Narrative Asset Pricing: Interpretable Systematic Risk Factors from News Text*

Leland Bybee
Yale University, USA

Bryan Kelly
Yale University, AQR Capital Management, and NBER, USA

Yinan Su
Johns Hopkins University, USA

May 3, 2023

Abstract

We estimate a narrative factor pricing model from news text of *The Wall Street Journal*. Our empirical method integrates topic modeling (LDA), latent factor analysis (IPCA), and variable selection (group lasso). Narrative factors achieve higher out-of-sample Sharpe ratios and smaller pricing errors than standard characteristic-based factor models and predict future investment opportunities in a manner consistent with the ICAPM. We derive an interpretation of the estimated risk factors from narratives in the underlying article text. (*JEL* C38, C52, G11, G12)

*We are grateful for comments and suggestions from Tarun Ramadorai (the editor) and two anonymous referees; several discussants, including Hui Chen, Diego García, Ryan Israelsen, Ben Matthies, Maximilian Rohrer, Qian Yang, and Dexin Zhou; and audience participants at EFA, Future of Financial Information, GSU-RFS, Holden Memorial Conference, JHU Carey, Kepos Capital, MFA, News and Finance Conference, Northeastern Finance Conference, UConn Finance Conference, and Yale SOM. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR. Send correspondence to Bryan Kelly, bryan.kelly@yale.edu.

A central premise of asset pricing is that differences in expected returns stem from differences in risk exposures, but what are the fundamental risks that investors care about? According to Merton's (1973) Intertemporal Capital Asset Pricing Model (ICAPM), risk is tied to news about "state variables" that track investors' wealth and forecast changes in future investment opportunities. Because the state variables determine optimal current consumption, their shocks constitute fundamental risks for the investor/consumer, and an asset's covariances with these shocks dictate the asset's risk premium in equilibrium.

The identity of potential ICAPM state variables has remained largely conceptual because of the limited success of empirical efforts to isolate interpretable risk factors. Some attempts propose macroeconomic variables as proxies for ICAPM state variables.[1] The main competing modeling framework uses statistical factor models based on characteristic-sorted stock portfolios. Statistical factor models tend to perform better than empirical ICAPM models in explaining covariances and risk premiums of "anomaly" portfolios and other assets, but have the drawback of being detached from interpretable fundamentals.[2]

In the ICAPM theory, state variables summarize an investor's information set and vary as she acquires unexpected new information. The state variables consist of the marketable (e.g., the market portfolio) and nonmarketable (e.g., human capital and real estate) portions of wealth, as well as forward-looking expectations about future investment opportunities. Naturally, a key challenge is measuring investors' assessments of nonmarketable wealth and future investment opportunities. It is common in existing work to infer investor expectations via predictive vector autoregression (VAR) from numerical macroeconomic data under the premise of rational expectations.

Our paper has two primary contributions. First, we attempt to narrow the gap between ICAPM theory and empirics by introducing additional data from news text. We are motivated by two potential advantages of using business news data in place of more standard macroeconomic data. News is released more continuously and is likely more timely than low frequency numerical macroeconomic data.[3] This presents an opportunity to measure covariances between assets and macroeco-

[1]For example, Chen, Roll, and Ross (1986), Cochrane (1996), Bali and Engle (2010), and Rossi and Timmermann (2015) use macroeconomic indicators, such as industrial production, investment, and inflation, to proxy for the state variables.

[2]Examples include Fama and French (1996), Fama and French (2016), Hou, Xue, and Zhang (2015), Kelly, Pruitt, and Su (2019), and Lettau and Pelger (2020).

[3]For example, Kelly, Manela, and Moreira (2021) show that WSJ text successfully forecasts and nowcasts official macroeconomic data releases.

nomic shocks more accurately (by using daily rather than monthly or quarterly data) and more synchronously (news text arrival is more likely to be concurrent with updates to the market's information set). Second, information in news text enjoys the richness of narrative; that is, it is derived from the sophisticated process of human understanding of complex contexts. News articles may contain information that is more accessible to investors because the hard cognitive work—inferring causes of business events and predicting their subsequent effects—is partially done by the journalists. Presumably, such information is in high demand by investors, thus business news outlets have an incentive to produce information relevant to market participants on a timely basis. Hence, we hypothesize that narratives (a) can proxy for shocks to marketable and nonmarketable wealth that might be poorly measured by market returns (i.e., the critique of Roll, 1977) and (b) have a meaningful forward-looking component that helps forecast future investment opportunities and can thus help proxy for latent ICAPM state variables.

These potential benefits of using news text are accompanied by a number of empirical challenges as well. News outlets face incentives to produce articles that are sensationalized or biased, or irrelevant for asset pricing (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010). This can obscure the information content of news that is useful for modeling asset prices. Beyond biased and irrelevant news, the inherent intricacies of natural language present a challenge to extracting and quantifying information encoded in news text.

Our second contribution is proposing a method for incorporating news text into an ICAPM factor pricing model that addresses the challenges of working with text data. Our empirical approach has three main components. First, we winnow news down to a set of articles that have a comparatively high likelihood of relevance for asset pricing. While many news sources are available for potential analysis, we choose to focus on *The Wall Street Journal* (WSJ) given its specialization in business news. Of course, the WSJ also produces nonbusiness news, therefore we follow Bybee et al. (2021, henceforth BKMX) and further filter out articles appearing in sections other than the three core business sections ("Section One," "Marketplace," and "Money and Investing") as well as articles with subject tags corresponding to predominantly nonbusiness content, such as sports, leisure, and arts.

The crux of the empirical problem is distilling a parsimonious set of risk factors (and eventually a pricing kernel) from the vast amount of textual data. Our empirical approach tackles this problem using the textual dimension reduction technique of latent Dirichlet allocation (LDA) to automatically group terms (unigrams and bigrams) into interpretable narrative themes based on their co-occurrences in news articles (following the LDA analysis in BKMX).[4] The LDA model consists of 180 topics and is chosen as the statistically optimal specification according to a Bayes factor criterion. Many of these topics correspond to important investor issues, such as "Recession," "U.S. Senate," "Economic growth," and "Federal Reserve." LDA estimates the attention that WSJ allocates to narratives for each topic on a given day, and we use the time series of allocations as candidate state variables in an ICAPM. LDA also estimates the term composition for each topic, thereby providing an interpretation of the state variables.

The third component of our empirical design estimates a mapping from the 180 narrative attention series into a small number of common asset pricing risk factors using instrumented principal component analysis (IPCA; Kelly, Pruitt, and Su, 2017). IPCA is estimated from an economically motivated criterion. It searches for tradable mimicking portfolios of the state variables that best fit realized individual stock returns, much like the two-step regression approach of Fama and MacBeth (1973). However, Fama-MacBeth estimates mimicking portfolios of *observable* state variables. IPCA is a dimension reduction method that upgrades the Fama-MacBeth logic to infer a small number of latent state variables from a large set of candidates (such as the 180 narratives). We introduce a new penalized version of IPCA called Sparse IPCA that selectively excludes irrelevant or especially noisy narratives before performing IPCA's usual dimension-reduced risk factor estimation.

Once the narrative factor model is estimated, we demonstrate its performance as a factor pricing model. It achieves lower out-of-sample pricing errors for 78 anomaly portfolios compared to the five Fama-French factors plus momentum (which we refer to as the "FFC6" model). The narrative factor model delivers an out-of-sample mean-variance efficient (MVE) portfolio with a Sharpe ratio of 1.3, compared to a Sharpe ratio of 0.8 for the FFC6 MVE portfolio. This result is remarkable in that narrative factors are formed based on stock covariances with narratives and no other firm characteristic data.

We then demonstrate that estimated narrative risk factors are consistent with the ICAPM framework. In particular, the narrative MVE portfolio predicts future market returns, consump-

[4] For example, terms like "economic downturn," "steep decline," "hardest hit," and "steep drop," show up together in a narrative. The narrative label is manually assigned by summarizing the common theme displayed in the automatically grouped topic terms. This example is labeled as the "Recession" topic.

tion growth, and a host of other macroeconomic indexes. The signs of the predictive relationships align with ICAPM restrictions; that is, the MVE has a positive risk premium, indicating its association with "good news," and indeed the MVE positively predicts procyclical indexes (like market returns and GDP growth) and negatively predicts countercyclical indicators (like credit spreads and unemployment).

Lastly, we derive an interpretation of the risks extracted from news narratives using the estimated model. The MVE state variable is theoretically important: it is an estimate of the univariate pricing kernel that determines asset risk premiums. The estimated MVE is a linear combination of around a dozen narratives selected from the 180 candidate topics. Our interpretation approach links these weights to narratives to quantify how changes in topic attention affect the model-implied MVE portfolio. The "Recession" topic has the largest negative impact on the model-implied MVE, while "Record high" and "Optimism" are the leading positive narratives. Through the model we can trace narrative impacts on the MVE not only to topics but also to specific articles (headlines like "Consumer Confidence Slides on Fears of Layoffs" and "Home Building Continues Recovery as Demand Rises" induce large changes in the MVE) or individual terms (see the term clouds in Figure 8). We also use this approach to identify the articles associated with the largest market swings over the past three decades. These articles relate to concrete issues like the chances of a double-dip recession and ripple effects from the European debt crisis.

We contribute to the new and promising area of research using text to understand asset markets (see the surveys of Loughran and Mcdonald, 2016; Gentzkow, Kelly, and Taddy, 2019). Early papers in this area include Tetlock (2007, 2011) and García (2013).

Closer to our paper, Liu and Matthies (2022) show that news text predicts consumption growth over low frequencies and construct a news-based index to proxy for the pricing kernel in a long-run risks framework. In contrast, our narratives are estimated with little explicit human input (other than selecting the news data source), and covariances between stocks and narratives leverage higher frequency daily data. Engle et al. (2020) identify climate change risk by tracking the fluctuations in WSJ attention to climate change news and propose a dynamic trading strategy to hedge climate risk. Our empirical methodology is a high-dimensional multivariate generalization of their hedging portfolio construction.

Ke, Kelly, and Xiu (2020), Manela and Moreira (2017), and Kelly, Manela, and Moreira (2021) develop textual machine learning methods to predict stock returns, volatility, and macroeconomic activity. Jeon, McCurdy, and Zhao (2021) attribute news as a source of jumps in stock returns. BKMX analyze a topic model for WSJ text and analyze its role in a macroeconomic VAR, and Lopez-Lira (2020) conducts a topic analysis of 10-K text to measure firms' risk exposures.

In terms of statistical methods, Sparse IPCA combines the selection and shrinkage functions of (group) lasso (Tibshirani, 1996; Yuan and Lin, 2006) with latent factor analysis via IPCA (Kelly, Pruitt, and Su, 2019). It is similar to the Sparse Principal Components Analysis (SPCA) (Zou, Hastie, and Tibshirani, 2012), which imposes lasso-type regularization on factor loadings. Pelger and Xiong (2021) impose hard-thresholding on factor loadings and emphasize the improved interpretability from sparse estimates. Sparse IPCA selects instruments according to their effects on factor loadings, in contrast to the two above that select factor loadings themselves. As an extension to Fama-MacBeth with regularization, our work is also related to Bryzgalova (2015).

1 Model

Figure 1 summarizes the data-generating process for innovations in narrative attention (z), stock excess returns (r), as well as the term frequencies of individual news articles (w). The figure has three parts: The main novelty is connecting state variables to news narratives (x to z). The empirical goal is to estimate this relationship in order to understand the fundamental risks from the perspective of news narratives. The return generating process (f to r) is the canonical latent factor model with time-varying factor loadings. We estimate the narrative-based f for factor pricing tests. The text data-generating process (θ to w, in the gray box) follows LDA and is treated as a stand-alone data preprocessing step to prepare narrative innovations (z). Throughout the paper, τ indexes days, and t indexes months. Since z, x, and r are innovations, they can be accumulated from daily to monthly frequencies.[5]

[5] The discrete-time state variable is a linear approximation of the continuous-time ICAPM. Write the instantaneous state variable as X_s, which the discrete-time state variable that we use is X's change: $x_\tau = \int_{\text{during day } \tau} \mathrm{d}X_s$, and $x_t = \int_{\text{during month } t} \mathrm{d}X_s$.

Figure 1: Illustration of the data-generating process

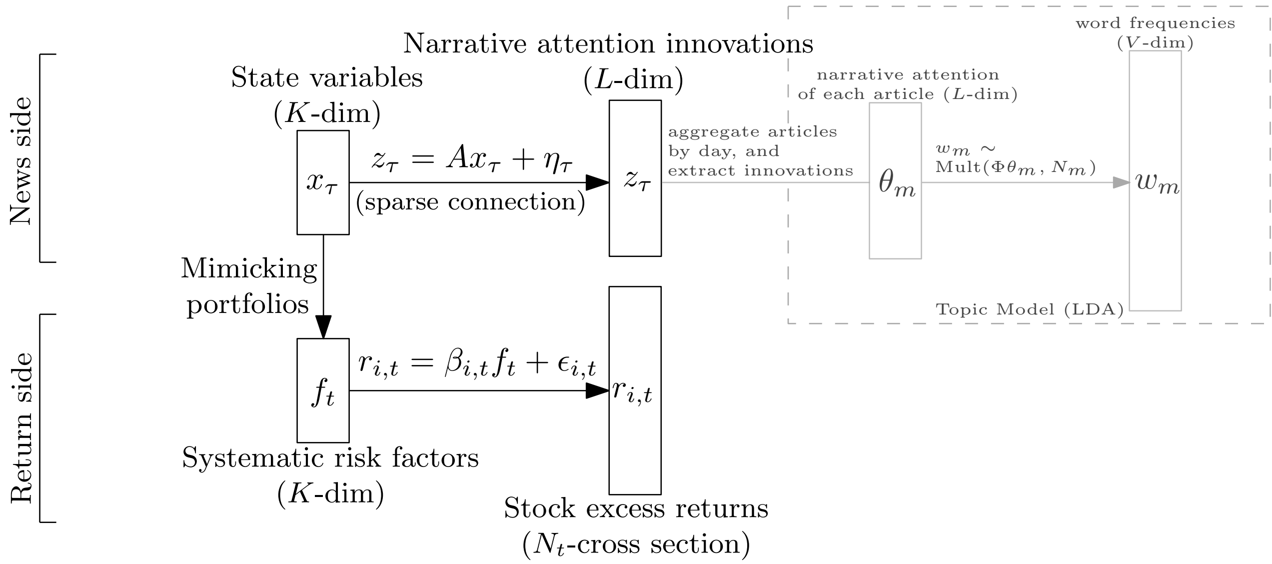

τ indexes days, t for months, i for stocks, and m for articles. L is the number of narratives, and V is the size of the vocabulary.

1.1 The news side of the model

Let x_τ ($K \times 1$ vector) be the ICAPM state variables. It contains the growth in the wealth portfolio (both the marketable and nonmarketable parts) and revisions to expectations about future investment opportunities. We do not specify an individual entry of x_τ to be the market factor as a proxy for wealth, but instead allow all news narratives to contribute to it in a unified fashion (estimated from data). This helps circumvent Roll's critique if, for example, the human capital component of wealth is not well proxied by the market factor but reflected in news narratives.[6]

Let z_τ be the innovation of each narrative's attention on day τ arranged in an $L \times 1$ vector with L being the number of narrative topics. We assume z_τ is related to the K ICAPM state variables x_τ via an $L \times K$ matrix A:

$$z_\tau = Ax_\tau + \eta_\tau, \tag{1}$$

where η_τ represents the part of narrative innovations that is irrelevant to the state variables and the asset market in general.

[6]The empirical results turn out consistent with this intuition: the factor space almost perfectly spans the market factor (see Section 6.1), and the narrative-based state variables forecast changes in the labor market (see Figure 6).

Equation (1) relates news narratives to asset pricing state variables and is the key to our model. We do not stipulate the causal mechanism for how the two are linked. Our interpretation is that agents read the news to form assessments of the current state, which in turn guides their asset pricing and consumption-saving decisions.

Numerical representations of text data are typically high-dimensional. The model is designed to accommodate this feature by allowing the number of narratives L to be much larger than the number of asset pricing state variables K. The A matrix summarizes which narratives matter for asset prices and by how much, and thus it plays a central role in deriving narrative interpretations of risks factor from our estimated model (discussed in detail in Section 6). Our estimator also supports row-wise sparsity in A, which allows the estimator to zero-out the effect of narratives that are entirely irrelevant to asset return dynamics.

We construct narrative attention shocks (z_τ) from LDA estimates of the WSJ from BKMX. We treat the LDA estimation as a stand-alone data processing step to deliver a numerical representation of the raw news text. Each narrative attention series is accompanied by an estimate for the term composition for that narrative. Through these estimates, we can eventually trace our risk factors back to influential individual articles or individual terms. This plays a valuable role in the interpretation analysis of Section 6.

The LDA model is becoming a standard tool in textual analysis for finance and economics. We briefly describe its structure to precisely define narrative attention and refer interested readers to Blei, Ng, and Jordan (2003) for further details. It begins from a "bag-of-words" representation of the raw text. The bag-of-words dimensionality is large because the WSJ vocabulary size (denoted as V) is over 18,000. LDA searches for a tractable thematic summary of the text with much lower dimensionality than V. To do so, it imposes a factor structure on term counts and nests this in a count distribution. In particular, LDA assumes the V-dimensional vector of term counts for a given article m, denoted as w_m, is distributed according to a multinomial distribution:

$$w_m \sim \text{Mult}(\Phi\theta_m, N_m), \tag{2}$$

where N_m is the sum of all term counts for article m (which governs the scale of the multinomial distribution). According to (2), expected term counts are summarized by a comparatively low dimension set of parameters, θ_m and $\Phi = [\phi_1, ..., \phi_L]'$. The l^{th} "topic" is a V-dimensional parameter vector ϕ_l, where $\phi_{l,v} \geq 0$ for all v and $\sum_v \phi_{l,v} = 1$. That is, a news topic is a probability distribution that defines the term composition of the topic. Terms with especially high probabilities in ϕ_l convey the topic's thematic content. The model's dimension reduction is achieved by setting the total number of topics L much smaller than the size of the vocabulary.

Finally, narrative attention corresponds to article-specific parameter vector $\theta_m = (\theta_{m,1}, ..., \theta_{m,L})'$. It is also a probability vector and describes article m's allocation of attention across topics. We aggregate article attention to a daily attention measure by averaging attention for all articles on day τ (denoted as $m \in \tau$) weighted by the number of terms in each article:

$$\theta_\tau = \frac{1}{\sum_{m\in\tau} N_m} \sum_{m\in\tau} N_m \theta_m.$$

The vector θ_τ is our main narrative attention time series. We calculate attention shocks, z_τ, as innovations in attention level θ_τ.

1.2 The return side of the model

Next, we describe the model component for the cross-section of returns. Let f_τ be the projection of the state variables, x_τ, onto the space of excess returns. We refer to the state-mimicking portfolios f_τ as systematic risk factors. The projection residual, $\nu_\tau := x_\tau - f_\tau$, is orthogonal to all excess returns. Therefore, we can rearrange model (1) so that

$$z_\tau = Af_\tau + g_\tau, \tag{3}$$

where $g_\tau := A\nu_\tau + \eta_\tau$ is the composite residual of z_τ. We have assumed both ν_t and η_τ are uncorrelated with excess returns, hence so is g_τ.

The cross-section of excess stock returns follows a latent factor structure:

$$r_{i,t+1} = \beta_{i,t} f_{t+1} + \epsilon_{i,t+1}. \tag{4}$$

ICAPM theory implies that the intercept is zero (i.e., there is no α) and that $\epsilon_{i,t}$ is mean zero and orthogonal to f_t.

Risk exposures, $\beta_{i,t}$, are allowed to vary over time as the firm evolves and economic conditions change. We assume $\beta_{i,t}$ is a function of instruments that provide guidance about the asset's risk exposures and also assume $\mathbb{C}\mathrm{ov}_t(f_{t+1}) = \Sigma_{\mathrm{ff}}$. These serve as identifying assumptions that enable our eventual IPCA estimator to recover estimates of time-varying exposures (see Kelly, Pruitt, and Su, 2019). Based on the hypothesis that observed narrative shocks are related to factor risks, we assume the instruments include the $1 \times L$ vector of covariances between the asset's return and narrative attention shocks, $cov_{i,t} := \mathbb{C}\mathrm{ov}_t(r_{i,t+1}, z_{t+1})$. Together with the rest of the model structure, this implies that

$$cov_{i,t} = \beta_{i,t}\Sigma_{\mathrm{ff}}A^\top.$$

Inverting the above expression, we can express $\beta_{i,t}$ in terms of narrative covariances:

$$\beta_{i,t} = cov_{i,t}A\left(A^\top A\right)^{-1}\Sigma_{\mathrm{ff}}^{-1} := cov_{i,t}\widetilde{\Gamma}, \tag{5}$$

where the $L \times K$ matrix $\widetilde{\Gamma} := A\left(A^\top A\right)^{-1}\Sigma_{\mathrm{ff}}^{-1}$ (or, equivalently, $A = \widetilde{\Gamma}\left(\widetilde{\Gamma}^\top\widetilde{\Gamma}\right)^{-1}\Sigma_{\mathrm{ff}}^{-1}$) parameterizes the instrumental mapping from $cov_{i,t}$ to $\beta_{i,t}$. Equations (4) and (5) embed the return model in the IPCA framework, whose estimation we will discuss next.

2 Estimation Method

Next, we turn to the procedure for estimating the model described in Equations (1)–(5). The procedure takes the following steps:

1. For each stock i and month t, calculate covariances between $r_{i,\tau}$ and z_τ from daily data:

$$\widehat{cov}_{i,t} := \sum_\tau \kappa(\tau;t)r_{i,\tau}z_\tau^\top - \left(\sum_\tau \kappa(\tau;t)r_{i,\tau}\right)\left(\sum_\tau \kappa(\tau;t)z_\tau^\top\right), \tag{6}$$

   where $\widehat{cov}_{i,t}$ is a $1 \times L$ row vector and $\kappa(\tau;t)$ is an exponentially decaying weighting function (kernel) that ends before the last day of month t (see details in Internet Appendix B.1).

2. Append a constant to the covariances to form a set of $(L+1)$ instruments $c_{i,t} := [1, \widehat{cov}_{i,t}]$, which

is supplied to the IPCA model

$$r_{i,t+1} = c_{i,t}\Gamma f_{t+1} + e_{i,t+1}, \tag{7}$$

where Γ is $(L+1)\times K$ with rows indexed from 0 to L ($\Gamma := [\Gamma_0; \Gamma_1; \ldots; \Gamma_L]$). Estimate the state-mimicking portfolios $\{f_t\}$ and Γ with Sparse IPCA, defined by the optimization:

$$\min_{\Gamma,\{f_t\}} \frac{1}{2}\sum_{i,t\in\mathbb{S}} (r_{i,t} - c_{i,t-1}\Gamma f_t)^2 + \lambda N_{\mathbb{S}} \sum_{l=0}^{L} \sigma_l^{\mathrm{c}} \|\Gamma_l\|_2 + \sum_{t\in\mathbb{S}} \|f_t\|_2^2, \tag{8}$$

where $\mathbb{S}$ denotes a training set, $\|\Gamma_l\|_2 = \sqrt{\Gamma_{l,1}^2 + \cdots + \Gamma_{l,K}^2}$ is the Euclidean (L_2) norm for group lasso regularization, and λ is the regularization hyperparameter (whose tuning we discuss below). Then estimate Σ_{ff} as the sample covariance of f_t estimates. Internet Appendix B.2 describes the numerical method used to solve the optimization.

3. Wrap-up step. Given estimates for $\widetilde{\Gamma} := [\Gamma_1; \ldots; \Gamma_L]$ and Σ_{ff}, calculate the plug-in estimate for A based on $A = \widetilde{\Gamma}\left(\widetilde{\Gamma}^\top\widetilde{\Gamma}\right)^{-1}\Sigma_{\mathrm{ff}}^{-1}$ as implied by Equation (5), and in turn calculate the nontradable state variable estimates via $x_\tau = \left(A^\top A\right)^{-1} A^\top z_\tau$ as implied by Equation (1).

The first two steps are a generalization of the Fama-MacBeth two-step regression approach to constructing mimicking portfolios. In a simple case in which each narrative represents a stand-alone state variable, Fama-MacBeth estimates the mimicking portfolios f_t of the observed state variables (Giglio and Xiu, 2021). Our generalization extends to cases in which the state variables that we wish to mimic must be reduced and selected from a large set of noisy proxies. The estimated f_t are portfolios of individual stocks formed with narrative covariances as firm-level signals (see detailed expressions in Internet Appendix B.2). The third step recovers the map from estimated asset pricing factors back to text narratives. The map plays a central role in our model interpretation analysis.

As discussed in the introduction, some narratives are irrelevant for asset pricing, leaving the corresponding rows of A with zero entries. Γ inherits row-wise sparsity structure A. If narrative l is irrelevant, the lth row of both A and Γ will be zero. The Sparse IPCA estimator selects relevant narratives by inducing row-wise sparsity in Γ estimation, and is achieved through the group lasso penalty in the second term of (8).[7]

[7]See Yuan and Lin (2006) and Freyberger, Neuhierl, and Weber (2020) for group lasso and an application in asset

The third term in (8) is included for a technical reason. Regularizing f_t is necessary to properly induce sparsity in the Γ estimate by preventing the second term from making Γ infinitesimally small and f_t arbitrarily large. See Internet Appendix B.1 for further detail.

We take an economically motivated approach to tuning the regularization hyperparameter λ. A successful factor pricing model should span the global MVE frontier and attain the highest Sharpe ratio among all excess returns. For a sample panel $\mathbb{S}$, each λ value corresponds to a different estimate of $(\Gamma, \{f_t\})$. For each value of λ, we calculate the training sample Sharpe ratio of the model-implied MVE portfolio as $\mathrm{SR}(\lambda; \mathbb{S}) := \sqrt{\widehat{\mu}_{\mathrm{f}}^{\top} \widehat{\Sigma}_{\mathrm{ff}}^{-1} \widehat{\mu}_{\mathrm{f}}}$, where $\widehat{\mu}_{\mathrm{f}}$ and $\widehat{\Sigma}_{\mathrm{ff}}$ are the mean and variance of f_t in $\mathbb{S}$. We set λ to its Sharpe ratio maximizing value, $\lambda^*_{\mathbb{S}} := \arg\max_{\lambda} \mathrm{SR}(\lambda; \mathbb{S})$, and in turn parameters Γ, A estimated under $\lambda^*_{\mathbb{S}}$ are the tuned estimates, which we report and analyze in the next section.[8]

3 Empirical Results

3.1 Data

Daily LDA-based narrative attention estimates for WSJ come from BKMX. This is based on full article text from 1984:01–2017:06 after applying preliminary filters to remove articles about nonbusiness topics (for details, see BKMX and Internet Appendix A.1). The model uses 180 topics and is selected based on a Bayes factor criterion (details in Appendix A.3). BKMX manually assigns topic labels based on automatically generated keyword lists for each topic. Appendix A.4 visualizes the LDA model estimates including keywords (ϕ_l) and daily attention levels (θ_τ) for a subset of narratives that we will later show to be influential in our factor pricing model. Our timing convention defines θ_τ as topic attention for the WSJ edition published on the morning of calendar day $\tau + 1$, based on the view that this reflects market information accrued on day τ and thus synchronizes z_τ with asset returns $r_{i,\tau}$.[9] We construct shocks to narrative attention (z_τ) as the deviation of attention level (θ_τ)

pricing.

[8]Note that we select λ based on the training sample rather than a separate tuning/validation sample, because the tuning criterion (MVE Sharpe ratio) is distinct from the estimation criterion (factor model mean-squared error). This has the benefit of efficiency by including more data in the training set, but may positively bias in-sample model performance and negatively bias out-of-sample performance if it introduces in-sample overfit. In Internet Appendix C.3, we report empirical results based on standard "leave-one-out" cross-validation (LOOCV), which are essentially unchanged compared to our main analysis.

[9]To avoid the potential leak of ex post information after the start of month t into the ex ante portfolio weights of f_t, we calculate $\widehat{cov}_{i,t-1}$ in the window up to the second last day before the start of month t. See Internet Appendix B.1.

from its 5-day moving average: $z_\tau := \theta_\tau - \frac{1}{5}\sum_{j=1}^{5}\theta_{\tau-j}$.[10]

Stock return data are from CRSP for firms listed on NYSE, AMEX, and NASDAQ.[11] To match the span of our news data, we use daily stock returns from 1984:01 to 2016:12 to construct monthly covariances $\widehat{cov}_{i,t}$. After a 1-year burn-in period to prepare the earliest $\widehat{cov}_{i,t}$, the full sample spans 32 years from 1985:01 to 2016:12 and contains 1,850,401 stock-month observations for 15,831 unique stocks.

3.2 Full-sample estimates and the role of regularization

Our first empirical results report the overall model fit in the full sample and explore how estimates change as we vary the degree of regularization. Figure 2 reports the following four full-sample statistics along the path of λ values (for specifications of $K = 1$ to 6 state variables):

1. **Total** R^2, defined as $1 - \sum_{i,t}(r_{i,t} - c_{i,t-1}\Gamma f_t)^2 / \sum_{i,t} r_{i,t}^2$, reports the model fit in terms of the proportion of realized return variation explained by the factors. It restates the first term in the Sparse IPCA objective (8) as an R^2 for ease of interpretation.

2. **Factor MVE Sharpe ratio**, defined in annual terms as $\sqrt{12\mu_{\mathrm{f}}^\top \Sigma_{\mathrm{ff}}^{-1}\mu_{\mathrm{f}}}$, and the maximization target for selecting λ.

3. **Number of selected narratives** is the number of rows of $\widetilde{\Gamma}$ with nonzero elements.

4. L^2 **Norm of each row of** Γ, defined as $\|\Gamma_l\|_2$, measures each instrument's marginal contribution to the K factor loadings ($\beta_{i,t}$).

The horizontal axes in Figure 2 begin on the left with unregularized IPCA ($\lambda = 0$). As λ increases, the estimation becomes more heavily penalized, which reduces the model R^2 (panel 1), shrinks the estimates of Γ toward zero (panel 4), and reduces the number of selected narratives (panel 3). Panel 2 shows that the model-implied MVE Sharpe ratio is nonmonotonic in λ. When λ is near zero, the model overfits the in-sample estimation objective, mechanically inflating the R^2. The overfit in R^2 belies poor model performance and mis-estimation of the systematic risk factors. Overfit is revealed in the suboptimal MVE Sharpe ratio (which is not mechanically inflated by the estimation

[10] In Internet Appendix C.5, we verify that our empirical results are robust to using other shock definitions.

[11] We restrict the sample to firms with nonmissing Fama-French characteristics (Fama and French, 2016) to match the sample underlying the benchmark factors.

Figure 2: Full-sample estimates with varying λ values

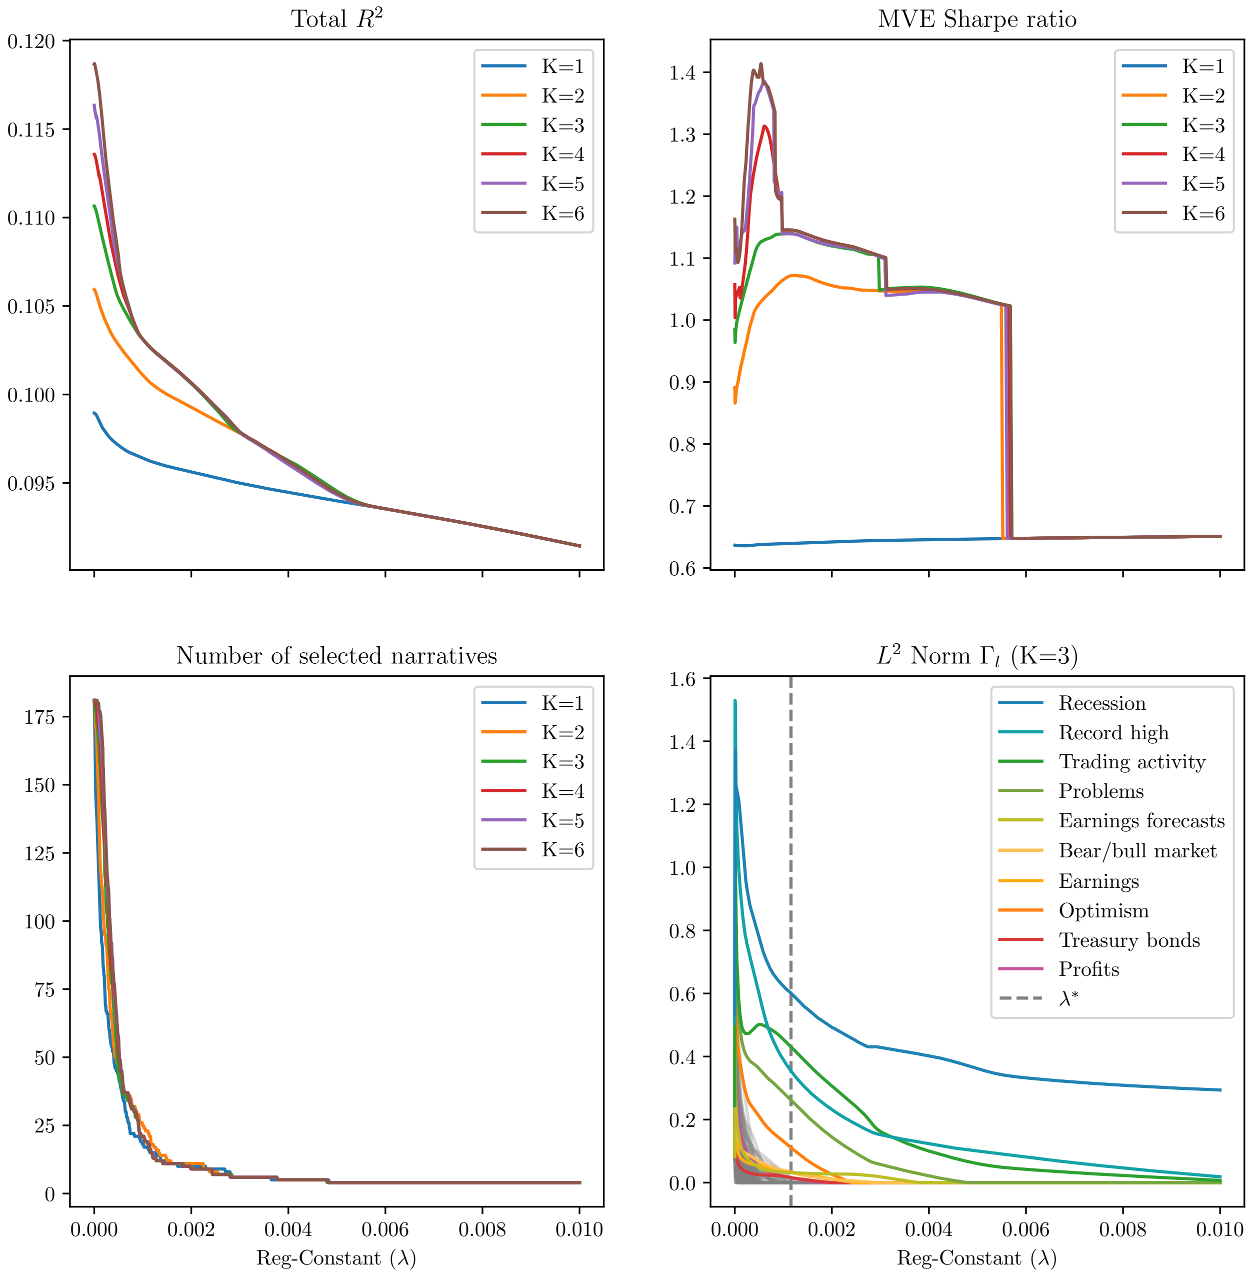

objective). A small amount of regularization begins to combat the effects of overfit and the MVE Sharpe ratio initially rises. This increase is evidence that regularization is effective in improving systematic risk estimation. Eventually the MVE Sharpe ratio peaks, corresponding to a balance point at which the model achieves a good in-sample fit without being overfit. Raising λ beyond this point leads to a deterioration in both R^2 and Sharpe ratio, indicating the region in which the model is too heavily regularized and thus underfit. Next, comparing across specification choices for the number of latent states K, we find large improvements in MVE Sharpe up to $K = 5$. At the Sharpe ratio maximizing penalty, we also see large improvements in R^2 up to $K = 3$. Further increasing K above 3 still pushes up the full-sample MVE Sharpe ratio, although the improvement becomes sensitive to the regularization constant. To be conservative, we use $K = 3$ as our main specification going forward (see plots of the estimated factors in Internet Appendix C.1). Panel 4 reports the sensitivity of Γ estimates along the λ path in the benchmark case of $K = 3$. The optimal hyperparameter is represented by the vertical dashed line, which is at the peak of the Sharpe ratio curve (the green "$K = 3$" curve in panel 2). As λ increases, each instrument's contribution to $\beta_{i,t}$ (measured by $\|\Gamma_l\|_2$) first shrinks toward zero, and eventually drops to zero for sufficiently high λ. This demonstrates both the shrinkage and selection effects of group lasso and Sparse IPCA. We color and label the ten most influential narratives in terms of $\|\Gamma_l\|_2$, with plots for the remaining narratives in gray. The top narratives in the legend are closely related to business and economics in general (see their keyword lists in Figure A.2). To further demonstrate that the selection method can effectively distinguish relevant and irrelevant instruments, we conduct an experiment with simulated placebo narratives. In addition to the 180 observed narratives (z_τ), we introduce random placebo narratives to "confuse" the estimator. The results in Internet Appendix C.2 show that Sparse IPCA effectively filters out the placebo narratives that we know for certain are irrelevant. Moreover, the estimates of the remaining selected narratives are largely unaffected by interference from the irrelevant ones.

4 Asset Pricing Performance

This section evaluates the asset pricing performance of narrative-based systematic risk factors. Factor estimates that successfully mimic the true unobservable systematic risks should deliver small pricing errors (α) for test assets and their MVE combination should achieve a high Sharpe ratio. Our benchmark factor models include the CAPM, the Fama-French three-factor model ("FF3"), the Fama-French five-factor model ("FF5"), and the Fama-French-Carhart six-factor model ("FFC6"). We report the performance of the narrative factor model with up to six factors ("NF1" to "NF6"), to match the number of factors in the benchmark models.

4.1 Cross-sectional pricing performance

Table 1: Cross-sectional pricing results

A. 78 anomaly portfolios as test assets					*B. 25 size/bm double sorts as test assets*				
Factors	avg $\lvert\hat{\alpha}_a\rvert$	avg $\lvert t(\hat{\alpha}_a)\rvert$	$\frac{\#\lvert t(\hat{\alpha}_a)\rvert>1.96}{\#\text{test assets}}$	*GRS*	Factors	avg $\lvert\hat{\alpha}_a\rvert$	avg $\lvert t(\hat{\alpha}_a)\rvert$	$\frac{\#\lvert t(\hat{\alpha}_a)\rvert>1.96}{\#\text{test assets}}$	*GRS*
Mkt	1.11	2.64	0.54	8.56	Mkt	0.42	3.03	0.84	11.61
FF3	0.97	2.65	0.54	8.50	FF3	0.32	3.89	0.84	15.12
FF5	1.18	3.20	0.68	7.48	FF5	0.27	3.31	0.72	13.06
FFC6	1.27	3.43	0.74	7.41	FFC6	0.28	3.40	0.76	12.68
NF1	1.29	3.03	0.73	8.78	NF1	0.35	2.33	0.64	6.56
NF2	0.97	2.72	0.54	7.98	NF2	0.16	1.14	0.16	5.29
NF3	0.84	2.62	0.54	7.84	NF3	0.25	1.95	0.56	5.72
NF4	0.92	2.81	0.55	7.53	NF4	0.23	1.75	0.44	5.58
NF5	0.91	2.78	0.60	7.43	NF5	0.23	1.78	0.48	5.48
NF6	0.96	2.89	0.63	7.38	NF6	0.25	1.91	0.48	5.56

Statistics: (1) avg $|\hat{\alpha}_a|$, cross-sectional average absolute value of the intercepts (a indexes anomaly test assets); (2) avg $|t(\hat{\alpha}_a)|$, the cross-sectional average absolute value of the intercept t-statistics; (3) $\frac{\#|t(\hat{\alpha}_a)|>1.96}{\#\text{test assets}}$, the proportion of intercepts statistically significantly different from zero; and (4) GRS, the GRS statistic for the joint test of all intercepts are zero.

Table 1 reports the cross-sectional asset pricing results with respect to two sets of test assets. We follow the standard empirical procedure in the factor pricing literature, calculating alphas and their t-statistics from time series regressions of test asset returns on factors over the full sample. We also calculate the GRS test statistic for the joint significance of alphas among all test assets as a model comparison metric. The test assets in panel A are 78 anomaly portfolios constructed as long-short portfolios of 78 characteristics used in Gu, Kelly, and Xiu (2020), including standard anomaly characteristics, such as idiosyncratic volatility, accruals, short-term reversal, and so forth.[12] Test assets in panel B are 25 double-sorted portfolios based on size and book-to-market from Kenneth French's data library. In general, the narrative factor models yield smaller and less significant pricing errors than the benchmarks. For example, anomaly portfolios pricing errors from the NF6 model are on average 24% smaller than those from the FFC6 model, with smaller average alpha t-statistic magnitudes (2.9 for NF6 vs. 3.4 for FFC6), and fewer significant pricing errors (63% for NFC6 vs. 74% for FFC6). We see similar patterns among size and value portfolios, where NF6 pricing errors are 11% small than those from FFC6, average t-statistics are 1.9 for NF6 versus 3.4 for FFC6, and

[12] In detail, the test assets are managed portfolios defined as $r_{a,t} := \sum_i char_{a,i,t-1} r_{i,t}$, where $char_{a,i,t-1}$ is the rank-standardized characteristic a of stock i at time t.

only 48% of alphas are significant for NF6 versus 76% for FFC6 (despite the fact that Fama-French factors are essentially designed to price size and value portfolios). Internet Appendix Figure C.3 displays each anomaly's pricing errors under FFC6 and NF3 for comparison. For example, NF3 is effective in pricing anomalies based on realized risk measures, such as betting-against beta (BAB) and idiosyncratic volatility. Short-term reversal remains the strongest anomaly for both models. This is understandable as we do not expect the ICAPM mechanism to explain fleeting mispricings that are driven by short-term market frictions.

4.2 Out-of-sample MVE Sharpe ratio

Next, we evaluate the narrative factor models in terms of their mean-variance efficiency. We construct out-of-sample factors and their MVE as tradable portfolios that mimic narrative state variables. This analysis focuses on out-of-sample performance to avoid Sharpe ratio inflation for the MVE portfolios arising from any potential in-sample overfit. The out-of-sample factor estimate for month $t+1$ is

$$f_{t+1}^{\text{OOS}} = \left[\left(\sum_i \hat{\beta}_{i,t} \hat{\beta}_{i,t}^{\top} \right) + 2\mathbb{I}_K \right]^{-1} \sum_i \hat{\beta}_{i,t} r_{i,t+1}$$

where $\hat{\beta}_{i,t} = c_{i,t}\widehat{\Gamma}$ and $\widehat{\Gamma}$ is estimated in a training sample that ends before month t (for further details on the f_t expression, see Internet Appendix B.2). We combine the out-of-sample factor estimates into an out-of-sample MVE portfolio according to

$$f_{t+1}^{\text{MVE,OOS}} = \widehat{\mu}_{\text{f}}^{\top} \widehat{\Sigma}_{\text{ff}}^{-1} f_{t+1}^{\text{OOS}},$$

where $\widehat{\mu}_{\text{f}}$ and $\widehat{\Sigma}_{\text{ff}}$ are the estimated factor means and covariances based on the in-sample factor realization in the training sample. We form out-of-sample f_{t+1}^{MVE} series with expanding training samples. The first out-of-sample realization corresponds to January 2000 (the initial training sample ends in December 1999). To reduce computation, we retrain the model once each December and apply these estimates to construct the out-of-sample MVE portfolio for each of the next 12 months. The result is a 16-year out-of-sample evaluation period for the MVE portfolios. We construct the out-of-sample MVE portfolio of benchmark Fama-French factors using the same expanding training scheme for their means and covariances. Figure 3 plots $\|\Gamma_l\|_2$ across the expanding training samples

Figure 3: L^2 norm of Γ in expanding training samples

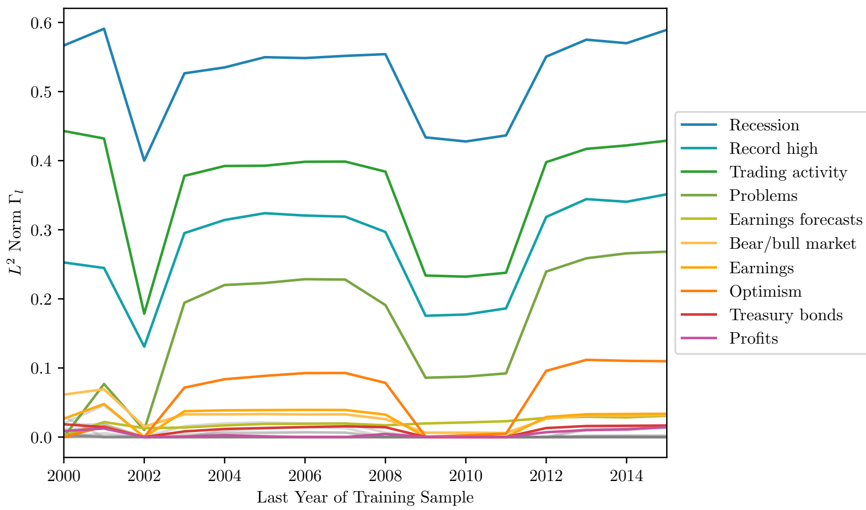

The horizontal axis is the end of the expanding training samples. The start is always 1985.

Table 2: Out-of-sample Sharpe ratios of MVE portfolios

A. Narrative factor model							
		K					
Tuning	Statistics	1	2	3	4	5	6
$\lambda = \lambda^*_{\mathbb{S}}$	Sharpe ratio	0.48	1.00	1.10	1.26	1.32	1.31
	# narratives	2.9	4.9	12.1	39.1	43.4	61.8
$\lambda = 0$	Sharpe ratio	0.44	0.66	0.73	0.73	0.79	0.91
	# narratives	All 180 narratives, no selection					
B. Benchmark factors							
		Mkt	SMB	HML	RMW	CMA	UMD
	Sharpe ratio	0.25	0.13	0.36	0.82	0.90	0.67

Panel A reports out-of-sample MVE portfolio Sharpe ratios for narrative factor models with up to six factors. The rows $\lambda = \lambda^*_{\mathbb{S}}$ correspond to Sparse IPCA and the "# narratives" reports the number of selected narratives (averaged over all training samples). The rows $\lambda = 0$ correspond to unregularized IPCA and thus all narratives are included without penalization. In panel B, each column reports the out-of-sample MVE portfolio Sharpe ratio using all factors up to that column (e.g., the column labeled "RMW" reports the MVE combination of Mkt, SMB, HML, and RMW). The out-of-sample period is from January 2001 to December 2016.

to illustrate how each topic l's contribution changes as we vary the training sample. Note that the identity and magnitude of estimates for each topic are relatively stable over time. The stability of Γ estimates provides the foundation for constructing robust out-of-sample factor portfolios. Panel A of Table 2 reports out-of-sample Sharpe ratios for narrative factor model MVE portfolios. We

Figure 4: Out-of-sample MVE Sharpe ratio versus λ

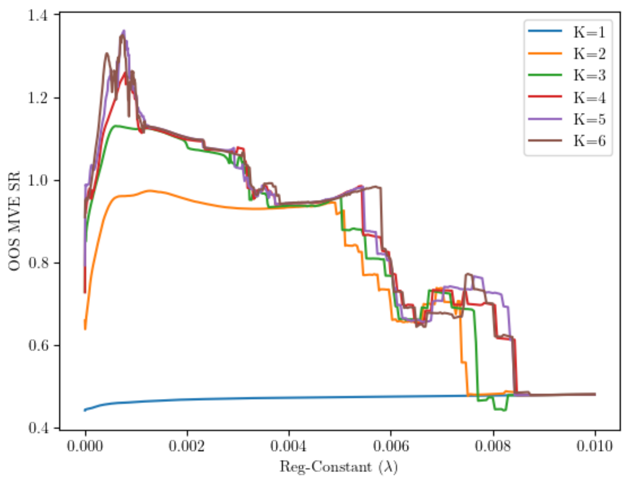

Out-of-sample MVE Sharpe ratios for different regularizing constant (λ) values.

present models estimated under the tuned regularization constant ($\lambda = \lambda^*_{\mathbb{S}}$) and the unregularized ($\lambda = 0$) for comparison. The "# narratives" row reports the number of selected narratives averaged across the expanding training samples. In panel B, each column reports the out-of-sample MVE portfolio Sharpe ratio using all factors up to that column (e.g., the column labeled "RMW" reports the MVE combination of Mkt, SMB, HML, and RMW). The MVE Sharpe ratio of narrative factors dominates that of the benchmark factor for each value of K. At $K = 5$, both the narrative model and the benchmark model achieve their highest out-of-sample MVE Sharpe ratios. In this case, the narrative Sharpe ratio of 1.3 achieves an improvement of roughly 50% over the benchmark. Panel A also shows that as K increases, the number of narratives selected also increases, indicating that a larger factor dimension gives the model more capacity to incorporate narratives and better matching asset price behavior out-of-sample. The results also show the large gains in out-of-sample model performance due to regularization. Without this, the narrative model MVE is much closer to the benchmark (i.e., the row labeled $\lambda = 0$). Figure 4 is another illustration of the benefits of Sparse IPCA regularization. It shows out-of-sample MVE Sharpe ratios for the full range of λ values. The steep increase in Sharpe ratio on the left side of the plot directly quantifies the benefit of regularization. Furthermore, this plot is very similar to its in-sample analog Figure 2, panel 2, demonstrating the efficacy of our in-sample tuning scheme for identifying λ values that optimize out-of-sample model performance. Finally, Figure 5 plots MVE portfolio cumulative returns over the evaluation sample.

Figure 5: Out-of-sample MVE cumulative returns

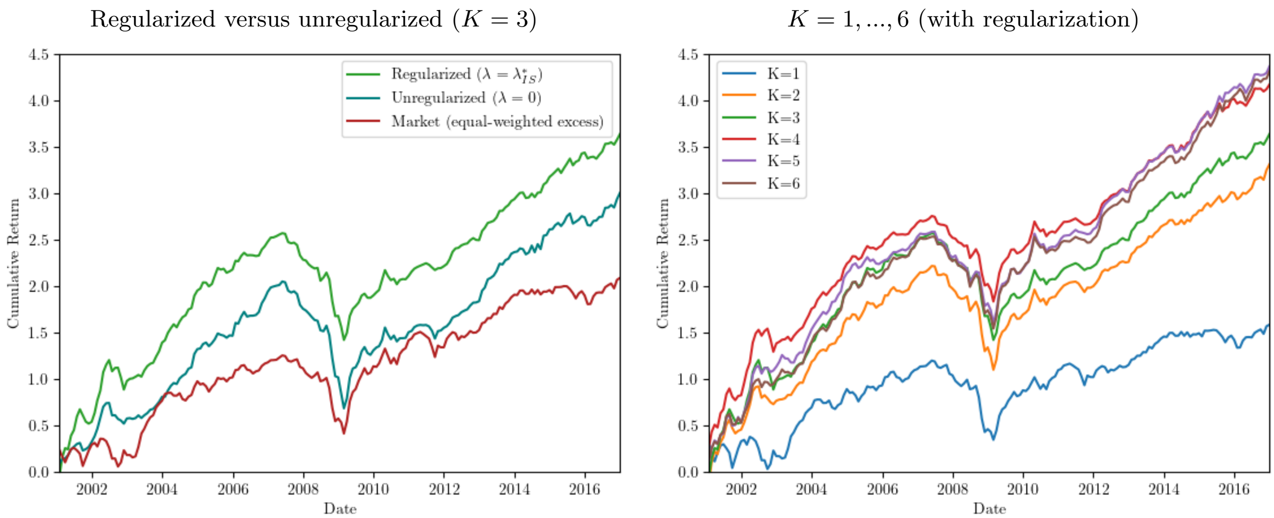

Cumulative returns of the OOS f^{MVE} in different specifications. The left panel fixes $K = 3$ and plots the benchmark specification in green, the unregularized in blue, and the market factor in red for comparison. The right panel always uses regularization but varies $K = 1 \sim 6$. In both figures, the return sequences are standardized with the same standard deviations as the market return.

The first panel compares the MVE return with and without regularization holding $K = 3$ fixed (and compared to the equal-weighted excess market return). The second panel compares the MVE for different specification choices of K. We standardize all return series to have the same volatility over the evaluation sample. The figures show that the investment performance of the narrative MVE portfolio is not concentrated in a particular period or driven by a particular event. While statistical horse races and anomaly pricing tests favor narrative factors, their interpretability and ICAPM properties (discussed in detail below) are further economic reasons to favor the narrative-based model over the benchmarks. As emphasized by Cochrane (2009), "It is probably not a good idea to evaluate economically interesting models with statistical horse races against models that use portfolio returns as factors. . . . one should always ask of a factor model, 'what is the compelling economic story that restricts the range of factors used?' ... If the purpose of the model is not just to predict asset prices but also to explain them, this puts an additional burden on economic motivation of the risk factors."

Table 3: Sharpe ratios of MVE portfolios combining FFC6 and NF

	K (number of NF added)						
Specification	0	1	2	3	4	5	6
NF		0.48	1.00	1.10	1.26	1.32	1.31
NF + Mkt	0.25	0.48	1.00	1.08	1.26	1.32	1.31
NF + FF3	0.36	0.41	0.60	0.68	1.02	1.17	0.98
NF + FF5	0.90	0.89	0.94	0.92	1.01	1.01	1.08
NF + FFC6	0.67	0.65	0.76	0.80	0.87	0.92	1.19

The first row (only NF, no FFC6) repeats the numbers from Table 2, panel A, Line "$\lambda = \lambda^*_{\mathbb{S}}$." The first column (only FFC6, no NF) repeats the numbers from Table 2, panel B. The rest of the entries are new results combining K NF's and a subset of factors from FFC6.

4.3 Robustness

We conduct additional sensitivity analyses for narrative factor model performance. First, we examine the effect of replacing news-based topics with a set of 129 numerical macroeconomic data series from FRED-MD as candidates for the latent state variables. The results are reported in panel B of Internet Appendix Table C.3 (which can be compared to our main results restated in panel A of Table C.3 for ease of reference). Model performance deteriorates when we replace narrative data with FRED-MD data. The out-of-sample MVE Sharpe ratio drops from 1.3 to 0.7. Second, we conduct robustness checks for our construction of narrative attention innovations. Instead of the 5-day trailing moving average ($z_\tau := \theta_\tau - \frac{1}{5}\sum_{\iota=1}^{5}\theta_{\tau-\iota}$), the alternatives include 1-day, 3-day, and 20-day moving averages (Table C.3, panel C). In these cases, the out-of-sample MVE Sharpe ratios remain well above one and the cross-sectional pricing errors are as good as in our main analysis. In the third exercise, we examine the characteristics of stocks with high and low exposures to the narrative factors. We regress narrative β's on the traditional firm characteristics across the stock-month panel. The correspondence is strikingly high (the R^2 is around 40%), although the two are from drastically different information sources, namely, textual news versus CRSP/Compustat. Stocks with high β's to the narrative-based pricing kernel, f_t^{MVE} is associated with anomalous characteristics known to be related to high average returns. For example, they tend to have smaller sizes, higher momentum, lower market beta, and higher dividend-to-price ratios. See Internet Appendix C.6 for the regression results and more detailed analysis. Lastly, instead of running horse races between separate MVE portfolios based on either narrative factors or characteristics-sorted portfolios, we analyze MVE portfolios that combine the two sets of factors in a joint model. To do so, we reestimate each narrative factor model (with $K = 1$ to 6) controlling for various combinations of the benchmark factors (more details in Internet Appendix C.7). We find that adding narrative factors to the benchmark factors produces a large economic improvement in the out-of-sample MVE Sharpe ratio (see, e.g., the last row of Table 3). The reverse is not true. If we begin from the NF6 specification and gradually introduce the benchmark FFC6 factors into the model, the out-of-sample MVE Sharpe ratio either does not change or degrades (see the last column of Table 3).

5 Narrative Factors and the ICAPM

We introduce news narratives in a factor pricing model in hopes of bringing empirical asset pricing closer to its theoretical underpinnings of macroeconomic risk. We adopt the ICAPM framework by hypothesizing that narratives have a meaningful forward-looking component that helps forecast future investment opportunities and track investor wealth. Given the narrative-based state variables estimated above, we loop back to test this hypothesis by examining the forecasting power of our estimated state variables, particularly the MVE combination of the narrative attention measures (which is an estimate of the model's pricing kernel). We run long-term predictive regressions of different macroeconomic variables onto the estimated MVE state variable (x_t^{MVE}). The macroeconomic variables include market return, inflation, interest rates, credit spreads, and growth in consumption, GDP, employment, payroll, and housing from FRED-MD (McCracken and Ng, 2016). For each forecast target, we predict the cumulative changes over different horizons (h):

$$\sum_{s=1}^{h} \psi_{t+s} = b_h \left(x_t^{\mathrm{MVE}} / \mathrm{std}(x_t^{\mathrm{MVE}}) \right) + \mathrm{error}_t^{(h)}. \tag{9}$$

On the left-hand side, ψ_t denotes the one-month change in a macroeconomic variable, such as $\psi_t =$ GDP growth$_t$, and the summation takes the cumulative change in the future horizon of h months. We standardize x_t^{MVE} so that the coefficient b_h can be interpreted as the effect per one-standard-deviation change in the state variable. The ICAPM not only implies the existence of predictive relationships between state variables and future outcomes but also imposes sign restrictions. A pricing kernel (x^{MVE}) should be procyclical, meaning it should rise with "good news," positively correlate with contemporaneous consumption growth, and positively predict future investment opportunities as well

Figure 6: News-based pricing kernel (x^{MVE}) forecasts future macroeconomic variables

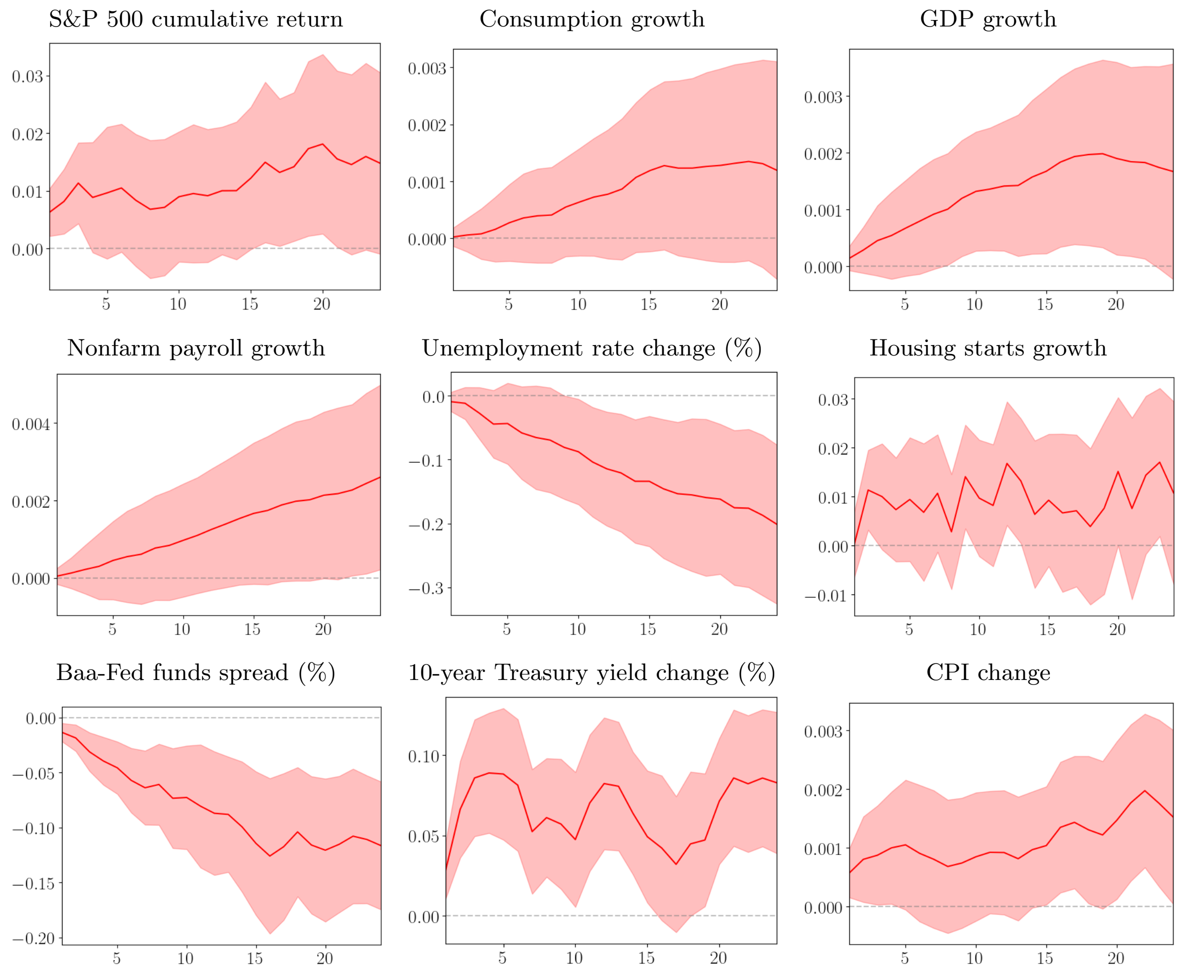

Each figure reports the estimation results of predictive regression (9) for a macroeconomic variable. The red line plots coefficient (b_h) estimates, and the shaded band represents the 90% confidence interval. We use the Newey-West standard errors computed with h lags to account for the autocorrelation.

as future cash flow of investor wealth. Our working hypothesis is thus that x_t^{MVE} positively predicts market returns, inflation, interest rates, and growth in consumption, GDP, payroll, and housing, and that it negatively predicts credit spreads and changes in unemployment.[13] Figure 6 reports predictive regression results. Each panel corresponds to a different prediction target, the horizontal axis is the prediction horizon h from 1 to 24 months, the vertical axis is the estimated coefficient b_h, and the shaded area marks the 90% confidence interval. The sign of the predictive associations between x_t^{MVE} with future economic outcomes is consistent with the ICAPM in every regression, though the result is

[13]Our analysis is similar to that of Maio and Santa-Clara (2012) and Liu and Matthies (2022), who analyze state variables constructed by other means.

not statistically significant in all cases. The first panel implies that a one-standard-deviation increase in x_t^{MVE} is associated with a 1.2-percentage-point increase in market value over the subsequent 24 months (the estimate is borderline insignificant at most horizons). The most statistically significant effects correspond to longer-term predictions of unemployment, credit spreads, and 10-year interest rates, and to a lesser extent of GDP growth, payroll growth, and inflation. The positive association with future consumption growth (which is though insignificant as in the second panel) supports the view that the news-based pricing kernel is compatible with long-run consumption risk as in Bansal and Yaron (2004) and Liu and Matthies (2022).[14] In the fourth and fifth panels, the positive prediction for nonfarm payroll and negative prediction for unemployment rate suggest that our state variables track (in part) the human capital component of investor wealth from anticipated labor income.[15] That x_t^{MVE} positively predicts housing development and negatively predicts credit spread suggests that narrative states track other nonequity components of investor wealth as well. Additionally, we employ the VAR framework to inspect the prediction properties of narrative state variables above and beyond other standard macroeconomic variables. We estimate VAR(3) systems of five variables: the narrative attention levels (θ_t^{MVE}), the S&P 500 index, the Federal Funds rate, employment, and industrial production. We experiment with two VAR specifications by switching the order of θ_t^{MVE} and the S&P 500 index. They correspond to different specifications of the unidentified structural shocks. The impulse response functions show θ_t^{MVE} shock has a positive and significant impact on market values over the next year. The effect persists beyond 2 years, though it becomes statistically insignificant at this horizon. This is true even with the conservative specification where the contemporaneous effect of θ_t^{MVE} on the S&P 500 index is forced to zero. Internet Appendix C.8 reports the details.

6 Narrative Interpretation of the Asset Pricing Model

Our estimated factor pricing model explicitly links news narratives to state variables, allowing us to trace factor behavior back to the narrative attention data. Furthermore, the estimated LDA model provides a link between narrative attention and the underlying primitive text data. Taken together,

[14] We report contemporaneous correlations of narrative state variables with macroeconomic variables in Table C.1. It shows x^{MVE} is positively correlated with contemporaneous consumption growth, supporting the general rationale of consumption CAPM models.

[15] The literature that focuses on labor income risks in asset pricing includes Jagannathan and Wang (1996), Julliard (2007), and Liu (2021).

the estimated factor model and LDA model provide a vehicle for interpreting asset pricing risk factors by relating them directly to human-readable news.

6.1 Interpretation method

Our interpretation method quantifies the impact of news on state variables. Imagine a hypothetical state of the world s, where a new piece of text is reflected as a narrative attention shock, $z(s)$. The news shock can be at the level of daily narrative attention, a single news article, or even an individual term count. Equation (1) maps narrative attention shocks to state variables according to

$$x(s) = (A^\top A)^{-1} A^\top z(s) := I_{z\to x}^\top z(s), \tag{10}$$

where the $L \times K$ matrix $I_{z\to x} := A(A^\top A)^{-1}$ summarizes how L narratives affect the K states. The latent factor structure of our model (because of the concomitant rotational indeterminacy) makes the identity of $x(s)$ ambiguous. However, the mapping between $z(s)$ on the MVE combination of state variables

$$x^{\mathrm{MVE}} := b^{\mathrm{MVE}} x, \quad \text{with} \quad b^{\mathrm{MVE}} := \mu_{\mathrm{f}}^\top \Sigma_{\mathrm{ff}}^{-1}$$

is invariant and thus unambiguous.[16] The MVE state variable is special in that it represents the univariate pricing kernel, which according to the theory is the sole source of cross-sectional variation in expected returns. Therefore, we study the news impulse for x^{MVE} as

$$x^{\mathrm{MVE}}(s) = b^{\mathrm{MVE}} (A^\top A)^{-1} A^\top z(s) := I_{z\to \mathrm{MVE}}^\top z(s), \tag{11}$$

where the $L \times 1$ "impact vector" $I_{z\to\mathrm{MVE}}$ summarizes the impact of each narrative on the pricing kernel. We also trace the MVE impulse one step further from the narrative level to the term level. Imagine now that in hypothetical state s, a new piece of text is reflected as a change in term frequencies, $\Delta w(s)$ (a $V \times 1$ vector with V the vocabulary size). The LDA model maps V-dimensional term frequencies into L-dimensional narrative attention according to $z(s) = \left(\Phi^\top \Phi\right)^{-1} \Phi^\top \Delta w(s)$. The induced narrative-level innovations then impact the MVE according to (11). Combined, the mapping

[16] x^{MVE} is defined such that its mimicking portfolio is the tradable MVE combination of systematic factors ($f^{\mathrm{MVE}} := b^{\mathrm{MVE}} f$). The rotational indeterminacy refers to that the model is invariant if β, f, x are changed to $\beta R^{-1}, Rf, Rx$, for some $K \times K$ invertible matrix R. However, $f^{\mathrm{MVE}}, x^{\mathrm{MVE}}$ are invariant to the rotation matrix R.

from term shocks to state variable response is

$$x^{\mathrm{MVE}}(s)=\underbrace{b^{\mathrm{MVE}}(A^\top A)^{-1}A^\top\left(\Phi^\top\Phi^\top\right)^{-1}\Phi^\top}_{:=I^\top_{w\to\mathrm{MVE}}}\Delta w(s). \tag{12}$$

Notice that the term-level impact vector $I_{w\to\mathrm{MVE}}$ is a linear combination of the L columns of Φ. Hence the term-level effects ($I_{w\to\mathrm{MVE}}$) describe a composite narrative that combines the basis narratives (ϕ_l's) according to their impacts on x^{MVE}. Unlike Φ, which consists of only nonnegative entries, $I_{w\to\mathrm{MVE}}$ has both positive and negative entries because the pricing kernel x^{MVE} can be affected in both directions. Besides the MVE combination of latent factors, we can also interpret the observed market factor, Mkt_t, given that a combination of the latent factors almost perfectly spans Mkt_t: the R^2 of projection $Mkt_t = b^{\mathrm{Mkt}} f_t + \mathrm{errors}_t$ is 97.6%. Replacing the combination weights b^{MVE} by b^{Mkt} in (11) and (12), respectively, we similarly define impact vectors, $I_{z\to\mathrm{Mkt}}$ and $I_{w\to\mathrm{Mkt}}$, that trace the impact on the market factor originating from narrative-level and term-level shocks.[17]

6.2 Interpretation results

We report impact vectors from full-sample model estimates for the $K = 3$ specification. Figure 7, panel 1, plots the impact of each selected narrative on the pricing kernel ($I_{z\to\mathrm{MVE}}$). The "Recession" narrative (whose keywords include "economic downturn," "weak economy," and "economic slump") has the most negative impact on the pricing kernel. Thus, the model infers that an increase in "Recession" narrative attention associates with an increase in the marginal utility of consumption. It also implies that assets whose returns positively correlate with "Recession" attention provide hedging benefits and hence carry high valuations (low risk premiums). The narratives "Record High" (whose keywords include "highest level," "pent-up demand," and "remain strong") and "Optimism" (whose keywords include "remain optimistic," "express confidence," and "positive sign") have the largest positive impact on x^{MVE}. The next three panels in Figure 7 report the impacts on the three individual risk factors comprising x_t (the three columns of $I_{z\to x}$). While the three impacts are not rotationally identified, they show that some narratives like "Trading Activity" (keywords including "market action" and "volume total") and "Problems" (keywords including "major problem" and

[17] We run full-sample projection of Mkt_t onto the estimated f_t to construct b^{Mkt} and use full-sample estimates $\widehat{\mu}_{\mathrm{f}}, \widehat{\Sigma}_{\mathrm{ff}}$ for b^{MVE}. Besides Mkt_t, one can follow the projection method to build impact vectors for other series as well.

Figure 7: Risk factor interpretation at the narrative level

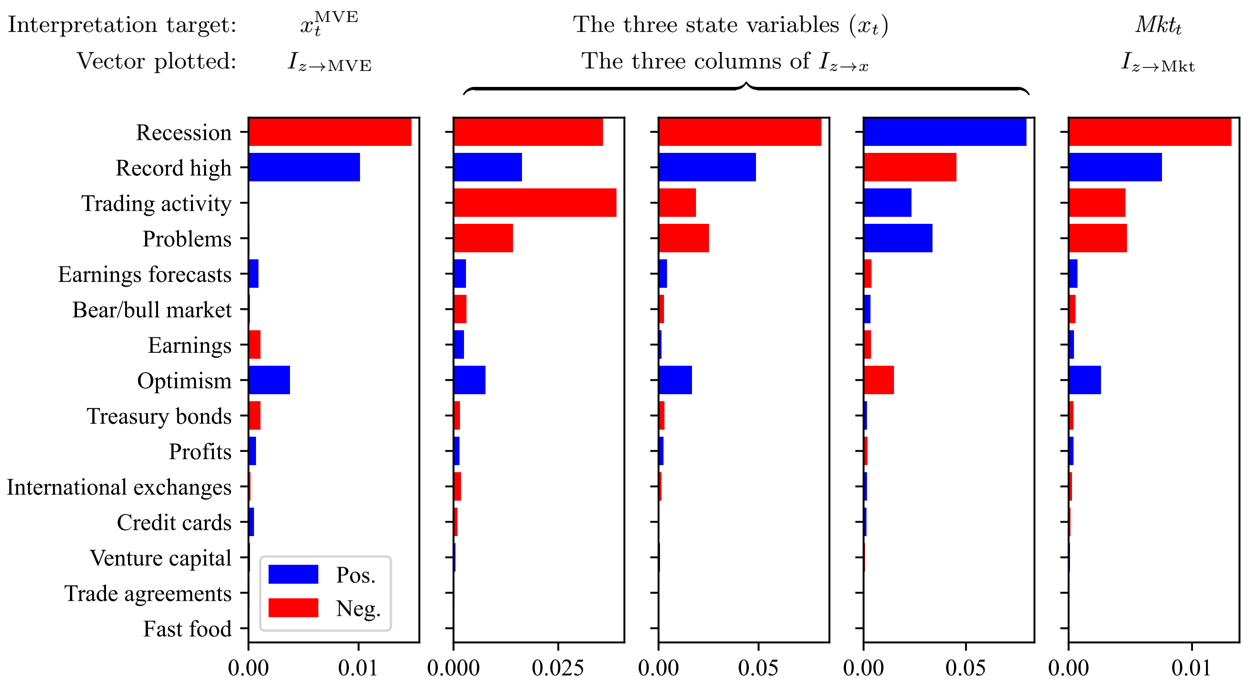

This figure plots the narrative-level impact vectors. The length of the bars represents the absolute values of corresponding entries, with red for negative impact and blue for positive.

"debacle") can have large yet opposing impacts on individual factors that net out to a small impact on the MVE. In the last panel in Figure 7, we see the market factor's impact vector has similar signs with the MVE in terms of "Recession," "Record high," and "Optimism." The differences are in the additional loadings on the narratives that describe financial market activeness, such as "Trading activity," "Bear/bull market," and "Problems." As these narratives are orthogonal to x^{MVE}, they add risk but contribute little expected return, dragging Mkt_t away from multivariate mean-variance efficiency. Moving from narratives to individual terms, Figure 8 reports the model's term-level impact vector, $I_{w\to\mathrm{MVE}}$. The dimension of this vector is over 18,000, so we display the terms' impact on x^{MVE} in term of clouds, where the magnitude of a term's impact is proportional to the size of the term in the cloud. Terms with negative and positive impacts are reported in separate clouds. The term clouds most reflect keywords from the "Recession" and the "Record high" narratives. Although the construction is without human input on the terms' semantic meanings, the term clouds show two distinct extremes in the language used to describe the current market condition and the investment outlook. It is understandable that these terms could be related to

Figure 8: MVE state variable (x^{MVE}) interpretation at the term level

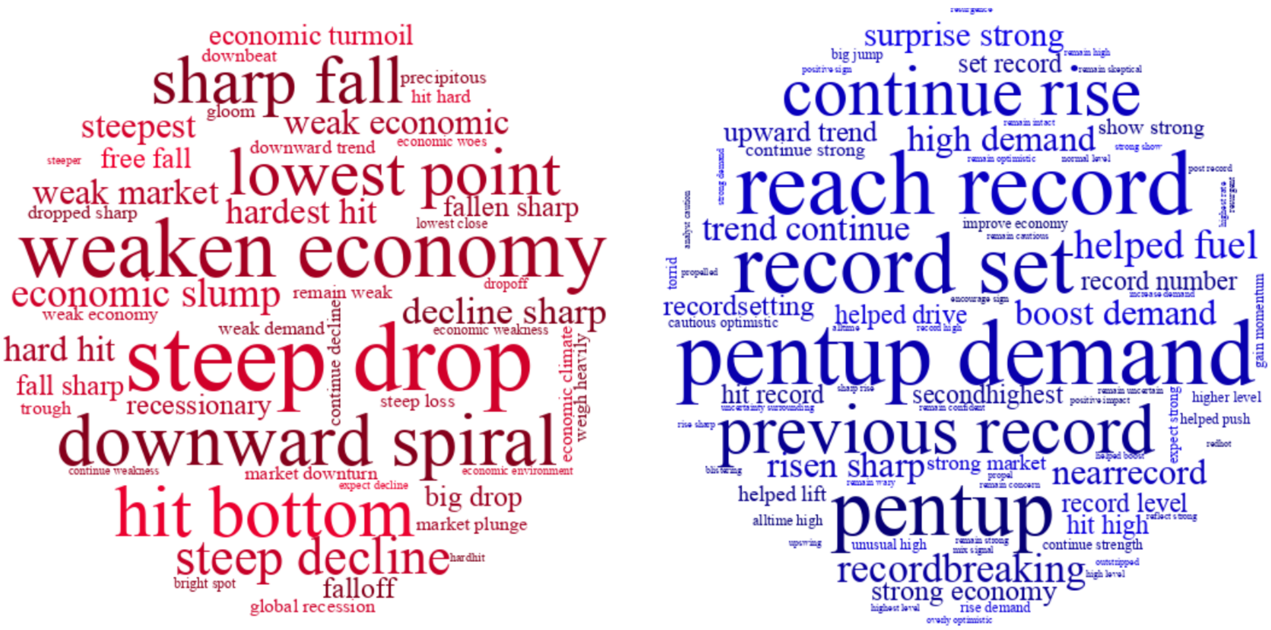

Negative | Positive

The figure illustrates each term's impact on the MVE state variable ($I_{w \to \mathrm{MVE}}$). A term's font size corresponds to the absolute magnitude of its impact. The construction method is detailed in Section 6.1.

concurrent adjustments to consumption and hence the pricing kernel. Table 4 describes the detailed contents of a few example narratives. It reports headlines of articles with the highest attention to "Recession," "Record High," and "Trading Activity" narratives. LDA identifies these articles by their disproportionate usage of terms prominent to each narrative (highlighted in yellow). "Recession" articles cover various aspects of the economic outlook, ranging from durable goods consumption to labor market activities. And, while the keyword list for "Record High" in the Internet Appendix might appear to indicate a focus on asset price movements, we see that its articles are in fact about fluctuations in the real economy. Both of these examples are consistent with the ICAPM mechanism linking the pricing kernel to real investment opportunities, as the "Recession" and "Record High" narratives have the largest negative and positive impacts on x^{MVE}, respectively. As a contrasting example, articles closest to the "Trading Activity" narrative mostly provide an ex post account of the prior day's asset returns, and interestingly this narrative has essentially zero net impact on the pricing kernel.

Table 4: Example news articles of the three most prominent narratives

	Recession
1993-05-07	Auto Registrations Continued to Slump In Europe Last Month
2001-04-25	Consumer Confidence Slides on Fears of Layoffs
2009-02-19	U.S. News: Housing Starts Hit Lowest Level In Half-Century
2011-08-02	World News: Manufacturing Slowdown Adds to Gloom on Economy
2016-07-08	World News: U.K. Consumer Sentiment Takes Dive
	Record high
1989-07-05	Japan Vehicle Sales Rise
1994-07-01	Purchasing Managers In U.K. Survey Report Rise for June Orders
1995-02-27	Hiring Outlook For Second Quarter Appears Vigorous
2006-01-12	Wall Street Bonuses Hit a Record in 2005
2016-07-20	U.S. News: Home Building Continues Recovery as Demand Rises
	Trading activity
1993-12-30	Industrials Rise A Bit to Record; Bonds Decline
1994-10-20	Profit News Helps Boost Stock Prices — Indexes Gain Ground Despite Weakness Of Bonds and Dollar
1996-06-21	Nasdaq Sinks Amid Sell-Off Of Tech Stocks
1997-12-09	Blue Chips Fall As Dollar's Rise Causes Concern
1998-04-21	Drug Stocks Resume Gains; Blue Chips Fall

This table reports example articles with the highest attention to the three narratives, respectively (titles only; see text bodies in Internet Appendix D.1). The shade of the yellow highlighting on each term reflects $\phi_{v,l}$, the term's probability for the corresponding narrative.

6.3 Narrative retrieval

We can use the impact vectors from the model to interpret daily fluctuations in the market factor. The model-implied impact on the market factor of an article m with narrative attention innovation $z(m)$ is $I_{z\to\mathrm{Mkt}}^{\top} z(m)$.[18] Figure 9 plots the observed market return series adjusted by the conditional volatility. We focus on the days with extreme values in order to understand the types of news associated with large swings in the market factor. For each of those days, we retrieve the article whose model-implied impact on the market, $I_{z\to\mathrm{Mkt}}^{\top} z(m)$, is the greatest. [19] The retrieved articles in Figure 9 provide a human-readable account of influential events that triggered investor concerns and moved the market factor. For instance, the articles capture the Black Monday in October 1987, the onset of the global financial crisis in 2007, and the worries that the looming Brexit would slow

[18] We calculate $z(m)$ similar to z_τ: $z(m) := \theta_m - \frac{1}{5}\sum_{\iota=1}^{5} \theta_{\tau_m-\iota}$, where $\theta_m \in \Delta^V$ is article m's topic attention levels calculated from LDA; the summation is over the attention levels in the 5 days before the date of article m.

[19] The method can pick more than one influential article to provide a more thorough narrative account of the day. In principle, this approach to narrative retrieval can be potentially applied in real time to "translate" from textual news to quantitative price changes.

Figure 9: Event retrieval for market returns

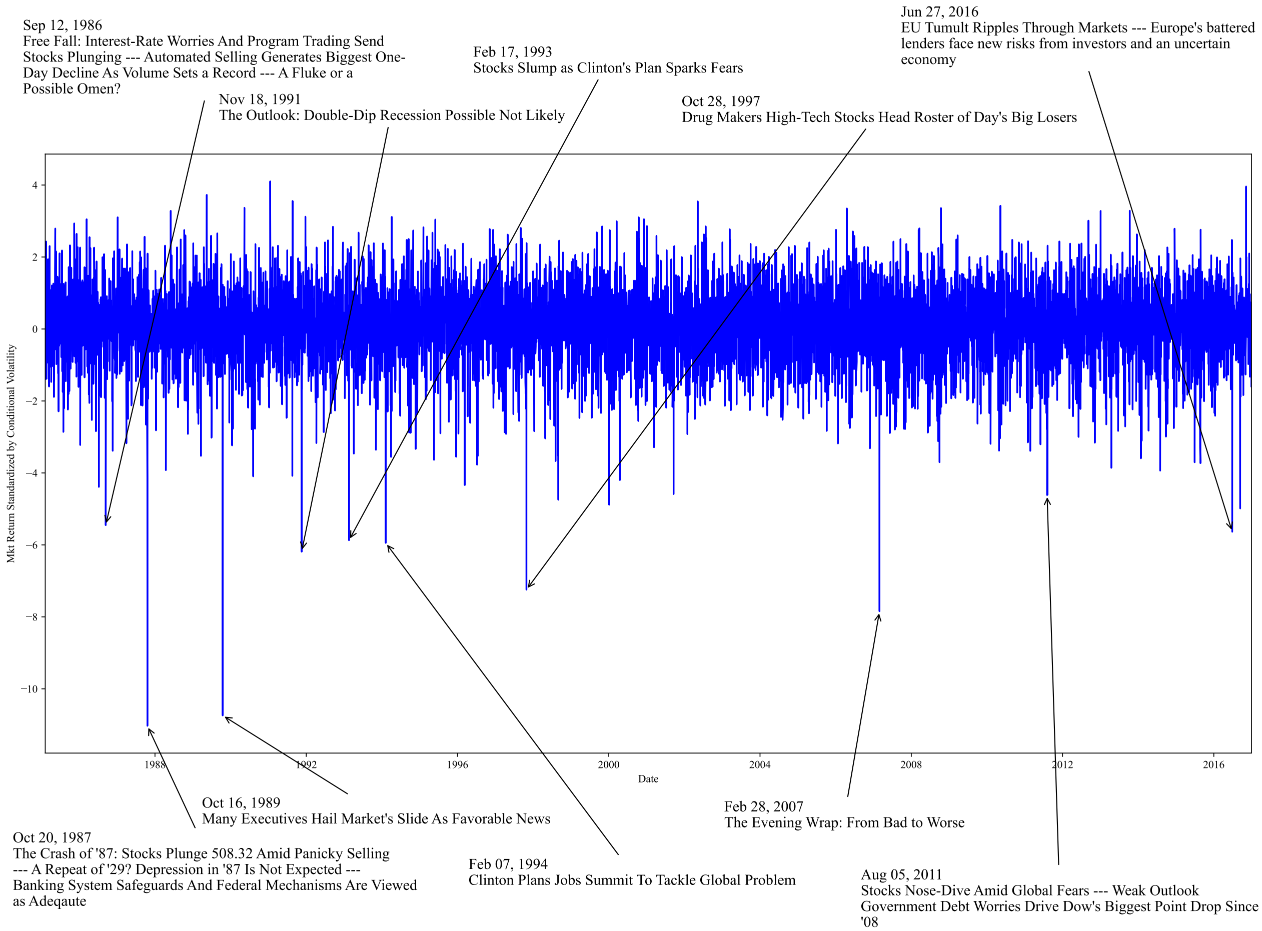

Blue curve: Daily market returns adjusted by conditional volatility, $\frac{f_\tau^{\text{Mkt}}}{\sigma_\tau^{\text{Mkt}}}$, where σ_τ^{Mkt} is EWMA volatility. We tag large spikes with retrieved news events. The article titles are reported in this figure, and the text bodies are in Internet Appendix D.2.

Europe's recovery from the debt crisis in 2016. The news articles connect market fluctuations with investor's forward-looking concerns about economic fundamentals. For example, the event in 1986 is linked to worries about interest rates. Fiscal and labor market policy concerns during the Clinton administration are retrieved in 1993 and 1994. To show more detail, Internet Appendix D.2 reports excerpts of the retrieved articles and displays how the machine "reads" quantitative contents from the textual articles. We color highlight each term according to its impact on the market return ($I_{w\to\text{Mkt}}$), with red for negative and blue for positive impacts. Terms "weak outlook," "global economic slowdown," "government debt worries," "EU Tumult Ripples," etc., are picked up by the model and assigned negative impacts on the market factor. These terms belong to the "Recession" and "Problems" narratives, which are the leading negative drivers of the market as reported in the last panel of Figure 7. Aggregating these terms in an article gives rise to the article's overall impact on the market return.

7 Conclusion

Investors' macroeconomic risk assessments are central to theoretical asset pricing but are empirically difficult to measure. We propose a novel method for estimating ICAPM state variables and asset pricing factors from business news narratives in *The Wall Street Journal*. Our method fuses natural language processing models with empirical asset pricing models. First, we use an LDA topic model to distill a comparatively low-dimensional set of 180 narratives from the much higher-dimensional raw term counts. Next, we use a form of IPCA that generalizes the Fama-MacBeth two-step procedure to deal with the situation in which state variables are not directly observed and must be reduced and selected from many potential proxies (in the form of narrative attention in our application). We devise a Sparse IPCA implementation that uses a group lasso penalty to filter out narratives that are irrelevant for describing the behavior of asset prices, and we find that this regularization approach significantly improves the model's out-of-sample asset pricing performance. Quantitatively, the narrative factors explain returns on anomaly portfolios with smaller pricing errors than benchmark characteristic-based factor models. The narrative factors are portfolios formed with only textual news data, a source of conditioning information that is completely different from the vast literature that uses firm characteristics. The MVE combination of narrative factors achieves higher out-of-sample Sharpe ratios than Fama-French and momentum factors. The narrative factors are empirically consistent with the ICAPM mechanism that state variables are priced sources of risk because they forecast future investment opportunities and components of nonmarketable wealth. Integrating news text in an asset pricing model affords concrete interpretations of the estimated risk factors. Among the narratives studied, we find that attention to the "Recession" narrative has the largest negative impact on the pricing kernel. We show that the underlying raw text associated with the "Recession" narrative covers a variety of events related to real macroeconomic activity. We view news text as a promising source of information for quantitative models of asset prices. Concepts like public information and investor expectations are central to both rational and behavioral models but are difficult to proxy for with traditional data. This paper demonstrates a set of tools for extracting signals from news text that can quantitatively match asset pricing phenomena, while aligning with theoretical restrictions, such as those from the ICAPM.

References

Baker, S. R., N. Bloom, and S. J. Davis. 2016. Measuring economic policy uncertainty. *Q. J. Econ.* 131:1593–636.

Bali, T. G., and R. Engle. 2010. The intertemporal capital asset pricing model with dynamic conditional correlations. *Journal of Monetary Economics* 57:377–90.

Bansal, R., and A. Yaron. 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–509.

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

Bryzgalova, S. 2015. Spurious factors in linear asset pricing models. *Working Paper, Stanford* .

Bybee, L., B. T. Kelly, A. Manela, and D. Xiu. 2021. Business news and business cycles. *Working Paper, Yale University* .

Chen, N.-F., R. Roll, and S. Ross. 1986. Economic forces and the stock market. *Journal of Business* 59:383–403.

Cochrane, J. H. 1996. A cross-sectional test of an investment-based asset pricing model. *Journal of Political Economy* 104:572–621.

———. 2009. *Asset pricing, revised edition.* Princeton, NJ: Princeton University Press.

Cong, L. W., T. Liang, and X. Zhang. 2019. Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. *Working Paper, Cornell University* .

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel. 2020. Hedging climate change news. *Review of Financial Studies* 33:1184–216.

Fama, E. F., and K. R. French. 1996. Multifactor explanations of asset pricing anomalies. *Journal of Finance* 51:55–84.

———. 2016. Dissecting anomalies with a five-factor model. *Review of Financial Studies* 29:69–103.

Fama, E. F., and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81:607–36.

Freyberger, J., A. Neuhierl, and M. Weber. 2020. Dissecting characteristics nonparametrically. *Review of Financial Studies* 33:2326–77.

García, D. 2013. Sentiment during recessions. *Journal of Finance* 68:1267–300.

Gentzkow, M., B. Kelly, and M. Taddy. 2019. Text as data. *Journal of Economic Literature* 57:535–74.

Gentzkow, M., and J. M. Shapiro. 2010. What drives media slant? Evidence from US daily newspapers. *Econometrica* 78:35–71.

Giglio, S., and D. Xiu. 2021. Asset pricing with omitted factors. *Journal of Political Economy* 129:1947–90.

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

Hou, K., C. Xue, and L. Zhang. 2015. Digesting anomalies: An investment approach. *Review of Financial Studies* 28:650–705.

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

Jeon, Y., T. H. McCurdy, and X. Zhao. 2021. News as sources of jumps in stock returns: Evidence from 21 million news articles for 9000 companies. *Journal of Financial Economics* 145:1–17.

Julliard, C. 2007. Labor income risk and asset returns. *Working Paper, London School of Economics* .

Ke, Z. T., B. T. Kelly, and D. Xiu. 2020. Predicting returns with text data. *Working Paper, Harvard University* .

Kelly, B., A. Manela, and A. Moreira. 2021. Text selection. *Journal of Business & Economic Statistics* 39:859–79.

Kelly, B. T., S. Pruitt, and Y. Su. 2017. Instrumented principal component analysis. *Working Paper, Yale SOM* .

———. 2019. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics* 134:501–24.

Lettau, M., and M. Pelger. 2020. Factors that fit the time series and cross-section of stock returns. *Review of Financial Studies* 33:2274–325.

Liu, Y. 2021. Labor-based asset pricing. *Working Paper, University of Rochester* .

Liu, Y., and B. Matthies. 2022. Long run risk: Is it there? *Journal of Finance* 77:1587–633.

Lopez-Lira, A. 2020. Risk factors that matter: Textual analysis of risk disclosures for the cross-section of returns. *Working Paper, University of Florida* .

Loughran, T., and B. Mcdonald. 2016. Textual analysis in accounting and finance: A survey. *Journal of Accounting Research* 54:1187–230.

Maio, P., and P. Santa-Clara. 2012. Multifactor models and their consistency with the ICAPM. *Journal of Financial Economics* 106:586–613.

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 & Economic Statistics* 34:574–89.

Merton, R. C. 1973. An intertemporal capital asset pricing model. *Econometrica* 41:867–87.

Mullainathan, S., and A. Shleifer. 2005. The market for news. *American Economic Review* 95:1031–53.

Pelger, M., and R. Xiong. 2021. Interpretable sparse proximate factors for large dimensions. *Journal of Business & Economic Statistics* 40:1642–64.

Roll, R. 1977. A critique of the asset pricing theory's tests part i: On past and potential testability of the theory. *Journal of Financial Economics* 4:129–76.

Rossi, A. G., and A. Timmermann. 2015. Modeling covariance risk in Merton's ICAPM. *Review of Financial Studies* 28:1428–61.

Steyvers, M., and T. Griffiths. 2007. Probabilistic topic models. *Handbook of Latent Semantic Analysis* 427:424–40.

Stock, J. H., and M. W. Watson. 2001. Vector autoregressions. *Journal of Economic Perspectives* 15:101–15.

Tetlock, P. C. 2007. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance* 62:1139–68.

———. 2011. All the news that's fit to reprint: Do investors react to stale information? *Review of Financial Studies* 24:1481–512.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58:267–88.

Yang, Y., and H. Zou. 2015. A fast unified algorithm for solving group-lasso penalize learning problems. *Statistics and Computing* 25:1129–41.

Yuan, M., and Y. Lin. 2006. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68:49–67.

Zou, H., T. Hastie, and R. Tibshirani. 2012. Sparse principal component analysis. *Journal of Computational and Graphical Statistics* 15:265–86.

Internet Appendix

Narrative Asset Pricing:

Interpretable Systematic Risk Factors from News Text

Leland Bybee
Yale University

Bryan Kelly
Yale University, AQR Capital Management, and NBER

Yinan Su
Johns Hopkins University

A From *WSJ* Archive to Narrative Innovations (z_t), Details

A.1 Textual Data Cleaning and Constructing the Document-Term Matrix

We conduct data processing steps in the following order:

1. Remove all articles prior to January 1984 and after June 2017 (data purchased at the beginning of July 2017).
2. Replace all nonalphabetical characters with an empty string and set the remaining characters to lower-case.
3. Parse article text into a white-space-separated word list retaining the article's word ordering. Exclude single-letter words.
4. Exclude articles with page-citation tags corresponding to any sections other than A, B, C, or missing.
5. Exclude articles corresponding to weekends.
6. Exclude articles with subject tags associated with obviously noneconomic content, such as sports. List of exclusions available from authors on request.
7. Exclude articles with certain headline patterns (such as those associated with data tables or those corresponding to regular sports, leisure, or books columns). List of exclusions available from authors on request.
8. Concatenate articles with the same accession-number as these are chained articles.
9. Exclude articles with less than 100 words.
10. Remove common "stop" words and URL-based terms. List of exclusions is standard but available from authors on request.
11. Lastly, we conduct light lemmatizing of derivative words. The following rules are applied in the order given. In each case, the stemming is only applied if the multiple terms reduce to the same stem.

(a) Replace trailing "sses" with "ss"
(b) Replace trailing "ies" with "y"
(c) Remove trailing "s"
(d) Remove trailing "ly"
(e) Remove trailing "ed." Replace remaining trailing "ed" with "e"
(f) Replace trailing "ing" with "e." For remaining trailing "ing" that follow a pair of identical consonants, remove "ing" and one consonant. Remove remaining trailing "ing."
(g) Remove words with fewer than three letters.

12. From the resultant unigrams, generate the set of bigrams as all pairs of (ordered) adjacent unigrams.
13. Exclude terms (both unigrams and bigrams) appearing in less than 0.1% of articles. The unique set of terms is the corpus vocabulary. Each column of the DTM corresponds to an element of the vocabulary.
14. Convert an article's word list into a vector of counts for each term in the vocabulary. This vector is the row of the DTM corresponding to the article.

A.2 Topic Model Estimation

We estimate θ and ϕ via the Gibbs sampling procedure proposed by Steyvers and Griffiths (2007). This procedure uses an equivalent form to the DGP given in Equation (2), while introducing an intermediate parameter $y_{m,i}$ corresponding to the topic assignment for each word.

$$\omega_{m,i} \sim \text{Mult}(\phi_{y_{m,i}}, 1), \qquad\qquad y_{m,i} \sim \text{Mult}(\theta_m, 1), \tag{13}$$

where $\omega_{m,i}$ is the observed word assignment of the i'th word in article m. The Gibbs sampler generates $\{\theta_m\}, \{\phi_l\}$ as well as the intermediate $\{y_{m,i}\}$ such that,

$$\theta_{m,l} = \frac{\sum_{i=1}^{Len_m} \mathbb{I}(y_{m,i} = l)}{Len_m}, \qquad \phi_{v,l} = \frac{\sum_{m=1}^{M} \sum_{i=1}^{Len_m} \mathbb{I}(\omega_{m,i} = v)\mathbb{I}(y_{m,i} = l)}{\sum_{v=1}^{V} \sum_{m=1}^{M} \sum_{i=1}^{Len_m} \mathbb{I}(\omega_{m,i} = v)\mathbb{I}(y_{m,i} = l)}.^{20,21} \tag{14}$$

While the topic model is estimated using article-level data, we aggregate news attention at the daily frequency. The daily attention level θ_l ($L \times 1$ vector) is formed as

$$\theta_{\tau,l} = \frac{\sum_{m \in \mathcal{M}_\tau} \sum_{i=1}^{Len_m} \mathbb{I}(y_{m,i} = l)}{\sum_{m \in \mathcal{M}_\tau} Len_m} \tag{15}$$

[21] The denominator in $\phi_{v,l}$'s formula can be equivalently written as $\sum_{m=1}^{M} \sum_{i=1}^{Len_m} \mathbb{I}(y_{m,i} = l)$.

[21] Unlike the standard Gibbs sampling procedure from Steyvers and Griffiths (2007), we do not incorporate the prior terms in our estimates of $\hat{\theta}_{m,l}$ and $\hat{\phi}_{l,v}$. See Bybee et al. (2021) for a fuller discussion of this point.

where $\mathcal{M}_\tau$ is the set of articles published on the next morning of calendar day τ.

A.3 Choosing the Number of Topics for LDA

LDA requires specifying the number of topics (L) when estimating the model. To make this decision, we estimate a series of models with various L's from 50 to 250 in increments of 10. We then choose the L that maximized the Bayes factor, the ratio of posterior probabilities for the specified model to the null model. As an additional robustness check, we perform a ten-fold cross-validation, where we iteratively refit the model on a subset of training data and then evaluate it on a held-out validation sample. Figure A.1 reports the results for both methods and suggests that 170–180 is an optimal range.

Figure A.1: Choosing the Number of Topics for LDA

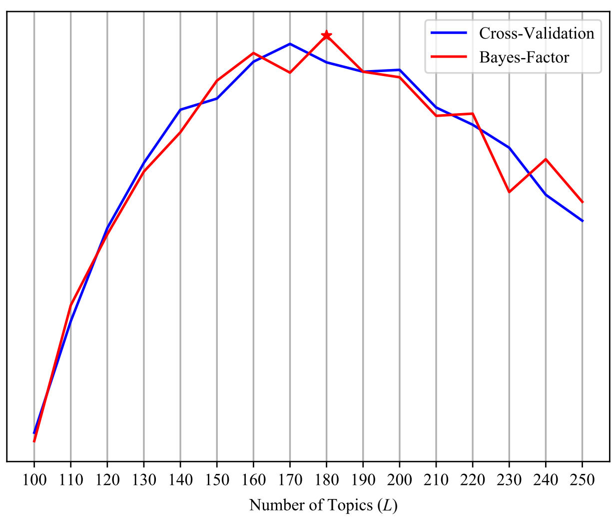

A.4 Visualization of the Selected Narratives

Figure A.2 plots $\theta_{\tau,l}$ time series for the few relevant narratives. The words below each plot are those with large $\phi_{v,l}$. They help display the readable content of the narrative. The red curves visualized the low-frequency variation using the HP filter. Visualization of all other narratives can be found on structureofnews.com.

Figure A.2: Narrative Attention and Keywords

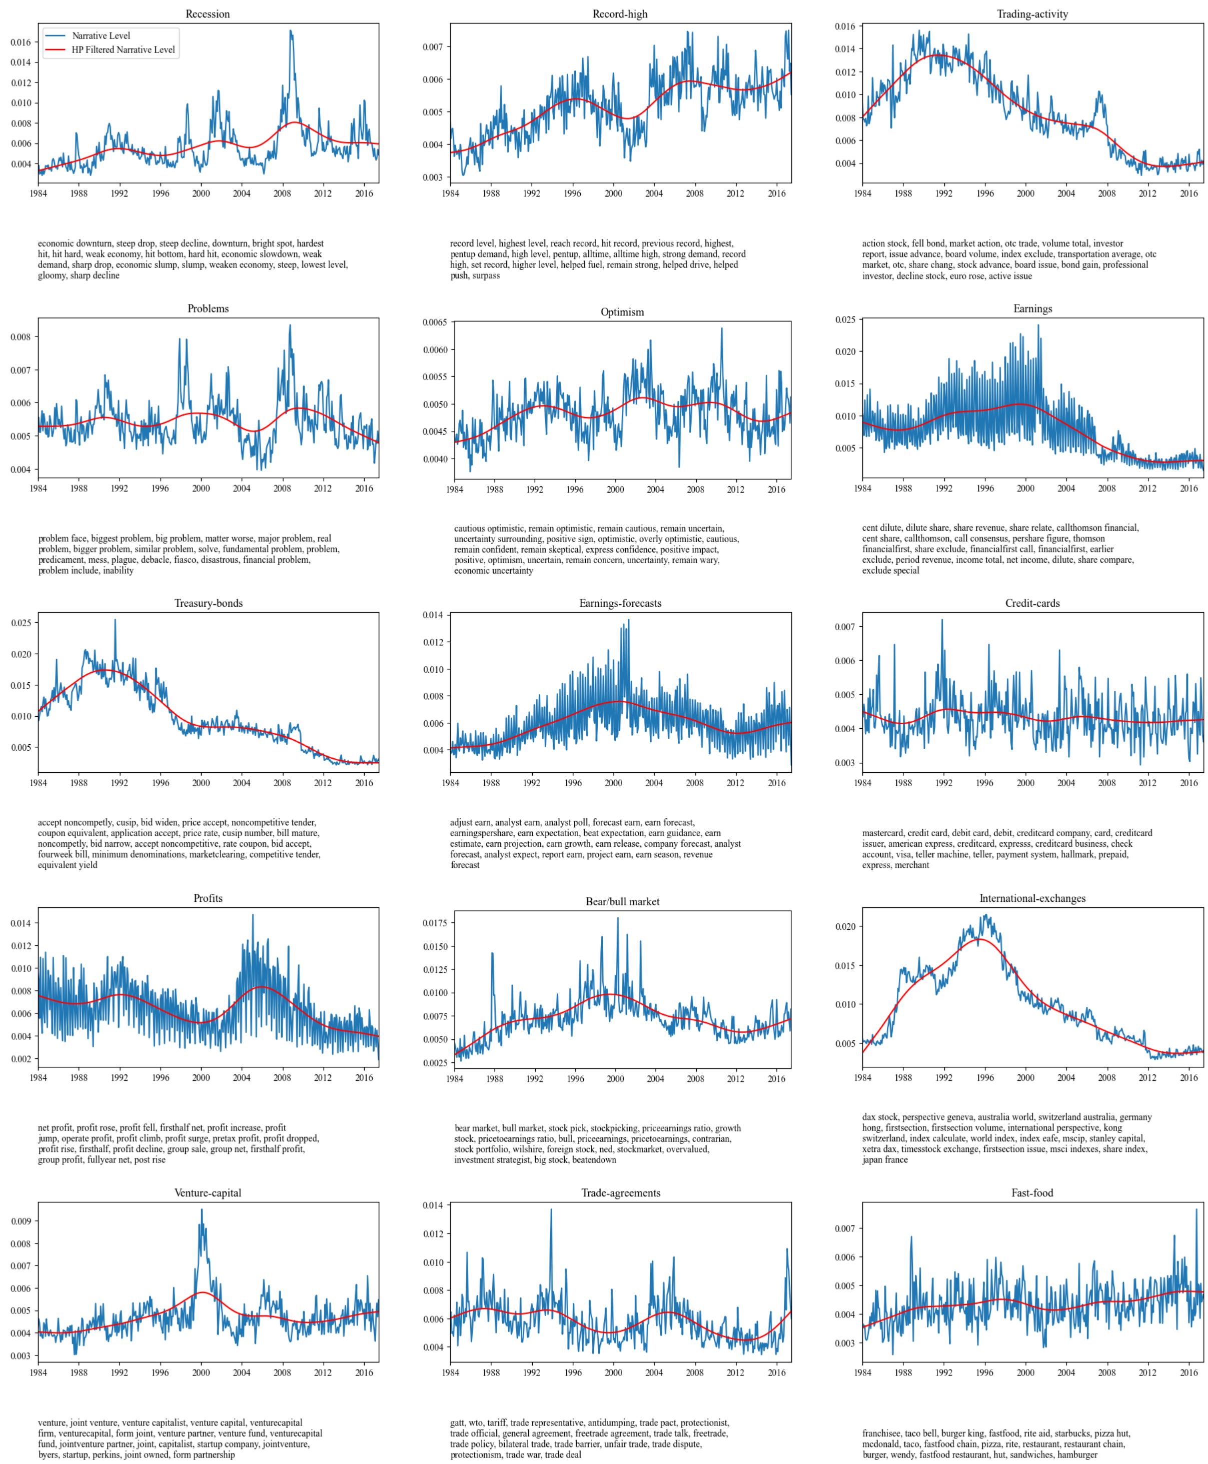

Note: Blue curve: narrative attention level ($\theta_{l,t}$) at the monthly frequency for selected narratives. Red curve: Trend component of the blue curve extracted by the monthly HP-filter. Word list below: terms (unigrams and bigrams) with high $\phi_{v,l}$.

B Additional Details about the Estimation Method

B.1 Technical Details of Steps 1–3

In step 1, we use exponentially decaying weights to estimate the realized narrative covariances. The kernel function is $\kappa(\tau; t) := \xi^{(t-t_\tau)}/(\sum_{\tau<\tau_t} \xi^{(t-t_\tau)})$ for $\tau < \tau_t$ and 0 otherwise, where $\xi := 0.99$, τ_t is the last day of month t, and t_τ is the month of day τ. Next, we provide additional details for a few terms in the target function (8), which is repeated here for ease of reference:

$$\min_{\Gamma,\{f_t\}} \frac{1}{2} \sum_{i,t\in\mathbb{S}} (r_{i,t} - c_{i,t-1}\Gamma f_t)^2 + \lambda N_{\mathbb{S}} \sum_{l=0}^{L} \sigma_l^{\mathrm{c}} \|\Gamma_l\|_2 + \sum_{t\in\mathbb{S}} \|f_t\|_2^2,$$

We extend the linear instrumental mapping to allow for a constant term. To wrap it in the matrix form, we append a constant 1 to the L covariances and stack $1 \times K$ parameters Γ_0 on top of $\widetilde{\Gamma}$. That is, $c_{i,t} := [1, \widehat{cov}_{i,t}]$, $\Gamma = [\Gamma_0; \widetilde{\Gamma}]$, and equivalently $c_{i,t}\Gamma = \Gamma_0 + \widehat{cov}_{i,t}\widetilde{\Gamma}$. Once Sparse IPCA estimates both $\{f\}$ and Γ, the wrap-up step $\langle 3 \rangle$ cuts out $\widetilde{\Gamma}$ from the estimated Γ, from which matrix A is backed out. The formula of backing out A in terms of $\widetilde{\Gamma}$ is $A = \widetilde{\Gamma}\left(\widetilde{\Gamma}^\top\widetilde{\Gamma}\right)^{-1}\Sigma_{\mathrm{ff}}^{-1}$. This is reversed from the definition $\widetilde{\Gamma} := A\left(A^\top A\right)^{-1}\Sigma_{\mathrm{ff}}^{-1}$. To do that, first, $\widetilde{\Gamma}\Sigma_{\mathrm{ff}} = A\left(A^\top A\right)^{-1}$. Then, $\Sigma_{\mathrm{ff}}\widetilde{\Gamma}^\top\widetilde{\Gamma}\Sigma_{\mathrm{ff}} = \left(A^\top A\right)^{-1}$. Plug that into the first expression, $A = \widetilde{\Gamma}\Sigma_{\mathrm{ff}}\left(A^\top A\right)$, then we have the formula of A in terms of $\widetilde{\Gamma}$. We design Sparse IPCA with the row-wise group lasso regularization, $\|\Gamma_l\|_2$ (where $\|\cdot\|_2$ is the Euclidean norm, a.k.a. L_2 norm, such that $\|\Gamma_l\|_2 = \sqrt{\Gamma_{l,1}^2 + \Gamma_{l,2}^2 + \cdots + \Gamma_{l,K}^2}$), instead of the simpler element-wise lasso that additively penalizes the absolute value of each element, $\sum_k |\Gamma_{l,k}|$. It is not meaningful to distinguish the particular factor an instrument matters to, since the K individual factors are rotationally unidentified anyway. Instead, we regularize the norm of Γ_l without distinguishing the direction in which it deviates from $\mathbf{0}_{1\times K}$. Scaling constant σ_l^{c} is the standard deviation of $\widehat{cov}_{l,i,t-1}$ across $\mathbb{S}$ (with σ_0^{c} assigned as 1). The purpose of σ_l^{c} is to place the strength of regularization of each narrative on the same scale. We are effectively regularizing the panel-wise standard deviation of $\|c_{l,i,t-1}\Gamma_l\|_2$, which is the part of $\beta_{i,t}$ variation contributed by instrument l. This adjustment with σ_l^{c} indeed follows the conventional procedure in lasso regressions, where regressors are standardized in order to bring the coefficients to the same level such that the coefficients are subject to the same strength of regularization. We do not want to simply standardize the regressors (in our case $c_{l,i,t-1}$), as it will change the magnitude of the coefficients (in our case Γ). Instead, the strength of the penalty is adjusted by σ_l^c directly, such that the magnitude (and the interpretation) of Γ is preserved. Constant $N_{\mathbb{S}}$ is the number of $\{i,t\}$ observations in the estimation sample panel ($\mathbb{S}$); It rescales λ such that the magnitude of the regularization term keeps up with the model fit term as the sample size changes. The rescaling is necessary with expanding window based OOS construction (appearing in Section 4.2). Otherwise, the effective strength of the same λ value would be varied across samples with different sizes. The third term, $\sum_{t\in\mathbb{S}} \|f_t\|_2^2$, is added for a technical reason. Notice the model fit term (the first term) is invariant if we shrink Γ and expand f_t by the same multiple. Therefore, without the third term's restriction on f_t, the minimization will return an infinitesimal Γ that bypasses its penalties. The third term is to balance the shrinkage effect on Γ, by also regularizing the sum of squared of f_t.

B.2 Numerical Minimization Algorithm of Sparse IPCA

Minimization problem (8) is solved numerically with an Alternating Regularized Least Squares (ARLS) algorithm. It alternates between minimizing over Γ, while holding $\{f_t\}$ fixed and minimizing over $\{f_t\}$, while holding Γ fixed. The ARLS algorithm can be seen as a special case of the Block Coordinate Descent algorithm, with the two parameters as the two "blocks." The process is terminated when the joint target function's descent is small (or when the first-order condition is satisfied) within a numerical tolerance. This algorithm is similar to the unregularized IPCA's alternating least squares (ALS) method, except that the two subproblems become regularized least squares. In particular, the Γ subproblem is a group lasso regression on the $\{i,t\}$ panel:

$$\min_{\Gamma} \frac{1}{2} \sum_{\{i,t\}\in\mathbb{S}} \left(r_{i,t} - \left(c_{i,t-1} \otimes f_t^{\top}\right) \mathrm{vect}\left(\Gamma\right)\right)^2 + \lambda N_{\mathbb{S}} \sum_{l=0}^{L} \sigma_l^c \left\|\Gamma_l\right\|_2 ,$$

where $c_{i,t-1}\otimes f_t^{\top}$ constitutes the $(L+1)K$-variate regressor, and the regression coefficients are $\mathrm{vect}\left(\Gamma\right)$ ($LK \times 1$ vector), the vectorization of Γ that transposes and stacks up the rows of Γ. We solve the group lasso regression numerically using Yang and Zou's (2015) algorithm. The $\{f_t\}$ subproblem is done period by period by solving the cross-sectional ridge regression

$$\min_{f_t} \frac{1}{2} \left\|r_t - C_{t-1}\Gamma f_t\right\|_2^2 + \left\|f_t\right\|_2^2 ,$$

with the analytical solution

$$f_t = \left(\Gamma^\top C_{t-1}^\top C_{t-1}\Gamma + 2\mathbb{I}_K\right)^{-1} \Gamma^\top C_{t-1}^\top r_t, \tag{16}$$

where r_t ($N_t \times 1$ vector) is the cross-section of $r_{i,t}$ at time t. Similarly, C_{t-1} ($N_t \times (L+1)$ matrix) is $c_{i,t-1}$'s stacked up. Notice, according to this formula, the sample estimates of f_t are tradeable portfolio returns constructed from individual assets. As long as C_{t-1} and Γ are estimated with data observable before the start of month t, the portfolio weights are ex ante available. Both the two subproblems easily adapt to unbalanced panels, hence so does Sparse IPCA overall. As a result, Sparse IPCA is applicable to data sets with a large panel and missing entries like stock returns.

C Additional Empirical Results

C.1 Factor Estimation Results

Figure C.1 visualizes the full-sample estimates of f_t and f_t^{MVE} at $K = 3$. Table C.1 reports the

Figure C.1: Factor Time Series

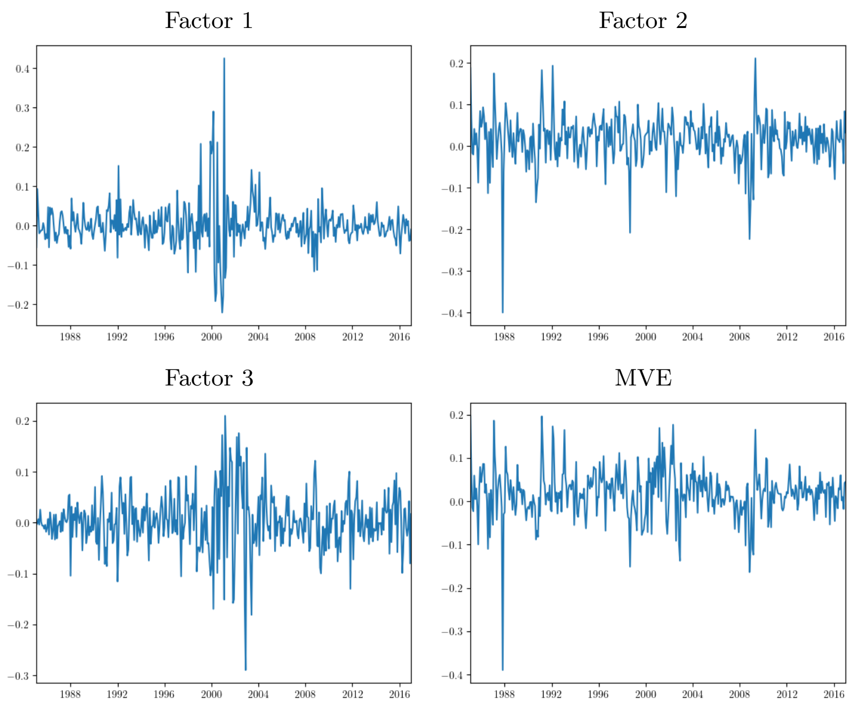

factor's correlation with macroeconomic variables and canonical financial time series at the monthly frequency.

C.2 Instrument Selection Robustness of Sparse IPCA

To further justify that Sparse IPCA is effective in distinguishing relevant and irrelevant instruments, we conduct an experiment with simulated placebo instruments. In addition to the 180 observed narratives (z_τ), we randomly generated an equal amount of placebos to "confuse" the estimation. In detail, we generate each placebo $z_{l,\tau}$ as an i.i.d. normal sequence that matches the times series variance of a corresponding real $z_{l,\tau}$ sequence. We examine whether the narratives that we know for sure are irrelevant can be successfully filtered out. Figure C.2 reports the results of this experiment at the benchmark of $K = 3$. The y-axis plots each instrument's $\max_l^\lambda$ (the maximum λ at which

Table C.1: Narrative Factor's Correlations with Macroeconomic Variables and Financial Time Series

	Factor 1	Factor 2	Factor 3	MVE
GDP growth	0.05	0.09	0.08	0.11
Consumption growth	0.03	0.06	0.05	0.07
10-year Treasury yield change	0.17	0.23	0.22	0.25
CPI change	0.00	-0.00	-0.00	-0.01
Nonfarm payroll growth	-0.05	-0.05	-0.05	-0.05
Unemployment rate change	0.05	0.03	0.04	0.01
Housing starts growth	0.05	0.06	0.06	0.06
Baa-Fed funds spread	0.06	0.04	0.05	0.00
S&P 500 returns	0.35	0.47	0.45	0.48
VIX	-0.53	-0.63	-0.61	-0.60
Mkt	0.30	0.41	0.41	0.40
SMB	0.16	0.16	0.15	0.17
HML	-0.10	-0.09	-0.09	-0.07
RMW	-0.15	-0.19	-0.19	-0.18
CMA	-0.11	-0.15	-0.14	-0.14
MOM	-0.09	-0.15	-0.15	-0.14
Risk free rate	-0.07	-0.09	-0.08	-0.09

narrative l is still selected) in log scale. The black dots mark the real instruments' $\max_l^\lambda$ in the original estimation with only the real narratives (i.e., $L = 180$). The blue and red bars are for real and placebo instruments, respectively, estimated under the specification with placebos (i.e., $L = 360$). They are sorted according to $\max_l^\lambda$ from left to right. The gray horizontal line represents the tuned $\lambda^*_{\mathbb{S}}$ under the specification with placebos ($L = 360$), so that only the bars that are higher than the gray line are eventually selected. The figure shows all of such bars are blue (real). That means none of the 180 placebos is selected under $\lambda^*_{\mathbb{S}}$. The vertical differences between the blue bars and black dots represents the difference in selection before and after introducing placebos. We see vertical differences are narrow, at least for the more relevant ones (the ones on the left with high $\max_l^\lambda$). This implies that Sparse IPCA can effectively filter out irrelevant instruments, and the estimates are largely unaffected by the interference of irrelevant ones.

C.3 λ Tuning Based on Leave-one-out Cross-validation (LOOCV)

We provide the results with the λ tuning based on the leave-one-out cross-validation (LOOCV) method, following the discussion in footnote 8.

Figure C.2: Placebo Test

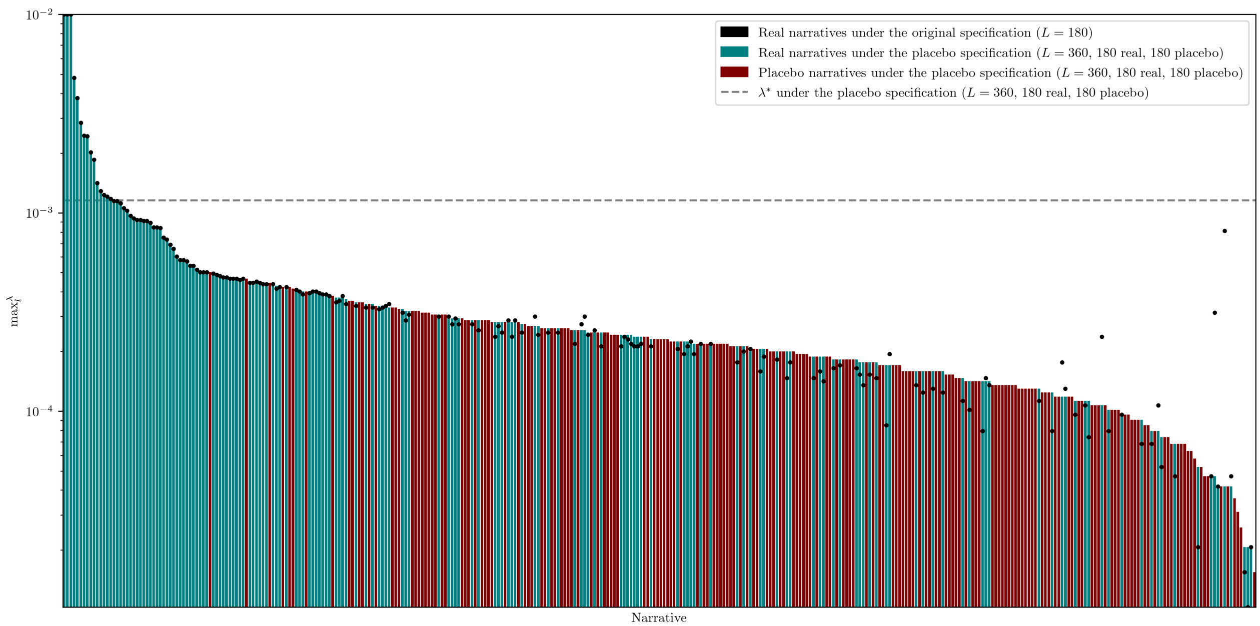

Note: Vertical axis: $\max_l^\lambda$, the maximum λ at which a narrative is still selected. Bars: $\max_l^\lambda$ of 180 real (in green) + 180 placebo (in red) narratives, sorted by $\max_l^\lambda$. Black dots: $\max_l^\lambda$ of 180 real narratives in the original specification without placebos. Gray horizontal line: the tuned $\lambda^*_{\mathbb{S}}$ in the specification with placebos ($L = 360$).

Method: The key difference is three stages of sample partitioning rather than two as in the baseline specification. Within each expanding in-sample $\mathbb{S}$, we further partition it into training samples and validation samples following the LOOCV method. In detail, for each month $t \in \mathbb{S}$, form an "leave-one-out" training sample $\mathbb{S} \setminus \{t\}$. Estimate parameters ($\Gamma$, μ, Σ) in the training sample with different λ values. Form f_t^{MVE} in month t that has been left out. Stitch these f_t^{MVE} together over all the t's in $\mathbb{S}$ to form the testing sequence of investment returns and evaluate the sequence's Sharpe ratio as $\mathrm{SR}^{\mathrm{LOOCV}}(\lambda; \mathbb{S})$. Tune the regularization constant according to $\lambda_{\mathbb{S}}^{*\mathrm{LOOCV}} := \arg\max_\lambda \mathrm{SR}^{\mathrm{LOOCV}}(\lambda; \mathbb{S})$. Once $\lambda_{\mathbb{S}}^{*\mathrm{LOOCV}}$ is tuned from $\mathbb{S}$, the rest of the OOS formation is the same as before: estimate parameters Γ, μ, Σ under $\lambda_{\mathbb{S}}^{*\mathrm{LOOCV}}$ in $\mathbb{S}$, bring the parameters to subsequent months to form the OOS MVE portfolio returns. The reported Sharpe ratios are still OOS, the only difference is it is from the cross-validated $\lambda_{\mathbb{S}}^{*\mathrm{LOOCV}}$ rather than $\lambda^*_{\mathbb{S}}$. We also conduct a joint λ and K tuning based on the LOOCV method. The process is largely the same as above except for also tuning K at the place where λ is tuned. That is, for each in-sample $\mathbb{S}$, we find the pair (λ, K) that yields the highest cross-validated f_t^{MVE} Sharpe ratio. The f_t^{MVE} sequence is formed in the same fashion as above by stitching the leave-one-out t's over $\mathbb{S}$.

Results: Table C.2 adds the results with the LOOCV tuning to the results with standard tuning already reported in Table 2 for comparison. We find $\lambda_{\mathbb{S}}^{*\mathrm{LOOCV}}$ tend to be slightly smaller than $\lambda_{\mathbb{S}}^*$, resulting in greater numbers of selected instruments. However, Sharpe ratio performance does not show any consistent improvement. The Sharpe ratio results are very close. We suggest the reason is the simpler main configuration ($\lambda_{\mathbb{S}}^*$) is already good enough as it does not suffer in-sample overfitting. We also find the joint λ, K tuning is able to achieve the ex post highest level of Sharpe ratio of 1.30, suggesting that the tuning method is also able to make the discrete choice of K.

Table C.2: Sharpe ratios of MVE portfolios

				K			
λ selection method	Statistics	1	2	3	4	5	6
$\lambda = \lambda_{\mathbb{S}}^*$	Sharpe ratio	0.48	1.00	1.10	1.26	1.32	1.31
(copy Table 2)	# narratives	2.9	4.9	12.1	39.1	43.4	61.8
$\lambda = \lambda_{\mathbb{S}}^{*\mathrm{LOOCV}}$	Sharpe ratio	0.48	0.95	1.11	1.21	1.20	1.30
	# Narratives	2.9	7.6	19.3	44.1	49.1	76.8
λ, K joint	Sharpe ratio				1.30		

C.4 Details on Anomaly Portfolios' Pricing Errors

Figure C.3 shows the pricing errors (α) of the 78 characteristic-sorted portfolios under the two factor models, NF3 and FFC6. The anomaly portfolios are listed from top to bottom in the order of $|\alpha|$ under the NF model. The figure shows the signs and the magnitudes of the α's across the tests assets tend to be correlated between FFC6 and NF3. The 1-month momentum portfolio remains the strongest anomaly not explained by both models. It produces large negative α's under both models. This anomaly (also known as the short-term reversal strategy with the sign switched) is known for being an illiquidity-driven phenomenon unrelated to risk exposure and supported by limits to arbitrage. This makes sense as we do not expect the ICAPM mechanism to explain fleeting mispricings that are driven by short-term market frictions. A few other anomalies based on accounting ratios are also not well explained by NF, for example, Depreciation/PP&E, Asset growth, and Cash holdings. Betting-against beta (BAB, denoted as "Beta" in Figure C.3) is among the more interesting portfolios whose expected return is well explained by NF. Both the BAB factor and NF sort stocks on realized risk exposures. While BAB measures exposure to the market factor, our model is more circumspect in how it quantifies systematic risk exposure. We observe that other anomalies that have a risk exposure flavor (beta squared, return volatility, and idiosyncratic return volatility) are also well explained by NF.

C.5 Asset Pricing Robustness under Different Specifications

Table C.3 repeats the asset pricing tests reported in Tables 1 and 2 for factor models estimated with different specifications. The specifications are different from the benchmark reported in the main text only in terms of the instruments supplied to the estimation procedure. Panel A repeats the benchmark specification. Panel B abandons the news narrative-based approach. It uses a set of 129 macroeconomic series from the FRED monthly data series as z_t inputs in calculating $c_{i,t}$. The FRED series are processed following the transformations recommended by the data set documentation (McCracken and Ng, 2016). Panel C changes the benchmark specification for narrative attention innovation calculated against the 5-day moving average, $z_\tau := \theta_\tau - \frac{1}{5}\sum_{\iota=1}^{5}\theta_{\tau-\iota}$, to three different specifications, including "Daily Innovation": $z_\tau := \theta_\tau - \theta_{\tau-1}$; "3-Day Moving-average Innovation": $z_\tau := \theta_\tau - \frac{1}{3}\sum_{\iota=1}^{3}\theta_{\tau-\iota}$; "20-Day Moving-average Innovation": $z_\tau := \theta_\tau - \frac{1}{20}\sum_{\iota=1}^{20}\theta_{\tau-\iota}$.

C.6 Stock Characteristics and the Exposure to Narrative Factors

We examine the characteristics of the stocks with high and low β's under the narrative factor model. We regress $\beta_{i,t}$ (which is instrumented by narrative covariances) on the 78 traditional characteristics across the stock-month panel. Table C.4 reports the regression results. The narrative β's have a strikingly high correspondence with traditional characteristics (the R^2 is around 40%), although the two are from drastically different information sources, namely, textual news covariances rather than "anomaly" characteristics computed from CRSP/Compustat. The signs of the MVE beta's coefficients conform to the existing knowledge on characteristics-based expected return patterns (aka anomalies). In particular, we find a negative coefficient for Size (consistent with the size anomaly), a negative coefficient for Beta (consistent with BAB), and a positive coefficient for Momentum (consistent with the momentum). The usual measure for value, Book-to-market, is insignificant,

Figure C.3: Model Implied α Comparison Between FFC6 and NF3

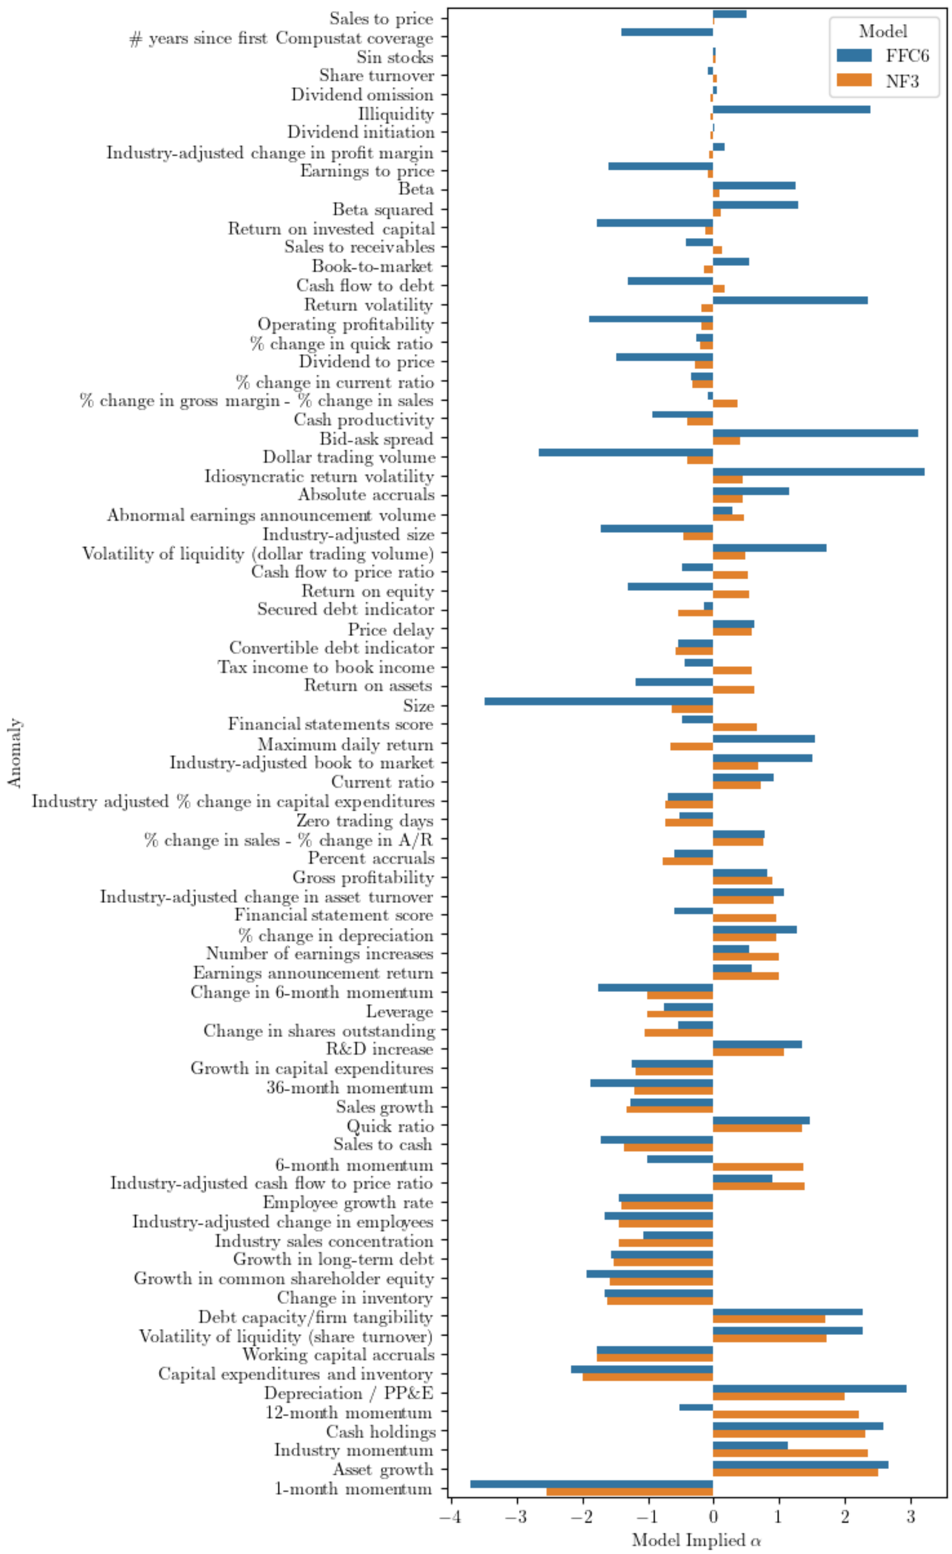

Table C.3: Asset Pricing Performance under Different Specifications

Specification	Statistic	1	2	3	4	5	6
	A. Benchmark specification (repeated from Tables 1 and 2)						
5-Day MA Innov.	OOS Sharpe Ratio	0.48	1.00	1.10	1.26	1.32	1.31
	Avg. Abs. t-Stat. — Anomaly LS	1.29	0.97	0.84	0.92	0.91	0.96
	Avg. Abs. t-Stat. — Size/BM	0.35	0.16	0.25	0.23	0.23	0.25
	B. FRED-MD series rather than news narratives						
FRED	OOS Sharpe Ratio	0.46	0.65	0.67	0.63	0.68	0.70
	Avg. Abs. t-Stat. — Anomaly LS	1.27	0.96	0.95	0.94	0.98	0.99
	Avg. Abs. t-Stat. — Size/BM	0.37	0.20	0.16	0.16	0.15	0.16
	C. Alternative innovation specifications						
Daily Innov.	OOS Sharpe Ratio	0.51	1.08	1.05	1.21	1.12	1.11
	Avg. Abs. t-Stat. — Anomaly LS	1.31	0.96	0.89	0.90	0.87	0.88
	Avg. Abs. t-Stat. — Size/BM	0.34	0.15	0.19	0.21	0.21	0.20
3-Day MA Innov.	OOS Sharpe Ratio	0.49	0.98	1.08	1.22	1.24	1.19
	Avg. Abs. t-Stat. — Anomaly LS	1.30	1.01	0.87	0.92	0.90	0.91
	Avg. Abs. t-Stat. — Size/BM	0.35	0.15	0.23	0.22	0.22	0.22
20-Day MA Innov.	OOS Sharpe Ratio	0.46	0.93	1.00	0.85	1.42	1.44
	Avg. Abs. t-Stat. — Anomaly LS	1.28	0.97	0.87	0.82	0.94	0.94
	Avg. Abs. t-Stat. — Size/BM	0.37	0.17	0.25	0.28	0.30	0.30

See the table explanation in Appendix C.5 and the interpretation of the results in Section 4.3.

while Dividend-to-price is positive and significant, consistent with the value premium. To understand the economic magnitude of these coefficients, let us take the coefficient of MVE β on Size (-0.29) as an example. The characteristics are cross-sectionally rank-standardized from -0.5 to 0.5 at each t. That mean, moving from the smallest to the biggest firm, the "Size" characteristic increases by 1, so the exposure (β) to MVE should drop by 0.29. The implied cross-sectional difference of expected return between the smallest and biggest stocks is $0.29 \times 22.35\%/12 = 0.54\%$ per month (where 22.35% is the annual factor premium of f_t^{MVE}). It matches the magnitude of the small size premium.

Table C.4: Stock Characteristics and the Exposure to Narrative Factors

β with respect to	Factor 1	Factor 2	Factor 3	f^{MVE}
Absolute accruals	-0.02 *	-0.02 ***	0.05 ***	0.01 ***
Working capital accruals	-0.08 ***	-0.06 ***	0.16 ***	0.05 ***
Abnormal earnings announcement volume	-0.02 ***	0.02 ***	-0.05 ***	-0.01 ***
# years since first Compustat coverage	-0.12 ***	-0.18 ***	0.46 ***	0.12 ***
Asset growth	-0.04 **	0.01	-0.03	-0.00
Bid-ask spread	0.10 ***	0.08 ***	-0.20 ***	-0.05 ***
Beta	0.47 ***	0.62 ***	-1.54 ***	-0.38 ***
Beta squared	-0.30 **	0.06	-0.13	-0.01
Book-to-market	-0.10 ***	0.03	-0.07	-0.01
Industry-adjusted book to market	0.07 ***	0.00	-0.02	-0.01
Cash holdings	0.05 ***	0.04 ***	-0.13 ***	-0.04 ***

Cash flow to debt	-0.05 **	0.10 ***	-0.22 ***	-0.04 ***
Cash productivity	-0.03	-0.02 *	0.08 ***	0.03 ***
Cash flow to price ratio	-0.08 ***	-0.11 ***	0.30 ***	0.08 ***
Industry-adjusted cash flow to price ratio	0.07 ***	0.02 **	-0.09 ***	-0.03 ***
Industry-adjusted change in asset turnover	0.00	0.03 ***	-0.08 ***	-0.02 ***
Change in shares outstanding	-0.01	0.00	-0.00	-0.00
Industry-adjusted change in employees	-0.05 ***	-0.01	0.02	0.01
Change in inventory	-0.04 ***	0.02 ***	-0.05 ***	-0.01 **
Change in 6-month momentum	-0.02 **	-0.01 *	0.03 **	0.01 **
Industry-adjusted change in profit margin	0.01	-0.01 *	0.02	0.00
Convertible debt indicator	0.05 **	0.05 ***	-0.12 ***	-0.03 ***
Current ratio	-0.07 *	0.07 ***	-0.16 ***	-0.03 **
Depreciation / PP&E	0.08 ***	0.08 ***	-0.22 ***	-0.06 ***
Dividend initiation	0.20 ***	0.10 ***	-0.23 ***	-0.06 ***
Dividend omission	-0.00	-0.07 ***	0.17 ***	0.04 ***
Dollar trading volume	0.18 ***	0.03	-0.12	-0.05 **
Dividend to price	-0.20 ***	-0.16 ***	0.41 ***	0.10 ***
Earnings announcement return	0.01	0.01 ***	-0.02 ***	-0.00 **
Growth in common shareholder equity	-0.01	-0.03 ***	0.07 ***	0.02 ***
Earnings to price	0.09 ***	0.01	-0.05 *	-0.02 ***
Gross profitability	-0.07 ***	0.00	-0.00	0.00
Growth in capital expenditures	0.04 ***	-0.00	0.00	-0.00
Industry sales concentration	-0.01	0.02	-0.02	0.00
Employee growth rate	0.05 ***	-0.00	0.00	-0.00
Idiosyncratic return volatility	0.31 ***	0.08 ***	-0.22 ***	-0.06 ***
Illiquidity	0.04	-0.02	0.11 *	0.04 **
Industry momentum	0.00	-0.01 **	0.03 **	0.01 *
Capital expenditures and inventory	0.10 ***	-0.02 **	0.06 **	0.01
Leverage	-0.13 ***	0.02	-0.02	0.01
Growth in long-term debt	-0.05 ***	-0.01	0.02	0.01
Maximum daily return	0.04 ***	-0.01	0.01	-0.00
12-month momentum	0.03 ***	-0.10 ***	0.25 ***	0.06 ***
1-month momentum	-0.00	-0.03 ***	0.08 ***	0.02 ***
36-month momentum	0.06 ***	-0.05 ***	0.10 ***	0.02 ***
6-month momentum	0.01	-0.01	0.03	0.01
Financial statement score	-0.12 ***	0.01	-0.01	0.00
Industry-adjusted size	-0.05 **	-0.02	0.03	0.01
Size	-0.33 ***	0.61 ***	-1.41 ***	-0.29 ***
Number of earnings increases	0.02 **	0.02 ***	-0.05 ***	-0.01 ***
Operating profitability	-0.01	0.03 *	-0.07 *	-0.02
Industry adjusted % change in capital expenditures	-0.03 ***	-0.01 ***	0.04 ***	0.01 ***
% change in current ratio	-0.01	0.01	-0.03	-0.01
% change in depreciation	0.00	-0.03 ***	0.07 ***	0.02 ***
% change in gross margin - % change in sales	0.05 ***	0.01 ***	-0.05 ***	-0.02 ***
% change in quick ratio	0.03	-0.03 **	0.08 **	0.02 **
% change in sales - % change in A/R	0.02 ***	-0.01	0.01	0.00
Percent accruals	0.04 **	-0.05 ***	0.10 ***	0.02 **
Price delay	-0.06 ***	0.03 ***	-0.07 ***	-0.01 ***
Financial statements score	-0.05 ***	-0.03 ***	0.08 ***	0.03 ***
Quick ratio	0.04	-0.02	0.06	0.01
R&D increase	-0.01	0.03 ***	-0.07 **	-0.02 **
Return volatility	-0.07 ***	0.05 ***	-0.11 ***	-0.02 ***
Return on assets	0.01	-0.05 ***	0.13 ***	0.03 ***
Return on equity	-0.06 ***	0.01	-0.02	-0.00
Return on invested capital	-0.03	-0.01	0.02	0.00
Sales to cash	0.01	-0.07 ***	0.18 ***	0.04 ***
Sales to receivables	-0.00	-0.01	0.01	-0.00
Secured debt indicator	0.07 ***	0.03 ***	-0.07 ***	-0.02 ***
Sales growth	0.02 *	-0.03 ***	0.07 ***	0.01 **

Sin stocks	-0.03	-0.07	0.18	0.04
Sales to price	-0.08 **	0.07 ***	-0.17 ***	-0.04 ***
Volatility of liquidity (dollar trading volume)	0.11 ***	-0.15 ***	0.30 ***	0.05 ***
Volatility of liquidity (share turnover)	-0.05 ***	0.05 ***	-0.08 ***	-0.01
Debt capacity/firm tangibility	0.01	0.03 **	-0.07 **	-0.02 **
Tax income to book income	-0.01	-0.02 ***	0.05 ***	0.01 **
Share turnover	-0.02	0.06 ***	-0.12 ***	-0.02 **
R^2	20.85	44.61	43.69	38.96

C.7 Investment Performance Combining NF and FFC6

We form MVE portfolios that combine narrative factors and FFC6 portfolios together and show adding narrative factors improves the Sharpe ratio investment performance of the characteristics-sorted portfolios. To connect FFC portfolios with our narrative selection and tuning method, the detailed procedure goes as the following. Each entry in Table 3 represents the combination of K NFs and a particular subset of the FFC6 portfolios. At each in-sample $\mathbb{S}$, for each λ value, the same procedure estimates narrative factors $\{f_t\}$ as well as parameter Γ. Calculate the (annualized) Sharpe ratio of the MVE portfolio of the combination of narrative factors and FFC portfolios as $\mathrm{SR}^{\mathrm{combo}}(\lambda; \mathbb{S})$. In-sample tuning picks $\lambda^*_{\mathrm{combo}} := \arg\max_\lambda \mathrm{SR}^{\mathrm{combo}}(\lambda; \mathbb{S})$. The rest of the OOS portfolio formation is once again the same as before using the in-sample estimated MVE weights under $\lambda^*_{\mathrm{combo}}$, except that the weights are with respect to the combination of K NFs plus a subset of the FFC6 portfolios.

C.8 Details of the VAR Analysis

We analyze the prediction properties of narrative state variables above and beyond other standard macroeconomic predictors in a canonical VAR framework. Through the classical macroeconomic time-series tool of VAR, we can control for expectations of future outcomes summarized in macroeconomic covariates. This provides a conservative test for the predictive effects of narrative states above and beyond standard macroeconomic variables. In particular, we estimate two five-variate VAR(3) systems with the same set of five variables but in different orders. In the first, the variables are, in order, the MVE combination of the narrative attention levels (θ_t^{MVE}), the log level of the S&P 500 index, the Federal Funds rate, log employment, and log industrial production. In the second, we switch the order of the first variables by having log S&P 500 first and θ_t^{MVE} second, keeping the rest unchanged. The structural VAR shocks are orthogonalized with a Cholesky factorization such that the shock's coefficient matrix is lower triangular (c.f. Stock and Watson, 2001), hence the ordering of the variables matter for the result. This identification condition restricts the contemporaneous impulse response of any variable ψ to the shock of another variable ϕ as zero if ϕ is ordered after ψ in the VAR. The analysis follows Baker, Bloom, and Davis (2016), and we update their data to fit our sample window.

Figure C.4 plots the impulse response function (IRF) of future equity market values to a one-standard-deviation shock in θ_t^{MVE} (in red). For comparison, we show the IRF of future equity market values to a shock in current market value (in yellow). The two panels correspond to the results attained switching the order of θ^{MVE} and log S&P. The figures show the θ_t^{MVE} shock has a

Figure C.4: S&P 500's Impulse Response Functions

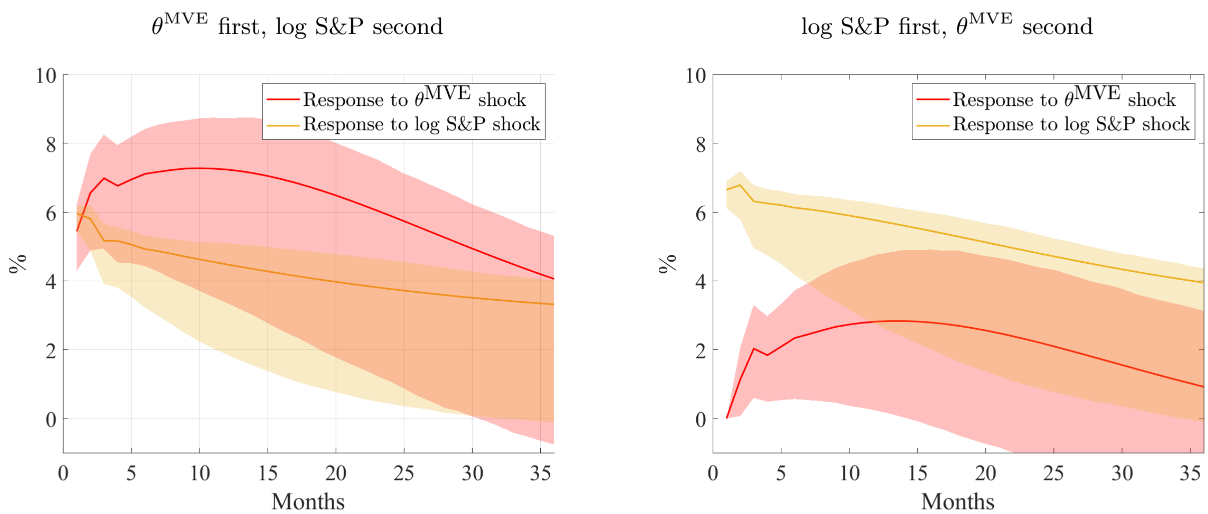

Note: Impulse response functions (IRF) of five-variate three-order VARs. Variable ordering in the left panel: θ^{MVE}, log S&P 500, Fed funds rate, log employment, and log industrial production. The right panel switches the ordering of the first two variables. Red curve: IRF of the logarithm of S&P 500 to standardized shocks of θ^{MVE}. Yellow curve: IRF of the logarithm of S&P 500 to standardized shocks of the logarithm of S&P 500 itself. Shaded band: 90% confidence intervals.

large positive and highly significant impact on market values over the next year. The effect persists beyond 2 years, though it becomes statistically insignificant at this horizon. In the right panel, by listing the market index in front of θ^{MVE}, we force a zero contemporaneous response of the market to the θ^{MVE} impulse. Despite this additional conservatism in assessing the impact of θ^{MVE}, the intertemporal effect from θ^{MVE} on future market value remains positive and significant in the first year but reverts over longer horizons.

D Raw News Text Content

This appendix section provides the text bodies of the news article titles that appear in Table 4 and Figure 9.

D.1 Text Bodies of the Articles in Table 4

The yellow highlighting follows the same rule as that in Table 4. The shades of the yellow highlighter on each term reflect $\phi_{v,l}$, the term's probability for the corresponding narrative. For each article, we provide the excerpt of the first 160 words to save space.

Date: Headline/Text Body
Recession
1993-05-07: Auto Registrations Continued to Slump In Europe Last Month BONN Newcar registrations continued to plunge across Europe last month as a recession in most markets kept consumers away from auto showrooms. In April car registrations in the European Community were off 18.3 from a year earlier according to provisional figures released by the European Automobile Manufacturers Association an industry lobbying group in Brussels. The sharpest declines came in Denmark where 35 fewer cars were registered than in April 1992 and Spain where registrations fell 30.2. The declines were larger than many analysts had expected and bolster the view that Europes auto industry is facing its leanest year in recent memory. These are very hefty declines and will certainly force a lot of us to reexamine our estimates said Bob Barber auto analyst at James Capel Co. in London. Among Europes five biggest car markets Italy Spain Britain France and Germany only Britain is showing signs of life.
2001-04-25: Consumer Confidence Slides on Fears of Layoffs WASHINGTON Consumer confidence is sliding again after stabilizing in March as jobloss fears threaten to undermine what has been surprisingly resilient consumer spending. The Conference Board said its index of consumer confidence fell to 109.2 in April from 116.9 in March. The index is back to its February level which was the lowest since 1996. Consumers gloomier assessment of their present situation accounted for most of the drop. The presentsituation index fell to 155.6 its lowest level since 1997 compared with 167.5 the previous month. The index of expectations slid to 78.2 from 83.1 but remains above its February low of 70.7. The confidence index based on the responses of several thousand households to a monthly questionnaire has fallen 23 since September mostly because consumers have been more pessimistic about the future as layoff announcements have mounted energy costs have risen and stock prices have fallen.
2009-02-19: U.S. News: Housing Starts Hit Lowest Level In Half-Century Housing starts plunged to new lows in January as a large number of vacant homes tight mortgage financing and a deepening recession created the worst housing market in a halfcentury. Meanwhile a report on industrial production showed that a broad collection of manufacturers cut back in January as falling sales and mounting inventories forced them to reduce worker hours and shut down factories. The two reports echoed what has been an emerging theme in the past few months Makers of goods and homes are slashing production as fast as they can to match falling consumer and business demand. There is nothing in these reports that says we are remotely close to turning around said Nigel Gault an economist at forecasting firm IHS Global Insight. Housing starts fell 16.8 in January from a month earlier to a seasonally adjusted annual rate of 466000 units the lowest in at least 50 years according to the Commerce Department.
2011-08-02: World News: Manufacturing Slowdown Adds to Gloom on Economy

LONDON The U.K. manufacturing sector posted an unexpected contraction in July falling to its lowest level in more than two years while activity at eurozone factories slowed to a nearstandstill. The July data released Monday suggested a poor start to the third quarter and damped hopes for a rebound. The U.K. manufacturing purchasing managers index fell to 49.1 in July from 51.4 in June Markit Economics and the Chartered Institute of Purchasing and Supply said. Markit Economics final eurozone manufacturing purchasing managers index fell to 50.4 in July from 52 in June. A reading below 50 indicates activity is contracting. The last time the sector contracted in the U.K. was June 2009 when Britain was still in recession. Eurozone new orders a forwardlooking indicator of activity fell to a reading of 47.6 the lowest since June 2009.

2016-07-08: World News: U.K. Consumer Sentiment Takes Dive
LONDON British consumer confidence suffered its steepest fall in more than two decades after voters decided to take the U.K. out of the European Union an ominous sign that could foreshadow a broader economic slowdown. A longrunning barometer of consumer confidence published by market researchers GfK U.K. Ltd. recorded an 8point fall in early July the biggest monthly drop since 1994 according to results published on Friday. The index fell to minus 9 from minus 1 in June. The survey of 2002 people conducted June 30 to July 5 is the first gauge of household sentiment published since the June 23 referendum. It suggests some consumers have been shaken by the political and market turmoil sparked by the vote including a steep drop in the pound and may rein in spending as uncertainty persists in the months ahead. The Bank of England is bracing for such a slowdown.

Record high

1989-07-05: Japan Vehicle Sales Rise
TOKYO Sales of cars trucks and buses in Japan climbed 15.5 in June from a year earlier to 508319 units the Japan Automobile Dealers Association said. The total was a record for the month surpassing the previous high of 439966 units set in June last year. The brisk June sales were the latest sign of the strength of the domestic auto market which has seen demand surging in recent months. In May and April for instance sales renewed the record for these months. In March they set an alltime high totaling 683299 units. This is totally unexpected one association official said of the June sales. Everybody here is surprised. We didnt think sales would remain so strong for so long.

1994-07-01: Purchasing Managers In U.K. Survey Report Rise for June Orders
LONDON Britains purchasing managers index rose to a record in June the latest monthly survey from the Chartered Institute of Purchasing Supply shows. The index rose from 59.2 in May to 61.4 in June its highest level ever and the fifth month in a row that purchasing managers have reported an upsurge in manufacturing activity. There was significant growth in manufacturing activity during the month overtaking previous record levels and prices were forced up as suppliers failed to meet the increase in demand. The institute said the June index was boosted by record rises in new orders and employment and a strong surge in output. Order books improved across all U.K. industries and regions in June. Increased demand in the domestic market was led by sales promotions and seasonal factors and was supported by a recovery in exports.

1995-02-27: Hiring Outlook For Second Quarter Appears Vigorous
MILWAUKEE Hiring activity in the second quarter should be at the strongest pace since mid1989 a quarterly joboutlook survey suggests. The survey by Manpower Inc. indicates that 23 more employers will be increasing their work forces during the second quarter than cutting jobs. That net hiring gain would be the largest since the third quarter of 1989 and is 3 higher than the projected net hiring increase in last years second quarter. In the first quarter the net hiring increase was 10. About 15000 U.S. employers were surveyed by Manpower a leading temporaryhelp concern. Mitchell S. Fromstein Manpowers chief executive contended that the strong hiring activity projected for the second quarter merely reflects a continuation of heavy hiring last year. The recent increase in the unemployment rate is more related to an increase in people seeking work than to an economic change in hiring trends Mr. Fromstein maintained. Hiring plans are still on an upward track.

2006-01-12: Wall Street Bonuses Hit a Record in 2005
NEW YORK Wall Streets collective bonuses climbed to a projected record of 21.5 billion last year as firms revenue grew according to the New York state comptrollers office. Comptroller Alan Hevesi said 2005s bonus tally was 2 billion more than the old record which was set in 2000. In 2004 Wall Street bonuses came to an estimated 18.6 billion. Last years average bonus was pegged at 125500 also a record Mr. Hevesi said. Revenue at Wall Street firms rose 45 through the first three quarters of 2005 climbing to the highest level since 2000 the year when the stock market peaked Mr. Hevesis office said. The mergers and acquisitions business accounted for most of the surge.

2016-07-20: U.S. News: Home Building Continues Recovery as Demand Rises

WASHINGTON Home building in the U.S. rebounded in June a sign demand for housing continues to firm heading into the second half of the year. Housing starts rose 4.8 from a month earlier to a seasonally adjusted annual rate of 1.189 million in June the Commerce Department said on Tuesday. Home building continues to gradually recover from the housing bust that accompanied the great recession said PNC chief economist Stuart Hoffman. Demand for new singlefamily homes is slowly but steadily improving. That rising demand has led to concerns about the low inventory of new and existing homes on the market which is pushing up prices and could weigh on further expansion. But Tuesdays report showed an estimated 1.015 million homes under construction in June the highest level since February 2008. Junes uptick was driven by a jump in starts in the West and the Northeast two of the pricier regions in the country.

Trading activity

1993-12-30: Industrials Rise A Bit to Record; Bonds Decline
The Dow Jones Industrial Average crept to a thirdstraight record. Bond prices fell and the dollar rose. The industrial average added a scant 0.56 point to 3794.33. Standard Poors 500 stock index fell 0.36 to 470.58 and the Nasdaq Composite Index rose 3.92 to 768.48. The industrial average climbed in early trading nearly cracking the 3800 level but then spent most of the day in negative territory until just before the close. Investors were greeted early by some positive economic news the Commerce Departments index of leading indicators rose 0.5 in November and existinghome sales jumped a betterthanexpected 2.9. The Dow Jones Transportation Average declined after hitting a record on Tuesday. But Larry Rice chief market strategist at Josephthal Lyon Ross wasnt surprised that the average slipped. The history of this market lately is that you get marginal new highs in the averages and then they back off he said.

1994-10-20: Profit News Helps Boost Stock Prices — Indexes Gain Ground Despite Weakness Of Bonds and Dollar
Stock prices moved higher on the strength of robust earnings shrugging off declining bond prices and a weak dollar. The Dow Jones Industrial Average rose 18.50 to 3936.04 marching closer to its record high of 3978.36. The bluechip indicator which was up more than 30 points late in the session has gained an impressive 160.48 points or 4.2 in the past nine sessions. Other indexes have failed to keep pace with the Dow industrials recent climb but yesterday they also marched forward. The Standard Poors 500 stock index jumped 2.62 to 470.28 the New York Stock Exchange Composite Index gained 1.07 to 258.32 and the Nasdaq Composite Index bolstered by strong technology earnings rose 5.81 to 770.62. Analysts said another round of solid thirdquarter earnings highlighted by AMRs report yesterday helped drive stock prices higher.

1996-06-21: Nasdaq Sinks Amid Sell-Off Of Tech Stocks
Bluechip stocks continued to benefit from an investor exodus from stocks of emergingtechnology companies that again sank the Nasdaq Composite Index. Bond prices held steady and the dollar strengthened. The Dow Jones Industrial Average wavered throughout the day then surged just before the closing bell. The Dow ended up 11.08 points at 5659.43 its second straight gain after falling in six of the previous seven sessions. Broad market indexes were mixed. Standard Poors 500 stock index edged up 0.14 to 662.10. The New York Stock Exchange Composite Index slipped 0.36 to 354.96. But investors spooked by more warnings about disappointing secondquarter earnings rushed to unload smaller stocks. The Nasdaq Composite plunged more than 20 points early in the day before recovering to a drop of 11.93 points to 1167.34. With the end of the second quarter a week away money managers are dumping the highest fliers in their portfolios along with shares of any company announcing bad news.

1997-12-09: Blue Chips Fall As Dollar's Rise Causes Concern
Bluechip stocks broke a sixday winning streak with a decline prompted in part by concerns over the strengthening dollar. Bonds fell. The Dow Jones Industrial Average dropped 38.29 points or 0.47 to 8110.84 its first decline after six sessions that lifted the average 354.35 points. The Dow Jones industrials fell more than the SP 500 which dropped only 1.42 or 0.14 to 982.37. In part that was because Dow industrials component CocaCola dropped 2 12 to 63 916 as some analysts lowered their 1998 estimates for the company citing the negative currency translation impact of the strengthening dollar on foreign sales. Boeing another component of the Dow average dropped 1516 to 51 716. It said the Asian economic crisis could cause some airlines to request a total of 60 delivery delays over the next three years.

1998-04-21: Drug Stocks Resume Gains; Blue Chips Fall

Drug and technology stocks soared financial and economically sensitive stocks swooned and the stock market finished mixed. The dollar also finished mixed and bonds declined. Pulling back from its Friday record the Dow Jones Industrial Average lost 25.66 to close at 9141.84. But Standard Poors 500 stock index and the Nasdaq Composite Index both bettered their Friday records. The SP 500 gained just 0.93 to 1123.65 but the technologystockheavy Nasdaq surged 20.54 or 1.1 to close at 1887.14. After slipping last week on disappointing earnings announcements drug stocks resumed their rise with news that Pfizers Viagra impotence pill is a huge seller and that Eli Lillys Evista may prevent breast cancer. Tech stocks particularly Internetrelated shares have regained momentum following recent favorable earnings news. KTel International which has announced that it will sell compact disks and other recordings over the Internet rose 12 1516 to 41 58 it traded at 6 58 earlier this month.

D.2 News Text of Headlines in Figure 9

We highlight each word according to its impact on the market return ($I_{z\to\mathrm{Mkt}}$), with red for negative and blue for positive impacts. The shades of the highlighting reflects the absolute magnitude.

Date (Market Return, %): Headline/Text Body
1986-09-12 (-4.42): Free Fall: Interest-Rate Worries And Program Trading Send Stocks Plunging — Automated Selling Generates Biggest One-Day Decline As Volume Sets a Record — A Fluke or a Possible Omen? The stock market showed its explosive new character as never before as prices plunged on huge volume yesterday and the Dow Jones Industrial Average fell a record 86.61 points. The selloff was triggered by bad news on interest rates. But it picked up momentum as waves of computerdriven Wall Street trading strategies increased the pressure. Such wide unpredictable price swings have become almost commonplace this year. At such high levels investors have to get used to the fact that stocks have taken on the trading characteristics of commodities which have long been known for swift wide swings says Leon Cooperman the head of research at Goldman Sachs Co. The avalanche of selling came just a few days after the Dow Jones industrials climbed to a record 1919.71 Sept. 4 and only two months after the industrials previous record drop of 61.87 points on July 7.
1987-10-20 (-17.44): The Crash of '87: Stocks Plunge 508.32 Amid Panicky Selling — A Repeat of '29? Depression in '87 Is Not Expected — Banking System Safeguards And Federal Mechanisms Are Viewed as Adeqaute Can it happen again On Oct. 28 1929 the stock market fell 12.8 ushering in the Great Depression. While the market plunged 22.6 yesterday economists generally dont expect another depression. I dont think the economy looks like it did in 1929 says George Stigler the winner of the 1982 Nobel Memorial Prize in Economics and a University of Chicago economics professor. The most violent and urgent of factors in the great crash was the collapse of the banking system. That cant happen anymore because of the Federal Deposit Insurance Corp. and additional safeguards. Mr. Stigler like other economists stresses that todays financial system and economic policy mechanisms provide considerably more protection against the type of cascading economic collapse that crippled the nation during the Depression which lasted from 1929 to 1933. During that period the value of the nations output contracted by more than 50 and unemployment rates rose to nearly 25.
1989-10-16 (-5.52): Many Executives Hail Market's Slide As Favorable News HOT SPRINGS Va. Many of the nations highestranking executives saluted Fridays market plunge as an overdue comeuppance for speculators and takeover players. Assuming that the market doesnt head into a bottomless free fall some executives think Fridays action could prove a harbinger of good news as a sign that the leveraged buyout and takeover frenzy of recent years may be abating. This is a reaction to artificial LBO valuations rather than to any fundamentals said John Young chairman of HewlettPackard Co. whose shares dropped 3.125 to 48.125. If we get rid of a lot of that nonsense it will be a big plus. A few of the executives here for the fall meeting of the Business Council a group that meets to discuss national issues were only too happy to personalize their criticism.
1991-11-18 (-3.55): The Outlook: Double-Dip Recession Possible Not Likely

NEW YORK With recoveries like this who needs a recession Last weeks gloomy news from the drop in retail sales to the jump in joblessinsurance claims to Fridays stockmarket plunge hardly inspires confidence that the economy is indeed recovering and will avoid a doubledip recession. Much will depend of course on what occurs in coming weeks in Washington. For now no one knows if the economys modest thirdquarter rise was merely an uptick in a longrunning slump or the start of a sustained recovery. But the evidence on balance still points to the latter eventuality. Doubledip recessions are rare but they do occur. The economy rose briefly amid the yearlong recession of 196970. And many analysts regard the sixmonth recession of 1980 and the 16month recession of 198182 as really a huge doubledipping slump interrupted by a year of economic growth.

1993-02-17 (-2.71): Stocks Slump as Clinton's Plan Sparks Fears
Frightened by the prospect of higher taxes that could choke off the budding economic recovery investors sent the stock market tumbling in a broadbased selloff. Stock prices plunged from the opening bell as investors gave President Clintons economic plans an initial and strong vote of no confidence. Although stock prices recovered slightly from their lows of the day the Dow Jones Industrial Average ended down 82.94 points or 2.44 to 3309.49. It was the biggest oneday point decline since Nov. 15 1991. Standard Poors 500 stock index fell 10.67 or 2.40 to 433.91. Small stocks fared even worse. The Nasdaq Composite Index which has surged in recent months surrendered 25.15 points or 3.64 to 665.39 its worst decline since Oct. 26 1987. Bond investors stayed cool with the Treasurys benchmark 30year bond losing less than a quarter of a point or less than 2.50 for each 1000 face amount. Shorterterm Treasury issues rose modestly.

1994-02-07 (-2.32): Clinton Plans Jobs Summit To Tackle Global Problem
WASHINGTON AP President Clinton will hold an international jobs summit in Detroit next month to tackle the global problem of persistently high unemployment the White House announced. The March 1415 conference will bring together economic labor finance and industry ministers from the Group of Seven industrialized democracies the U.S. Canada Germany Italy Japan France and Britain. The conference will send a message that we intend to confront the challenge of job creation and unemployment not retreat to the economic structures of yesterday the White House announcement said. During a G7 meeting last July in Tokyo Mr. Clinton announced his intention to convene such a conference. He said then the G7 officials would search for the causes and possible answers for this stubbornly high unemployment. The president originally hoped to hold the meeting last fall but it was pushed into 1994 by the crush of other items on his firstyear agenda.

1997-10-28 (-6.58): Drug Makers High-Tech Stocks Head Roster of Day's Big Losers
NEW YORK In a day that saw the largest trading volume ever on the New York Stock Exchange the 30 stocks in the Dow Jones Industrial Average lost a total of 129 billion in market capitalization. From the Dows peak Aug. 6 when the average closed at 8259 and the market capitalization stood at 1.94 trillion the 30 stocks in the industrial average have surrendered 264 billion. Yesterdays Big Board volume totaled 685496330 breaking the previous record of 683800820 set Jan. 23. The outstanding losers in a session chockablock with big losses came from two groups healthcare stocks including both healthcare providers and drug makers and hightechnology stocks. Among the healthcare stocks the biggest loser was Oxford Health Plans which plunged 42 78 to 25 78 after the company said it would post a thirdquarter loss despite expectations that the company would show a profit for the period.

2007-02-28 (-3.43): The Evening Wrap: From Bad to Worse
The worlds stock markets took a dive today beginning with a sharp plunge in Chinas stock market that pulsed through global trading floors and culminated with one of the worst days for U.S. markets in recent memory. The Dow Jones Industrial Average posted a staggering loss of 416.02 points or 3.3 to end at 12216.24. For anyone caught in the turmoil the close of the session couldnt come soon enough. Blue chips began the session deeply lower and continued to step down until about an hour before the end of the session. Then the industrials plunged in a heartbeat. Weaker by about 280 points the index suddenly was down more than 500 points its deepest intraday swoon since the markets reopened after days of inactivity following the Sept. 11 2001 terrorist attacks.

2011-08-05 (-5.04): Stocks Nose-Dive Amid Global Fears — Weak Outlook Government Debt Worries Drive Dow's Biggest Point Drop Since '08

Stocks spiraled downward Thursday as investors buckled under the strain of the global economic slowdown and the failure of policy makers to stabilize financial markets. The selling began in Europe and continued in the U.S. where stocks plunged from the opening bell. The Dow Jones Industrial Average posted its worst point drop since the financial crisis in December 2008 falling 512.76 points or 4.31 to 11383.68. Oil and other commodities were also hammered. Even gold was a safe haven no more as prices fell. Asian markets slid on Friday morning with benchmark indexes in Tokyo Australia South Korea and Hong Kong all falling more than 3 by midday. It was an absolute bloodbath said John Richards head of strategy at RBS Global Banking Markets. There was no one single catalyst for the downdraft traders said. Rather it reflected multiple concerns that have mounted over the past month and came to a head this week. Worries about a U.S.

2016-06-27 (-3.70): EU Tumult Ripples Through Markets — Europe's battered lenders face new risks from investors and an uncertain economy

Just a few years ago Europes banks managed to stagger out of crisis brought on by the Continents debt woes. Britains looming exit from the European Union analysts and investors fear could push them back in. A wide swath of European financial institutions are at risk Hobbled behemoths like Deutsche Bank AG and Credit Suisse Group AG that are limping through difficult turnarounds clusters of regional banks pressured by negative interest rates and banks across Europes weak periphery that are reeling under piles of bad loans. Still wounded from the eurozone debt crisis European banks need investor confidence and steady economic growth to prosper. Brexit risks both. Perhaps most acutely Britains breakaway calls into question the durability of the European Union and the euro. All of a sudden the prospects of Europes political framework disintegrating at its core considered farfetched just days ago have edged up.
