The Review of Financial Studies

News and Asset Pricing: A High-Frequency Anatomy of the SDF

Saketh Aleti
Duke University, USA

Tim Bollerslev
Duke University, USA

Utilizing real-time newswire data, together with a robustly estimated intraday stochastic discount factor (SDF), we identify and quantify the economic news that is priced. News related to monetary policy and finance on average accounts for most of the variation in the SDF, followed by news about international affairs and macroeconomic data. We also document nontrivial temporal variation in the relative importance of the news, along with marked differences in the estimated news risk premiums in the "factor zoo." To further highlight the economic mechanisms at work, we associate the different news effects with interest rate, growth, and risk premium shocks. (*JEL* C58, G12, G14)

Received: February 1, 2023; Editorial decision: March 2, 2024
Editor: Itay Goldstein.
Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Asset pricing theory states that in a frictionless financial market a unique stochastic discount factor (SDF), or pricing kernel, exists that prices all risky assets (e.g. Back 2010). But what economic news drives large changes in the SDF, and how is specific news topics priced? We seek to provide new insights into these key questions by empirically linking large intraday changes or "jumps" in a robustly estimated, high-frequency SDF to directly observable real-time economic news and corresponding news topics. We find that news related to monetary policy and finance typically accounts for the largest portion of the variation in the SDF and the tangency portfolio risk premium, followed

We are grateful to Itay Goldstein (editor) and two anonymous referees for their most helpful and constructive comments. We would also like to thank Yacine Aït-Sahalia (discussant), Anna Cieslak, Jia Li, Andrew Patton, and George Tauchen and participants at the December 2022 EC2 Conference in Paris on "Econometrics of High-Frequency Data and Factor Models" and the January 2023 NBER conference in Boston on "Big Data and Securities Markets," along with the participants at various other seminar presentations, for their useful comments and suggestions. We also gratefully acknowledge financial support from the Carlsberg Foundation that helped with purchasing the Dow Jones Newswire data. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Saketh Aleti, saketh.aleti@duke.edu.

The Review of Financial Studies 38 (2025) 712–759

https://doi.org/10.1093/rfs/hhae019 Advance Access publication May 16, 2024

by news about international affairs and macroeconomic data. Building on the same mimicking portfolio approach underlying these results, we further document marked differences in the way in which the different news topics, and the compensation thereof, manifest in the "factor zoo." Further elaborating on the economic mechanisms at work, we find that the different news topics tend to affect different fundamental economic shocks. Perhaps surprisingly, we also find that a large fraction of the monetary policy news that is priced is not covered by the traditional monetary policy calendars.

Our estimation of the high-frequency SDF closely follows the minimax adversarial estimation approach recently advocated by Chen, Pelger, and Zhu (2024), including the use of neural networks for flexibly approximating the unknown functional form of the SDF. This approach is directly motivated by the seminal work of Hansen and Jagannathan (1997) and the result that the approximating SDF closest to the true SDF in a mean-square-error sense may be obtained by minimizing the largest possible squared pricing errors. In contrast to Chen, Pelger, and Zhu (2024), however, who base their estimation of a monthly SDF on monthly individual stock returns, we deliberately rely on a zoo of high-frequency factors recently constructed by Aleti (2024) for spanning the intraday SDF. Using factor portfolios instead of individual stocks allows us to succinctly incorporate the many pricing anomalies documented in the existing asset pricing literature, while simultaneously affording a greater degree of robustness to market microstructure noise in the high-frequency setting. It also heeds the concerns of Kozak and Nagel (2022), who explicitly caution against the use of too few factors to span the SDF.

The idea of associating volatility in the SDF with specific news topics is perhaps most closely related to the recent work by Bybee, Kelly, and Su (2023), who combine textual analysis of daily news articles with latent factor analysis to indirectly infer a set of systematic narrative news risk factors and a corresponding univariate pricing kernel. However, instead of first estimating different news factors, we directly link large changes, or "jumps," in our estimate of the high-frequency SDF to specific news topics.[1] Our work is also related to the emerging literature on "narrative asset pricing" that more generally includes Bybee et al. (Forthcoming), Ke, Kelly, and Xiu (2021), Pettenuzzo, Sabbatucci, and Timmermann (2020), and Mai and Pukthuanthong (2021), all of whom seek to infer cross-sectional and time-series predictability from the themes and sentiments in daily news articles.

The specific news topics that we rely on are based on the hand-constructed categorizations previously developed by Baker et al. (2019) and

[1] Our identification of the "jumps" in the estimated tangency portfolio formally relies on infill asymptotic arguments adapted from the high-frequency financial econometrics literature (see, e.g., Aït-Sahalia and Jacod 2014), together with the thresholding technique of Mancini (2001). Bajgrowicz, Scaillet, and Treccani (2016) have argued that this approach may falsely identify rapid bursts in volatility as "jumps." We purposely do not try to distinguish between these two alternative theoretical-based explanations for the "large" intraday changes in the SDF. We simply refer to both as "SDF jumps" in the sequel.

Baker et al. (2021) in their news-based explanations for aggregate stock market volatility and large market moves, respectively.[2] To allow for a more aggregated bird's-eye view of the resultant variance decompositions of the SDF and the estimated topic news risk premiums, we further combine some of these previously defined news topics into a smaller set of easy-to-interpret so-called "metatopics." By focusing on the SDF and the tangency portfolio returns, as opposed to the returns on the aggregate market portfolio underlying the previous work by Baker et al. (2021) and others, we thus obtain a more accurate picture of the news that truly matters from a systematic risk and pricing perspective.

Importantly, and in contrast to a substantial portion of existing work that relies on lower-frequency return and textual data, we explicitly rely on real-time news via the *Dow Jones Newswires Archive*—a machine-readable collection of articles from the Dow Jones Company that is commonly used by active market participants—in conjunction with our high-frequency estimate of the SDF. Using this precisely timed news data substantially sharpens our empirical analyses and inference by allowing us to filter the news to narrow time windows around the precisely timed abrupt changes in the SDF, thereby more cleanly identifying the cause(s) for the changes, and in turn allowing for more accurate analyses of the variation and risk premiums associated with specific news categories.[3] By comparison, the use of daily news articles and coarser daily returns invariably blurs the distinction between jumps and smoother price moves, the latter of which is more likely associated with the progressive assimilation of the constant stream of different news that arrives throughout the day.[4]

To help further elucidate the economic mechanisms at work, we rely on the approach of Cieslak and Schrimpf (2019) and the high-frequency comovements of the SDF and the term structure of interest rates to classify the jumps in the SDF into four types of primitive economic shocks: economic growth shocks, financial risk premium shocks, and short- and long-rate shocks. We then connect these economic jump classifications with the news associated with the jumps and estimate the corresponding risk premiums. Interestingly, we find that much of the risk premiums associated with *Monetary policy and*

[2] To allow for more recent news trends, we further augmented these previously defined news topics with a few additional key terms, most notably related to the COVID-19 pandemic. For completeness, we also include a few missing terms from the "automatic" news topic categorizations in the studies by Bybee et al. (Forthcoming) and Bybee, Kelly, and Su (2023).

[3] In a similar vein, Chinco, Clark-Joseph, and Ye (2018) have recently sought to link the emergence and disappearance of high-frequency cross-sectional predictive relationships to precisely timed news stories about firm fundamentals. Jeon, McCurdy, and Zhao (2022) similarly associate precisely timed firm-specific news with daily "jumps" in a large cross-section of individual stock returns, while Christensen, Timmermann, and Veliyev (2022) document jumps in after-hours prices immediately following earnings announcements.

[4] Our focus on intraday variation implicitly assumes that the most important news for the pricing of U.S. assets occurs during regular U.S. trading hours. This, of course, does not rule out international news, only international news that is released outside U.S. market hours.

finance news has been driven by growth and long-rate shocks, consistent with the critical role that forward guidance has played over our sample period, while the risk premiums for news associated with *Macroeconomic data*, *Labor*, and *Regulation* are, not surprisingly, dominated by growth shocks. Further underscoring the advantages of our all-encompassing text-based news classification procedures vis-à-vis the existing literature on stock market reaction to monetary policy that has mostly focused on FOMC and other scheduled central bank communications (e.g., Bernanke and Kuttner 2005; Lucca and Moench 2015), we find that a preponderance of the jumps in the SDF attributed to our *Monetary policy and finance* metatopic stem from unscheduled central bank communications and speeches by Fed governors or other Fed officials (see also the analysis in Cieslak, Morse, and Vissing-Jorgensen 2019).

The advantage of using high-frequency data for more accurately identifying and estimating news announcement effects has previously been emphasized in the literature. Fair (2002), in particular, explicitly points to the weak identification afforded by the use of "coarse" daily data as the reason for the apparent lack of a clear news-based explanation for many of the largest aggregate stock market changes reported in the widely cited study by Cutler, Poterba, and Summers (1989). Relatedly, there is now also a large existing, mostly empirically oriented, literature pertaining to macroeconomic news announcement effects and the way in which announcement surprises affect the intraday returns on specific assets and/or asset classes (see, e.g., Andersen et al. 2003, 2007; Faust et al. 2007; Lee and Mykland 2008; Evans 2011; Lahaye, Laurent, and Neely 2011; Lee 2012, along with the many other references therein). A more recent growing literature emphasizes the advantages of the use of high-frequency data for obtaining more reliable identification of various economic mechanisms through heteroscedasticity and the idea that return volatilities tend to be higher shortly after macroeconomic and other scheduled news announcements (see, e.g., Bollerslev, Li, and Xue 2018; Nakamura and Steinsson 2018; Bianchi et al. 2023, among others). The present paper adds to both of these strands of literature by formally characterizing the economic significance and relative importance of *all* different types of news, including unscheduled announcements, through the decomposition of the tangency portfolio risk premium into separate news risk premiums associated with a set of well-defined news topics. As such, the paper also speaks directly to the broader ongoing debate about whether macroeconomic risk matters for asset pricing, highlighted in the recent review article by Brunnermeir et al. (2021). It also helps clarify possible economic causes behind "nonmarket" jumps as a source of systematic risk and risk premiums recently emphasized in a series of papers, including Aït-Sahalia, Jacod, and Xiu (2023), Chabi-Yo, Huggenberger, and Weigert (2022), Jacod, Lin, and Todorov (2022), and Li, Todorov, and Zhang (2024).

1. Estimating the High-Frequency SDF

Our estimation of the high-frequency stochastic discount factor (SDF) is based on the same general adversarial method of moments procedure recently advocated by Chen, Pelger, and Zhu (2024) (henceforth CPZ) in their estimation of a monthly SDF. This approach is directly motivated by the theoretical results in Hansen and Jagannathan (1997), formally establishing that the approximate SDF closest to the true SDF in a mean-square-error sense may be obtained by minimizing the largest possible squared pricing errors. A similar approach has also previously been proposed by Bansal, Hsieh, and Viswanathan (1993).

1.1 Methodology

To fix ideas, consider the canonical conditional moment restrictions implied by the standard no-arbitrage condition,

$$E_t[M_{t+1} R^e_{i,t+1}] = 0, \tag{1}$$

where $R^e_{i,t+1}$ denotes the excess return on asset i from time t to $t+1$, and M_{t+1} refers to the SDF over that same time interval. Projecting the SDF onto the space of returns, it readily follows that

$$M_{t+1} = 1 - w_t^{\intercal} R^e_{t+1}, \tag{2}$$

where $w_t^{\intercal} R^e_{t+1}$ equals the return on the tangency portfolio, defined by

$$w_t = \left(E_t[R^{e,\intercal}_{t+1} R^e_{t+1}]\right)^{-1} E_t[R^e_{t+1}]. \tag{3}$$

While this straightforward solution for the SDF involves a simple function of estimable quantities, it is challenging to implement in practice due to the difficulties in accurately estimating the conditional risk premiums $E_t[R^e_{t+1}]$, and the inversion of the potentially high-dimensional second-moment matrix $E_t[R^{e,\intercal}_{t+1} R^e_{t+1}]$.

Instead, we proceed by estimating the weights for the tangency portfolio as a function of time t information. In particular, define:

$$w_t \equiv f_w(I_t; \theta_w), \tag{4}$$

where $I_t \in \mathbb{R}^K$ refers to the time t information set, and $f_w : \mathbb{R}^K \to \mathbb{R}^N$ is a known function of I_t parameterized by θ_w. Our estimate of the tangency portfolio,

$$F_{t+1} = f_w(I_t; \theta_w) R_{t+1}, \tag{5}$$

is then constructed by choosing the θ_w parameters such that the implied SDF minimizes the resultant conditional alphas for some deliberately chosen set of test assets.

Consistently estimating the conditional alphas over fixed time intervals is, of course, impossible without additional assumptions (see, e.g., Merton 1980).

Accordingly, we rely on a set instrumented unconditional moment conditions directly implied by the conditional no-arbitrage restrictions in Equation (1). That is,

$$E[M_{t+1} R^e_{t+1} g_t] = 0, \tag{6}$$

where the g_t vector of instruments is determined by the $f_g : \mathbb{R}^K \to \mathbb{R}^{N \times N_g}$ function parameterized by θ_g,

$$g_t \equiv f_g(I_t; \theta_g). \tag{7}$$

The $R^e_{t+1} g_t(I_t)$ term in (6) is readily interpreted as the return on a set of N_g portfolios, with the equation stipulating that the implied SDF is able to price all of these portfolios, or equivalently that no alpha can be found by trading on the I_t information set.

The θ_g parameters that determine the optimal set of instruments and the θ_w parameters that determine the optimal weights are, of course, both unknown and must be estimated. Motivated by the aforementioned results in Hansen and Jagannathan (1997), we jointly estimate these parameters based on the following minimax objective for the weights and instruments:[5]

$$\min_w \max_g \frac{1}{N_g} \sum_{i=1}^{N_g} \left\| \underbrace{\mathbb{E}\left[\overbrace{(1 - w_t R^e_{t+1})}^{M_{t+1}} R^e_{t+1} g_{t,i} \right]}_{\text{Conditional Error } \alpha_{g,i}} \right\|^2 . \tag{8}$$

The inner maximization problem, pertaining to the instruments, may naturally be interpreted as that of an arbitrageur trying to pick the θ_g parameters for the function $f_g(I_t; \theta_g)$ so as to maximize the resultant portfolios' conditional alphas, say $\alpha_{g,i}$. The outer minimization problem, pertaining to the weights, may in turn be interpreted as that of an asset pricer trying to pick the θ_w parameters for the function $f_w(I_t; \theta_w)$ such that the implied SDF minimizes the conditional alphas for the specific set of portfolios, $R^e_{t+1} g_{t,i}$.

1.2 Practical implementation and functional approximations

The population expectations defining the minimax problem in (8) are not directly observable. To obtain a practically feasible version, we replace the expected conditional alphas with their full-sample averages,

$$\hat{\alpha}_{g,i} = \frac{1}{T} \sum_t (1 - \hat{w}_t R^e_{t+1}) R^e_{t+1} \hat{g}_{t,i}, \tag{9}$$

[5] In addition to providing the closest approximation to the true SDF in a mean-square-error sense, Chernozhukov et al. (2020) have recently shown that under additional regularity conditions, the SDF estimated by this minimax procedure will formally converge to the true SDF at an almost parametric rate.

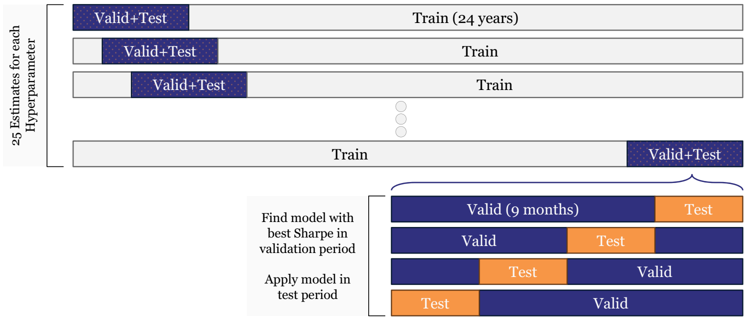

Figure 1
SDF cross-validation estimation procedure
This figure illustrates the folded "out-of-sample" cross-validation procedure that we rely on for estimating the high-frequency SDF. The procedure consists of first estimating the SDF for each set of hyperparameters, dropping one year from the full 25-year sample at a time. We then validate the resultant estimates on a quarterly basis, dropping one quarter at a time.

resulting in the corresponding minimax problem:

$$\min_{w} \max_{g} \frac{1}{N_g} \sum_{i=1}^{N_g} \hat{\alpha}_{g,i}^2. \tag{10}$$

The practical implementation of this problem still necessitates a choice for the $f_g(I_t;\theta_g)$ and $f_w(I_t;\theta_w)$ instrument and weight functions. Again closely following CPZ, building on Bybee, Kelly, and Su (2023) and techniques from the recent Machine Learning (ML) literature, we rely on neural networks for flexibly approximating both of these functions. Neural network-based approximations have also previously been used in the literature for the nonparametric estimation of the SDF in more traditional method-of-moments-based settings by Bansal and Viswanathan (1993) and Chen and Ludvigson (2009), among others.

Additionally, to help mitigate problems with overfitting, we employ a modified K-fold cross-validation procedure, as illustrated in Figure 1. More precisely, we loop over each calendar year in our full 25-year sample and generate an out-of-sample estimate of the tangency portfolio for that year by minimizing the loss function for the remaining 24 years; we repeat this procedure for each hyperparameter set in our grid. Armed with the results from this outer loop, we then determine the best estimate for each quarter through an additional folding process, in which one quarter serves as the test period and the other three quarters are used for validation. Following CPZ, we rely on the Sharpe ratio of the estimated tangency portfolio as our validation metric for this inner loop. This K-fold cross-validation approach has the advantage that it allows the SDF model parameters and hyperparameters to vary, while at the same time reducing overfitting and maintaining computational feasibility.[6] A related cross-validation procedure has also recently been applied by Kozak, Nagel, and Santosh (2020). In contrast to that latter study, however, which analyzes the out-of-sample performance of their SDF estimates across a range of hyperparameters, we purposely seek to obtain a singular SDF estimate based on an "optimal" set of hyperparameters within each quarter. Further details concerning the explicit functional forms for the weights and instruments along with the hyperparameter grid that we employ are provided in Appendix A and the Internet Appendix.

1.3 Return data and conditioning information

The set of assets underlying our estimation consists of the 272 high-frequency portfolios recently constructed by Aleti (2024). All of the portfolio returns are sampled at a 15-minute frequency and cover the sample period from January 2, 1996, to December 31, 2020, for a total of 169,965 intraday 15-minute return observations.[7] We will refer to the returns on these portfolios as $Z_t \in \mathbb{R}^{272}$ in the sequel. The 272 portfolios comprise 218 factor portfolios following Chen and Zimmermann (2021) and Jensen, Kelly, and Pedersen (2023), 48 industry portfolios, plus the commonly used six Fama-French portfolios (FF6). The long-short, or net-zero investment, portfolios are all directly compatible with the excess return format in the basic no-arbitrage condition in Equation (1). Correspondingly, for the market portfolio and the 48 industry portfolios, we subtract the risk-free rate to obtain the relevant excess returns.[8] Taken as a whole, these portfolios serve as a powerful set of span and test assets. The 218 high-frequency factor portfolios, in particular, effectively capture the many "anomalies" highlighted in the asset pricing literature, while the industry portfolios account for well-documented industry-specific effects. The inclusion of the FF6 portfolios further ensures that the estimated SDF will be able to price the Fama-French workhorse factors.

Our use of portfolio returns to span the high-frequency SDF contrasts with CPZ and several other recent studies that rely on individual stock returns for

6 We thank an anonymous referee for suggesting the use of a folding procedure to obtain "out-of-sample" estimates. We also note that the temporal dependencies in the conditioning variables used in the $f_g(I_t;\theta_g)$ and $f_w(I_t;\theta_w)$ functions may introduce a subtle look-ahead bias. However, setting aside a sufficiently long initial sample purely for training and validation purposes to allow for a traditional rolling out-of-sample estimation would substantially limit the sample horizon available for the empirical analysis.

7 As discussed in more detail in Aleti (2024), the use of a coarse 15-minute sampling frequency effectively mitigates the impact of market microstructure noise. As an aside, the present context also facilitates the practical estimation compared to the use of finer, say 5-minute returns, which would be prohibitively more expensive from a computational perspective.

8 We proxy the risk-free rate by the daily returns for the 1-month Treasury-bill rate from Kenneth French's website, "distributing" the returns equally across each of the within-day 15-minute intervals. This theoretically motivated excess return adjustment, of course, has virtually no effect on any of our estimates, especially considering that the risk-free rate is close to zero in our sample.

spanning the SDF at lower daily or monthly frequencies.[9] Our motivation for doing so is threefold. First, portfolios are always well-defined, ensuring a balanced panel. Second, portfolios better represent systematic risk relative to stocks. And third, portfolio returns are generally much less susceptible to market microstructure noise, thereby allowing for more meaningful use of higher-frequency data. On the other hand, our use of factor returns and univariate portfolio sorts instead of individual stock returns implicitly rules out higher-order interactions between the returns and firm-specific characteristics. Such interaction effects have recently been found to be valuable for better understanding variation in asset returns, in both the time-series (Gu, Kelly, and Xiu 2020) and cross-sectional (Chen, Pelger, and Zhu 2024; Bryzgalova, Pelger, and Zhu Forthcoming) dimensions. Indeed, it is possible that extending the set of basis assets beyond the "zoo" of univariate factors via more complex interactions among the characteristics may similarly improve the performance of the estimated SDF. However, since our high-dimensional, high-frequency estimation already pushes the boundaries of our computational constraints, we do not pursue that here.

In terms of the conditioning set that we use for incorporating important economic information, we again closely follow CPZ and rely on data drawn from three different sources. The first data set, *FRED-MD*, consists of 126 monthly macroeconomic variables, as further discussed in McCracken and Ng (2016). The second data set consists of the cross-sectional medians of the 153 firm characteristics recomputed on a monthly basis using the characteristic data from Jensen, Kelly, and Pedersen (2023).[10] The third, and last, data set consists of the eight popular monthly equity risk premium predictor variables highlighted in the oft-cited study by Welch and Goyal (2008). We further transform each of these individual time series as necessary to render them stationary and follow with a rolling standardization. We also lag all of the series by 1 month to ensure that the combined I_t information set is actually available at time t. Additional details concerning these transformations and the interpolations that we use in the construction of the combined high-frequency data set are provided in the Internet Appendix.

1.4 Tangency portfolio estimates

With our methodology and data defined, we proceed with the estimation—see the Appendix A for additional details—and now discuss our results. Here, the

[9] Other recent lower-frequency SDF estimation procedures that explicitly rely on large cross-sections of individual stocks include the so-called "agnostic" approach of Pukthuanthong, Roll, and Wang (2020) and Kim and Korajczyk (2021), and the Bayesian approach of Kozak, Nagel, and Santosh (2020). Bryzgalova, Pelger, and Zhu (Forthcoming) similarly rely on large numbers of individual stocks in their construction of managed portfolios to span the SDF.

[10] Unlike CPZ, who rely on their own choice of 46 firm characteristics, we rely on the 153 characteristics produced by Jensen, Kelly, and Pedersen (2023). This choice is primarily motivated by data availability for our sample period.

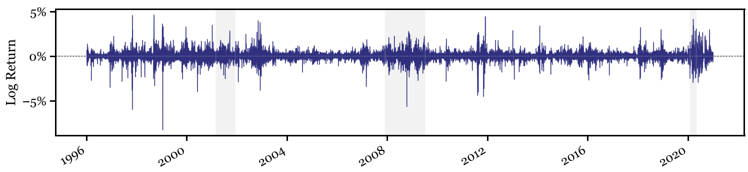

Figure 2
Tangency portfolio returns
This figure reports the 15-minute intraday and overnight returns on our estimated tangency portfolio, rescaled to have the same realized volatility as the Fama-French market portfolio. NBER-defined recessions are highlighted in gray.

key expression is Equation (5), which formally defines the tangency portfolio as a weighted combination of the span assets. Armed with our estimated weight function and the high-frequency portfolio returns, we write our estimated tangency portfolio as

$$\hat{F}_{t+1} = 1 - \hat{f}_w(I_t; \hat{\theta}_w) Z_{t+1}. \tag{11}$$

Figure 2 displays the resultant SDF returns, where for ease of comparison we have rescaled the returns to have the same full-sample realized volatility as the Fama-French market portfolio.

The general features of the high-frequency SDF returns closely mirror those of most other high-frequency return series. The high-frequency tangency portfolio returns are, not surprisingly, on average also positively correlated with the high-frequency returns on the aggregate market portfolio, with the full-sample correlation equal to 23.9%.[11] Meanwhile, as further detailed in Appendix A.4, the high-frequency SDF does a much better job of explaining the cross-sectional variation in the 272 high-frequency test portfolios than both the CAPM and the FF6 model; the full-sample cross-sectional R^2 on the SDF is about 52.7% compared to 16.9% for the CAPM and 33.6% for the FF6. Lastly, the annualized Sharpe ratio of our estimated tangency portfolio is 0.98, economically plausible and markedly higher than that of the Fama-French market portfolio which achieves a Sharpe of 0.41.

Further illustrating the key features of the estimated SDF, Figure 3 reports the annualized realized volatilities computed from the summation of the 15-minute intraday and overnight squared returns. For ease of interpretation, the figure shows the backward-looking exponentially weighted moving average based on a half-life of 30 days. In line with the extensive literature on time-varying financial market volatility (see, e.g., Bollerslev et al. 2018, and the many references therein), the volatility of the SDF clearly varies over time in a strongly positively autocorrelated fashion. Regressing the daily realized

11 The magnitude of this correlation is about double that found by Chen, Pelger, and Zhu (2024). As evidenced by the additional results reported in the Internet Appendix, there is also a fair amount of variation in this correlation over time, as there is in the correlations with the representative factor cluster portfolios.

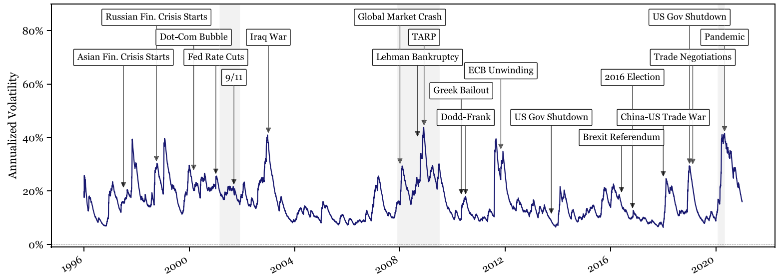

Figure 3
Tangency portfolio volatility
The figure plots the annualized realized volatility of the estimated tangency portfolio. For visual clarity, we smooth the underlying realized variance series using a 30-day half-life EWMA. The tangency portfolio itself is rescaled to have the same unconditional realized volatility as the Fama-French market portfolio.

volatility of the SDF on the lagged daily, weekly, and monthly realized SDF volatilities, as in the popular HAR model of Corsi (2009), also results in a fairly high R^2 of 32.1%, mirroring the strong degree of return volatility predictability typically observed for the aggregate market portfolio.

As indicated by the specific events annotated in Figure 3, periods associated with high economic uncertainty and/or financial crises clearly tend to be accompanied by relatively high SDF volatility. Corroborating this visual impression, the correlations between the monthly realized SDF volatility and the monthly *Financial*, *Macro*, and *Real uncertainty* indices from Jurado, Ludvigson, and Ng (2015) equal 59.5%, 47.9%, and 43.3%, respectively. This again directly mirrors existing evidence for the market portfolio, and the tendency for realized volatility to increase during periods of market "stress" (see, e.g., Banulescu et al. 2016). This connection has also previously been used by Manela and Moreira (2017) and Baker et al. (2019) in the construction of news-based aggregate market volatility indexes.

To help more clearly delineate these linkages and the economic news that is actually priced, we further decompose the high-frequency SDF into separate continuous and more abrupt jump components.

1.5 Tangency portfolio jumps

Interpreting the estimated 15-minute tangency portfolio returns $\hat{F}_t$ defined in (11) as the discrete-time realization of some true underlying continuous-time log price process, we rely on techniques from high-frequency financial econometrics to separate the continuous and discontinuous moves in this portfolio. Intuitively, if the increment in $\hat{F}_t$ over a given 15-minute interval is "too large" in an absolute value sense to have likely been generated by a Gaussian process with a local variance commensurate with the usual variation over the interval, we classify the increment as a "jump." A more formal discussion of this thresholding procedure, including the estimation of the local variance and the construction of the threshold, is given in Appendix B.

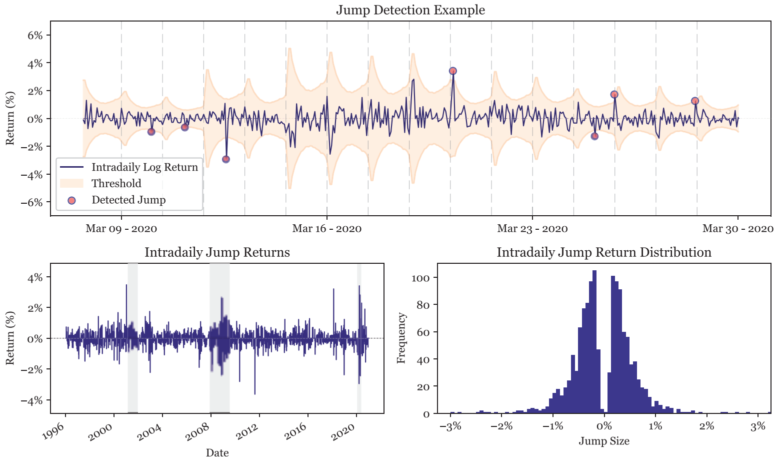

Figure 4
Jump identification example and jump returns
The top subplot shows the intradaily returns on the estimated tangency portfolio for the last 3 weeks of March 2020, together with the specific jump threshold that we rely on and the jump returns that exceed the threshold marked (in red). The full-sample jump returns for the estimated tangency portfolio, again rescaled to have the same realized volatility as the Fama-French market portfolio, are plotted as a time series in the lower-left subplot and as a histogram in the lower-right subplot.

To illustrate the basic idea, the top panel in Figure 4 plots the intraday high-frequency returns on the estimated tangency portfolio for the last 3 weeks of March 2020, coincident with the global onset of the COVID-19 pandemic, together with our corresponding time-varying jump thresholds. As indicated by the red dots in the figure, our procedure, not surprisingly, identifies several SDF jumps during this 3-week period. As will be discussed further below, all of these jumps may also naturally be linked to specific news about the severity of the pandemic and/or statements and actions by the Federal Reserve and other policymakers intended to help mitigate its economic impact.[12]

Using this same thresholding technique, the bottom-two panels in Figure 4 show the time-series and size distribution of the identified intraday jumps in the high-frequency SDF. On average, there are 52 intraday jumps per year, with a maximum of 70 jumps in 2010 and a minimum of 40 jumps in 2001. The second panel helps further visualize the jump magnitudes, showing that the jumps tend to be fat-tailed.[13] In contrast to the jump returns, the continuous returns are far smaller. They are generally very difficult to associate with observable economic information or specific news. Instead, we deliberately exploit the

[12] Figure B.1 in Appendix B provides additional illustrative examples of SDF jumps during other time periods that may similarly be linked to specific economic news events.

[13] The gap in the center of the distribution is simply an artifact of the thresholding technique, which inevitably cannot identify “small” jumps.

more abrupt changes in the SDF, as represented by the jump returns, to more clearly delineate the news that is actually priced by investors.

Also, our jump detection procedure leans on infill asymptotic arguments, which assume that the empirically observed returns are sampled at an increasingly higher frequency. This assumption motivates and necessitates our use of high-frequency data. As further detailed in the Internet Appendix, relying on a similar threshold-based procedure while instead using lower-frequency daily (close-to-close) returns delivers far noisier results compared to our high-frequency approach. To quantify the degree of misclassification, we perform a jump detection exercise, attempting to classify days with/without a jump, using said daily returns along with our 15-minute returns. We find that both sets of classifications agree upon 69.6% of the days in our sample being classified as nonjump days and 3.2% of days being classified as jump days. However, a nontrivial 15.5% of the days are classified as jump days by the 15-minute returns but as nonjump days by the daily returns, while the converse holds for the remaining 11.7% days.[14] The findings for open-to-close returns are qualitatively similar suggesting that the use of high-frequency data is critical for effective jump detection.

As a precursor to our textual analyses and to help further underscore the advantages of the use of high-frequency returns for pinpointing important news, the Internet Appendix also reports the correlation between the daily SDF returns on days with prescheduled macroeconomic news based on the Bloomberg Economic Calendar and the 15-minute SDF returns that span the exact time of the news release. One might naturally expect that on said days the correlation between the daily and the news-specific high-frequency returns should be relatively high. However, the correlation between the daily and the high-frequency returns only equals 0.434, implying there exists a substantial amount of additional daily variation in the SDF beyond that observed right around the time of the precisely identified economic news releases that clouds the lower frequency returns.[15] In other words, the use of automatic text-based news attribution procedures to explain "large" daily SDF returns, even on days with known important economic news, is likely to be hampered by the considerable amount of additional news that typically occurs throughout most days.[16]

---

14 Supposing that the 15-minute classifications are the ground truth, this amounts to a Type I error rate of 11.7/(11.7 + 69.6) = 14.4%, while the Type II error rate is as high as 15.5/(15.5 + 3.2) = 82.9%.

15 Excluding the overnight portion of the daily return and focusing on open-to-close returns only slightly increases the correlation to 0.494.

16 Corroborating this thesis, the Internet Appendix reports the results from applying the exact same news attribution scheme that we rely on for the high-frequency jump returns, as discussed in more detail below, on days with well-defined important economic news. Case in point, on FOMC announcement days the average "topic weight"—introduced in Section 2.2—for the topic *Monetary policy* is only around 20%, while it is almost 60% using the high-frequency identification scheme on said days. Similar results hold true for other well-defined news events.

Correspondingly, with hundreds of news articles typically being published in the 17.5-hour time interval between the close of the market on one day and the opening of the market on the following day, the approach that we rely on would not be tenable, let alone reasonable, for "explaining" the overnight SDF returns. We thus ignore the overnight portion of the returns in all of our news-based attributions. This, of course, does not rule out the pricing of international news *per se*, only news that is always published outside regular U.S. market hours. Indeed, anticipating our results, news concerning *International affairs*, as defined in the next section, emerges as our overall second most important news topic.

Next, we turn to a more detailed description of the news data and our approach for linking the SDF jumps with the different news topics.

2. Linking Systematic Jumps with News

Our approach for linking the jumps in the tangency portfolio with news is based on the counts of keywords in several million precisely timed newswire articles. We begin by briefly detailing the news data, followed by a discussion of the way in which we group key news article terms into prespecified *news topics*. Aggregating the mentions of each news topic over 15-minute intervals, we in turn "explain" each of the SDF jumps with the dominant news topics over the relevant time interval.

2.1 News data

Our primary data set consists of the *Dow Jones Newswires Archive*, a machine-readable collection of articles from the Dow Jones' real-time news feeds. Retrieving all articles from January 1, 1996, to December 31, 2020, leaves us with a total of 50,879,472 articles. Each of these articles consists of a headline, a body text, a subject/product/company code, along with additional identifiers. We deliberately exclude articles that are seemingly irrelevant to investment decisions, such as those about sports, entertainment, and lifestyle. We also deem articles that simply state open/close prices and various technical indicators as being irrelevant for capturing innovations in the state variables that drive the SDF. Similarly, we remove certain articles about company-specific news, which is arguably idiosyncratic and hence should not affect the SDF. However, consistent with the idea that certain company news and the news pertaining to systematically important firms may have economy-wide implications (e.g., Patton and Verardo 2012; Savor and Wilson 2016), we deliberately design our filters to retain firm news that is deemed to be relevant more broadly. A more in-depth discussion of all the exclusion filters that we rely on, including examples of company-specific news that is filtered out and company news that is retained, is provided in the Internet Appendix.

All in all, after applying the above filters, we are left with a total of 5,167,880 relevant articles over the full sample period. As the resultant article counts in

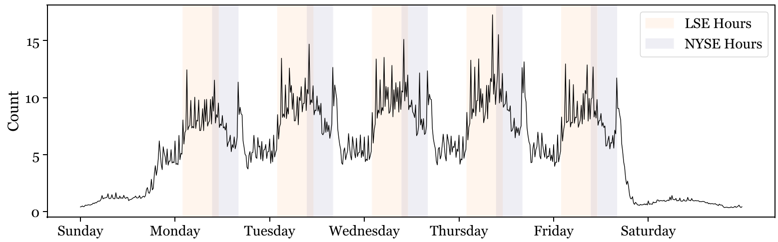

Figure 5
Article counts
The figure shows the average number of news articles in our filtered data set for each 15-minute interval and day of the week. The shaded areas represent London Stock Exchange and New York Stock Exchange market hours.

Figure 5 show, the news articles tend to be posted on weekdays at the start of European market hours until slightly after American markets close. As further shown by the additional summary statistics provided in the Internet Appendix, even though there has been a decline in the average number of "relevant" articles being posted per month from the beginning to the end of the sample (in part because of changes and updates to the Dow Jones data set itself), there are typically still several hundred investment-related articles being posted on a daily basis throughout the entire sample period.

2.2 Extracting news topics

To determine which topics are prevalent in the news, we begin by assigning topic counts for each article in our filtered data set of relevant articles. These topic counts are computed as the number of key terms associated with each topic in a given article. For example, the topic *Monetary policy* and its associated key terms naturally include *Federal Reserve*, *money supply*, *open market operations*, and *Fed funds rate*, among others. Counting the mentions of these key terms provides us with a direct measure of the prevalence of *Monetary policy* as a topic in a given news article. Repeating this procedure across all articles in turn reveals what topic dominates the news at any given point in time. A similar automatic approach of aggregating news topic counts has also recently been employed by Bybee, Kelly, and Su (2023), albeit over coarser daily time intervals.

Our topic counts are primarily based on the topics and key terms previously defined by Baker et al. (2019). These topics and terms were all "hand-selected" with the explicit purpose of studying stock market volatility. They extend the news topics and associated terms previously used by Baker, Bloom, and Davis (2016) and Davis (2017) for measuring economic policy uncertainty. However, the former topic list is still not entirely up-to-date, requiring additional modifications to accommodate more recent trends in the news, most notably the COVID-19 pandemic. Additionally, to ensure that our list of key terms is comprehensive, we further augment the lists with any missing key terms from the aforementioned Bybee, Kelly, and Su (2023) study. Altogether, this leaves us with a set of 44 handpicked news topics with an average of 24.7 key terms each. The Internet Appendix comprehensively lists all the topics and associated key terms.

By contrast, Bybee et al. (Forthcoming) and Larsen and Thorsrud (2019) both rely on Latent Dirichlet Allocation (LDA), a generative Bayesian model proposed by Blei, Ng, and Jordan (2003), for defining key terms and topics. However, LDA and other unsupervised machine learning procedures often produce word clusters that are difficult to interpret and/or have seemingly little to do with news about the state of the economy, invariably requiring some additional hand-cleaning or, in downstream applications, regularization to avoid fitting on irrelevant topics. To illustrate these challenges more concretely in the present context, the Internet Appendix reports the results from running LDA on the Dow Jones news data. As these additional results show, many of the "machine-learned" topics produced by this automatic procedure tend to be "noisy" and difficult to interpret from an economic perspective.

Having defined the topics and corresponding key terms, we calculate the topic counts for each article. To do so, we begin by combining the headline and body text into a single block of text. Next, we preprocess this text data using standard transformations from the Natural Language Processing (NLP) literature, consisting of the following steps: (a) convert all text to lowercase letters, (b) remove any stop words,[17] (c) delete multiple spaces and line breaks, (d) remove nonalphanumeric characters, and (e) stem and lemmatize each word.[18] We then tokenize the text and extract n-grams, or groups of n-adjacent words.[19] Finally, we count the number of n-grams that appear in each topic's key term list for each article, producing the requisite topic counts. This now fairly standard type of approach for automatic text processing also underlies the related work by Ke, Kelly, and Xiu (2021) and Bybee et al. (Forthcoming).

Although uniquely identified by the key terms, many of the news topics determined by the above-defined counts are fairly specific, yet also inherently related. Hence, to provide a more broad-based view on the type of news that matters, we further combine the 44 more detailed news topics into a smaller set of 8 "metatopics." The compositions of most of these metatopics are fairly self-explanatory. For instance, our *Monetary policy and finance* metatopic naturally combines the previously defined *Monetary policy*, *Other financial indicators*, *Financial regulation*, *Interest rates*, and *Inflation* topics into a single topic. As another example, our *Commodities and energy* metatopic combines the *Commodity Markets* and *Energy markets* topics into a single metatopic.

---

17 Some examples of stop words are *the*, *is*, and *are*. We obtain our list of stop words from the *Natural Language Toolkit* (NLTK), a Python library developed by Bird, Klein, and Loper (2009).

18 Lemmatization entails converting the words such as *taxes* and *taxation* to their root word *tax*.

19 As an example, the text *conduct monetary policy* would be tokenized into {*conduct*, *monetary*, and *policy*}. The set of unigrams is simply the set itself. The set of bigrams is {*conduct monetary*, *Monetary policy*}. The set of trigrams is {*conduct monetary policy*}. Since our list of key terms consists of unigrams, bigrams, and trigrams, these are also the only n-grams that we extract.

Of course, not all of the 44 more detailed news topics are as easily categorized and combined, and as such some of the original topics appear in more than one of the 8 metatopics, while some do not appear in any metatopic.[20] Appendix C.1 provides an exact description of the relevant metatopic definitions.

To help shed more light on the potential overlap between topics and more concretely motivate our metatopics, the Internet Appendix reports the topic clusters obtained by applying hierarchical agglomerative clustering and K-means clustering to our news data, and the document topic-count matrix in particular. As these additional results show, the clusters produced by both commonly used automatic procedures tend to be "noisy" and difficult to interpret from an economic perspective. However, they still highlight certain topics that tend to be mentioned jointly in the news. For instance, the hierarchical approach reveals that *Monetary policy* and *Interest rates* link together, along with the pairs *Energy markets* and *Commodity markets*, as well as *Government spending, deficits, and debt* and *Taxes* to name a few. In sum, while the machine-learned clusters are not directly useful, they are broadly congruent with and further support our hand-tailored groupings and metatopics.

2.3 Linking topics with SDF jumps

As discussed in Sections 1.3 and 1.5, to help alleviate concerns about the impact of market microstructure noise, we purposely rely on a "coarse" 15-minute sampling frequency for the estimation of the jumps in the SDF. Accordingly, to link the estimated SDF jumps with the news, we begin by aggregating the topic counts for all of the relevant news articles across the same 15-minute time intervals used in identifying the jumps. Although the SDF, in theory, should respond immediately upon the release of new economic information that investors care about, this 15-minute temporal aggregation of the topic counts simultaneously serves to allow for more gradual incorporation of the news, consistent with the idea that it might take market participants some time, however brief, to fully digest and interpret the news, in turn resulting in "gradual" jumps as first hypothesized by Barndorff-Nielsen et al. (2009) (see also the recent discussion in Bollerlsev 2022).

To further justify this choice, we impute finer 1-minute SDF returns based on the estimated tangency portfolio weights together with the 1-minute factor returns. Figure 6 plots the resultant average realized variance of the SDF at this finer 1-minute frequency around the time of the prescheduled news announcements on the Bloomberg Economic Calendar. Per the more detailed discussion in the Internet Appendix, an average of 217 such "Bloomberg events" happen per year. In addition to centering the realized variance at

[20] More precisely, 7 of the 37 news topics that define our 8 metatopics are repeated twice, while 11 of the 44 original news topics are not included in any metatopic.

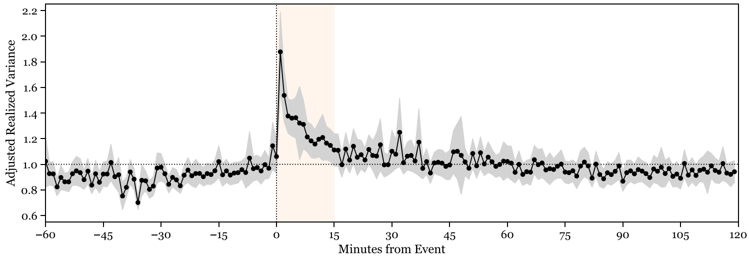

Figure 6
News event time variation
The figure shows the average time-of-day adjusted realized variance centered at Bloomberg news announcement times. The time-of-day adjustment is computed using the average realized variance across days for each particular time of day. Minute zero on the *x*-axis corresponds to the Bloomberg calendar time of the event. The gray area represents 95% confidence intervals.

the exact news announcement times, to control for the well-known U-shaped pattern in the intraday return variation, we also adjust the realized variance for a time-of-day effect.[21] Accordingly, if the news does not affect the SDF, this normalized realized variation measure should remain steady around one. Meanwhile, as the figure shows, the normalized average realized variation remains statistically significantly above unity for about 15 minutes following the event time. In other words, the new information associated with these precisely timed Bloomberg news announcements is typically absorbed by the tangency portfolio within 15 minutes. Of course, some news events may take longer to be fully digested, while others may be absorbed at an even faster speed, but 15 minutes appears as a reasonably balanced choice.[22]

To allow for slow-moving variation, or trends, in the importance of different types of news over the full 1996–2020 sample period, we further demean the aggregated 15-minute topic counts based on a backward-looking 30-day moving average, deleting all demeaned topic counts below zero over each interval. We then sort these demeaned aggregated topic counts to determine

21 More formally,

$$AdjRV(c)=\left(\sum_{t,i}\left(\Delta_i^{1min}F_t\right)^2\cdot TOD_i\cdot\mathbb{1}_{\left[Event_{t,i}=c\right]}\right)\Big/\left(\sum_{t,i}\mathbb{1}_{\left[Event_{t,i}=c\right]}\right),$$

where $\Delta_i^{1min}F_t$ refers to the 1-minute tangency portfolio return on day t at time i, constructed by multiplying the estimated weights from our main analysis by the 1-minute returns on the underlying spanning factors, and the TOD_i time-of-day adjustment variable is simply defined as:

$$TOD_i=\frac{1}{\#(\Delta_{1min})}\sum_t\left(\Delta_i^{1min}F_t\right)^2.$$

22 The Internet Appendix provides additional robustness checks pertaining to other choices of sampling frequencies and event windows.

the dominant, or primary, news topics for each of the 15-minute intervals, in turn associating each of the jumps in the tangency portfolio with a weighted average of the top news topics for the particular time interval containing the jump.[23] If a 15-minute time interval with a jump does not have any news articles that contain at least one key term associated with one of our news topics, we simply associate that jump with the topic *None*.

More specifically, let $DTC^{+}_{k,t,i} \equiv DTC_{k,t,i} \cdot \mathbb{1}_{[DTC^{k}_{t,i}>0]}$ denote the demeaned topic count for topic k on day t during time interval i. We then define the "soft" topic weights for time (t,i) based on the top K^* topics as:

$$TW_{k,t,i} = \frac{DTC^{+}_{k,t,i} \cdot \mathbb{1}_{[rank(DTC^{+}_{k,t,i}) \leq K^*]}}{\sum_k DTC^{+}_{k,t,i} \cdot \mathbb{1}_{[rank(DTC^{+}_{k,t,i}) \leq K^*]}}. \tag{12}$$

Fixing $K^* = 1$ would simply associate each jump with the dominant topic for each time interval. However, "distributing" the topic weights over the top K^* topics may allow for a more meaningful set of sparse topic weights. At the same time, using "too large" a value for K^* is likely to result in the inclusion of some topics that may have little or nothing to do with the specific news event(s) that caused the jump. Based on the additional analysis reported in the Internet Appendix and in an effort to balance this bias-variance trade-off, we fix $K^* \equiv 3$ for our main empirical analyses reported below.[24]

3. What Drives Systematic Risk?

We begin our analysis by assessing which of the news topics are associated with most of the jumps in the SDF, followed by an assessment of how much each topic contributes to the overall jump variation. We also present additional results based on aggregating the more detailed news topics into our more broadly defined, and easier-to-interpret, metatopics.

3.1 Why does the SDF jump?

To help assess the relative importance of the different news topics for explaining the jumps in the SDF, we begin by summarizing the unconditional and conditional topic frequencies. Specifically, we define the Topic Unconditional Frequency (TUF) of a given topic as the frequency by which it appears as one of the top K^* topics across all the 15-minute intervals in the sample,

$$\text{TUF}(k) \equiv \sum_{t,i} TW_{k,t,i} \Big/ \left(\sum_{t,i} 1 \right). \tag{13}$$

[23] The Internet Appendix provides an illustrative example for the 13:15-13:30 time interval on January 3, 2001, for which *Monetary policy* naturally emerges as the dominant news topic. The procedure is also robust to the choice of the window size we use for demeaning.

[24] Additional results for other values of K^* reported in that same appendix show that our main results remain robust across a range of K^*s, including $K^* \equiv 1$.

Table 1
News and jumps

Topic	TCF	TUF	Diff	t(Diff)
	Topic frequency (%)			
Monetary policy	7.22	14.74	7.53	11.34
U.S. politics	6.40	10.15	3.76	6.84
Energy markets	5.57	8.07	2.50	4.76
Middle East	4.56	6.16	1.60	3.51
Russia	3.12	4.67	1.55	3.81
National security	3.63	4.14	0.51	1.45
Labor markets	2.65	4.12	1.46	4.07
Inflation	2.66	3.96	1.30	4.17
Real estate markets	2.24	3.92	1.68	4.92
Taxes	2.82	3.54	0.72	2.20
None	25.20	3.42	−21.81	−43.32
Broad quantity indicators	2.13	3.17	1.04%	3.43
Trade	2.59	3.02	0.43	1.50
China	6.26	2.73	−3.53	−11.08
Commodity markets	2.36	2.62	0.26	0.97
Consumer spending and sentiment	1.65	2.52	0.88	3.40
Natural disasters	2.58	2.34	−0.23	−0.81
Financial regulation	1.52	2.25	0.73	2.89
Elections and political governance	1.55	2.10	0.55	2.11
Gov. spending, deficits, and debt	1.76	1.93	0.17	0.70

The table reports the unconditional and conditional frequency counts for the different news topics, given by the $TUF(k)$ and $TCF(k)$ statistics formally defined in the main text. The last two columns report the differences in the frequencies along with the t-statistics for testing whether the differences are statistically significant. The table only reports the top-20 news topics sorted by their conditional frequencies.

We similarly define and calculate the Topic Conditional Frequency (TCF) as the same average over the jump intervals only,

$$\text{TCF}(k) \equiv \sum_{t,i}\left(TW_{k,t,i} \times \underbrace{\mathbb{1}_{|F_{t,i}|\geq\alpha\sqrt{\tau_i BV_t}\,\Delta_n^{\varpi}}}_{\text{Jump in Tangency Portfolio}} \right) \Big/ \left(\sum_{t,i} \mathbb{1}_{|F_{t,i}|\geq\alpha\sqrt{\tau_i BV_t}\,\Delta_n^{\varpi}} \right). \tag{14}$$

This latter statistic represents the probability that topic k is one of the dominant news topics over an interval conditional on a jump being detected. Thus, if the jumps are statistically independent of the news, the $TUF(k)$ and $TCF(k)$ frequencies for any given topic k should be the same. Hence, by comparing the differences in the two frequencies, we obtain a simple assessment as to which of the different news topics primarily appear to be associated with systematic jumps.

Table 1 reports the resultant frequencies, along with the t-statistics for testing whether the differences in the $TUF(k)$ and $TCF(k)$ frequencies are statistically significant.[25] For brevity, we only include the top-20 topics based on the conditional frequency counts.[26] As the table shows, most of

25 The t-statistics are conveniently calculated from regressions of the form $TW_{k,t,i} = \beta_0 + \beta_1 \cdot \mathbb{1}_{|F_{t,i}|\geq\alpha\sqrt{\tau_i BV_t}\,\Delta_n^{\varpi}}$, and the significance of the β_1 coefficient associated with the indicator for the jumps.

26 As discussed in Appendix B, these results rely on a jump threshold parameter of $\alpha = 3.0$. The conditional frequencies for values of α in excess of 3.0, and the corresponding t-statistics for testing the differences in

the differences are strongly statistically significant. The largest difference manifests for *Monetary policy*, which appears with a frequency of 7.22% unconditionally compared to a conditional frequency of 14.74%. This finding is directly in line with Baker et al. (2021), who report that news associated with monetary policy seemingly triggered many of the largest (in an absolute value sense) daily stock market returns over the past half-century. It also corroborates the extensive historical analysis of intraday stock market jumps in Johnson, Medeiros, and Paye (2022), which suggests that monetary policy news has become increasingly important in recent decades. Of course, the importance of Fed policy for understanding asset markets has also long been emphasized in the macroeconomics literature – see, for example, the early studies by Kuttner (2001); Rigobon and Sack (2004); Bernanke and Kuttner (2005), along with the more recent review by Rogers, Scotti, and Wright (2014).

Although *Monetary policy* stands out as the overall most important news topic, *U.S. politics*, *Energy markets*, *Middle East*, and *Labor markets*, also all exhibit highly statistically significant, although smaller, differences in their $TUF(k)$ and $TCF(k)$ frequencies. Interestingly, the topic *None*, which again refers to intervals without a primary topic, stands out as by far the most frequent topic unconditionally, being associated with 25.20% of all 15-minute intervals in the sample. However, conditional on a jump being detected, only 3.42% of the intervals lack a primary topic. In other words, the SDF almost never jumps without identifiable economic news. This finding of a news-based explanation for most of the jumps in the SDF also accords with previous studies that have successfully linked various high-frequency jumps in aggregate equity index and individual stock returns to precisely timed marketwide and company-specific news releases (see, e.g., Lee and Mykland 2008; Lee 2012; Jeon, McCurdy, and Zhao 2022, and the many other references therein). In contrast to most previous studies, however, which have sought to link jumps with prescheduled news announcements, as will be discussed further below, we also find a very important role for unscheduled news.

Going one step further, Figure 7 displays the $TCF(k)$ conditional topic frequencies computed on a disaggregated annual basis. Not surprisingly, *Monetary policy* consistently ranks as one of the most frequent news topics for explaining the occurrence of jumps throughout the sample. The frequency of the *Monetary policy* topic appears particularly high during the 2001 recession, the 2008-2009 Great Recession, and from 2012 to 2016 and the time of the European Sovereign Debt Crisis, the Taper Tantrum, and the latter rounds of Quantitative Easing. Meanwhile, the topic *U.S. politics* tends to play a relatively more important role in explaining the jumps in the years following US presidential elections, and especially so during the first year of the Trump presidency. The frequency of SDF jumps related to news about *Energy markets*

$TUF(k)$ and $TCF(k)$, reported in the Internet Appendix, are generally very similar to the results based on $\alpha = 3.0$ reported in Table 1.

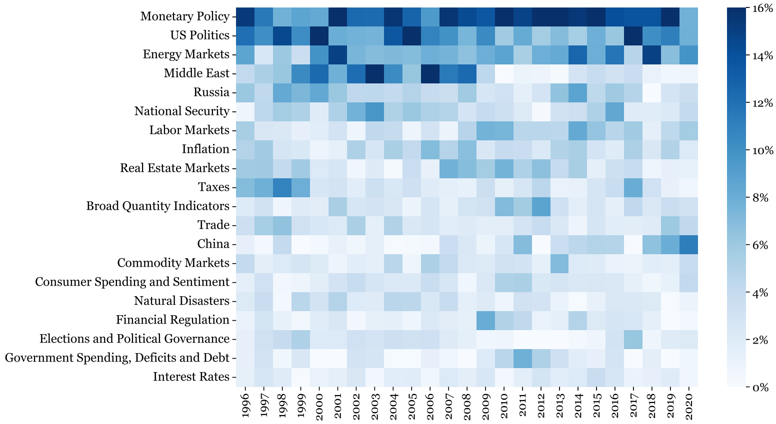

Figure 7
Conditional topic frequencies
The heatmap shows the fraction of jumps associated with each topic across all detected jumps for each year in the 1996–2020 sample. The display is limited to the top-20 topics by frequency over the full sample.

seemingly peaked in 2018, coincident with the steep decline in oil prices and a series of OPEC announcements, while both *Middle East*- and *National security*-related jumps appeared relatively more frequent after 9/11 and near the beginning of the Iraq War in 2003. The relative frequency of jumps attributed to news about *Russia*, unsurprisingly, peaked around the time of the Russian Financial Crisis in 1998, while *Disease* naturally stands out as the overall most frequent news topic for explaining the jumps in the SDF at the height of the COVID-19 pandemic in 2020.

To more concretely convey the specific news stories that actually matter, we perform a more granular analysis by extracting the news headlines for a sample of the SDF jumps associated with our top-three news topics: *Monetary policy*, *U.S. politics*, and *Energy markets*. Figure 8 displays all of the relevant detected jumps together with a select set of news headlines.[27] Mirroring the analogous display pertaining to the largest (in an absolute value sense) daily aggregate market returns in Bybee, Kelly, and Su (2023), the economic news stories underlying the SDF jumps, although topically related, are seemingly also quite diverse. For instance, for *Monetary policy* the underlying causes of the jumps include unscheduled announcements, official statements from Fed officials, ECB news, and various other Fed-related news. For *U.S. politics*, the headlines tend to be about Congress and elections. For *Energy Markets*, the news often has to do with OPEC announcements, and to a lesser extent, oil-related conflicts in the Middle East and statements by the International Energy Agency. Our text-based analysis of the news data

27 Similar headline displays for each of the three individual news topics are given in the Internet Appendix.

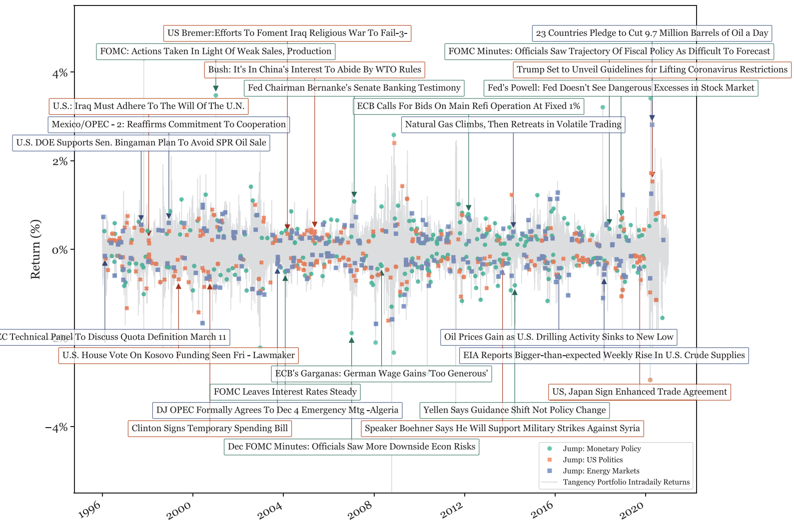

Figure 8
Select topic headlines
The figure displays the intradaily 15-minute returns for the estimated tangency portfolio in gray, together with all the jumps associated with *Monetary policy*, *US politics*, and *Energy markets*. The select headlines are drawn from articles that were published in the same time interval as the jumps; the jumps themselves are marked as of a particular topic if the corresponding topic weight exceeds 50%. Not all jumps are annotated in order to prevent overlapping headlines.

succinctly categorizes all of these disperse news stories into a set of well-defined news topics.

3.2 SDF jump variance decompositions

In addition to the frequency counts discussed in the previous section, it is instructive to consider how much each of the different news topics contributes to the total SDF jump variation.[28] To do so, we calculate the sum of all the squared jumps multiplied by their TW topic weights and divide that sum by the total sum of the squared jumps. Since the topic weights are normalized to sum to one, the resultant variance contributions naturally sum to 100%. Seeing that the risk-return relationship directly links the variation in the tangency portfolio to its expected return, these calculations therefore also provide a simple first indicative answer as to which news topics account for most of the compensation for systematic jump risks.

The year-by-year contributions for the top-20 topics fairly closely mirror the annual conditional jump frequencies displayed in Figure 7, and we defer a summary of these results to the Internet Appendix. Instead, to afford a more

[28] Relatedly, a long list of studies seeks to identify news and/or economic variables associated with the variation in aggregate stock market volatility at lower, typically monthly, frequencies (see e.g., the early oft-cited studies by Cutler, Poterba, and Summers 1989; Schwert 1989).

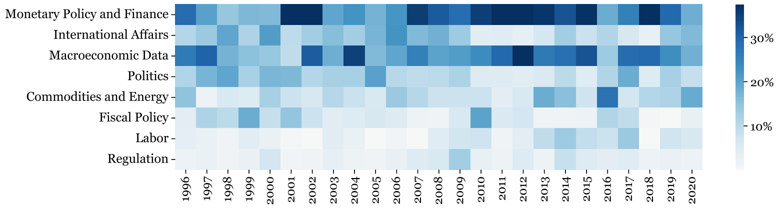

Figure 9
Metatopic variance contributions
The heatmap displays the fraction of jump variation associated with each metatopic relative to the total jump variation for each year in the 1996–2020 sample.

general picture of the news that drives the variation in the SDF, Figure 9 shows the variance decompositions for our broader metatopics. The results reaffirm that most of the jump variation may be attributed to news about *Monetary policy and finance* broadly defined. Interestingly, the metatopic *International affairs*, which comprises news about *China*, the *Middle East*, *North Korea*, and *Russia*, stands out as the overall second most important category of news for explaining the jump variation, although arguably less so for the second half of the sample. By comparison, *Macroeconomic data* on balance appears relatively more important during the second half of the sample, and especially so at the start of the European debt crisis in 2010, and more recently in 2018 during the height of the U.S.-China trade war. The patterns for the *Politics* and the *Commodities and energy* metatopics again fairly closely mirror the patterns previously seen for the conditional jump frequencies for the *U.S. politics* and *Energy markets* topics in Figure 7. Note, however, that this close coherence between the results for some of the topic jump frequencies and the related metatopics is not merely by construction, as the magnitudes of the jumps also figure importantly in the jump variance decompositions. Lastly, foreshadowing our news risk premium estimates discussed next, the remaining three metatopics each account for relatively little of the SDF jump variation.

4. News Risk Premiums

The news risk premiums associated with a particular news topic is naturally defined as the return an investor would be willing to sacrifice to orthogonalize her/his portfolio with respect to the systematic variation stemming from news associated with said topic. Accordingly, we estimate the news risk premiums based on a mimicking portfolio approach, in which we quantify the compensation for exposure to the systematic jumps linked with each of the news topics.

4.1 News risk premium estimation

To begin, define the set of nontradable *topic jump factors*,

$$F_t^{J,k} \equiv F_t^J \cdot TW_{k,t}, \tag{15}$$

where F_t^J denotes the time t tangency portfolio jump return, and $TW_{k,t}$ refers to the fraction of the time t jump "explained" by topic k. As such, if one were able to orthogonalize a portfolio with respect to $F_t^{J,k}$, one would effectively neutralize the portfolio's exposure to any systematic jumps associated with news topic k. Correspondingly, the risk premium associated with topic jump factor $F_t^{J,k}$, say $\lambda_t^{J,k}$, may be interpreted as the jump risk premium for news topic k. Of course, topic jump factors are nontradable. We therefore rely on a standard mimicking portfolio approach to identify and estimate their risk premiums.

In particular, it readily follows from standard asset pricing arguments that

$$\lambda_t^{J,k} = \beta_t^{J,k} \lambda_t^J, \tag{16}$$

where $\beta_t^{J,k}$ denotes the exposure of the topic jump factor to jumps in the tangency portfolio, and λ_t^J refers to the risk premium on tangency portfolio jumps. Our estimation of the jump betas $\beta_t^{J,k}$ is based on "jump regressions" of $F_t^{J,k}$ on F_t^J (see Li, Todorov, and Tauchen 2017, for a more formal discussion of jump regressions). Our approach for estimating the SDF jump risk premium λ_t^J essentially follows the continuous-time Fama-MacBeth approach recently developed by Aït-Sahalia, Jacod, and Xiu (2023).

Empirically, for estimating λ_t^J, we rely on the same 272 high-frequency portfolios used for spanning the SDF as our "test assets." We reestimate the continuous and jump betas for all of these test assets with respect to the SDF on a rolling monthly basis using 1-month and 1-year backward-looking windows, respectively.[29] Armed with the monthly beta estimates, we then estimate separate monthly continuous and jump risk premiums $(\lambda_t^C, \lambda_t^J)$ for the SDF using what effectively amounts to a standard cross-sectional Fama-MacBeth regression approach. Finally, we obtain the requisite risk premium estimate for topic k by averaging $\hat{\beta}_t^{J,k} \hat{\lambda}_t^J$ across all of the months in the sample. A more extensive theoretical discussion of the approach and the practical implementation details are provided in the Internet Appendix.

The above estimation procedure relies critically on the jumps in the SDF for identifying the news topics that are priced. The resultant news risk premium estimates are thus jump-risk specific, ignoring any "continuous" news-related risk compensation. To obtain news risk premium estimates that also take into account the possible compensation manifest in the diffusive variation, we assume that the continuous and jump returns are equally exposed to the news so that the continuous beta for topic factor k equals the jump beta for that same topic factor, or $\beta_t^{C,k} = \beta_t^{J,k}$. In other words, as discussed more formally in the Internet Appendix, the relative contributions of each topic toward "small" and "large" moves in the SDF are assumed to be the same. This assumption in turn

[29] Our use of a longer estimation window for the jump betas to account for the additional estimation error uncertainty mirrors Aït-Sahalia, Jacod, and Xiu (2023) and Aleti (2024).

allows for the estimation of a combined risk premium for news topic k based on the relation:

$$\beta_t^{C,k}\lambda_t^C + \beta_t^{J,k}\lambda_t^J = \beta_t^{J,k}(\lambda_t^C + \lambda_t^J) = \beta_t^{J,k}\lambda_t, \tag{17}$$

where λ_t denotes the total risk premium on the tangency portfolio. This risk premium may, of course, be estimated directly by averaging the SDF tangency portfolio returns over some nontrivial time interval. The resultant topic news risk premium estimates obtained by averaging $\hat{\beta}_t^{J,k}\hat{\lambda}_t$ across all of the months in the sample naturally exceed the previously defined estimates based on averaging $\hat{\beta}_t^{J,k}\hat{\lambda}_t^J$. To differentiate these more inclusive estimates from the earlier jump risk premium estimates, we will refer to the latter as the "total" news risk premium in the sequel. We also rely on the same type of calculations for quantifying the analogous more broadly defined metatopic news risk premiums.

The assumption that $\beta_t^{C,k} = \beta_t^{J,k}$ underlying our total news risk premium estimates is not directly testable since, again, we cannot reliably determine the news that drives the "small" moves in the SDF. However, estimated jump betas generally tend to be fairly close to their continuous counterparts (see, e.g., Zhang et al. 2022). Indeed, estimating the jump and continuous betas with respect to the tangency portfolio for our 272 test/span assets over the full 25-year sample period, the correlation between the two different betas equals 88.9%.

4.2 News risk premium estimates

We begin our discussion by considering the individual news topic risk premiums. Table 2 reports the resultant estimates, together with the t-statistics in parentheses for testing whether the premiums are statistically significantly different from zero.[30] Since the topic risk factors are constructed from the jumps in the tangency portfolio, the magnitudes of the estimated premiums are directly proportional to the average return on the tangency portfolio. Hence, to facilitate the interpretation of the results, we normalize the return on the tangency portfolio to 100% per annum, so that the numbers directly reveal the fraction associated with each of the different news topics. The table shows both the jump and the total news topic risk premiums for the top-20 most important news topics sorted by their jump risk premiums.

Consistent with our earlier findings pertaining to the news that causes the SDF to jump, *Monetary policy* commands the largest (normalized) jump risk premium of 18.55% ($t = 2.00$). In other words, close to one-fifth of the return earned on the tangency portfolio may be attributed to exposure to *Monetary policy* related jump risk. Moreover, if we assume that the continuous and jump

[30] These t-statistics implicitly treat the SDF as given. Explicitly accounting for the first-step estimation error in the SDF would entail a nontrivial extension of the continuous-time Fama-MacBeth approach of Aït-Sahalia, Jacod, and Xiu (2023) formally underlying the inference.

Table 2
Topic risk premiums

Topic	Jump		Total	
Monetary policy	18.55%	(2.00)	21.01%	(4.29)
Energy markets	7.14%	(1.72)	8.38%	(4.58)
U.S. politics	6.75%	(1.92)	9.27%	(4.79)
Middle East	6.73%	(3.04)	4.40%	(3.41)
Labor markets	4.47%	(2.44)	3.73%	(4.58)
Russia	3.87%	(1.63)	3.69%	(3.01)
Government spending, deficits, and debt	3.48%	(3.64)	1.30%	(2.01)
Real estate markets	3.41%	(1.90)	3.55%	(3.95)
Elections and political governance	3.04%	(2.39)	3.43%	(4.22)
Natural disasters	3.04%	(3.64)	1.57%	(3.43)
Inflation	2.79%	(1.21)	4.70%	(3.10)
Litigation matters	2.76%	(2.32)	1.80%	(3.46)
North Korea	2.44%	(2.32)	1.47%	(3.03)
Taxes	2.29%	(1.46)	4.62%	(4.33)
Commodity markets	2.16%	(2.21)	2.35%	(4.38)
Entitlement and welfare programs	1.96%	(1.71)	0.95%	(2.05)
Financial regulation	1.88%	(1.35)	1.48%	(2.55)
Broad quantity indicators	1.30%	(1.04)	2.84%	(3.23)
National security	1.28%	(0.91)	2.91%	(3.18)
Consumer spending and sentiment	1.02%	(1.10)	1.89%	(3.58)

The table reports the estimated jump and total risk premiums for each topic factor, with the corresponding t-statistics in parentheses. The estimates are computed from the estimated tangency portfolio rescaled to a return of 100% per annum over the full sample.

returns are similarly exposed to the news, up to 21.01% ($t=4.29$) of the return on the SDF may be traced to news about *Monetary policy*. This echoes several other recent studies which suggest that much of the equity risk premium is earned around the time of FOMC announcements (e.g., Savor and Wilson 2013; Lucca and Moench 2015; Cieslak, Morse, and Vissing-Jorgensen 2019). By comparison, risks associated with news about *Energy markets* and *U.S. politics* account for 8.38% and 9.27% of the overall return on the tangency portfolio, respectively.

The relative importance of the different topic risk premiums also adheres fairly closely, although not perfectly, to the ordering of the jump frequencies and variance decompositions discussed earlier. This, of course, is not surprising as one would naturally expect that topics that account for most of the jumps and the jump variation in the SDF also demand the largest compensation. It is also noteworthy that even though the estimated risk premiums for many of the less important topics are quite small, most of the estimates are still statistically significant at conventional levels.

Turning to the estimates for our more broadly defined metatopics reported in Table 3, the results imply that more than a quarter of the return on the tangency portfolio stems from risk related to news about *Monetary policy and finance*.[31] This, of course, is entirely in line with the more nuanced results in Table 2, and the finding that *Monetary policy* alone accounts for slightly

[31] Since some of the news topics appear in more than one metatopic, the total risk premium estimates reported in Table 3 based on the topic estimates in Table 2 do not exactly add up to 100%.

Table 3
Metatopic risk premiums

Metatopic	Jump		Total	
Monetary policy and finance	24.45%	(1.85)	28.97%	(4.07)
Macroeconomic data	14.19%	(1.61)	21.31%	(4.28)
International affairs	12.53%	(2.24)	12.60%	(4.28)
Politics	9.79%	(2.23)	12.70%	(5.04)
Commodities and energy	9.30%	(1.92)	10.73%	(4.89)
Fiscal policy	7.73%	(2.75)	6.87%	(4.04)
Labor	5.58%	(2.61)	4.30%	(4.58)
Regulation	4.34%	(2.15)	2.84%	(3.26)

The table reports the estimated annualized jump and total risk premia for each of the metatopic factors, constructed by summing the topic factors for all the component topics within each metatopic. The estimates are computed from the tangency portfolio rescaled to a return of 100% per annum over the full sample. The t-statistics are reported in parentheses.

more than one-fifth of the total risk premium. Table 3 also shows that news about *Macroeconomic data*, broadly defined, plays a very important role in explaining the total risk premium. This result adheres with a long list of prior studies documenting large (in an absolute value sense) stock market returns in response to macroeconomic news announcements – see, e.g., the early studies by Pearce and Roley (1985); French and Roll (1986); Andersen et al. (2003, 2007), along with the more recent work by Gürkaynak, Kısacıkoğlu, and Wright (2020), among others.[32] As previously noted, news about *International affairs*, *Politics*, and *Commodities and energy* also all account for a nontrivial portion of the total risk premium. Even though the remaining three metatopics on average appear somewhat less important, the corresponding total news risk premium estimates are still strongly statistically significant when judged by their individual t-statistics.[33]

The metatopic risk premium estimates reported in Table 3 are all based on the full-sample averages of the monthly $\hat{\beta}_t^{J,k'} \hat{\lambda}_t^J$ estimates. This obviously masks any temporal variation in the compensation for exposure to the different news topics. Meanwhile, the metatopic variance contributions previously discussed in Figure 9 clearly suggest that the relative importance of the different types of news varies over time. To this end, Figure 10 plots the cumulative daily returns for the top-three (based on their total risk premium estimates) metatopic mimicking portfolios. To allow for easier interpretation and make the returns comparable in magnitude to those of the market portfolio, we rescale the tangency portfolio returns to 10% per annum. In addition, to help better

[32] A theoretical explanation for the importance of macroeconomic announcements, rooted in revealed preference theory and the idea that the announcements provide important information about the prospects of future economic growth, has also recently been developed by Ai and Bansal (2018). Relatedly, Ai et al. (2022) propose an equilibrium-based model for the cross-sectional pricing of FOMC announcements.

[33] This significance also easily "survive" a standard Bonferonni-type correction for multiple testing at the 5% level. The precision of the total risk premium estimates is driven by the fact that the tangency portfolio, which has a high Sharpe ratio, is used to construct the mimicking portfolios, while the topic jump risk premium estimates tend to be relatively more imprecise due to the estimation error associated with the continuous-time Fama-MacBeth procedure.

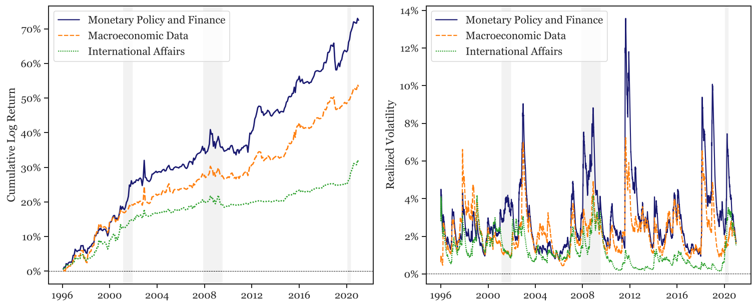

Figure 10
Metatopic risk premiums over time
The left subplot shows the cumulative returns on the *Monetary policy and finance*, *Macroeconomic data*, and *International affairs* metatopic-mimicking portfolios computed from the tangency portfolio returns rescaled to 10% per annum over the full sample. The right subplot shows a 30-day EWMA of the realized volatility of the same three portfolios. The shaded regions correspond to NBER-defined recessions.

understand the time-varying risks embedded in the portfolios, the right subplot shows the realized volatility of the same three portfolios (computed as a 30-day EWMA).

Looking at the first subplot, the cumulative returns on the *Monetary policy and finance* mimicking portfolio steadily increased over most of the sample, reinforcing the idea that investors generally put a high premium on the risks associated with said news. Meanwhile, there also appears to be a substantial amount of variation in the corresponding return volatility, with notable peaks in 2003, when the market surged, in 2007–2008 associated with the Global Financial Crisis, in 2012 and the time of the European Sovereign Debt Crisis, and from 2017 to the end of the sample manifesting both political and COVID-19-induced uncertainty. In line with a traditional risk-return trade-off relationship, these periods of heightened economic and financial market uncertainty are generally also associated with a higher premium on *Monetary policy and finance* related news.

The cumulative returns and the volatility for the *Macroeconomic data*-mimicking portfolio, not surprisingly, evidence fairly similar overall patterns, underscoring the heightened importance of new economic data and indicators at various points in time. Putting these results further into perspective of the existing literature, a number of studies have previously documented that the stock market tends to react differently to macroeconomic news announcements in recessions compared to expansions (see, e.g., McQueen and Roley 1993; Boyd, Hu, and Jagannathan 2005; Andersen et al. 2007). More recently, Law, Song, and Yaron (2020) have argued that the reaction to macroeconomic news is closely tied to investors' expectations about the likelihood of the Fed tightening its policies, while Schmeling and Wagner (2024) and Gardner, Scotti, and Vega (2022) emphasize the importance of tone and sentiment, respectively. The temporal variations in the returns and volatility observed for the *Macroeconomic data*-mimicking portfolio further corroborate these ideas.

Turning to the results for the *International affairs* mimicking portfolio, much of the return was evidently earned during the first few years of the sample, coincident with the Asian Financial Crisis, the Russian Financial Crisis, and the immediate aftermath of 9/11. The run-up around 2003, along with the moderate returns throughout most of the 2000s, may naturally be ascribed to the War on Terror and the 2003 invasion of Iraq. The return and volatility also increased in 2020 at the start of the global pandemic when investors again put a higher premium on international news. Similar plots for the five other metatopic mimicking portfolios are included in Appendix C.2. Of these, perhaps the ones for *Commodities and energy* and *Politics* stand out with their own most clearly distinct patterns, highlighting the different economic news that investors are most concerned about, and thus carry the largest news risk premium, at different points in time.

4.3 News risk premiums in the factor zoo

The significance and economic motivation for the myriad of risk factors proposed in the asset pricing literature continues to be an area of active debate (see, e.g., Harvey, Liu, and Zhu 2016; Hou, Xue, and Zhang 2020; Jensen, Kelly, and Pedersen 2023). Relatedly, the success of the Fama-French model has been attributed to the ability of the corresponding characteristic-based factors to capture certain macroeconomic risks and innovations in state variables naturally associated with changes in the investment opportunity set (see, e.g., Vassalou 2003; Vassalou and Xing 2004; Petkova 2006; Aretz, Bartram, and Pope 2010, among others). As an alternative look at the risks actually priced by the factor zoo, we now decompose factor risk premiums into components associated with each of our news topics.

To understand the basic approach that we use for estimating the factor news risk premiums, consider a zero-cost investment portfolio, or long-short factor j. We write the time-t risk premium on the SDF tangency portfolio as λ_t, with the contributions stemming from news topic k denoted by λ_t^k. The time-t risk premium for factor j may then naturally be decomposed as:

$$\mu_t^j = \beta_t^j \lambda_t = \beta_t^j \cdot \sum_k \lambda_t^k, \tag{18}$$

where β_t^j denotes the usual factor loading, or exposure, of the factor with respect to the tangency portfolio and the equality of λ_t and $\sum_k \lambda_t^k$ is guaranteed by the inclusion of the news topic *None*. The factor loading β_t^j is readily estimated by a standard time-series regression of the factor returns on the tangency portfolio returns. In the results reported below, we rely on the high-frequency 15-minute returns over rolling monthly windows for this estimation. Following the discussion in connection with Equation (17) in Section 4.1, we similarly estimate the λ_t^k's on a rolling monthly basis by $\hat{\beta}_t^{J,k} \hat{\lambda}_t$. Combining the resultant $\hat{\beta}_t^j$ and $\hat{\lambda}_t^k$ estimates, our full-sample factor-specific topic risk premium for factor j and news topic k is in turn obtained by averaging $\hat{\beta}_t^j \hat{\beta}_t^{J,k} \hat{\lambda}_t$ over all of the months in the sample. This approach closely mirrors a traditional mimicking portfolio approach for the estimation of factor risk premium, except for the rescaling by $\hat{\beta}_t^{J,k}$. This additional rescaling stems from Equation (17), and the decomposition of the total risk premium on the tangency portfolio into the different news topics.

To help more concisely convey the results, rather than report the estimates for all of the individual factors, we focus on a set of thirteen representative factor cluster portfolios. Our definitions of the factor clusters follow that of Jensen, Kelly, and Pedersen (2023) and Aleti (2024), with the returns on the factor cluster portfolios constructed as the average returns on the factors within each of the clusters, rescaled to match the volatility of the Fama-French market portfolio. Table 4 reports the resultant full-sample metatopic risk premium estimates.[34] For comparison, we also include the news risk premium contributions for the Fama-French market portfolio in the first row of the table.

The relative importance of the different topics for the risk premium on the market portfolio fairly closely mirrors the total metatopic risk premium estimates for the tangency portfolio previously reported in Table 3. Interestingly, however, compared to the estimates for the tangency portfolio, news about *International affairs* commands a relatively larger market risk premium. In parallel to the findings for the market portfolio, news about *Monetary policy and finance* is also important for explaining the risk premium on many of the factors, with the estimates for the Profitability, Quality, Profit Growth, and Debt Issuance factors all statistically significant at the usual 5% level. This finding is naturally explained by shocks to interest rates affecting the characteristics underlying these factors. News about *International affairs* and *Macroeconomic data* likewise appear to be important for many of the factors, the aforementioned factors included, supporting the idea that the premiums are compensating for exposures to broader economic conditions and systematic risks embedded in the factors. Consistent with the relatively lower risk premium estimates for the *Fiscal policy*, *Labor*, and *Regulation* metatopics for both the tangency and the market portfolio, these same news topics are generally also less important for explaining the factor risk premiums. Perhaps not surprisingly, the risk premiums for the Seasonality and Skewness factor cluster portfolios are the least affected by any of the economic news, indirectly suggesting that these factors may be capturing mispricing rather than systematic risks (see also Keloharju, Linnainmaa, and Nyberg 2021).

[34] To help preserve space, we defer the results for the individual topic risk premiums to the Internet Appendix.

Table 4
Metatopic risk premia for factor cluster portfolios

	Monetary policy and finance	International affairs	Macroeconomic data	Politics	Commodities and energy	Fiscal policy	Labor	Regulation
Market	1.22	1.26	0.64	0.70	0.54	0.37	0.23	0.19
	(1.48)	(2.40)	(2.53)	(2.39)	(2.67)	(2.16)	(2.42)	(1.93)
Value	−1.34	−1.04	−0.62	−0.81	−0.56	−0.47	−0.26	−0.20
	(−2.65)	(−2.86)	(−2.38)	(−2.66)	(−2.46)	(−2.35)	(−2.97)	(−2.42)
Investment	−0.51	−0.52	−0.34	−0.46	−0.33	−0.27	−0.18	−0.12
	(−0.94)	(−1.44)	(−1.29)	(−1.77)	(−1.59)	(−1.37)	(−2.27)	(−1.68)
Low risk	−0.35	−0.45	−0.39	−0.53	−0.15	−0.30	−0.07	−0.16
	(−0.57)	(−1.13)	(−1.40)	(−1.63)	(−0.75)	(−1.47)	(−0.83)	(−1.58)
Profitability	0.90	0.82	0.29	0.28	0.40	0.19	0.20	0.03
	(2.07)	(2.42)	(1.47)	(1.16)	(2.54)	(1.39)	(2.71)	(0.44)
Quality	1.64	1.16	0.56	0.63	0.66	0.35	0.29	0.09
	(2.44)	(2.75)	(2.49)	(2.01)	(3.04)	(2.94)	(3.04)	(0.79)
Leverage	1.01	0.91	0.56	0.75	0.42	0.41	0.20	0.19
	(1.80)	(2.40)	(1.86)	(2.42)	(2.03)	(1.81)	(2.57)	(2.43)
Momentum	1.08	0.70	0.37	0.40	0.49	0.33	0.20	0.06
	(1.86)	(1.94)	(1.50)	(1.40)	(2.33)	(2.12)	(2.58)	(0.69)
Size	−0.71	−0.84	−0.37	−0.47	−0.40	−0.35	−0.21	−0.08
	(−1.82)	(−2.73)	(−1.79)	(−2.21)	(−2.90)	(−2.32)	(−3.65)	(−1.44)
Profit growth	1.40	0.74	0.49	0.38	0.41	0.45	0.19	0.06
	(2.44)	(2.04)	(2.19)	(1.32)	(2.23)	(2.90)	(2.42)	(0.73)
Accruals	0.55	0.46	0.40	0.57	0.34	0.20	0.07	0.08
	(1.15)	(1.41)	(2.18)	(2.16)	(2.05)	(1.61)	(1.23)	(1.32)
Debt issuance	1.60	1.21	0.69	0.77	0.51	0.55	0.23	0.18
	(3.26)	(3.33)	(3.52)	(3.51)	(4.02)	(3.53)	(3.65)	(2.61)
Skewness	−0.26	−0.17	−0.17	−0.32	−0.18	−0.19	−0.09	−0.13
	(−0.52)	(−0.56)	(−0.75)	(−1.33)	(−1.03)	(−1.14)	(−1.16)	(−1.98)
Seasonality	−0.41	−0.56	−0.32	−0.46	−0.22	−0.21	−0.11	−0.15
	(−0.76)	(−1.70)	(−1.69)	(−1.97)	(−1.50)	(−1.68)	(−1.62)	(−2.15)

The table reports the estimated percentage annualized returns for each of the thirteen representative factor cluster portfolios that are are earned from exposure to each of the eight metatopics. The definitions of the factor clusters follow Jensen, Kelly, and Pedersen (2023) and Aleti (2024), while the cluster portfolios are constructed as the average of the factor returns within each cluster; these cluster portfolios have also been rescaled to match the realized volatility of the market portfolio. The first row reports the corresponding risk premia contributions for the Fama-French market portfolio. The t-statistics, computed using the realized volatilities of the mimicking portfolios, are reported in parentheses.

5. Economic Mechanisms and Monetary Policy News

To help further elucidate the economic mechanisms at work and what sets the different news topics apart, it is instructive to associate the jumps in the SDF with a set of more easily interpretable and readily identifiable economic shocks. The composition of these shocks could possibly also be used to help differentiate between existing asset pricing models, and in the development of new more empirically realistic models. In a similar vein, we demonstrate that restricting the analysis of monetary policy news to FOMC and other scheduled central bank communications misses a substantial portion of the news that is actually priced.

5.1 Economic news classifications

We rely on the methodology proposed by Cieslak and Schrimpf (2019), to designate all the intraday jumps in the SDF as growth, risk premiums, short-rate, or long-rate shocks, based on the high-frequency comovements of the SDF and interest rates, together with the volatility of yields across different maturities. More specifically, we begin by calculating the realized covariances between the SDF returns and the yields in 45-minute windows surrounding each jump event. We rely on finer 1-minute returns on the SDF, imputed from the estimated 15-minute tangency portfolio as discussed in Section 2.3, together with 1-minute returns on 2-year Treasury-bond futures to do so.[35]

We define all the SDF jumps for which the stock-yield covariance is positive as either growth or risk premium shocks, while we consider the jumps where the stock-yield covariance around the time of the jump is negative as being associated with interest rate shocks.[36] Intuitively, a positive/negative growth shock naturally results in higher/lower future cash flows, thereby affecting equity valuations and yields in the same direction. Along the same lines, a positive/negative risk premium shock is associated with a higher/lower degree of risk aversion, or lower/higher risk appetite, and a flight-to-safety similarly causing stock returns and yields to covary positively. Conversely, negative high-frequency stock-yield covariances naturally arise from positive/negative interest rate shocks that cause the discounted present values of future cash flows and thus the values of equities to fall/rise. Again, closely following the conceptual framework of Cieslak and Schrimpf (2019) and the idea that risk premium shocks manifest more strongly over longer horizons, we further categorize the positive covariance shocks into either growth or risk premium shocks, depending on whether the realized variance of the 2-year Treasury-bond futures returns is higher (growth) or lower (risk premium)

[35] The high-frequency data for the Treasury-bond futures are obtained from Tick Data Inc.

[36] This approach closely follows that of Cieslak and Schrimpf (2019), except we rely on the high-frequency return on the tangency portfolio in place of the market portfolio, and 45-minute windows around the jump returns instead of the exact timestamps of the events.

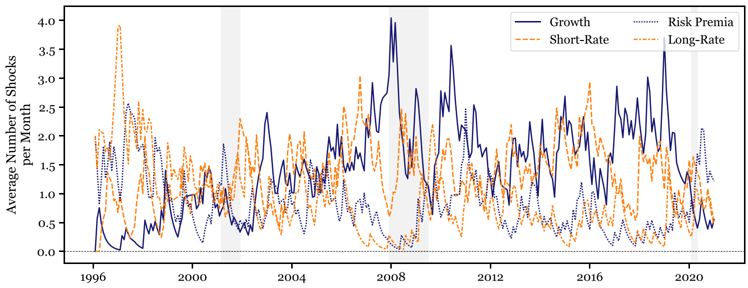

Figure 11
Average jump count by classification
The figure shows the average monthly (3-month EWMA) number of growth, risk premium, short-rate, and long-rate shocks associated with the SDF jumps, based on the shock classification scheme discussed in the main text.

than the realized variance of the 10-year Treasury-bond futures returns. Our categorization into so-called "short- and long-term interest rate shocks" is based on the same realized volatility comparisons.

This easy-to-implement classification scheme is invariably somewhat stylized. Nonetheless, applying the approach to all the SDF jumps observed throughout the sample period, results in a total of 427 growth, 243 risk premiums, 385 short-rate, and 260 long-rate economic shocks. As evidenced by the monthly averages depicted in Figure 11, growth and short-rate shocks tend to dominate during most of the sample, while the relative importance of risk premiums and long-rate shocks generally increase during periods of heightened economic and financial uncertainty.

Next, to help illuminate the news that tends to affect the different economic shocks, we combine the above classification scheme with our textual-based news topic classification. We begin, by determining the "dominating" metatopic for each of the SDF jumps, as the one with the largest topic weight over the relevant jump time interval. In the case of a tie, we simply split the allocation accordingly. Also, in the few instances without metatopic-relevant news, we simply omit that jump from the sample. Armed with this additional classification, we connect the metatopics for each of the jumps with the corresponding economic shocks.

The top panel in Table 5 reports the resultant joint frequencies over all of the jumps in the sample, together with the corresponding 95% confidence intervals. Not surprisingly, looking at the jumps associated with the *Monetary policy and finance* metatopic, we see only short-rate shocks manifest at a statistically significant higher rate, as judged by the average row frequency falling below the corresponding confidence interval. On the other hand, looking at the average frequencies across each of the rows for the *Macroeconomic data*, *Politics*, *Commodities and energy*, and *Labor* metatopics they are all associated with significantly higher proportions of growth shocks.

Table 5
Metatopic and economic shock classifications

A. Joint frequencies (% of jumps)

	Growth	Short-rate	Risk premiums	Long-rate
Monetary policy and finance	6.75	8.29	4.25	4.42
	(5.52, 7.98)	(6.91, 9.67)	(3.23, 5.27)	(3.41, 5.44)
International affairs	5.83	5.58	3.92	3.76
	(4.65, 7.01)	(4.45, 6.72)	(2.98, 4.87)	(2.77, 4.74)
Macroeconomic data	6.68	5.56	3.60	4.07
	(5.41, 7.94)	(4.44, 6.68)	(2.67, 4.52)	(3.09, 5.06)
Politics	5.39	4.36	2.36	3.09
	(4.19, 6.59)	(3.28, 5.43)	(1.55, 3.18)	(2.20, 3.97)
Commodities and energy	4.04	3.27	2.21	1.76
	(3.03, 5.06)	(2.33, 4.21)	(1.45, 2.97)	(1.06, 2.46)
Fiscal policy	1.58	1.60	0.95	1.74
	(0.92, 2.25)	(0.92, 2.27)	(0.43, 1.46)	(1.03, 2.46)
Labor	1.07	0.27	0.40	0.08
	(0.61, 1.53)	(0.01, 0.53)	(0.13, 0.67)	(−0.03, 0.19)
Regulation	1.12	0.84	0.56	0.59
	(0.58, 1.67)	(0.37, 1.30)	(0.17, 0.94)	(0.20, 0.97)

B. Total risk premiums (Annualized %)

	Growth	Short-rate	Risk premiums	Long-rate
Monetary policy and finance	10.24	5.45	4.98	8.30
	(2.68)	(2.77)	(3.26)	(2.88)
International affairs	2.13	2.80	3.39	4.28
	(1.80)	(3.47)	(4.44)	(4.58)
Macroeconomic data	8.16	5.11	4.06	3.98
	(3.50)	(3.86)	(3.09)	(2.22)
Politics	3.33	2.90	3.14	3.34
	(3.10)	(4.26)	(4.17)	(4.38)
Commodities and energy	2.33	2.47	2.29	3.64
	(2.98)	(3.62)	(3.71)	(4.16)
Fiscal policy	1.37	1.40	1.89	2.21
	(2.40)	(4.53)	(2.59)	(2.42)
Labor	2.50	0.75	0.49	0.55
	(4.25)	(2.66)	(3.12)	(2.94)
Regulation	1.40	0.33	0.35	0.75
	(2.54)	(1.51)	(3.53)	(2.29)

Panel A reports the percent of jumps in the tangency portfolio for the metatopic-economic-shock two-way classification scheme discussed in the main text. The 95% confidence intervals reported in parentheses are computed by bootstrap. Panel B reports the estimated total risk premia, with t statistics in parentheses, for the same two-way classification scheme based on the mimicking portfolio approach discussed in the main text. All of the risk premium estimates are reported in annualized percentage form.

In an effort to more concretely quantify the economic significance of these differences, we further construct “metatopic-shock” factors by interacting the metatopic weights with indicator variables for each of the economic shock classifications. In parallel to the news risk premium estimation detailed in Section 4.1, we then estimate the total risk premiums for each of the metatopic-economic-shock classifications, in essence decomposing the total risk premium estimates previously reported in Table 3 into additive economic shock premiums. To facilitate interpretation, and in parallel with the earlier results, we again report the estimates as a fraction of the total return on the tangency portfolio.

Looking at the bottom panel of Table 5, we see the risk premium estimates broadly align with the joint frequencies reported in the top panel.

However, there are also some notable differences. For instance, despite the relatively lower occurrence of long-rate shocks driven by *Monetary policy and finance* news, the risk premium for such shocks is the second largest of all the premiums, with the premium for growth shocks associated with the same type of news being the largest. In other words, although short-rate shocks are the most common type of economic shocks affected by *Monetary policy and finance*-related news, investors put a higher premium on long-rate shocks, possibly because of concerns about forward guidance and/or unconventional policy interventions. For *Macroeconomic data*, *Labor*, and *Regulation*, the majority of the risk premiums also seem to be driven by growth shocks, consistent with the idea that news related to these metatopics convey information about the current state of the economy and/or future economic conditions. On the other hand, news about *International affairs*, *Commodities and energy*, and *Fiscal policy* draw the largest compensation from long-rate shocks, although the differences across the premiums are relatively small in magnitude.

To more explicitly clarify the actual news that drives the different economic premiums, the Internet Appendix provides example headlines for each dominant metatopic and shock classification pair, as well as a sample of headlines for all of the different shock classification pairs. As these specific examples show, information about the current or future state of the economy or specific economic events often drive growth shocks. Meanwhile, short- and long-rate shocks are naturally associated with news about financial markets and interest rates stemming directly from monetary policy, or information about government spending and/or economic policies. Risk premium shocks often arise from news that affects uncertainty when broadly defined. Consistent with prior work on monetary policy-driven uncertainty (see, e.g., Bekaert, Hoerova, and Lo Duca 2013; Husted, Rogers, and Sun 2020), communications by Federal Reserve officials also sometimes generate risk premium shocks.

5.2 A closer look at monetary-policy-related news

Our main empirical analyses have established that the largest portion of the topic-based variation in the SDF stems from news associated with the *Monetary policy* topic. In contrast to much of the existing literature concerned with the role and financial market impact of monetary policy, which has typically focused on FOMC and other scheduled central bank announcements (e.g., Bernanke and Kuttner 2005; Lucca and Moench 2015), our text-based analyses incorporate the possible effects of nonscheduled monetary-related news. Extending the discussion of the type of economic shocks that drive our results, we show that such off-calendar monetary news accounts for a substantial fraction of the *Monetary policy* news risk premium. This echoes recent work by Bianchi, Ludvigson, and Ma (2023) and Cieslak and McMahon (2023), who have similarly emphasized the role of nontraditional channels for communicating monetary policy by the Fed.

We begin by constructing a traditional "monetary policy calendar" as a benchmark for events known by the literature. Following previous work, as further detailed in the Internet Appendix, this calendar naturally comprises FOMC Decisions and Statements, FOMC Minutes, Economic Data Reports and Projections, as exemplified by the Beige Book and similar reports, and Other Fed Events, including scheduled meetings and various press conferences. We then compute the proportion of the jump variation in the SDF that is explained by each of these event categories. Since some of the categories have overlapping events, we also consider an all-encompassing category that we refer to as "Any," marking intervals with any of the above specified events.

As the more detailed results reported in the Internet Appendix show, jumps associated with this combined category explain 18.2% of the total jump variation in the SDF. Restricting the set of jumps to those where the topic weights on the *Monetary policy* topic exceed 25% increases that proportion to 31.5%. Further limiting the set of jumps to those where the Fed or FOMC was explicitly mentioned in one or more of the headlines over the relevant 15-minute time intervals further increases that proportion to 47.1%. On the one hand, these findings show that nearly half the jump variation around this subset of jumps, purposefully constructed to capture Fed/FOMC-induced variation, can be attributed to traditional calendar events. On the other hand, this means that more than half of the systematically important Fed and FOMC-related news events is not captured by the traditional monetary policy calendars.

Going one step further, to more precisely quantify the economic significance of these "missing" events, we decompose the topic risk premiums of the *Monetary policy* topic depending on whether or not the underlying news is included in our monetary policy calendar. To do so, we break our *Monetary policy* topic risk factor F_t^{MP} into two separate additive factors for calendar and noncalendar news, that is $F_t^{Any} = F_t^{MP} \times \mathbb{1}_{[t \in Calendar]}$ and $F_t^{NotAny} = F_t^{MP} \times \left(1 - \mathbb{1}_{[t \in Calendar]}\right)$. Estimating the corresponding risk premiums as in Section 4.2, we find that the total topic risk premium for *Monetary policy*, previously estimated to be 21.01% ($t = 4.29$), decomposes into a total topic risk premium for the *Any* factor equal to 4.90% ($t = 3.73$) and a premium for the *NotAny* factor equal to 16.11% ($t = 3.69$). In other words, the monetary policy calendars employed in a number of previous studies likely miss many systematically important monetary news events that are priced by investors.

To alleviate concerns that much of this "missing" risk premium arises from events that are essentially unrelated to the Fed or FOMC, we construct a *FedMentioned* topic factor based on the interaction between the *Monetary policy* topic factor and an indicator variable for whether the Fed or FOMC was explicitly mentioned in one or more of the news headlines associated with the relevant SDF jump. The total risk premium for this factor is 14.27% ($t = 3.73$), indicating that Fed or FOMC news is indeed driving much of the *Monetary policy* topic risk premium. Yet, in light of the risk premium estimated for the *Any* topic factor, it appears that much of this very same news is missing from the traditional monetary policy calendars, further highlighting the advantages of our high-frequency text-based approach for successfully identifying the news that matters.[37]

6. Conclusion

We exploit high-frequency data and real-time economic news to provide a novel characterization of systematic financial market risks. Our approach sidesteps the need for an explicit model for the stochastic discount factor, relying instead on a large panel of high-frequency portfolio returns combined with a minimax method of moments approach to robustly recover the tangency portfolio under minimal assumptions. By directly linking the jumps in the estimated tangency portfolio with the textual information in a comprehensive collection of precisely timed news articles, we are able to explicitly identify the type of news that matters to investors. Grouping the news articles into intuitive and interpretable categories of news topics, we find that *Monetary policy*, *U.S. politics*, and *Energy markets* stand out as the overall most important news topics for explaining the variation in the tangency portfolio returns.

To further address the economic significance of the news, we employ a mimicking portfolio approach, allowing us to decompose the risk premium on the tangency portfolio into separate components associated with each of the different news topics. Consistent with other recent studies emphasizing the importance of FOMC announcements for explaining the return on the market portfolio, news about *Monetary policy* again emerges as the overall most salient news topic, explaining more than 30% of the tangency portfolio risk premium. Further combining the news articles into more broadly defined metatopics, we find that news related to *Monetary policy and finance* explains close to 40% of the tangency portfolio returns. Importantly, this includes a myriad of nonscheduled monetary news beyond FOMC announcements and other scheduled central bank communications, that our text-based news classification procedure identifies as being significant. Extending our procedure to allow for the decomposition of the news risk premiums on other assets, we also shed new light on the type of news that accounts for the risk premiums earned by the zoo of different factors proposed in the asset pricing literature. In parallel to the results for the tangency portfolio, news related to *Monetary policy and finance* again emerges as the overall most important news topic for explaining most of the factor risk premiums. At the same time, however, our results also reveal quite distinct news risk premium contributions for different factor cluster portfolios. Based on the joint high-frequency response of the SDF and short- and long-term interest rates, we further attribute the different channels behind

[37] The Internet Appendix also provides a series of concrete examples of this “missing” noncalendar news for which the *Monetary policy* topic weight exceeds 25% and the Fed or FOMC is explicitly mentioned in the headlines.

the news to economic growth, risk premium, and monetary policy shocks broadly defined.

The approach developed here could similarly be used to help illuminate the news that drives the systematic risks and returns for other financial assets. It may also be used for risk management purposes to help devise portfolios immune to certain types of news. Relatedly, the results could also be used in the development of active factor timing strategies based on specific types of scheduled news announcements. Going one step further, the detailed characterization of the news that is priced and the underlying economic shocks and mechanisms at work, could possibly also be used to help differentiate between competing asset pricing models. Following Liu and Matthies (2022), it is also possible that measures of news sentiments could be employed in conjunction with the asset-pricing-relevant news topics detailed here to help predict future economic conditions. We leave further work in these directions for future research.

Code Availability: The replication code and data are available in the Harvard Dataverse at doi:10.7910/DVN/TTKT4A.

Appendix

A SDF Estimation Details

This appendix provides additional details about the practical implementation of the key minimax optimization problem in Equation (10) that underlies our estimation of the high-frequency SDF.

A.1 Functional Forms and Hyperparameters. We rely on the functional forms implicitly defined by:

$$h_t^w \equiv LSTM(I_t; \theta_{w,1}),$$

$$h_t^g \equiv LSTM(I_t; \theta_{g,1}),$$

$$w_t \equiv f_w(I_t; \theta_w) = FFN\left(h_t^w; \theta_{w,2}\right),$$

$$g_t \equiv f_g(I_t; \theta_g) = FFN\left(h_t^g; \theta_{g,2}\right),$$

where the LSTMs denote long short-term memory neural networks with parameters $\theta_{w,1}$ and $\theta_{g,1}$, respectively, and the FFNs denote feedforward neural networks with parameters $\theta_{w,2}$ and $\theta_{g,2}$, respectively (see, e.g., Hochreiter and Schmidhuber 1997, for more formal definitions of the LSTM and FFN type networks). Intuitively, the LSTMs serve to condense the high-dimensional information in I_t into the lower-dimensional state variables, h_t^w and h_t^g, by recursively updating their past values with the relevant new information. The FFNs in turn use the constructed state variables to determine the weights and the instruments. We further restrict the weights for each asset to the $[-1,1]$ interval via "tanh" final layer activations for both networks. Accordingly, the FFN for the weights maps from $\mathbb{R}^{\dim(h^w)}$ to $[-1,1]_N$, while the FFN for the instruments maps from $\mathbb{R}^{\dim(h^g)}$ to $[-1,1]_{N\times N_g}$.

Table A.1 lists the hyperparameters used to define the weight and instrument functions, f_w and f_g.

The variable column indicates whether the specific hyperparameter is being defined for f_w or for f_g while the neural net column refers to the LSTM and FFN components of these two

Table A.1
Hyperparameters for main model

Variable	Neural net	Hyperparameter	Hyperparameter value
Weights	LSTM	Activation	Tanh
Weights	FFN	Layer structure	{[], [2], [4, 2], [8, 4, 2]}
Weights	FFN	Intermediate layer activations	ReLu
Weights	FFN	Final layer activation	Tanh
Weights	FFN	Dropout fraction	5%
Instruments	LSTM	Activation	Tanh
Instruments	FFN	Layer structure	{[], [2], [4, 2], [8, 4, 2]}
Instruments	FFN	Intermediate layer activations	ReLu
Instruments	FFN	Final layer activation	Tanh
Instruments	FFN	Dropout fraction	5%
Instruments	FFN	Output dimension	1000
Both	LSTM	State variable dimension	{2,4,8}

This table reports the hyperparameter grid for our SDF estimation procedure. Our estimation relies on two neural networks, one for estimation of the weights and another for estimation of the instruments. Each neural network itself consists of two other neural network, an LSTM that generates a set of state variables and a FNN that takes nonlinear transformations thereof to form the output. Each of these neural networks then accepts a set of hyperparameters. The choices for each are given in the final column.

variable networks. We purposely strive for simplicity when defining the hyperparameters, generally using the same values for both the weights and instruments. Additionally, to reduce the number of hyperparameters that need to be tuned, we refer to Chen, Pelger, and Zhu (2024) when setting default values; for instance, the intermediate layers are set to use so-called "ReLU activations" and dropout fractions of 5%, matching CPZ. However, for the output dimension of the instruments, we deviate from CPZ and set this value to $1,000$. This choice still reflects a fairly large set of generated assets while, critically, respecting computational constraints. Lastly, there are three hyperparameters to tune as marked by asterisks in the table: the state variable dimensions and the FFN layer structures used in the two networks. To reduce the computational burden, we jointly tune the state variable dimension for the weights and instruments. In total, this implies a grid of $4\times4\times3=36$ hyperparameters. The tuning itself is done through cross-validation as discussed in the main text.

A.2 Sharpe Ratio Bound. We also deliberately bound the Sharpe ratios for each of the individual SDF estimates to lie between 0.4 and 1.5. These particular bounds are directly motivated by Kozak, Nagel, and Santosh (2020) and the empirical analyses therein demonstrating that the imposition of similar priors on the Sharpe ratio results in SDF estimates with improved out-of-sample cross-sectional explanatory power. Our individually estimated SDFs also easily converge to tangency portfolios with Sharpe ratios within these bounds.[38] We enforce this constraint in the optimization procedure using the method of Lagrange multipliers, basically adding a penalty to the objective function defined by (9) and (10):

$$\left\|\frac{1}{N_g\cdot T}\sum_t \hat{\alpha}_{g,i}^2\right\|_2 + \underbrace{c\cdot\max\{S_{lower}-\hat{S},0,\hat{S}-S_{upper}\}}_{\text{Sharpe Penalty}}. \tag{A.1}$$

Scaling the penalty to $c=10\%$ is sufficient to enforce the constraint in practice, as evidenced by all of our fitted models achieving a Sharpe ratio within the $S_{lower}=0.4$ to $S_{upper}=1.5$ bounds. In other words, the penalty essentially works as a hard constraint.

[38] Note, that as a formality, by reducing the estimation error, and in turn the in-sample return volatility, the in-sample Sharpe ratio for our ensemble average SDF estimate may actually exceed the upper bound.

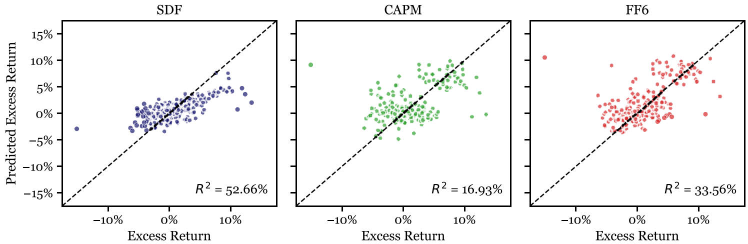

Figure A.1
Predicted versus actual returns
The three subplots report the predicted versus actual realized returns for each of our 272 test assets based on our estimated SDF, the CAPM, and the FF6 model. The predicted returns are computed using full-sample time-series regressions between each factor model and each of the test assets. The outlier seen on the far left of all three plots is the Fama-French coal industry portfolio, corresponding to SIC codes 1200–1299, which has seen abnormally low returns over our sample. All returns are reported in annualized form. The annotations give the uncentered R^2 between the predicted and the realized returns.

A.3 Numerical Optimization. With our hyperparameters and functional approximations defined above, we proceed to solve the minimax objective in Equation (A.1) in three steps. We begin by determining the initial set of weights $\{w_t\}$ that approximately minimize the unconditional and conditional alphas. By targeting both the monthly and full sample alphas, we obtain a sensible first guess for the tangency portfolio that should work reasonably well both conditionally and unconditionally. In total, this step consists of 1,024 iterations, optimized using the well-known "Adam" algorithm (Kingma and Ba 2015). In the second step, we then spend 64 iterations maximizing the objective function by updating the instruments $\{g_t\}$; like CPZ, we find that this step convergences quite rapidly and hence use the lower number of iterations. The final third step then spends 1,024 iterations minimizing the objective function by updating the weights $\{w_t\}$ using the instruments from the second step. Unlike the first step, this last step thus targets the conditional alphas implied by the instruments. Moreover, since the first step converges quite rapidly, we employ a "learning rate" of 0.01 to form our initial guess and then use a learning rate of 0.001 for the second and third steps. Finally, we repeat the last two maximization and minimization steps until convergence. Based on experimentation, we find that repeating the last two steps for a total of five iterations is generally sufficient for convergence.

A.4 Cross-Sectional Pricing. To assess the cross-sectional pricing ability afforded by our estimate of the SDF, we compute the predicted returns for each of our test assets and compare the predictions to the actual realized returns. As a reference, we also compute the analogous predictions for the CAPM using the Fama-French market portfolio, and the FF6 model comprising the five Fama-French factors together with Momentum. The predicted returns themselves are computed, for simplicity, using full-sample time-series regressions where the intercept naturally corresponds to the alpha. However, we obtain similar results if we instead form ex post mean-variance efficient portfolios for each of the factor models and instead rely on the usual no-arbitrage moment condition.

In any case, the resultant predictions reported in Figure A.1, not surprisingly, show that our SDF performs admirably compared to both the CAPM and the FF6, attaining a much higher R^2 between the predicted and the actual realized returns. As a natural consequence of our adversarial method of moments approach, the largest pricing errors for the SDF are also noticeably smaller than the largest errors for the CAPM and the FF6 model, both of which substantially misprice certain assets.

B SDF Jump Identification

In this subsection, we clarify our jump identification procedure. To begin, recall that the tangency portfolio is defined as a weighted combination of the span assets. Hence, under the maintained assumptions that all of the span assets follow no-arbitrage Itô semimartingale processes and that the weights $w_t \equiv f_w(I_t;\theta_w)$ form a bounded predictable process, the tangency portfolio will itself be an Itô semimartingale. Succinctly expressing this process as

$$\int_0^t F_s ds \;=\; F_0 + \int_0^t \mu_s ds \;+\; \int_0^t \sigma_s dW_s \;+\; J_t, \qquad t \geq 0, \tag{B.2}$$

where μ_s defines the drift in the SDF, σ_s defines the SDF diffusive volatility, W_s is a standard Brownian motion, our goal is to identify the realizations of the J_t jump process that accounts for large discontinuous changes in the SDF. To do so, we rely on the now standard thresholding approach originally proposed by Mancini (2001).

In particular, following Bollerslev and Todorov (2011), we classify an SDF return $F_{t,i}$ in the ith intraday time-interval on day t as a jump if the following condition holds,

$$|F_{t,i}| \geq \alpha \sqrt{\tau_i BV_t}\, \Delta_n^{\varpi}, \tag{B.3}$$

where Δ_n denotes the sampling frequency corresponding to n intraday observations per day, and the τ_i time-of-day indicator, and the BV_t bipower variation measure (Barndorff-Nielsen and Shephard 2006) are defined as follows:

$$\tau_i = \left(\sum_t F_{t,i}^2\right) \Big/ \left(\sum_{t,j} F_{t,j}^2\right), \tag{B.4}$$

and

$$BV_t = \frac{\pi}{2} \cdot \frac{n}{n-1} \cdot \sum_{i=2}^{n} |F_{t,i}||F_{t,i-1}|. \tag{B.5}$$

Following now common choices in the literature (e.g., Todorov and Bollerslev 2010; Bollerslev and Todorov 2011; Aït-Sahalia, Jacod, and Xiu 2023), we further fix the two tuning parameters at $\alpha = 3.0$ and $\varpi = 0.49$. Our procedure thus effectively amounts to classifying an SDF return that exceeds 3.0 local standard deviations as a jump. The Internet Appendix reports additional robustness checks for larger, and more conservative, values of α. Our qualitative findings remain intact with respect to these other choices of thresholds.

B.1 Illustration of News-Driven SDF Jumps. Figure B.1 shows three concrete examples of high-frequency SDF jumps, along with headlines from the Dow Jones Newswires data during the

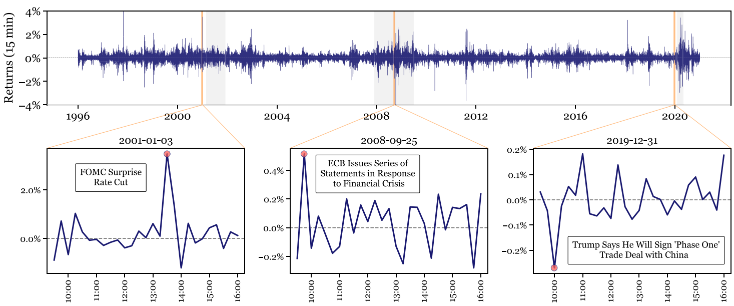

Figure B.1
Examples of news-driven jumps
The upper subplot shows the intradaily returns on the estimated tangency portfolio. The lower three subplots show 3 specific days with large jump returns readily associated with specific economic news.

Table C.1
Metatopic compositions

Metatopic	Associated topics
Monetary policy and finance	*Interest Rates; Other Financial Indicators; Financial Regulation; Monetary Policy; Inflation*
International affairs	*Middle East; Russia; China; North Korea*
Macroeconomic data	*Broad Broad Quantity Indicators; Inflation; Interest Rates; Other Financial Indicators; Labor Markets; Real Estate Markets; Trade; Business Investment and Sentiment; Consumer Spending and Sentiment*
Politics	*Elections and Political Governance; US Politics*
Commodities and energy	*Commodity Markets; Energy Markets*
Fiscal policy	*Taxes; Government Spending, Deficits and Debt; Entitlement and Welfare Programs*
Labor	*Labor Regulations; Labor Markets; Immigration*
Regulation	*Financial Regulation; Competition Policy; Intellectual Property Policy; Labor Regulations; Immigration; Energy and Environmental Regulation; Legal Reforms and Supreme Court; Housing and Land Management; Other Regulation*

This table reports the topics associated with each metatopic. Certain topics appear more than once.

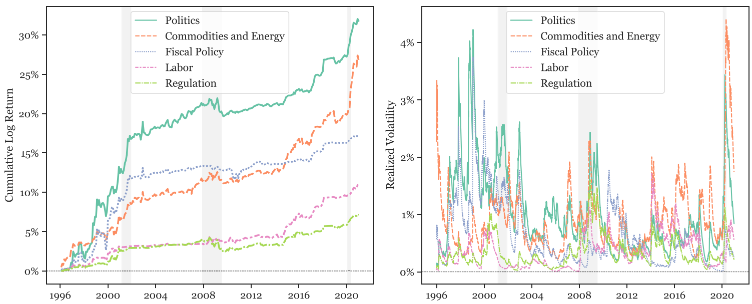

Figure C.1
Examples of news-driven jumps
The left subplot shows the cumulative returns on the *Politics*, *Commodities and energy*, *Fiscal policy*, *Labor*, and *Regulation* metatopic mimicking portfolios computed from the tangency portfolio returns rescaled to 10% per annum over the full sample. The right subplot shows a 30-day EWMA of the realized volatility of the same portfolios. The shaded regions correspond to NBER-defined recessions.

same 15-minute time intervals that contain the high-frequency jump returns. Each of these sharp changes is unsurprisingly coincident with significant economic news.

C Metatopics

C.1 Definitions. We define our metatopics by collecting groups of related topics. The metatopics *Macroeconomic data*, *Fiscal policy*, and *Regulation* are based on the topic categories previously defined by Baker et al. (2019). The remaining metatopics are based on our own definitions as detailed in Table C.1. The topics *Interest Rates*, *Labor Markets*, *Financial Regulation*, *Other Financial Indicators*, *Inflation*, *Labor Regulations*, and *Immigration* all appear twice. All other topics only appear once. As such, there are 37 associated topics in total, of which 30 are unique. Correspondingly, 14 of the 44 basic news topics, as detailed with their key terms in the Internet Appendix, are not associated with any of our eight metatopics.

C.2 Additional Risk Premium Estimates. Figure C.1 shows the time-series risk premium estimates for the additional five metatopics, not included in Figure 10 in the main text.

References

Ai, H., and R. Bansal. 2018. Risk preferences and the macroeconomic announcement premium. *Econometrica* 86:1383–430.

Ai, H., L. Jianyu, X. N. Pan, and L. Xu. 2022. Missing events in event studies: Identifying the effects of partially measured news surprises. *Journal of Financial Economics* 143:247–76.

Aït-Sahalia, Y., and J. Jacod. 2014. *High-frequency financial econometrics*. Princeton, NJ: Princeton University Press.

Aït-Sahalia, Y., J. Jacod, and D. Xiu. 2023. Continuous-Time Fama-MacBeth Regressions. Chicago Booth Research Paper No. 20-30.

Aleti, S. 2024. The high-frequency factor zoo. Working Paper, Duke University.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega. 2003. Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review* 93:38–62.

———. 2007. Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics* 73:251–77.

Aretz, K., S. M. Bartram, and P. F. Pope. 2010. Macroeconomic risks and characteristic-based factor models. *Journal of Banking and Finance* 34:1383–99.

Back, K. E. 2010. *Asset pricing and portfolio choice*. Oxford, UK: Oxford University Press.

Bajgrowicz, P., O. Scaillet, and A. Treccani. 2016. Jumps in high-frequency data: Spurious detections, dynamics, and news. *Management Science* 62:2198–217.

Baker, S., N. Bloom, S. J. Davis, and M. Sammon. 2021. What triggers stock market jumps? Working Paper, Northwestern University.

Baker, S. R., N. Bloom, and S. J. Davis. 2016. Measuring economic policy uncertainty. *Quarterly Journal of Economics* 131:1593–636.

Baker, S. R., N. Bloom, S. J. Davis, and K. J. Kost. 2019. Policy news and stock market volatility. Working Paper, Northwestern University. NBER Working Paper.

Bansal, R., D. A. Hsieh, and V. Viswanathan. 1993. A new approach to international arbitrage pricing. *Journal of Finance* 48:1719–47.

Bansal, R., and V. Viswanathan. 1993. No arbitrage and arbitrage pricing: A new approach. *Journal of Finance* 48:1231–62.

Banulescu, D., P. R. Hansen, Z. Huang, and M. Matei. 2016. Volatility during the financial crisis through the lens of high frequency data: A realized GARCH approach. University of North Carolina at Chapel Hill, Working Paper.

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard. 2009. Realized kernels in practice: Trades and quotes. *Econometrics Journal* 12:C1–C32.

Barndorff-Nielsen, O. E., and N. Shephard. 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4:1–30.

Bekaert, G., M. Hoerova, and M. Lo Duca. 2013. Risk, uncertainty and monetary policy. *Journal of Monetary Economics* 60:771–88.

Bernanke, B., and K. N. Kuttner. 2005. What explains the stock market's reaction to federal reserve policy? *Journal of Finance* 60:1221–57.

Bianchi, F., R. Gomez-Cram, H. Kung, and T. Kind. 2023. Threats to central bank independence: High-frequency identification with Twitter. *Journal of Monetary Economics* 135:37–54.

Bianchi, F., S. C. Ludvigson, and S. Ma. 2023. Monetary-based asset pricing: A mixed-frequency structural approach. Working Paper, Johns Hopkins University.

Bird, S., E. Klein, and E. Loper. 2009. *Processing raw text*. Sebastopol, CA: O'Reilly.

Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.

Bollerlsev, T. 2022. Realized semi(co)variation: Signs that all volatilities are not created equal. *Journal of Financial Econometrics* 20:219–52.

Bollerslev, T., B. Hood, J. Huss, and L. H. Pedersen. 2018. Risk everywhere: Modeling and managing volatility. *Review of Financial Studies* 31:2730–73.

Bollerslev, T., J. Li, and Y. Xue. 2018. Volume, volatility, and public news announcements. *Review of Economic Studies* 85:2005–41.

Bollerslev, T., and V. Todorov. 2011. Tails, fears, and risk premia. *Journal of Finance* 66:2165–211.

Boyd, J. H., J. Hu, and R. Jagannathan. 2005. The stock markets reaction to unemployment news: Why bad news is usually good for stocks. *Journal of Finance* 60:649–72.

Brunnermeir, M., E. Farhi, R. S. Koijen, A. Krishnamurthy, S. Ludvigson, H. Lustig, N. Stefan, and M. Piazzesi. 2021. Review article: Perspectives on the future of asset pricing. *Review of Financial Studies* 34:2126–60.

Bryzgalova, S., M. Pelger, and J. Zhu. Forthcoming. Forest through the trees: Building cross-sections of stock returns. *Journal of Finance*.

Bybee, L., B. T. Kelly, A. Manela, and D. Xiu. Forthcoming. Business news and business cycles. *Journal of Finance*.

Bybee, L., B. T. Kelly, and Y. Su. 2023. Narrative asset pricing: Interpretable systematic risk factors from news text. *Review of Financial Studies* 36:4759–4787.

Chabi-Yo, F., M. Huggenberger, and F. Weigert. 2022. Multivariate crash risk. *Journal of Financial Economics* 145:129–53.

Chen, A. Y., and T. Zimmermann. 2021. Open source cross-sectional asset pricing. Working Paper, Board of Governors of the Federal Reserve System.

Chen, L., M. Pelger, and J. Zhu. 2024. Deep learning in asset pricing. *Management Science* 70:714–50.

Chen, X., and S. C. Ludvigson. 2009. Land of addicts? An empirical investigation of habit-based asset pricing models. *Journal of Applied Econometrics* 24:1057–93.

Chernozhukov, V., W. Newey, R. Singh, and V. Syrgkanis. 2020. Adversarial estimation of riesz representers. Working Paper, MIT.

Chinco, A., A. D. Clark-Joseph, and M. Ye. 2018. Sparse signals in the cross-section of returns. *Journal of Finance* 74:449–92.

Christensen, K., A. Timmermann, and B. Veliyev. 2022. Prices moving at warp speed: Jumps after earnings announcements. Aarhus University, Working Paper.

Cieslak, A., and M. McMahon. 2023. Tough talk: The fed and the risk premium. Duke University, Working Paper.

Cieslak, A., A. Morse, and A. Vissing-Jorgensen. 2019. Stock returns over the FOMC cycle. *Journal of Finance* 74:2201–48.

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

Corsi, F. 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7:174–96.

Cutler, D. M., J. M. Poterba, and L. H. Summers. 1989. What moves stock prices? *Journal of Portfolio Management* 15:4–12.

Davis, S. J. 2017. Regulatory complexity and policy uncertainty: Headwinds of our own making. Working Paper, University of Chicago.

Evans, K. P. 2011. Intraday jumps and US macroeconomic news announcements. *Journal of Banking and Finance* 35:2511–27.

Fair, R. C. 2002. Events that shook the market. *Journal of Business* 75:713–31.

Faust, J., J. H. Rogers, S.-Y. B. Wang, and H. Wright, Jonathan. 2007. The high-frequency response of exchange rates and interest rates to macroeconomic annuncements. *Journal of Monetary Economics* 54:1051–68.

French, K. R., and R. Roll. 1986. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics* 17:5–26.

Gardner, B., C. Scotti, and C. Vega. 2022. Words speak as loudly as actions: Central bank communication and the response of equity prices to macroeconomic announcements. *Journal of Econometrics* 231:387–409.

Gu, S., B. Kelly, and D. Xiu. 2020. Empirical asset pricing via machine learning. *Review of Financial Studies* 33:2223–73.

Gürkaynak, R. S., B. Kısacıkoğlu, and J. H. Wright. 2020. Missing events in event studies: Identifying the effects of partially measured news surprises. *American Economic Review* 110:3871–912.

Hansen, L. P., and Jagannathan. 1997. Assessing specification errors in stochastic discount factor models. *Journal of Finance* 52:557–90.

Harvey, C. R., Y. Liu, and H. Zhu. 2016. ... and the cross-section of expected returns. *Review of Financial Studies* 29:5–68.

Hochreiter, S., and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9:1735–80.

Hou, K., C. Xue, and L. Zhang. 2020. Replicating anomalies. *Review of Financial Studies* 33:2019–2133.

Husted, L., J. Rogers, and B. Sun. 2020. Monetary policy uncertainty. *Journal of Monetary Economics* 115:20–36.

Jacod, J., H. Lin, and V. Todorov. 2022. Systematic jump risk. Working Paper, Northwestern University.

Jensen, T. I., B. T. Kelly, and L. H. Pedersen. 2023. Is there a replication crisis in finance? *The Journal of Finance* 78:2465–518.

Jeon, Y., T. H. McCurdy, and X. Zhao. 2022. News as sources of jumps in stock returns: Evidence from 21 million news articles for 9000 companies. *Journal of Financial Economics* 145:1–17.

Johnson, J. A., M. C. Medeiros, and B. S. Paye. 2022. Jumps in stock prices: New insights from old data. *Journal of Financial Markets* 60:100708–.

Jurado, K., S. C. Ludvigson, and S. Ng. 2015. Measuring uncertainty. *American Economic Review* 105:1177–216.

Ke, Z., B. T. Kelly, and D. Xiu. 2021. Predicting returns with text data. Working Paper, University of Chicago.

Keloharju, M., J. T. Linnainmaa, and P. Nyberg. 2021. Are return seasonalities due to risk or mispricing? *Journal of Financial Economics* 139:138–61.

Kim, S., and R. A. Korajczyk. 2021. Large sample estimators of the stochastic discount factor. Working Paper, Northwestern University.

Kingma, D. P., and J. Ba. 2015. Adam: A method for stochastic optimization. In *ICLR Conference Proceedings*.

Kozak, S., and S. Nagel. 2022. When do cross-sectional asset pricing factors span the stochastic discount factor. Working Paper, University of Chicago.

Kozak, S., S. Nagel, and S. Santosh. 2020. Shrinking the cross-section. *Journal of Financial Economics* 135:271–92.

Kuttner, K. N. 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics* 47:523–44.

Lahaye, J., S. Laurent, and C. J. Neely. 2011. Jumps, cojumps and macro announcements. *Journal of Applied Econometrics* 26:893–921.

Larsen, V., and L. A. Thorsrud. 2019. Business cycle narratives. Working Paper, Norges Bank.

Law, T.-H., D. Song, and A. Yaron. 2020. Fearing the fed: How wall street reads main street. Working Paper, Boston College.

Lee, S., and P. Mykland. 2008. Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies* 21:2535–63.

Lee, S. S. 2012. Jumps and information flow in financial markets. *Review of Financial Studies* 25:439–79.

Li, J., V. Todorov, and G. Tauchen. 2017. Jump regressions. *Econometrica* 85:173–95.

Li, J., V. Todorov, and Q. Zhang. 2024. Testing the dimensionality of policy shocks. *Review of Economics and Statistics* 106:470–82.

Liu, Y., and B. Matthies. 2022. Long-run riks: Is it there? *Journal of Finance* 75:1587–633.

Lucca, D. O., and E. Moench. 2015. The pre-FOMC announcement drift. *Journal of Finance* 70:329–71.

Mai, D., and K. Pukthuanthong. 2021. Economic narratives and market outcomes: A semi-supervised topic modeling approach. Working Paper, University of Missouri at Columbia.

Mancini, C. 2001. Disentangling the jumps of the diffusion in a geometric jumping Brownian motion. *Giornale dell'Istituto Italiano degli Attuari* 64:19–47.

Manela, A., and A. Moreira. 2017. News implied volatility and disaster concerns. *Journal of Financial Economics* 123:137–62.

McCracken, M. W., and S. Ng. 2016. FRED-MD: A monthly database for macroeconomic research. *Journal of Business and Economic Statistics* 34:574–89.

McQueen, G., and V. V. Roley. 1993. Stock prices, news and business conditions. *Review of Financial Studies* 6:683–707.

Merton, R. C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:323–61.

Nakamura, E., and J. Steinsson. 2018. High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics* 133:1283–330.

Patton, A. J., and M. Verardo. 2012. Does beta move with news? firm-specific information flows and learning about profitability. *Review of Financial Studies* 25:2789–839.

Pearce, D. K., and V. V. Roley. 1985. Stock prices and economic news. *Journal of Business* 58:49–67.

Petkova, R. 2006. Do the Fama-French factors proxy for innovations in predictive variables? *Journal of Finance* 61:581–612.

Pettenuzzo, D., R. Sabbatucci, and A. Timmermann. 2020. Cash flow news and stock price dynamics. *Journal of Finance* 75:2221–70.

Pukthuanthong, K., R. Roll, and J. Wang. 2020. An agnostic and practically useful estimator of the stochastic discount factor. Working Paper, University of Missouri.

Rigobon, R., and B. Sack. 2004. The impact of monetary policy on asset prices. *Journal of Monetary Economics* 51:1553–75.

Rogers, J. H., C. Scotti, and J. H. Wright. 2014. Evaluating asset-market effects of unconventional monetary policy: A multi-country review. *Economic Policy* 29:249–799.

Savor, P., and M. Wilson. 2013. How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis* 48:343–75.

———. 2016. Earnings announcements and systematic risk. *Journal of Finance* 71:83–138.

Schmeling, M., and C. Wagner. 2024. Does central bank tone move asset prices? *Journal of Financial and Quantitative Analysis* 1–48.

Schwert, G. W. 1989. Why does stock market volatility change over time? *Journal of Finance* 44:1115–53.

Todorov, V., and T. Bollerslev. 2010. Jumps and betas: A new framework for disentangling and estimating systematic risks. *Journal of Econometrics* 157:220–35.

Vassalou, M. 2003. News related to future GDP growth as a risk factor in equity returns. *Journal of Financial Economics* 68:47–73.

Vassalou, M., and Y. Xing. 2004. Default risk in equity returns. *Journal of Finance* 59:831–68.

Welch, I., and A. Goyal. 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21:1455–508.

Zhang, C., J. Li, V. Todorov, and G. Tauchen. 2022. Variation and efficiency of high-frequency betas. *Journal of Econometrics* 228:156–75.