. By optimally choosing the degree of flexibility in the model, which we accomplish via cross-validation, we obtain a model that exploits information from neighboring high-frequency periods, but allows the factor structure to vary over the trade day.

We apply our "Intraday PCA" to 15-minute returns on a collection of over 400 individual U.S. stock returns over the period 1996-2020. We find comprehensive evidence that the factor model of these returns varies over the trade day, and the improvement in explanatory power of Intraday PCA over standard PCA is significant at the 1% level. The improvement in fit compared with the famous Fama-French (1992) model (FF3), which is also significant at the 1% level, is larger than the improvement from using the FF3 compared with the workhorse capital asset pricing model (CAPM). Our analysis reveals that the gains from allowing for a time-varying factor structure mostly accrue in the opening hours of the trade day, when information accumulated in the overnight period (e.g., earnings announcements and pre-market macroeconomic announcements) is incorporated into asset prices.

We next propose a simple model of high frequency asset returns and information flows, and demonstrate that information flows can lead to changes in the factor structure of asset returns. The key mechanism in the model is that if prices are affected by different types of news (e.g., cashflow and discount rate news), arriving at different times, and if the different news types have different factor structures, then the factor structure of asset returns will naturally vary over the trade day. The model additionally predicts that changes in the factor structure will be greatest when the surprise component of the news is larger.

We test the predictions of the model in three distinct empirical analyses. Firstly, we sort days according to the surprise component of earnings announcements made during the previous overnight period. We find significant evidence that the changes in asset return factor structure are larger when earnings announcements are more surprising. Secondly, we examine Federal Open Market Committee (FOMC) announcements, which occur during the trade day rather than in the overnight period. Focusing only on FOMC days we see that the gains from allowing for a varying factor structure are greatest in ther periods immediately after 2pm, the typical FOMC announcement time. We use the "monetary policy surprise" measure of Bauer and Swanson (2023) to quantify the surprise component of a given FOMC announcement and find that the gains from the Intraday PCA model are larger for more surprising announcements. Finally, we consider four other important macroeconomic announcements and find the gains from allowing for intradaily variation in the factor structure are between two and nearly four times as great on announcement days than on average. All three analyses support our hypothesis that information flows lead to changes in the factor structure of asset returns.

References

Ahn, S. C. and Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, 81(3):1203–1227.

Aït-Sahalia, Y. and Xiu, D. (2019). Principal component analysis of high-frequency data. *Journal of the American Statistical Association*, 114(525):287–303.

Andersen, T. G. and Bollerslev, T. (1998). Deutsche mark–dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, 53(1):219–265.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Wu, J. (2005). A framework for exploring the macroeconomic determinants of systematic risk. *American Economic Review*, 95(2):398–404.

Andersen, T. G., Riva, R., Thyrsgaard, M., and Todorov, V. (2023). Intraday cross-sectional distributions of systematic risk. *Journal of Econometrics*, 235(2):1394–1418.

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

Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, 71(1):135–171.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2008). Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536.

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

Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155.

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

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

Boguth, O., Carlson, M., Fisher, A., and Simutin, M. (2016). Horizon effects in average returns: The role of slow information diffusion. *Review of Financial Studies*, 29(8):2241–2281.

Bollerslev, T., Engle, R. F., and Nelson, D. B. (1994). ARCH models. *Handbook of Econometrics, Vol. 4*, pages 2959–3038.

Bollerslev, T., Li, J., and Xue, Y. (2018). Volume, volatility, and public news announcements. *Review of Economic Studies*, 85(4):2005–2041.

Bollerslev, T., Li, S. Z., and Todorov, V. (2016). Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns. *Journal of Financial Economics*, 120(3):464–490.

Brennan, M. J., Jegadeesh, N., and Swaminathan, B. (1993). Investment analysis and the adjustment of stock prices to common information. *Review of Financial Studies*, 6(4):799–824.

Campbell, J. Y. and Vuolteenaho, T. (2004). Bad beta, good beta. *American Economic Review*, 94(5):1249–1275.

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

Chamberlain, G. and Rothschild, M. (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, pages 1281–1304.

Cieslak, A. and Schrimpf, A. (2019). Non-monetary news in central bank communication. *Journal of International Economics*, 118:293–315.

Cochrane, J. H. (2011). Discount rates. *Journal of Finance*, 66(4):1047–1108.

Collins, D. W. and Kothari, S. P. (1989). An analysis of intertemporal and cross-sectional determinants of earnings response coefficients. *Journal of Accounting and Economics*, 11(2-3):143–181.

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

Connor, G. and Korajczyk, R. A. (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics*, 15(3):373–394.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253.

Engle, R. (2009). *Anticipating correlations: A new paradigm for risk management*. Princeton University Press.

Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2):427–465.

Fama, E. F. and French, K. R. (2004). The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*, 18(3):25–46.

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

Fan, J., Ke, Z., Liao, Y., and Neuhierl, A. (2022). Structural deep learning in conditional asset pricing. *Available at SSRN 4117882*.

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

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

Giglio, S. and Xiu, D. (2021). Asset pricing with omitted factors. *Journal of Political Economy*, 129(7):1947–1990.

Gu, S., Kelly, B., and Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, 222(1):429–450.

Harris, L. (1986). A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 16(1):99–117.

Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H. (2009). *The elements of statistical learning: Data mining, inference, and prediction*, volume 2. Springer.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of finance*, 54(6):2143–2184.

Jagannathan, R. and Wang, Z. (1996). The conditional CAPM and the cross-section of expected returns. *Journal of Finance*, 51(1):3–53.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524.

Kormendi, R. and Lipe, R. (1987). Earnings innovations, earnings persistence, and stock returns. *Journal of Business*, pages 323–345.

Kothari, S. P. (2001). Capital markets research in accounting. *Journal of Accounting and Economics*, 31(1-3):105–231.

Lettau, M. and Pelger, M. (2020a). Estimating latent asset-pricing factors. *Journal of Econometrics*, 218(1):1–31.

Lettau, M. and Pelger, M. (2020b). Factors that fit the time series and cross-section of stock returns. *Review of Financial Studies*, 33(5):2274–2325.

Liao, Y. and Todorov, V. (2024). Changes in the span of systematic risk exposures. *Quantitative Economics*, forthcoming.

Lintner, J. (1965). Security prices, risk, and maximal gains from diversification. *Journal of Finance*, 20(4):587–615.

Livnat, J. and Mendenhall, R. R. (2006). Comparing the post–earnings announcement drift for surprises calculated from analyst and time series forecasts. *Journal of Accounting Research*, 44(1):177–205.

Lou, D., Polk, C., and Skouras, S. (2019). A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134(1):192–213.

Patton, A. J. and Verardo, M. (2012). Does beta move with news? firm-specific information flows and learning about profitability. *Review of Financial Studies*, 25(9):2789–2839.

Pelger, M. (2019). Large-dimensional factor modeling based on high-frequency observations. *Journal of Econometrics*, 208(1):23–42.

Pelger, M. (2020). Understanding systematic risk: A high-frequency approach. *Journal of Finance*, 75(4):2179–2220.

Pelger, M. and Xiong, R. (2022). State-varying factor models of large dimensions. *Journal of Business & Economic Statistics*, 40(3):1315–1333.

Rivers, D. and Vuong, Q. (2002). Model selection tests for nonlinear dynamic models. *Econometrics Journal*, 5(1):1–39.

Ross, S. (1976). The arbitrage pricing theory. *Journal of Economic Theory*, 13(3):341–360.

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

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442.

Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., Jordan, M. I., and Sastry, S. S. (2004). Kalman filtering with intermittent observations. *IEEE transactions on Automatic Control*, 49(9):1453–1464.

Su, L. and Wang, X. (2017). On time-varying factor models: Estimation and testing. *Journal of Econometrics*, 198(1):84–101.

Vuolteenaho, T. (2002). What drives firm-level stock returns? *Journal of Finance*, 57(1):233–264.

Watts, R. L. and Leftwich, R. W. (1977). The time series of annual accounting earnings. *Journal of Accounting Research*, pages 253–271.

Wood, R. A., McInish, T. H., and Ord, J. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, 40(3):723–739.

Appendix A An approximate leave-one-out R^2 metric

Recall that we estimate the factor loadings, $\hat{\boldsymbol{\beta}}$, in the estimation sample using the eigenvectors associated with the K largest eigenvalues of $\mathbf{R}^\top\mathbf{R}$, the second moment matrix of returns in the estimation sample. With those, we obtain the factor realizations by regressing returns on betas:

$$\underset{(K\times 1)}{\hat{\mathbf{f}}_t} = (\underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\underset{(N\times K)}{\hat{\boldsymbol{\beta}}})^{-1}\underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\underset{(N\times 1)}{\mathbf{R}_t} = \underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\underset{(N\times 1)}{\mathbf{R}_t} \tag{A.1}$$

where the last equality holds as we impose the normalization that $\hat{\boldsymbol{\beta}}^\top\hat{\boldsymbol{\beta}} = I_K$. From those we obtain the (infeasible) forecasts, $\hat{r}_{i,t}$, from the factor model:

$$\hat{r}_{i,t} = \hat{\boldsymbol{\beta}}_i^\top\hat{\mathbf{f}}_t \tag{A.2}$$

These are not true forecasts because $\hat{\mathbf{f}}_t$ uses, in part, the realized value of the target variable, $r_{i,t}$. One might instead consider using the estimation-sample mean of the factors, $\hat{\boldsymbol{\lambda}}$, rather than their out-of-sample realizations, leading to the forecast:

$$\bar{r}_{i,t} = \hat{\boldsymbol{\beta}}_i^\top\hat{\boldsymbol{\lambda}} \tag{A.3}$$

The "predictive" and "total" R^2 measures of Kelly et al. (2019) are based on these two forecasts:

$$R^2_{\text{total}} = 1 - \frac{\sum_i\sum_{(t,\tau)\in\mathcal{T}}(r_{i,t-1+\tau} - \hat{r}_{i,t-1+\tau})^2}{\sum_i\sum_{(t,\tau)\in\mathcal{T}} r^2_{i,t-1+\tau}} \tag{A.4}$$

$$R^2_{\text{pred}} = 1 - \frac{\sum_i\sum_{(t,\tau)\in\mathcal{T}}(r_{i,t-1+\tau} - \bar{r}_{i,t-1+\tau})^2}{\sum_i\sum_{(t,\tau)\in\mathcal{T}} r^2_{i,t-1+\tau}} \tag{A.5}$$

where $\mathcal{T}$ denotes the collection of days and intradaily observations in the out-of-sample period.

Note that the infeasible forecasts in equation (A.2) can be written as:

$$\hat{r}_{i,t} = \underset{(1\times K)}{\hat{\boldsymbol{\beta}}_i^\top}\underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\underset{(N\times 1)}{\mathbf{R}_t} = \hat{\boldsymbol{\beta}}_i^\top\sum_{j=1}^{N}\hat{\boldsymbol{\beta}}_j r_{j,t} = \hat{\boldsymbol{\beta}}_i^\top\hat{\boldsymbol{\beta}}_i r_{i,t} + \hat{\boldsymbol{\beta}}_i^\top\sum_{j\neq i}\hat{\boldsymbol{\beta}}_j r_{j,t} \tag{A.6}$$

Consider a forecast that replaces $r_{i,t}$ on the right-hand side with its in-sample mean, $\hat{\mu}_i$:

$$\tilde{r}_{i,t} = \hat{\boldsymbol{\beta}}_i^\top \hat{\boldsymbol{\beta}}_i \hat{\mu}_i + \hat{\boldsymbol{\beta}}_i^\top \sum_{j \neq i} \hat{\boldsymbol{\beta}}_j r_{j,t} \tag{A.7}$$

This is a proper forecast of $r_{i,t}$, being based on estimation sample estimates of factor loadings, and contemporaneous values only of the *other* variables in the panel, i.e. $r_{j,t}$ for $j \neq i$. This is an "approximate leave-one-out" (ALOO) forecast of variable i. It is particularly useful in applications with a large cross-sectional dimension, where a true leave-one-out forecast is computationally expensive, and in local PCA, where the bandwidth parameter must be tuned using a validation sample, increasing the computational burden. Below we show that one can obtain the ALOO forecast as a simple transformation of the original factors and infeasible forecasts.

It is convenient to write the above equation as:

$$\tilde{r}_{i,t} = \hat{r}_{i,t} - \hat{\boldsymbol{\beta}}_i^\top \hat{\boldsymbol{\beta}}_i \left(r_{i,t} - \hat{\mu}_i \right) \tag{A.8}$$

which is presented as equation (10) in the main paper. In matrix form, we then obtain the factors, infeasible forecasts, and ALOO forecasts:

$$\underset{(T\times K)}{\hat{\mathbf{F}}} = \underset{(T\times N)}{\mathbf{R}} \underset{(N\times K)}{\hat{\boldsymbol{\beta}}} \tag{A.9}$$

$$\underset{(T\times K)}{\hat{\mathbf{R}}} = \underset{(T\times K)}{\hat{\mathbf{F}}} \underset{(K\times N)}{\hat{\boldsymbol{\beta}}} = \underset{(T\times N)}{\mathbf{R}} \underset{(N\times K)}{\hat{\boldsymbol{\beta}}} \underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top} \tag{A.10}$$

$$\underset{(T\times N)}{\tilde{\mathbf{R}}} = \underset{(T\times N)}{\hat{\mathbf{R}}} - \Big(\underset{(T\times N)}{\mathbf{R}} - \underset{(T\times 1)}{\boldsymbol{\iota}} \underset{(1\times N)}{\hat{\boldsymbol{\mu}}^\top} \Big) \cdot \underbrace{\mathrm{diag}\{ \underset{(N\times K)}{\hat{\boldsymbol{\beta}}} \underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top} \}}_{(N\times N)} \tag{A.11}$$

where $\boldsymbol{\iota}$ is a vector of ones, and diag$\{A\}$ sets the off-diagonal elements of the matrix A to zero. Thus the ALOO forecasts, $\tilde{\mathbf{R}}$, can be obtained in closed form from the usual forecast, $\hat{\mathbf{R}}$, the matrix of factor loadings, $\hat{\boldsymbol{\beta}}$, and the vector of in-sample mean returns, $\hat{\boldsymbol{\mu}}$. No new estimation is required. The "approximate leave-one-out" R^2 metric is then computed as in equation (11) of the main paper.

For PCA applied to standardized data, as in our empirical application, a simple modification of the above derivations is required. Denote the standardized returns as:

$$
\begin{aligned}
x_{i,t} &= \left(r_{i,t} - \hat{\mu}_i\right) / \hat{\sigma}_i && \text{(A.12)} \\
\text{so} \quad \underset{(T\times N)}{\mathbf{X}} &= \Big(\underset{(T\times N)}{R} - \underset{(T\times 1)}{\boldsymbol{\iota}}\,\underset{(1\times N)}{\hat{\boldsymbol{\mu}}^{\top}}\Big)\,\underset{(N\times N)}{\hat{\mathbf{D}}^{-1}} && \text{(A.13)}
\end{aligned}
$$

where $\hat{\boldsymbol{\mu}}$ is the vector of estimation-sample means, and $\hat{\mathbf{D}}$ is a diagonal matrix of the estimation-sample standard deviations.

We estimate the factor loadings, $\hat{\beta}_i$, in the estimation sample using the eigenvectors associated with the K largest eigenvalues of $\mathbf{X}^{\top}\mathbf{X}$, again imposing the normalization that $\hat{\boldsymbol{\beta}}^{\top}\hat{\boldsymbol{\beta}} = I_K$. With those, we obtain the factor realizations:

$$
\underset{(K\times 1)}{\hat{\mathbf{f}}_t} = \underset{(K\times N)}{\hat{\boldsymbol{\beta}}^{\top}}\,\underset{(N\times 1)}{\mathbf{X}_t} \tag{A.14}
$$

and from those we obtain the forecasts from the factor model for the standardized returns and the original returns:

$$
\hat{r}_{i,t} = \hat{\mu}_i + \hat{\sigma}_i \hat{x}_{i,t} = \hat{\mu}_i + \hat{\sigma}_i \hat{\boldsymbol{\beta}}_i^{\top}\hat{\mathbf{f}}_t \tag{A.15}
$$

Analogous to equation (A.6), these infeasible forecasts can be written as:

$$
\hat{r}_{i,t} = \hat{\mu}_i + \hat{\sigma}_i \hat{\boldsymbol{\beta}}_i^{\top}\hat{\boldsymbol{\beta}}_i x_{it} + \hat{\sigma}_i \hat{\boldsymbol{\beta}}_i^{\top}\sum_{j\neq i}\hat{\boldsymbol{\beta}}_j x_{jt} \tag{A.16}
$$

We replace the realized value of $x_{i,t}$ with its estimation sample mean, which is zero by construction. Analogous to equations (A.7) and (A.8), we obtain the ALOO forecasts:

$$
\begin{aligned}
\tilde{r}_{i,t} &= \hat{\mu}_i + \hat{\sigma}_i \hat{\boldsymbol{\beta}}_i^{\top}\sum_{j\neq i}\hat{\boldsymbol{\beta}}_j x_{j,t} && \text{(A.17)} \\
&= \hat{r}_{i,t} - \hat{\sigma}_i \hat{\boldsymbol{\beta}}_i^{\top}\hat{\boldsymbol{\beta}}_i x_{i,t} && \text{(A.18)}
\end{aligned}
$$

Finally, we have the factors, original forecasts, and ALOO forecasts in matrix form:

$$\underset{(T\times K)}{\hat{\mathbf{F}}} = \underset{(T\times N)}{\mathbf{X}}\underset{(N\times K)}{\hat{\boldsymbol{\beta}}} \tag{A.19}$$

$$\underset{(T\times N)}{\hat{\mathbf{X}}} = \underset{(T\times K)}{\hat{\mathbf{F}}}\underset{(K\times N)}{\hat{\boldsymbol{\beta}}} \tag{A.20}$$

$$\underset{(T\times N)}{\hat{\mathbf{R}}} = \underset{(T\times 1)}{\boldsymbol{\iota}}\underset{(1\times N)}{\hat{\boldsymbol{\mu}}^\top} + \underset{(T\times N)}{\hat{\mathbf{X}}}\underset{(N\times N)}{\hat{\mathbf{D}}} \tag{A.21}$$

$$\underset{(T\times N)}{\tilde{\mathbf{X}}} = \underset{(T\times N)}{\hat{\mathbf{X}}} - \underset{(T\times N)}{\mathbf{X}} \cdot \underbrace{\operatorname{diag}\{\underset{(N\times K)}{\hat{\boldsymbol{\beta}}}\underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\}}_{(N\times N)} \tag{A.22}$$

$$\underset{(T\times N)}{\tilde{\mathbf{R}}} = \underset{(T\times N)}{\hat{\mathbf{R}}} - \underset{(T\times N)}{\mathbf{X}} \cdot \underbrace{\operatorname{diag}\{\underset{(N\times K)}{\hat{\boldsymbol{\beta}}}\underset{(K\times N)}{\hat{\boldsymbol{\beta}}^\top}\}}_{(N\times N)}\underset{(N\times N)}{\hat{\mathbf{D}}} \tag{A.23}$$

The ALOO forecasts, $\tilde{\mathbf{R}}$, can again be obtained in closed form from the usual forecast, $\hat{\mathbf{R}}$, the standardized data, $\mathbf{X}$, the matrix of factor loadings, $\hat{\boldsymbol{\beta}}$, and the diagonal matrix of standard deviations of the original data, $\hat{\mathbf{D}}$. The approximate leave-one-out R^2 is then computed as in equation (11) using the ALOO forecasts from equation (A.23).

Appendix B Details on the model in Section 3

Appendix B.1 Details on the Kalman filter

The underlying variable of interest in this problem is the $(N \times 1)$ vector of daily earnings growth, $\Delta \log \mathbf{X}_t$, for each of the firms (see equation 12). For reasons that become clear below, we define the state variable for this filtering problem, $\boldsymbol{\xi}_t$, as the most recent C values of this variable:

$$
\begin{aligned}
\Delta \log \mathbf{X}_t &= \mathbf{g} + \boldsymbol{\gamma} \mathbf{Z_t} + \mathbf{u}_t && \text{(A.24)} \\
\boldsymbol{\xi}_t &\equiv \operatorname{vec} \begin{bmatrix} \Delta \log \mathbf{X}_t^\top \\ \vdots \\ \Delta \log \mathbf{X}_{t-C+1}^\top \end{bmatrix} = F \boldsymbol{\xi}_{t-1} + \boldsymbol{\nu}_t && \text{(A.25)} \\
\text{where} \qquad F &= I_N \otimes F_1 && \text{(A.26)} \\
V[\boldsymbol{\nu}_t] &= Q_1 \otimes Q_2 && \text{(A.27)} \\
\text{and} \qquad F_1 &= \begin{bmatrix} \mathbf{0}_C^\top \\ I_{C-1}, \mathbf{0}_{C-1} \end{bmatrix} && \text{(A.28)} \\
Q_1 &= \boldsymbol{\gamma} \Sigma_Z \boldsymbol{\gamma}^\top + \sigma_u^2 I_N && \text{(A.29)} \\
Q_2 &= \mathbf{e}_1 \mathbf{e}_1^\top && \text{(A.30)}
\end{aligned}
$$

where I_N is the N-dimensional identity matrix, $\otimes$ is the Kronecker product, $\mathbf{0}_C$ is a $(C \times 1)$ vector of zeros, and $\mathbf{e}_1$ is a $(C \times 1)$ vector with a 1 in the first element and zeros elsewhere.

The measurement variable for this problem is complicated by the fact that earnings announcements only occur every $C > 1$ days. We follow Patton and Verardo (2012) and use the framework proposed by Sinopoli et al. (2004) for updating expectations when the signal is only observed intermittently. We proceed by defining a measurement variable that tracks, with error, earnings growth over the last C days:

$$
Y_{i,t} = \sum_{j=0}^{C-1} \Delta \log X_{i,t-j} + \eta_{i,t} \tag{A.31}
$$

The measurement error, $\eta_{i,t}$ is introduced to "mask" this signal in between earnings announcement dates. Specifically, on earnings announcment dates the variance of the measurement error is set to zero, and the earnings growth since the last announcement is perfectly observed. On other dates, the variance of this variable is set to a large value:

$$V[\eta_{i,t}] = \sigma_I^2(1 - A_{i,t}) \tag{A.32}$$

where $A_{i,t} = 1$ if day t is an announcement date for firm i, and zero otherwise, and $\sigma_I \to \infty$ is the variance of the measurement error on non-announcement days.

The measurement variable

$$\mathbf{Y}_t \equiv [Y_{1,t}, ..., Y_{N,t}]^\top \tag{A.33}$$

is linked to the state variable via:

$$\begin{aligned} \mathbf{Y}_t &= H^\top \boldsymbol{\xi}_t + \boldsymbol{\eta}_t && \text{(A.34)} \\ \text{where} \quad H^\top &= I_N \otimes \boldsymbol{\iota_C} && \text{(A.35)} \\ \text{and} \quad V[\boldsymbol{\eta}_t] &= \sigma_I(I_N - \text{diag}\{\mathbf{A}_t\}) && \text{(A.36)} \end{aligned}$$

where ι_C is a $(C \times 1)$ vector of ones, $\mathbf{A}_t$ is a $(N \times 1)$ vector of announcement indicators for each of the firms, and diag$\{\mathbf{a}\}$ is a diagonal matrix with the vector $\mathbf{a}$ on the diagonal.

With the above structure for the state variable, $\boldsymbol{\xi}_t$ and its connection to the measurement variable, $\mathbf{Y}_t$, standard Kalman filtering computations can be applied. The one-step-ahead forecasts and the nowcasts of the state variable can be obtained as follows:

$$\begin{aligned} \hat{\boldsymbol{\xi}}_{t|t-1} &= F\hat{\boldsymbol{\xi}}_{t-1|t-1} && \text{(A.37)} \\ \hat{\boldsymbol{\xi}}_{t|t} &= \hat{\boldsymbol{\xi}}_{t|t-1} + P_{t|t-1}H^\top \left(HP_{t|t-1}H^\top + \sigma_I(I_N - \text{diag}\{\mathbf{A}_t\})\right)^{-1} \left(\mathbf{Y}_t - H^\top\hat{\boldsymbol{\xi}}_{t|t-1}\right) && \text{(A.38)} \\ P_{t|t-1} &= FP_{t-1|t-1}F^\top + Q && \text{(A.39)} \\ P_{t|t} &= P_{t|t-1} - P_{t|t-1}H^\top \left(HP_{t|t-1}H^\top + \sigma_I^2(I_N - \text{diag}\{\mathbf{A}_t\})\right)^{-1} HP_{t|t-1} && \text{(A.40)} \end{aligned}$$

From equation (A.38) we can extract the nowcasts of the earnings growth of each firm between day $t-C+1$ and day t. By summing the growth rates since the last announcement day of a given firm (which can be no more than C days in the past) we obtain a nowcast of the log-level of the earnings of the firm. The first-difference of this nowcast is the "cashflow news" for this firm.

Appendix B.2 Parameter choices for the simulation

Parameter	Value	Description
T	3000	Number of trade days
M	21	Number of periods per day (including the overnight return)
N	100	Number of firms
C	10	Number of days between earnings announcements
D	5	Number of days between FOMC announcements
K	2	Number of common factors in earnings growth
π_{CF}	0.25	Diffusion of cashflow information
π_{DR}	0.25	Diffusion of discount rate information
$\sigma^2_{\varepsilon,O}$	0.25	Variance of noise trades as a proportion of cashflow news variance,
$\sigma^2_{\varepsilon,D}$	0.03	for the overnight period ($\sigma^2_{\varepsilon,O}$) and the intraday periods ($\sigma^2_{\varepsilon,D}$)
g	5/252	Growth rate of log-earnings
σ^2_u	0.1	Variance of unpredictable component of earnings growth
Σ_Z	$2I_N$	Covariance matrix of common factors in earnings growth
$\boldsymbol{\gamma}_1$	[2,3]	Coefficients on common earnings factors, for first N/2 firms
$\boldsymbol{\gamma}_2$	[2,-1]	Coefficients on common earnings factors, for second N/2 firms
σ^2_{rf}	0.25^2	Variance of the monetary policy surprise (MPS)
μ_{MPS}	-4	Mean (μ_{MPS}) and variance (σ^2_{MPS}) of the Normal distribution
σ^2_{MPS}	2	used to determine firm sensitivity to the MPS