New News is Bad News*

Paul Glasserman† Harry Mamaysky‡ Jimmy Qin§

March 10, 2025

Abstract

An increase in the novelty of news predicts negative stock market returns and negative macroeconomic outcomes over the next year. We quantify news novelty—changes in the distribution of news text—through an entropy measure, calculated using a recurrent neural network applied to a large news corpus. Entropy is a better out-of-sample predictor of market returns than a collection of standard measures. Cross-sectional entropy exposure carries a negative risk premium, suggesting that assets that positively covary with entropy hedge the aggregate risk associated with shifting news language. Entropy risk cannot be explained by existing long-short factors.

Keywords: entropy; natural language processing; news articles; empirical asset pricing
JEL Codes: E30, G12, G14, G17
Online Appendix: `https://sites.google.com/view/hmamaysky`

*We thank Jeremias Huber, Yu Yu (discussant), and talk participants from Columbia, Reichman, the SQA/CFA 7th Annual Data Science in Finance Conference, BlackRock, the 2023 Wolfe Research Global Quantitative and Macro Investment Conference, and CUNY for helpful comments. We thank the authors of van Binsbergen, Bryzgalova, Mukhopadhyay and Sharma (2024) for sharing their sentiment data.

†Columbia Business School, pg20@columbia.edu.

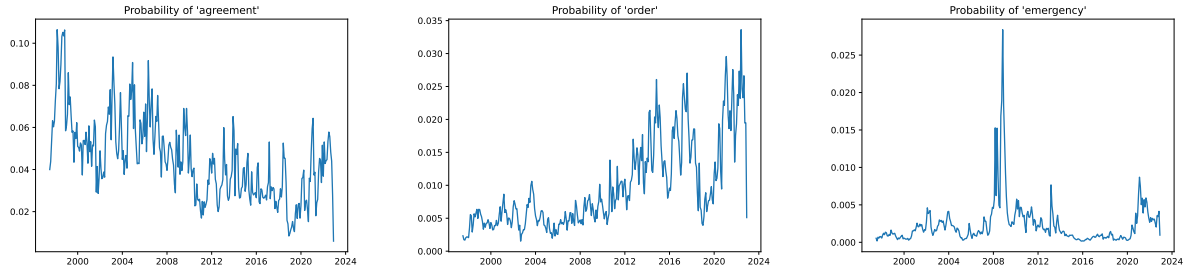
‡Columbia Business School, hm2646@columbia.edu.

§UT Dallas Naveen Jindal School of Management, jimmy.qin@utdallas.edu. Supported by the W. Edwards Deming Center for Quality, Productivity and Competitiveness at Columbia Business School.

1 Introduction

Several studies have documented that the sentiment of news text forecasts short-term changes in asset prices, with negative news forecasting negative returns. We find that a change in the distribution of news text—a measure of the novelty or unusualness of the news rather than its sentiment—forecasts negative market returns and macroeconomic outcomes over the subsequent year. Consistent with this pattern, we find that assets that positively covary with our measure carry a negative risk premium: investors accept lower compensation to hold assets that hedge the risk associated with a shift in the language of news.

To motivate the idea, consider the string of words "the securities and exchange commission issued an." What is the likelihood that the next word in the sentence is "agreement," "order," "emergency," or any other English language word? Using a recurrent neural network re-estimated in rolling windows over 1.6 million Reuters news articles spanning a period of 27 years, we can quantify exactly how this distribution evolves over time. For example, the likelihood that the next word is "agreement" has a decreasing time trend, while the likelihood that the next word is "order" has an increasing time trend. The likelihood that the subsequent word is "emergency" spikes during the financial crisis and again during the COVID-19 pandemic. Although in principle such shifts in language structure could be neutral, we find that they generally contain negative news about markets and the economy. This finding leads us to explore and subsequently uncover a rich structure of risk pricing associated with exposures to the changing distribution of news language.



The probability of "agreement", "order", and "emergency" following the string of words "the securities and exchange commission issued an." The distributional shift over time is estimated by a recurrent neural network.

We refer to our measure as entropy; it can also be seen as a measure of the cross entropy (or dissimilarity) between the current and past distributions of news text. Glasserman and Mamaysky (2019) used an entropy measure to forecast volatility, rather than the direction of returns. To estimate conditional word probabilities, they used the empirical

frequencies of word sequences in what is known as an n-gram method in the natural language processing (NLP) literature; see Jurafsky and Martin (2023). Exact phrases, like in the above example, occur rarely if they are more than a few words long, limiting the accuracy of the n-gram method.

Here we use a far more powerful recurrent neural network (RNN) model with a long short-term memory (LSTM) architecture (Mikolov et al. 2010; Hochreiter and Schmidhuber 1997) to estimate conditional word probabilities and thus to calculate entropy. Whereas in practice the n-gram approach conditions on only three or four words prior to the predicted word, an RNN model captures far more contextual information in assigning probabilities to words. We apply the RNN model with word embeddings that encode rich information about the relationships between words, adding to the power of the method. Unlike the latest and most complex large language models—like ChatGPT, Gemini, Claude, and their relatives—our RNN can be trained relatively quickly. This allows us to retrain our RNN monthly in rolling windows over our news data, and then use the most recent model to assign an entropy score to the next month's news articles, which reflects the degree to which the text in these articles deviates from the RNN model's conditional word distributions. We can then assess the in- *and* out-of-sample forecasting performance of our entropy measure.

Each month we calculate *ENT*, the change in our raw entropy measure over the previous 12 months. For simplicity, we often refer to *ENT* simply as *entropy* in the rest of the paper. In in-sample regressions, we find that a one standard deviation increase in *ENT* forecasts a 2.3% decline in cumulative S&P 500 returns over the next 12 months, even after controlling for a large number of alternative predictors. The effect is highly statistically significant.

To check out-of-sample predictability, we compare univariate forecasts for a large number of predictors; the out-of-sample R-squared for *ENT* is positive at multiple horizons but overwhelmingly negative for the other predictors. In particular, entropy is the best out-of-sample predictor of market returns when compared to a long list of variables proposed in prior studies, including the inverse of the cyclically adjusted price-to-earnings ratio of Campbell and Shiller (1988) *CAPE*, some of the variables tested in Welch and Goyal (2008), Campbell and Thompson (2008), and David and Veronesi (2022), as well as economic policy uncertainty (*EPU*, Baker et al. 2016), the *VIX* (Martin 2017), and the aggregate consumption to wealth ratio (*CAY*) of Lettau and Ludvigson (2001). We estimate our RNN model in rolling windows in order to run out-of-sample tests. Given the extensive literature devoted to identifying successful out-of-sample risk premium fore-

casters, the out-of-sample forecasting performance of *ENT* is remarkable.

If an increase in *ENT* forecasts a decline in the stock market, then an asset that positively covaries with *ENT* provides a hedge against this deterioration in the investment opportunity set. We, therefore, expect *ENT* to carry a negative risk premium: investors should accept lower expected returns on assets that covary positively with *ENT* because of the hedge they provide. This ICAPM-style (Merton 1973) argument provides the intuition underlying the sign consistency property proposed by Maio and Santa-Clara (2012). Through cross-sectional regressions that control for standard factors, we estimate a monthly risk premium for *ENT* that is indeed negative, ranging from -0.06% to -0.09% per unit standard deviation of cross-sectional entropy loading, depending on the choice of test assets. While the magnitude of this risk premium is in line with that of the other factors, none of the other cross-sectional pricing factors satisfy the Maio and Santa-Clara (2012) consistency property that a factor's market return forecasting direction and cross-sectional risk premium have the same sign.

The structure of our *ENT* measure allows its decomposition into two components, which we refer to as *news updates* and *model updates*. The first component measures the change in entropy, holding the distributional model of text fixed, and the second measures the change in the distributional model. The first component is thus more focused on the novelty of the text in the current month. We find that the out-of-sample predictive power of entropy stems primarily from the first component, which robustly forecasts future returns out-of-sample and carries a negative cross-sectional risk premium.

Entropy and its news innovation component forecast year-ahead macroeconomic outcomes. Increased entropy is associated with future increases in unemployment and the *VIX*, and with future declines in industrial production, inflation, interest rates, and corporate earnings. Consistent with an ICAPM-style hedging argument (see Boons 2016), market participants are willing to accept lower expected returns for securities that hedge entropy risk. That entropy also negatively forecasts aggregate market returns suggests that market dividends themselves provide a hedge against reduced consumption in future bad times forecasted by increased entropy.

To explore this further, we formulate a simple theoretical model that produces our main findings: an increase in entropy forecasts weaker economic conditions and a decline in the stock market. The model summarizes the state of the economy through a scalar process. This economic state variable influences the news produced by a media sector, with news text showing greater dispersion under poorer economic conditions; this feature is motivated by prior work finding greater dispersion of economic forecasts in recessions.

The economic state also influences the prices of risky assets in two ways: it affects the level of future dividends, and it enters into the stochastic discount factor used to price future dividends. We derive conditions on model parameters under which news entropy negatively covaries with future economic conditions and future stock returns.

We show that returns on the replicating portfolio for our entropy measure cannot be explained by the multitude of existing long-short factors (Harvey et al., 2016; Feng et al., 2020). Jensen et al. (2021) argue that most factor-based asset pricing results are replicable, and in doing so construct a database of over 150 long-short factor portfolios.[1] Using this database of over 150 long-short factors we show that our entropy-change replicating portfolio is among the long-short portfolios that are most poorly explained by existing factors. Furthermore, entropy appears largely unrelated to several measures of economic uncertainty (Baker et al., 2016; Bekaert et al., 2022; Jurado et al., 2015; Azzimonti, 2018; David and Veronesi, 2022) and contains information that is distinct from news sentiment. Finally, using a projection of entropy onto the dispersion of U.S. regional sentiment data from van Binsbergen et al. (2024), we extend our entropy measure back to 1970 and show that its macro and market forecasting power remains robust.

1.1 Literature Review

Our paper is related to the literature examining the impact of text similarity in financial markets. Tetlock (2011) measures news staleness as the similarity of news stories to prior news stories about the same firm, and finds that firms' stock returns are less responsive to stale news but that retail investors overreact to stale news. Cohen et al. (2020) show that firms with quarter-over-quarter changes in the language of their 10-Ks and 10-Qs experience significantly lower future stock returns relative to non-changers, which they interpret as market underreaction to the information content of corporate reports. Hoberg and Phillips (2016, 2018) construct firm peer groups using textual similarity of product descriptions contained in 10-K reports and show that there are cross-momentum effects within text-based peer groups.

The traditional definition of dissimilarity has focused on differences in word counts between small groups of related documents. Glasserman and Mamaysky (2019) extend this concept to measure how unusual an article's language is relative to the text of *all* prior articles. They show that, at the firm-level and in aggregate, unusual news forecasts increases in future realized and implied volatility. Calomiris and Mamaysky (2019) find

[1] Chen and Zimmerman (2021) make a similar point and also provide factor-replication code and data.

that this measure forecasts country-level returns. Prior measures of dissimilarity or informativeness suffer from several problems: word counts cannot account for article context or synonyms, while the n-gram approach of Glasserman and Mamaysky (2019) suffers from the sparsity issue that many feasible n-grams are never observed, especially for large n. Our entropy methodology offers a dramatic improvement over prior efforts by using word embeddings inside a deep-learning framework.

Our paper is related to a long literature in computer science and linguistics which aims to create probabilistic models for text using deep learning. Hochreiter and Schmidhuber (1997) introduce the long short-term memory (LSTM) architecture, which revolutionized the ability of recurrent neural networks to successfully model large text corpora (Charniak 2019). Pennington et al. (2014) introduces the GloVe model of word embeddings which map words into a high-dimensional space where vector operations reflect the semantic content of words. Transformer-based models, like ChatGPT (Radford et al. 2019) and its peers, have many more parameters than our model. Because of their size, these large language models are trained using *all* available text data up to the date of their release and their use for out-of-sample testing may be problematic.[2] Our approach, while restricted to a smaller model and data set, is trained in rolling windows, and makes forecasts using only data that would have been available to market participants in real-time, thus allowing for out-of-sample testing.

Several papers are concerned with aggregate measures of text or news flow. Jiang et al. (2019) show that an index of aggregate manager sentiment, based on the text of firms' 10-Ks, 10-Qs, and conference calls, negatively forecasts stock returns, suggesting stock investors overreact to the information content of management communication. Using a dictionary-based approach, Shapiro et al. (2022) construct an aggregate economic sentiment measure from articles in 16 major U.S. newspapers and show that positive economic news is associated with growth in future consumption, output, and real rates. van Binsbergen et al. (2024) construct a 170-year-long measure of U.S. economic sentiment at the national and state levels, and show that news coverage has become increasingly negative across all states in recent decades. Baker et al. (2016) introduce the economic policy uncertainty (EPU) measures which count the number of times terms like *uncertainty*, *economic*, and *legislation* (and other similar word combinations) appear in close proximity to each other in articles from 10 major U.S. newspapers. They then use a weighted average of the frequency of major news discussing economic policy uncertainty, expiring

[2]See Glasserman and Lin (2024), Sarkar and Vafa (2024), and Levy (2025) for a discussion of these issues.

tax provisions, and forecaster disagreement about government purchases to construct their measure. They show that *EPU* negatively forecasts investment, output, and employment, and at the firm-level *EPU* predicts "greater stock price volatility and reduced investment and employment" for firms most heavily exposed to government policy.

Brogaard and Detzel (2015) test the asset pricing implications of *EPU* and show that it positively forecasts market returns but carries a negative cross-sectional price of risk.[3] Based on these findings, *EPU* does not satisfy the Maio and Santa-Clara (2012) sign property, which suggests that if *EPU* positively forecasts future returns, it should have a positive cross-sectional price of risk. When controlling for other forecasting variables and for *ENT*, we also find that *EPU* forecasts returns positively and that it carries a negative risk premium, but the forecasting results are not consistent and don't hold out-of-sample, and the negative risk premium is not significant in all specifications. In contrast, *ENT* is a significant return forecaster in and out of sample, satisfies the ICAPM-style sign property, and has a significant price of risk across all specifications.

It is natural to contrast our entropy measure with measures of model uncertainty. Among others, Hansen and Sargent (2008), Aït-Sahalia et al. (2021), and Brenner and Izhakian (2018) link risk premia to uncertainty, distinguishing uncertainty about the true model and ambiguity about the true outcome probabilities from the market or economic volatility. An increase in *ENT* might suggest heightened uncertainty, but theory commonly associates uncertainty with a positive risk premium, and we find a negative risk premium for *ENT*. However, we formulate a simple theoretical model where entropy forecasts both negative economic outcomes *and* negative market returns.

The remainder of this paper proceeds as follows. Section 2 discusses the data we use. In Section 3, we define the entropy measure and compare several estimation methods. Section 4 provides an empirical assessment of the forecasting power of entropy for future market returns. Section 5 investigates the cross-sectional association between aggregate entropy and returns. Section 6 decomposes entropy into news and model innovation components, investigates the macroeconomic forecasting properties of entropy, and discusses potential channels for why entropy forecasts market returns. Section 7 performs some robustness checks and Section 8 concludes. We develop our theoretical model in the appendix (Section A). Technical details and supplementary results are included in an Online Appendix.

[3]Lin (2022) is a related paper that shows that another economic uncertainty index positively forecasts stock market volatility, but does not earn a significant price of covariance risk in a cross-section of 25 Fama-French size and book-to-market sorted portfolios, which is also a violation of the ICAPM sign property proposed in Maio and Santa-Clara (2012).

2 Data and Variable Construction

We construct our entropy measure using the Thomson Reuters News Feed Direct archive from January 1996 to December 2022. We removed articles that represent multiple rewrites of the same initial story, retaining only the first article in a given chain.[4] We only kept news articles that were written in English and discussed S&P500 companies.[5] We excluded articles with headlines containing certain keywords such as "shh margin trading", "nyse", "imbalance", "machine generated", and "research alert;" these keywords typically signal messages that are not really news articles or exhibit irregularities. For example, "machine generated" articles were produced using algorithms or automated processes, rather than being created manually by human writers and editors. As such, their content may not have the same level of context, nuance, or analysis as content created by human writers. We further discarded articles with fewer than 30 words. This filtering process leaves us with 1,642,517 articles. Figure 1 depicts the number of articles and average article length per month, respectively. The monthly article count increased from 1996 and peaked in the early 2000s, remaining steady until 2010 before gradually decreasing. The average number of words per article steadily increased from 1996, reaching an approximate steady-state level of 250 words per article in the early 2000s, with fluctuations observed after 2013.

We retrieve the Chicago Board Options Exchange's CBOE Volatility Index (*VIX*) from the Federal Reserve Economic Data (FRED) server. We write *VIX2* for the square of *VIX*. We include additional uncertainty measures from Bekaert et al. (2022) (*BEX*) and Jurado et al. (2015) (*JLN*), and the Azzimonti (2018) Partisan Conflict Index (*PCI*), all of which can be retrieved from the David and Veronesi (2022) review article data set. As a proxy for the risk-free rate, we use the market yield on U.S. Treasury securities at ten-year (*DGS10*) and two-year (*DGS2*) constant maturity, also obtained from FRED. We calculate the 2s-10s spread, *DGS10-2*, as *DGS10* minus *DGS2*. We include the monthly S&P 500 dividend yield (*DY*), monthly cyclically adjusted PE ratio for the S&P 500 (*CAPE*), and the consumption-wealth ratio (*CAY*).[6]

We also use daily and monthly returns of Fama and French (2015) five-factor model

[4]Thomson Reuters tracks articles by assigning a Primary News Access Code (PNAC). Articles that share the same PNAC are duplicates. This generally happens when there is an update to the coverage of the same event.

[5]A detailed description of the article selection methodology is in Glasserman et al. (2025). The methodology involves mapping Reuters company names to CRSP company names using fuzzy matching.

[6]The *CAY* variable obtained from Martin Lettau's website ends in 2019 Q3. We extrapolate it to the end of 2022 using the method described in Section A3 of the Online Appendix.

(*MKT*, *SMB*, *HML*, *RMW*, *CMA*) and the momentum factor (*UMD*).[7] *MKT* refers to the CRSP value-weighted index return net of the risk-free rate. We write *Return1* for the market return of the previous month, *Return12* for the cumulative market return of the previous 12 months excluding the most recent month, and *Return60* for the cumulative market return of the previous 60 months including the most recent month. Similarly, we construct *SMB60* (*HML60*, *RMW60*, *CMA60*, *UMD60*) as the *SMB* (*HML*, *RMW*, *CMA*, *UMD*) factor return of the previous 60 months. We use the last 60-month lagged factor returns in the forecasting regressions in Section 4 to convert factor returns to state variables as suggested by Maio and Santa-Clara (2012).[8] In Section 7, we also use factor returns obtained from Jensen et al. (2022).

3 The Entropy Measure

We build upon and extend the idea from Glasserman and Mamaysky (2019) to construct an entropy score for the unusualness of news text as a market signal. A text is considered unusual if it has a low likelihood relative to a model of language probability. This problem has been studied in the natural language processing literature on word prediction; see, in particular, Chapter 3 of Jurafsky and Martin (2023).

3.1 The Entropy of Text

Our goal is to estimate the probability of a new set of articles (an *evaluation text*) under a probability model $\mathbb{P}$ estimated from past articles (referred to as the *reference text* or *training corpus*). We can represent an evaluation text as a sequence of N words $w_1 w_2 \cdots w_N$. Its probability is given by the product of conditional probabilities

$$\mathbb{P}(w_1 \cdots w_N) = \prod_{k=1}^{N} \mathbb{P}(w_k | w_1 \cdots w_{k-1}), \tag{1}$$

[7] They are obtained from the Nasdaq Data Link `https://data.nasdaq.com/`, Robert Shiller's website `http://www.econ.yale.edu/~shiller/data.htm`, Martin Lettau's website `https://sites.google.com/view/martinlettau/data`, and Ken French's website `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/biography.html` respectively.

[8] Maio and Santa-Clara (2012) use 60-month lagged returns for two factors, but use a different construction for converting *SMB* and *HML* to state variables. For consistency, we use the 60-month lagged returns for all factors.

in which the first factor is $\mathbb{P}(w_1)$, the unconditional marginal probability of word w_1. The average negative log probability per word is then given by

$$-\frac{1}{N}\ln \mathbb{P}(w_1 \cdots w_N) = -\frac{1}{N}\sum_{k=1}^{N} \ln \mathbb{P}(w_k | w_1 \cdots w_{k-1}). \tag{2}$$

If $\mathbb{P}$ correctly models the word-generating process, and if this process is stationary and ergodic, then (2) converges to the entropy of the process as N increases (see Section 3.8 of Jurafsky and Martin 2023). For fixed N, (2) measures the cross-entropy between the empirical distribution of the evaluation text and the model $\mathbb{P}$. For brevity, we will refer to (2) as the entropy of the evaluation text $w_1 \cdots w_N$. High entropy means low probability and thus signals an unusual text, relative to $\mathbb{P}$.

To calculate entropy, we need a model or estimate for $\mathbb{P}$. The n-gram approach (often with $n \in \{4, 5\}$) truncates the conditioning on the right side of (2) to the $n-1$ words immediately preceding w_k and thus misses additional context provided by words that came earlier. The n-gram approach estimates conditional probabilities of the form $\mathbb{P}(w_k | w_{k-n+1} \cdots w_{k-1})$ as the ratio of the number of occurrences of the sequences $w_{k-n+1} \cdots w_{k-1} w_k$ and $w_{k-n+1} \cdots w_{k-1}$, with adjustments to the numerator and denominator for strings that were never previously observed, which we refer to as the *sparsity* issue. Choosing a larger n captures more conditioning information but makes it more likely that the conditioning sequence has rarely or never been observed before. The inability to condition on a large number of prior words leads to approximation errors in (2).

3.2 A Neural Network Approach

A neural network approach to modeling the probability $\mathbb{P}$ addresses the two shortcomings—approximation and sparsity—of the n-gram method. Recurrent Neural Networks (RNNs) are a type of neural network specifically designed for modeling sequential data, such as language. RNNs process input sequences one element at a time while maintaining an internal state that captures the context of the preceding elements in the sequence. This enables the network to capture complex dependencies between elements in the sequence, making them particularly useful for tasks like language modeling, where understanding the context of a word is critical for predicting the subsequent word in the sequence. We use an RNN with an LSTM (long short-term memory) architecture to incorporate additional conditioning information. Charniak (2019) provides an excellent introduction.

The above figure illustrates the evolution of the RNN model. The state of the RNN has

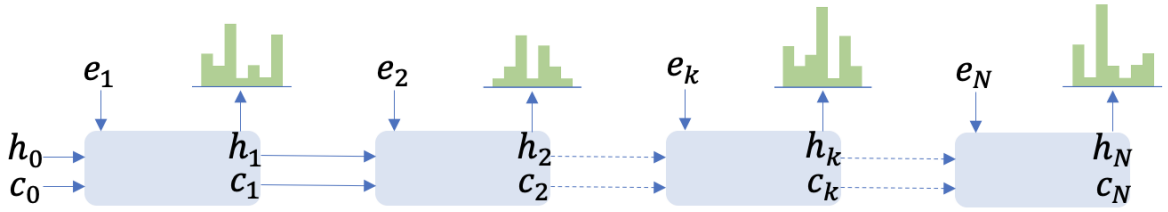


Illustration of RNN model with LSTM architecture. h contains the model's hidden state and c contains the long-term memory.

two components, h and c. The h component is a vector of numerical values summarizing the context of the input sequence processed up to that point. It serves as a summary of the past inputs seen by the RNN up to the current time step and is updated iteratively as the RNN processes each word in the input sequence. The cell state, denoted by c, is an internal memory that stores and propagates information across time steps to capture long-term dependencies; this is the distinguishing feature of the LSTM structure. It allows the RNN to learn and remember long-term dependencies by regulating the flow of information through a mechanism consisting of gates. The cell state is updated in parallel with the hidden state and is used to decide what information to retain, forget, or update.

The state of the RNN is updated after each word is read. Each word is represented through a word embedding (shown in the figure as the e), which is a vector of numerical values that encodes information about the meaning and use of the word. We use the well-known GloVe word embeddings (Pennington et al., 2014), which map each word to a 100-dimensional vector.[9]

At each step, the current state (h, c) determines a probability distribution over the model's vocabulary, which assigns a probability to each possible next word. The RNN thus embodies a model of $\mathbb{P}$. The training process chooses the parameters of the RNN to minimize (2) for articles that appear within the training period. Once the RNN is trained, we evaluate the probabilities of the form $\mathbb{P}(w_k|w_1 \cdots w_{k-1})$ by feeding the embeddings of $w_1, \ldots, w_{k-1}$ into the RNN and then evaluating the RNN's conditional distribution $\mathbb{P}(\cdot|w_1 \cdots w_{k-1})$ at each observed word w_k. Unlike the n-gram approach, the conditional probabilities returned from the RNN reflect information from a very large number of prior words in the document.

Each of the arrows in the LSTM figure corresponds to a combination of linear and nonlinear transformations involving a large number of parameters. These are described in more detail in Section A2 of the Online Appendix. Training the RNN means calibrating

[9]We obtain these from `https://nlp.stanford.edu/projects/glove/`.

these parameters to a reference or training text, a process we turn to next.

3.3 Model Training and Updating

Our neural network consists of an embedding (input) layer, an LSTM layer, and a fully connected (output) layer. The embedding layer maps words to their vector representations. The LSTM layer models temporal dependencies between words in a document. We set the dimension of the h and c vectors at 16. The LSTM layer, whose architecture is detailed in Section A2 of the Online Appendix, contains four W matrices which transform word embeddings into an internal state and have dimension 16×100, four U matrices which transform the internal state and have dimension 16×16, and four bias vectors of dimension 16 each. The total number of parameters in the LSTM layer is $4 \times (1600 + 256 + 16) = 7{,}488$. The output of the LSTM consists of a fully connected layer which maps the final hidden state h into a probability distribution over the vocabulary of 10,000 words. This layer contains a $10,000 \times 16$ matrix U_s which transforms the state h to 10,000 outputs and a 10,000-length bias vector for a total of 170,000 parameters. Thus the network has 177,488 parameters which must be trained.

All articles in the corpus are preprocessed by removing all punctuation strings and converting text to lowercase. We retain only the 10,000 most frequent words in the whole corpus, representing all other words as "UNK" for unknown words. We chunk each article into segments of 100 words.[10] Training the RNN uses mini-batches, which are subsets of the entire dataset. The model's likelihood function in (2) is evaluated on the mini-batch, and parameter updates occur using the average gradient from each mini-batch item. This process requires specifying a batch size and the number of epochs. The batch size is the number of 100-word segments in a mini-batch. A larger batch size typically results in more stable updates to the weights but requires more memory and may lead to longer training times. Conversely, a smaller batch size allows for faster training times but may cause more fluctuations in weight updates. We choose a batch size of 128, i.e., the model updates its parameters after processing each mini-batch of 128 samples, each of length 100 words. An epoch is a full pass through the training set, during which the network processes every sequence of words in the training set once. Increasing the number of epochs can lead to better convergence of the network weights and improved performance

[10]For segments containing fewer than 100 words (at the end of articles, for example), as a standard approach, we pad zeros to follow the actual words. When a batch of sequences is passed through the model, fixed-size length sequences allow the model to use vectorized operations and take advantage of hardware optimizations (i.e., parallel processing on GPUs). This results in faster training times and more efficient use of available computational resources.

on the validation set, but it also increases the risk of overfitting. We choose the number of epochs to be 50. These values for the batch size and number of epochs are standard (see Goodfellow et al. 2016; Keskar et al. 2016).

At the beginning of each epoch, the entire training dataset is randomly shuffled in 100-word chunks and divided into mini-batches without replacement.[11] During each epoch, the model processes one mini-batch at a time and updates its weights using backpropagation for each length-100 element of the mini-batch.[12] This process is repeated for all the mini-batches in the dataset until the entire dataset has been used for training in that epoch. Since the dataset is shuffled at the beginning of each epoch, the composition of the mini-batches changes from epoch to epoch, providing the model with a diverse set of samples to learn from. This shuffling means that, during training, the network state from 100-word sequence A that precedes 100-word sequence B in the same article will not carry over from A to B.

We initialize the training using the first six months of data (January 1996 – June 1996) to compute the entropy for all articles in the following month (July 1996). Write $m_{[t-6,t-1]}$ for the RNN model trained using articles in the six months prior to t, and write $m_{[t-6,t-1]}(t)$ for the equal-weighted average entropy score of articles in month t using the RNN model trained over the prior six months. We use this value to detect month t articles that seem unusual based on information known only prior to month t.

Suppose in month t we have n_t articles in the archive, and the i-th article, represented by $w_{ti1}w_{ti2}\cdots w_{tiN_{ti}}$, has length N_{ti}. By forming the sample counterpart of (2), the entropy of this article calculated using model $m_{[t-6,t-1]}$ is

$$H_{ti} = -\frac{1}{N_{ti}} \sum_{k=1}^{N_{ti}} \ln(p_{tik}^{x_{tik}}), \tag{3}$$

where p_{tik} is the output vector of length 10,000 when we feed $w_{ti1}w_{ti2}\cdots w_{ti(k-1)}$ into $m_{[t-6,t-1]}$, x_{tik} is the position of word w_{tik} in the vector, and $p_{tik}^{x_{tik}}$ is the estimated proba-

[11] If the dataset is not perfectly divisible by the batch size, only the last batch in an epoch will contain fewer samples than the specified batch size.

[12] We train the model using the categorical cross-entropy loss function and the Adam optimizer (Kingma and Ba, 2014). The optimization uses gradient descent with backpropagation: the loss function is computed for each step in the sequence of words in the training text, and gradients are backpropagated through the entire sequence, updating the weights of the model at each time step to reduce the loss function Werbos (1990).

bility of the k-th word being w_{tik} given the preceding words in the chunk. Thus,

$$m_{[t-6,t-1]}(t) = \frac{1}{n_t}\sum_{j=1}^{n_t} H_{tj} = -\frac{1}{n_t}\sum_{j=1}^{n_t}\frac{1}{N_{tj}}\sum_{k=1}^{N_{tj}} \ln(p_{tjk}^{x_{tjk}}). \tag{4}$$

Note that (4) is calculated for all articles in month t using the model $m_{[t-6,t-1]}$ which was estimated over the prior six months, and would thus have been know in month t. Unlike in model training–when the probabilities in (3) are evaluated only in the context of a 100-word chunk–all words in article j are used, in succession, to calculate $p_{tjk} \in \mathbb{R}^{10,000}$. The state of the RNN, i.e., (h, c), is then reset to zero before calculating the entropy of article $j + 1$.

From August 1996 onward, in order to calculate the entropies for all articles in month t, we first retrain the model $m_{[t-7,t-2]}$ to obtain an updated model $m_{[t-6,t-1]}$. The retraining starts from the parameter values obtained in $m_{[t-7,t-2]}$ and updates the values by running the parameter optimization on new text in month $t-1$ together with randomly sampled text from month $t-2$ to month $t-6$. In particular, we take all articles from month $t-1$, half of the articles randomly sampled from month $t-2$, one-quarter of the articles randomly sampled from month $t-3$, and so on until $\frac{1}{2^5}$ of the articles are randomly sampled from month $t-6$. Thus, in retraining the model based on the prior six months of news, we give more weight to texts from the most recent months. Each retraining happens with batches of 128 length-100 word sequences with 50 epochs. Our main variable of interest, monthly entropy or ENT_t, is the difference between the average entropy in month t and month $t-12$, using the most recently available model at each time, or

$$ENT_t \equiv m_{[t-6,t-1]}(t) – m_{[t-18,t-13]}(t-12) \tag{5}$$

Our first observation is from July 1997.

The major benefit of using a relatively small language model, with "only" 177,488 parameters, is that we are able to estimate it in rolling windows, which means that ENT_t would have been available to market participants in real-time.[13] Note that the time t rolling model still contains information about months before t through the influence of earlier articles on the model's parameters from prior training steps. Furthermore, ENT_t can be decomposed into a new information component and a model update component, a model feature which we explore in Section 6.

[13]Running on a GPU server with four NVIDIA T4 cards requires 21 hours to estimate the model in all rolling windows.

3.4 Entropy Overview

Table 1 shows summary statistics for the variables we use in our analysis, all measured at monthly frequency. *ENT* ranges from −0.140 to 0.188, with a mean close to 0 and a standard deviation of 0.061. The table also shows statistics for the Economic Policy Uncertainty (*EPU*) index of Baker et al. (2016) and the San Francisco Fed's News Sentiment (*SEN*) index introduced by Shapiro et al. (2022). *EPU* is a well-known measure that is plausibly related to *ENT*, and thus we pay special attention to *EPU* in our analysis. *SEN* is also a related news-based measure. The construction of the *EPU* and *SEN* indexes is described in Online Appendix Section A3. Figure 2 reports the average contemporaneous correlations between *ENT* and other control variables. *ENT* is not very correlated with any other control variables; its strongest correlation (0.393) is with *DY*. *ENT* has little correlation with *EPU* (0.112) and *SEN* (-0.275).

The time series of *ENT* and *EPU* are plotted in the top-left panel of Figure 3 as solid dark red and dashed light red lines, respectively. While *ENT* and *EPU* sometimes exhibit similar patterns (e.g., during the 2007–08 financial crisis), they often diverge. The time series of *ENT* and *SEN* are plotted in the top-right panel as solid dark red and dashed light red lines, respectively. The peaks in *ENT* mostly coincide with either peaks or troughs in the sentiment measure. For example, during the 2007–08 financial crisis period, *SEN* is very negative while *ENT* is high. The time series of *ENT* and *VIX* are plotted in the bottom-left panel as solid dark red and dashed light red lines, respectively. The peaks of *VIX* often, but not always, coincide with peaks of *ENT*, particularly during the financial crisis period. The bottom-right panel of Figure 3 plots *ENT* and 12-month ahead cumulative returns. There appears to be some tendency for large positive next-12-month returns to occur at times of low entropy, and for low future returns to occur at times of moderate to high entropy. We investigate this tendency further in Section 4.

Figure 3 shows that January 2007, November 2010, and October 2014 were all months with large *ENT* peaks (and, interestingly, all of these were also low *EPU* months). Table 2 shows sample articles with high entropies from each of these months. We saw in the example of Section 1 how the conditional probabilities of words can change over time. We now give examples of how these changes contribute to the peaks in *ENT* through the entropy calculated in (3) and (4):

- In January 2007, there were several high-entropy news articles related to Fannie Mae and Freddie Mac. In one such article, the phrase "lack of progress in reining in mortgage lenders Fannie Mae and Freddie Mac leaves the economy at" was

followed by the word "risk." But in January 2007, the word "risk" was perceived by the model to be unlikely in the context of this phrase, giving the sentence high entropy.[14] As illustrated in the middle panel of Figure 4, the conditional probability of "risk" increased dramatically through the mortgage crisis in late 2007 and the conservatorship of Fannie Mae and Freddie Mac by the U.S. government in 2008. The initial high entropy score thus signaled an important shift in the news.

This distributional shift can also be illustrated through the string of words "Fannie Mae and Freddie Mac growth will be" from another article in January 2007. The actual next word was "muted." The word clouds in Figure 5 compare the conditional distributions of the next word as calculated in December 2006 (left panel) and December 2008 (right panel). The onset of the global financial crisis in between these dates is clearly reflected in the word clouds, where the sizes of words reflect their relative likelihoods under the models estimated at each date.

- In late November 2010, there were high entropy articles discussing Ireland's financial strains. In one such article, the phrase "the Euro falls to a 5 week low on growing concerns Ireland will be forced to" was followed by the word "default." But "default" had low conditional probability under the October 2010 model, thus giving the sentence high entropy. The right panel of Figure 4 shows how the conditional probability of "default" changed over time, and became much higher by 2012 and again in 2014. In this example, as in the previous examples, the initial increase in entropy signaled a change in the distribution of text associated with economic developments.

- Lastly, the *ENT* peak in October 2014 we observed in Figure 3 coincides with the Fed's ending of QE3. Many moderate to high entropy articles from that month mentioned QE and its cessation.

4 Time Series of Returns

In this section, we present evidence that entropy forecasts future market returns, and we compare entropy with other predictors from the literature.

[14]The conditional probability is calculated by feeding the embeddings of the preceding string of words (e.g., "lack of progress in reining in mortgage lenders Fannie Mae and Freddie Mac leaves the economy at") into the RNN model from the prior month (i.e., December 2006) and then evaluating the probability at the target word (e.g., "risk").

4.1 In-Sample Regressions

We estimate a variety of time-series forecasting regressions of the form

$$R_{t+1,t+h} = \beta_0 + \beta_{ENT} ENT_t + \gamma^\top Control_t + \epsilon_{t+1,t+h} \tag{6}$$

where $R_{t+1,t+h}$ is the cumulative market return from month $t+1$ to month $t+h$ with $h \in \{3, 6, 9, 12\}$ and ENT_t is the entropy measure in month t. $Control_t$ contains sentiment (*SEN*), squared implied volatility (*VIX2*), interest rates (*DGS10* and *DGS10-2*), the dividend yield (*DY*), the consumption to wealth ratio (*CAY*), the market return of the previous month (*Return1*), the cumulative market return of the previous 12 months excluding the most recent month (*Return12*), the cumulative returns of the Fama-French five factors (Fama and French 2015) and momentum over the previous 60 months (*Return60*, *SMB60*, *HML60*, *RMW60*, *CMA60*, and *UMD60*). As explained in Section 2, using the 60-month lagged returns converts factors to state variables, as suggested in Maio and Santa-Clara (2012). We also use *1/CAPE*, the reciprocal of the cyclically adjusted price-to-earnings ratio, as a forecasting variable.

Column (1) of Panel A of Table 3 reports the coefficient estimates (and t-statistics) for equation (6) with $h = 12$, after scaling each coefficient estimate by the standard deviation of the independent variable.[15] A one standard deviation increase in entropy predicts a 2.387% decrease in the 12-month ahead cumulative market return. Column (2) replaces *ENT* in equation (6) with *EPU*. The results show that *EPU* positively predicts future market returns, which is consistent with Brogaard and Detzel (2015). Column (3) includes both *ENT* and *EPU* in a single regression model. *ENT* still forecasts future market returns, and the standardized coefficient for *ENT* (-2.276%) is nearly unchanged from column (1).

Panel B (C, D) reports the coefficient estimates (and t-statistics) for equation (6) with $h = 9$ ($h = 6$, $h = 3$). Column (1) shows that a one standard deviation increase in entropy predicts a 2.284% (1.450%, 1.221%) decrease in the 9-month (6-month, 3-month) ahead cumulative market return. Column (3) shows that a one standard deviation increase in entropy predicts a 2.242% (1.399%, 1.197%) decrease in the 9-month (6-month, 3-month) ahead cumulative market return, after controlling for both *ENT* and *EPU*. *ENT*

[15] Tables A1–A4 in the Online Appendix show the unscaled coefficient estimates for equation (6) for forecasting horizons of 12, 9, 6, and 3 months, respectively. These show that *CAY* forecasts future market returns negatively, whereas Lettau and Ludvigson 2001 find that *CAY* forecasts returns positively. We verified that, as a standalone forecaster, *CAY* forecasts return positively when restricted to the pre-2002 period but negatively in the full sample.

is significant in all cases using Newey and West (1987) standard errors with four lags.[16]

The full-sample analysis incorporates data from the COVID-19 period, which raises the concern that this time period may disproportionately impact our results. To address this, we replicate the analyses detailed in the first three columns of Table 3 using pre-COVID-19 data. In the pre-COVID-19 analysis, entropy is calculated up to the end of 2018 and used to predict returns through the end of 2019. Columns (4)–(6) show that *ENT* consistently and negatively forecasts 12-month (9-month, 6-month, 3-month) ahead cumulative returns: a one standard deviation increase in entropy is associated with a 2.699% (2.178%, 1.720%, 1.170%) negative market return. *EPU* does not significantly predict market returns once the COVID-19 period is dropped from the analysis, while *ENT* remains significant in column (6) at horizons of 6–12 months (and is marginally significant at 3 months, and in all cases in column 4).

4.2 Out-of-Sample Regressions

We saw in the previous section that an increase in entropy negatively forecasts future market returns in in-sample regressions. We now examine whether entropy provides out-of-sample predictability.

Let $R_{t+1,t+12}$ denote the cumulative return over months $t+1, \cdots, t+12$. For a fixed training window of length h, for each t, we run the univariate regression

$$R_{t-i-11,t-i} = \alpha_t + \beta_t ENT_{t-i-12} + \epsilon_{t-i-12}, \quad i = 1, 2, \ldots, h, \tag{7}$$

estimating $\hat{\alpha}_t$ and $\hat{\beta}_t$ from the h monthly observations in the training window (we discuss the choice of h momentarily). Then we predict the future return $R_{t+1,t+12}$ as $\hat{R}_{t+1,t+12} = \hat{\alpha}_t + \hat{\beta}_t ENT_t$. The out-of-sample R-squared is calculated as

$$1 - \frac{MSE(R_{t+1,t+12} - \hat{R}_{t+1,t+12})}{MSE(R_{t+1,t+12} - \frac{1}{h}\sum_{i=1}^{h} R_{t-i-11,t-i})}, \tag{8}$$

where the numerator is the mean squared error of predicting with the model in (7), and the denominator is the mean squared error based on using past historical means to forecast future returns. We are thus comparing the prediction accuracy of a model which uses ENT_t to make a conditional return forecast versus a model which just uses the rolling mean over the past h months. If predictions using entropy are better than the rolling mean

[16] In Section A4 of the Online Appendix we check other standard error methodologies, using closed form and bootstrap approaches, and find that Newey-West with four lags is an appropriate adjustment.

predictions, we will get a positive out-of-sample R-squared. If predicting with entropy is worse than using historical means, we will get a negative out-of-sample R-squared.

To estimate p-values for the out-of-sample R-squareds, we use bootstrap resampling of residuals. The idea is to construct a distribution of the R-squared test statistic under the null hypothesis that the β_t coefficient in (7) is zero. We first estimate (7) using the full sample, then in the nth bootstrap run, resample the residuals to construct a bootstrapped $R_{t-i-11,t-i}$ series by setting $\beta = 0$ while using the estimated α, and regress this bootstrapped series on ENT_{t-i-12} in rolling windows of length h to get estimated $\{\hat{\alpha}_t^{(n)}, \hat{\beta}_t^{(n)}\}$s under the no-predictability null, which are then used to form out-of-sample forecasts and to calculate the out-of-sample R-squared in (8). We resample residuals using the geometric block sampling approach of Politis and Romano (1994). The method preserves temporal dependencies in the residuals by selecting contiguous segments of data with varying block lengths. Details of the method can be found in Online Appendix Section A4. We use 1000 bootstrap iterations in computing the p-values, which indicate the fraction of simulated R-squareds which are larger than the observed ones.

Table 4 shows the out-of-sample R-squared of entropy as well as all other control variables we used in the in-sample analysis.[17] Each variable is used separately in univariate regressions of the form in (7). Each column represents a different choice of window length—$h \in \{12, 15, 18, 21, 24\}$ months—for estimation of the forecasting model in (7). In the first row, we see that entropy produces a positive out-of-sample R-squared in four out of the five estimation windows, and the out-of-sample R-squareds range from 0.024 (with $h = 15$ months) to 0.051 (with $h = 21$ months). In the second row, we see that *ENT_NEWS*, which captures the part of entropy due to news innovations while holding the text model fixed (we formally define *ENT_NEWS* in Section 6) produces a positive out-of-sample R-squareds in all five estimation windows, with the out-of-sample R-squareds ranging from 0.060 (with $h = 24$ months) to 0.096 (with $h = 18$ months). In fact, this news innovations measure outperforms *all* other predictors, most of which have negative R-squareds in all estimation windows. Only two other measures have non-negative out-of-sample R-squareds: *Return60* produces positive R-squareds in all estimation windows and *RMW60* produces positive R-squareds for estimation window lengths of 18, 21, and 24 months.

Table A7 in the Online Appendix shows the average of β_t from (7) for each of the training horizons used in the out-of-sample forecasting. The average forecasting coefficients associated with *ENT* are all negative and are highly statistically significant for

[17]To avoid look-ahead bias, we use the original *CAY* (through 2019 Q3) in the out-of-sample analysis.

horizons of 15 months and above. Though the β_t's vary over time, the average rolling entropy forecasting coefficient for future returns is consistent with the full sample results in Table 3.

Given the difficulty of out-of-sample market return forecasting (Welch and Goyal, 2008; Campbell and Thompson, 2008), the performance of *ENT* and *ENT_NEWS* in our out-of-sample tests is remarkable.[18]

5 Cross-Section of Returns

According to the argument in Maio and Santa-Clara (2012), the sign with which an ICAPM state variable forecasts aggregate market returns should match the cross-sectional risk price associated with that variable. If entropy negatively forecasts the return of the aggregate stock market, then high entropy is associated with an unfavorable investment opportunity set for investors. In this case, securities that do well during high entropy times provide a useful hedge, and should therefore earn a negative risk premium. In this section, we investigate how exposure to entropy is priced in the cross-section of stock returns.

According to the Merton (1973) ICAPM framework, in equilibrium, an asset's risk premium is determined by its conditional covariances with ex-post market returns and innovations of state variables of the form

$$\mathbb{E}_t[r_{i,t+1}] - r_{f,t+1} = \beta Cov_t(r_{i,t+1}, r_{m,t+1}) + \gamma Cov_t(r_{i,t+1}, \Delta x_{t+1}) \tag{9}$$

where $r_{i,t+1}$ is the return on asset i, $r_{f,t+1}$ is the risk-free rate (known at time t), $r_{m,t+1}$ is the market return, and x_t is a state variable that affects the investment opportunity set. The coefficients β and γ are risk premia. Equation (9)—which reproduces equation (8) of Maio and Santa-Clara (2012)— states that investors receive compensation for the covariance of stock returns with variables that impact future market returns. Based on the argument of Maio and Santa-Clara (2012), if Δx_{t+1} forecasts stock returns positively, then the ICAPM suggests that stocks that positively covary with Δx_{t+1} should earn positive risk premia because they add to investor risk. On the other hand, if, like change in entropy, Δx_{t+1} forecasts future negative stock returns, then stocks that covary positively with Δx_{t+1} hedge against unfavorable changes in the investment opportunity set and

[18]While our rolling LSTM models are completely out-of-sample, the GloVe word embeddings we use did not become available until October 2015. We provide empirical evidence that this is not a concern in Section A5 of the Online Appendix.

should earn negative risk premia. The ICAPM sign property of Maio and Santa-Clara (2012) can be summarized as:

$$\text{sgn}(\gamma) = \begin{cases} + & \text{if } \Delta x_{t+1} \text{ predicts market returns positively,} \\ - & \text{if } \Delta x_{t+1} \text{ predicts market returns negatively.} \end{cases} \tag{10}$$

We have seen that entropy forecasts a decline in market returns, so we expect it to carry a negative risk premium.

5.1 Factor-Mimicking Portfolios

Because ENT and EPU are not directly tradeable, we approximate them with factor-mimicking portfolios. In the first step, we extract innovations from the ENT and EPU time series by fitting an $AR(p)$ model. The number of lags in the AR model is determined by minimizing the Bayesian information criterion (BIC). The final models, estimated using monthly data, are $AR(12)$ for ENT and $AR(1)$ for EPU:

$$\begin{aligned} ENT_t &= \beta_{ENT,0} + \sum_{i=1}^{12} \beta_{ENT,i} ENT_{t-i} + \gamma_{ENT} MKT_{t-1} + \epsilon_{ENT,t}, \\ EPU_t &= \beta_{EPU,0} + \beta_{EPU,1} EPU_{t-1} + \gamma_{EPU} MKT_{t-1} + \epsilon_{EPU,t}. \end{aligned} \tag{11}$$

Following Breeden et al. (1989); Lamont et al. (2001); Ang et al. (2006), in the second step, we create the mimicking factor F_{ENT} (F_{EPU}) to track innovations in ENT (EPU) by estimating the coefficient b_{ENT} (b_{EPU}) in the following regressions

$$\begin{aligned} \hat{\epsilon}_{ENT,t} &= a_{ENT} + b_{ENT}^{\top} X_t + u_{ENT,t}, \\ \hat{\epsilon}_{EPU,t} &= a_{EPU} + b_{EPU}^{\top} X_t + u_{EPU,t}, \end{aligned} \tag{12}$$

where $\hat{\epsilon}_{ENT,t}$ and $\hat{\epsilon}_{EPU,t}$ denote the innovations of ENT and EPU from (11), and X_t denotes the monthly excess returns on a set of base assets. We estimate (12) over the full sample and construct factor mimicking portfolios as follows

$$\begin{aligned} F_{ENT,t} &= \hat{b}_{ENT}^{\top} X_t, \\ F_{EPU,t} &= \hat{b}_{EPU}^{\top} X_t, \end{aligned} \tag{13}$$

where $\hat{b}_{ENT}$ and $\hat{b}_{EPU}$ are the coefficient estimates from (12). Since X_t represents the excess returns of base assets portfolios, the coefficient vectors $\{b_{ENT}, b_{EPU}\}$ can be interpreted as

weights of a zero-cost portfolio with returns given by (13). While the regressions in (11) and (12) are run with monthly data, the factor replicating portfolios in (13) are available at a daily frequency.

We consider four sets of base assets: the 25 Fama-French portfolios sorted by size and momentum, size and book to market, size and investment, and size and profitability. In the risk premia tests of the next section, the base assets are selected to correspond to the test assets in the Fama and MacBeth (1973) analysis. Table 5 shows summaries of the Fama and French (2015) factor and momentum returns, as well as those of the *ENT* and *EPU* factor replicating portfolios using different base assets.

5.2 Factor Risk Premia

To estimate the risk premium associated with the entropy factor mimicking portfolio, we run a Fama and MacBeth (1973) analysis using *MKT* (market minus risk-free), *SMB*, *HML*, *RMW*, *CMA*, *UMD*, as well as F_{EPU} and F_{ENT} as the factors. These correspond to the five factor model of Fama and French (2015) augmented with momentum, and the *ENT* and *EPU* factor replicating portfolios. To properly reflect estimation risk from the first stage regressions, we use the following generalized method of moments (GMM) system in our estimation

$$\begin{aligned} \mathbb{E}[r_{i,t} - \alpha_i - \beta_i^\top f_t] &= 0, \quad i = 1, 2, \cdots, N \\ \mathbb{E}[(r_{i,t} - \alpha_i - \beta_i^\top f_t) f_t] &= 0_K, \quad i = 1, 2, \cdots, N \\ \mathbb{E}[\beta(r - \beta^\top \lambda)] &= 0_K, \end{aligned} \tag{14}$$

where $r_{i,t}$ is the excess return of test asset i on day t, β_i is a $K \times 1$ vector of factor loadings for test asset i, f_t is a $K \times 1$ vector of factors, and 0_K is a $K \times 1$ vector of zeros. In addition, $\beta = [\beta_1 \ \beta_2 \ \cdots \ \beta_N]$ is the $K \times N$ matrix of the N test asset betas, r is a vector of the N test asset average excess returns, and λ is a $K \times 1$ vector of factor risk premia.[19] The parameters of the system are $\{\alpha_1, \cdots, \alpha_N, \beta_1, \cdots, \beta_N, \lambda\}$. The first two moment conditions, which correspond to the first-stage Fama-MacBeth regressions, exactly identify α and β. The last moment condition represents the second stage regressions, with no constant, which pin down the prices of risk. See Cochrane 2009 page 241 for details.

Table 6 shows the risk premia from (14), but scaled by the standard deviations of

[19]The risk premia λ in (14) differ from β and γ in (9) in that they are premia on regression coefficients rather than covariances.

the first-stage betas to highlight the magnitudes of the effects.[20] Columns (1) – (4) show monthly risk premia using different sets of 25 Fama-French portfolios as the base and test assets: size-momentum, size-book/market, size-investment, and size-profitability. We see that on all four sets of test assets, F_{ENT} carries a negative and significant risk premium, with the monthly risk premium for a standard deviation increases in entropy beta ranging from -0.060% to -0.091%. The magnitude of the entropy risk premium is on par with that of the other factors, suggesting the risk compensation associated with hedging entropy risk is economically large.

In comparison, F_{EPU} carries a negative and significant risk premium in only two of the four tests, with a risk premium ranging from -0.008% to -0.083% per unit of standard deviation of EPU betas. The sign and the magnitude of the F_{EPU} risk premium are consistent with the results in Brogaard and Detzel (2015) (compare Panel B of their Table 5 against the unscaled risk premia for F_{EPU} in Table A8 of the Online Appendix). Among the five factors of Fama and French (2015), only MKT has a significant risk premium—ranging from 0.048% to 0.083%—for more than two out of four sets of test asset. The momentum factor carries a positive and significant risk premium in two tests: 0.083% for size and book-to-market portfolios and 0.064% for size and investment portfolios.

Column (5) of Table 6 shows the sign of the ICAPM property from (10). The determination of whether a factor forecasts year-ahead market returns positively or negatively is based on the forecasting regression in (6), shown in column (3) of Online Appendix Table A1. A "+" or "–" in column (5) of Table 6 corresponds to significant positive and negative coefficients at the 5% level, respectively; a blank indicates a lack of significance in the forecasting regression. The ICAPM sign property posits that any significant coefficients in columns (1) – (4) of Table 6 should have the same signs as the corresponding entry in column (5).[21] ENT is the only factor for which this property holds. CMA has positive signs in column (5) but only one out of the four tests shows a significant risk premium.

5.3 Which Securities Hedge Entropy Risk

We now investigate which types of securities are useful for entropy risk hedging. As test assets, we use decile portfolios from univariate sorts on size, book/market, operating profitability, and investment obtained from Ken French's website, and we estimate their

[20] Table A8 in the Online Appendix shows the raw risk premia.

[21] Recall that *UMD60* in Table A1 corresponds to *UMD* in Table 6, *Return60* corresponds to *MKT*, and so on. ENT and EPU are the state variable equivalents of F_{ENT} and F_{EPU}.

loadings on F_{ENT}. Each asset i's factor exposures are estimated using the full-sample time-series regression

$$r_{i,t} = \beta_{i,0} + \beta_{i,F_1} F_{1,t} + \cdots + \beta_{i,F_8} F_{8,t} + \epsilon_{i,t},$$

where $r_{i,t}$ is the excess return of asset i on day t, $F_{i,t}$ is the day-t return on factor i, and the indices $j = 1, \ldots, 8$ refer to *MKT*, *SMB*, *HML*, *RMW*, *CMA*, *UMD*, F_{EPU}, and F_{ENT}, respectively. For this analysis, F_{EPU} and F_{ENT} were constructed using 100 base assets in (12), corresponding to the 25 size-X sorted portfolios for $X \in \{$momentum, book/market, investment, profitability$\}$. Each β_{i,F_j} is asset i's loading on factor j. From this set of equations, we compute estimates $\hat{\beta}_{i,F_1}, \ldots, \hat{\beta}_{i,F_8}$ of the factor loadings for each asset i.

In Figure 6, we show the F_{ENT} loadings ($\hat{\beta}_{i,F_8}$) corresponding to each test asset. The first panel shows these loadings for size-sorted portfolios. The U-shaped pattern indicates that small and large companies hedge entropy risk, while medium-sized firms are negatively correlated with entropy innovations. Companies sorted on momentum (left panel, middle row) show that past losers hedge entropy risk while past winners do poorly during high entropy times, though the effect is non-monotonic. Companies sorted on book/market (right panel, middle row) show a similar pattern of exposures to the size-sorted portfolios, with growth firms (low book/market) and value firms (second highest book/market) having positive entropy betas, and all other portfolios having negative entropy betas.[22] In contrast, sorting on investment (left panel, bottom row) yields an inverted U-shape, indicating that very low and very high investment companies do poorly during times of high entropy, while medium-investment firms do well. The bottom right panel shows that companies with high operating profitability offer a hedge to entropy risk, while less profitable companies are entropy-risky. Given the typical high minus low long-short portfolio approaches, our finding of generally non-monotonic entropy exposure is notable. As a robustness check, Figure A5 in the Online Appendix repeats the analysis in each panel of Figure 6 but using F_{ENT} constructed from different base assets. The results are similar.

[22]We interpret the highest book/market decile as distressed, rather than value, stocks with extreme high book values relative to market values.

6 Channels

We've shown that entropy negatively and robustly forecasts aggregate market returns both in and out of sample. Furthermore, entropy is a priced cross-sectional risk and is the only factor in our analysis that satisfies the Maio and Santa-Clara (2012) ICAPM sign property between its forecasting direction for the market (negative) and its cross-sectional price of risk (also negative). In this section, we investigate some potential mechanisms underlying these findings. First, we show that entropy can be decomposed into a news innovation component and a model innovation component, and that the news innovation part plays a large role in our findings. Second, we study how entropy and its different components forecast macroeconomic outcomes. We then explain how these findings are consistent with the theoretical model we develop in the appendix.

6.1 Entropy Decomposition

We decompose the entropy measure from equation (5) into two parts

$$ENT_t = ENT_NEWS_t + ENT_MODEL_t. \tag{15}$$

The first part is

$$ENT_NEWS_t = m_{[t-18,t-13]}(t) – m_{[t-18,t-13]}(t-12)$$

which uses the text model from one year prior to month t, but applies it to month t news flow. This captures the change in the information content of news. The second part is

$$ENT_MODEL_t = m_{[t-6,t-1]}(t) – m_{[t-18,t-13]}(t)$$

which looks at the difference between month t entropy as seen through the lens of today's and the prior year's text models. In order to understand what drives the predictive power of ENT, we replace ENT_t by ENT_NEWS_t and ENT_MODEL_t, and repeat the time series prediction of 12-month ahead cumulative returns in-sample as in Section 4.1 and out-of-sample as in Section 4.2.

Online Appendix Table A5 columns (1) and (2) show the in-sample results, where the coefficient estimates are normalized by the standard deviation of the independent variables.[23] Column (1) indicates that a standard-deviation increase in ENT_NEWS is

[23]Online Appendix Table A6 shows the unscaled results.

associated with a -1.081% market return over the next 12 months. A standard-deviation increase in *ENT_MODEL*, column (2), is associated with a -1.595% next 12-month return. However, neither effect is significant. This suggests that the the large significant *ENT* coefficients in Table 3 rely on both the news and model components of entropy.

Table 4 presents the out-of-sample R-squareds for *ENT_NEWS* and *ENT_MODEL*. *ENT_NEWS* produces a positive out-of-sample R-squared irrespective of the model estimation window length, with the largest R-squared of 0.096 occurring at an estimation window length of 18 months. On the other hand, *ENT_MODEL* never produces a positive out-of-sample R-squared. The predictive power of *ENT* for future market returns comes from the change in the information content of news rather than from the year-over-year change in the model. We return to this finding in Section 6.2.

While *ENT_NEWS* is a very robust forecaster of year-ahead market returns, the direction of forecastability fluctuates over time. This can be seen by the negative, but insignificant average β_t coefficient from (7) for *ENT_NEWS* shown in Table A7 of the Online Appendix, as well as the negative but insignificant in-sample coefficient in Table A5. While the average direction of forecastability from *ENT_NEWS* to year-ahead market returns is negative, there are also times when this sign is positive, which renders the average impact of the effect only marginally significant.

Next, we explore whether *ENT_NEWS* and *ENT_MODEL* are priced cross-sectionally. We construct factor-mimicking portfolios for these two components of *ENT* using the procedure from Section 5.1. We run four tests. Each test uses one of the sets of 25 Fama-French portfolios (sorted on size-momentum, size-book-to-market, size-investment, or size-profitability) as both the base assets for the factor-mimicking portfolios and the test assets for pricing.

Table 7 shows that F_{ENT_NEWS} has a negative risk premium in three out of the four cases, ranging from -0.055% to -0.060% per month.[24] Though it is negative in all cases, the risk premium is significant for the book-to-market, investment, and profitability sorts, but not for momentum. The magnitudes of the risk premia are again large in comparison to those associated with the other factors, and are consistent with the results in Table 6, where we use F_{ENT} instead of F_{ENT_NEWS}. Column (5) shows that none of the factors satisfy the ICAPM property in (10) when the signs of the coefficients in the forecasting regression in (6) are obtained from column (1) of Table A6. However, as Tables A6 and A7 show, the general predictive sign of *ENT_NEWS* for future market returns is negative, though marginally significant.

[24]Table A9 of the Online Appendix shows the unscaled results.

Tables A10 (scaled coefficients) and A11 (unscaled) of the Online Appendix show that *ENT_MODEL* is not associated with a significant risk premium. And this part of entropy does not forecast future market returns in or out of sample. The evidence points to the market demanding a negative risk premium for securities that help hedge the part of entropy—*ENT_NEWS*, or news innovation—which is useful for forecasting market returns in our out-of-sample tests.

6.2 Entropy and Macro Outcomes

Having established that the predictive power of *ENT* stems largely from year-over-year news innovation, we now explore whether such changes in news flow predict economic fundamentals. We estimate time-series forecasting regressions of the form

$$Y_{t+12} = \beta_0 + \beta_{ENT} ENT_t + \gamma^\top Control_t + \beta_Y Y_t + \epsilon_t. \tag{16}$$

The dependent variable in our analysis is one of:[25]

Variable	Description	Units
EPU	economic policy uncertainty	standard units
UNRATE	unemployment rate	percent
INDPRO_YOY	year-over-year change in industrial production	percent
CPI_YOY	year-over-year change in the Consumer Price Index	percent
DGS10	10-year Treasury yield	percent
DGS10-2	10-year minus 2-year slope	percent
VIX	CBOE Volatility Index	percent
EPS	S&P500 last 12-month real earnings per share	dollars

For each dependent variable, we use the lagged value of entropy (ENT_t, ENT_NEWS_t, or ENT_MODEL_t), the lagged value of the dependent variable, and the lagged values of all other macro variables as controls.

Table 8 summarizes the results. Each column of Table 8 corresponds to one of the above dependent variables. The first row of the table shows the β_{ENT} coefficient from (16), scaled by the full-sample standard deviation of entropy, and the associated t-statistic. The second row shows the results of (16) where ENT_t is replaced with the news innovation part of entropy, ENT_NEWS_t. The third row shows the results for ENT_MODEL_t.[26]

[25]Earnings per share, in constant May 2023 dollars, are obtained from Robert Shiller's website.

[26]The full results reporting the scaled coefficients are in Online Appendix Tables A12–A14. The full results reporting the raw coefficients are in Online Appendix Tables A15–A17.

The table shows that the forecasting results for *ENT* and *ENT_NEWS* are very similar. An increase in either measure positively forecasts the unemployment rate and the *VIX* volatility index, and negatively forecasts industrial production, CPI, the level of interest rates, and the S&P 500 earnings per share. There is a weak positive forecasting relationship from these two entropy measures for *EPU*, and no relationship for the slope of 2s-10s curve. *ENT_MODEL* does not significantly forecast any of the macro outcomes variables, with the exception of the *VIX*, where the relationship is negative, not positive as with the other entropy measures. Overall, entropy and its news innovation component robustly forecast negative economic outcomes one year ahead.

6.3 Discussion

We have shown that high entropy significantly and negatively forecasts macroeconomic outcomes and market returns. Cross-sectional exposure to entropy shocks receives a significant and negative cross-sectional risk premium. Furthermore, as Figure A4 of the Online Appendix shows, daily market returns and shocks to entropy replicating portfolios are contemporaneously negatively correlated. In anticipation of future bad times, the market sells off on positive entropy innovation days. But the market does not sell off *enough* to generate high expected returns following positive entropy shocks.

One explanation for this finding is that the market underreacts to entropy news. This may be due to capacity constrained investors (Sims, 2003), slow-moving institutional capital (Gabaix and Koijen, 2021; Glasserman et al., 2025) along with limits to arbitrage (Gromb and Vayanos, 2010), or general asset-class momentum (Brooks et al., 2023). However, we favor a risk-based explanation.

In the appendix, we propose a simple theoretical model which explains our findings in a rational, frictionless setting. The general idea is that entropy is higher in bad times—because news text has higher dispersion under poorer economic conditions—and that bad times are persistent. An entropy news shock thus causes investors to revise down their beliefs about future dividends (which is consistent with the negative predictive coefficient of entropy for future *EPS* in Table 8). This causes stock prices to fall. However, the discount factor investors use to value future dividends rises, because a dollar of dividends becomes more valuable in bad times. Entropy shocks are thus associated with lower dividend expectations, lower prices, *and* lower discount rates. Details are in the appendix.

7 Robustness

We next check whether entropy is spanned by existing risk factors, whether the entropy risk premium can be explained by a news sentiment risk premium, whether other measures of uncertainty span entropy or can explain the forecasting power of entropy for market returns, and whether our main forecasting results hold over a longer time period.

7.1 Existing Pricing Factors

Entropy is distinctive because of its impressive in- and out-of-sample forecasting power for market returns, the fact that it is the only factor that satisfies the ICAPM sign property of (10), and its macroeconomic forecasting ability. However, it is possible that the information content of entropy is spanned by one of the multitude of existing factors in the literature. To check for this, we use the 153 value-weighted factors discussed in Jensen et al. (2022), whose daily returns are available at `https://jkpfactors.com/`. We augment this set of factors with the *ENT* and *EPU* factor replicating portfolios, F_{ENT} and F_{EPU}, using size-momentum as the base assets.[27] We regress each of the 155 factor return series on the remaining 154 factors using daily data.

In Figure 7 we show these R-squareds for each factor, ranked from low to high. Factors close to the lower left corner are not well explained by the existing factor zoo, while factors close to the upper right corner are largely spanned by existing, alternative factors. We label the top and bottom five factors based on this spanning R-squared measure and also indicate *ENT* via red text. F_{ENT}, the entropy mimicking portfolio with size-momentum as the base assets, has the seventh lowest R-squared among the 155 factors, showing that *ENT* conveys novel information relative to existing factors.

7.2 Economic Sentiment

Another well-known news-based variable is *SEN*, the San Francisco Fed's news sentiment index. The main return forecasting results in Table 3 show that entropy is a strong forecaster of year-ahead market returns even after controlling for the information content of *SEN*, though *SEN* is itself a significant and positive return forecaster in specifications which include *EPU* as a forecasting variable. However, Table 4 shows that *SEN* is a poor out-of-sample forecaster of market returns. Online Appendix Tables A18 (scaled coefficients) and A19 (unscaled) show the Fama-MacBeth risk premium analysis which

[27] The results for the other base assets are similar.

includes F_{SEN}, the *SEN* replicating portfolio constructed using the method described in Section 5.1, in addition to the other factor exposures. Exposure to *SEN* carries a positive and significant risk premium in three of the four tests but does not satisfy the ICAPM property because it is not a significant in-sample market return forecaster at the 5% level; though, as Table A1 shows, *SEN is* a significant and positive return forecaster at the 10% level. Tables A18 and A19 show that even in the presence of F_{SEN}, exposure to F_{ENT} continues to carry a negative (in all specifications) and significant (in three of the four specifications) risk premium, consistent with the results of Section 5.2.

7.3 Other Uncertainty Measures

Another concern is that entropy may be spanned by other uncertainty measures studied in the literature. To check for this, we utilize measures of macroeconomic uncertainty discussed in the survey article David and Veronesi (2022).[28] In particular, we consider the Bekaert et al. (2022) (*BEX*) uncertainty measure of time-varying risk aversion, the Jurado et al. (2015) (*JLN*) variable which provides econometric estimates of the conditional volatility of the purely unforecastable component of the future values of "hundreds of macroeconomic and financial indicators," the Azzimonti (2018) Partisan Conflict Index (*PCI*) which tracks the degree of disagreement among U.S. politicians at the federal level, as well as *ENT*, *EPU*, *SEN*, and *VIX*.

We regress each of the seven uncertainty measures on the other six uncertainty measures and calculate the resultant R-squared. In Figure 8 we show these R-squareds for each uncertainty measure, ranked from lowest to highest. Variables close to the lower left corner are not well explained by the remaining ones, while variables close to the upper right corner are largely spanned by the remaining measures. *ENT*, indicated by a red dot, has the lowest R-squared (0.13) among all considered uncertainty measures, while *EPU* has the highest R-squared (0.70). The second smallest is *PCI*, whose R-squared of 0.47 is still substantially greater than that of *ENT*. Our entropy measure is the one that is furthest removed from the information content of standard measures of uncertainty.

We expand the in-sample analysis of Section 4.1 to include multiple uncertainty measures along with the more standard controls used previously in Table 3. We estimate a variety of time-series forecasting regressions of the form

$$R_{t+1,t+12} = \beta_0 + \beta_{ENT} ENT_t + \eta^\top Uncertain_t + \gamma^\top Control_t + \epsilon_t \tag{17}$$

[28] All the data used in this part of the analysis are available from the survey article. The source of the data can be found in Footnote 3 in David and Veronesi (2022).

where $R_{t+1,t+12}$ is the cumulative market return from month $t+1$ to month $t+12$ and ENT_t is the entropy measure in month t. $Uncertain_t$ contains non-text-based uncertainty measures: *BEX*, *JLN*, *PCI*, and squared implied volatility (*VIX2*). $Control_t$ contains all the remaining variables from equation (6), as well as *EPU* and *SEN*.

Columns (1) and (2) of Table 9 report the scaled coefficient estimates (and t-statistics) for (17), without and with the uncertainty controls.[29] A one standard deviation increase in entropy predicts a 2.355%-2.218% decrease in the 12-month ahead market return. Columns (3) and (4) replace *ENT* in (17) with two text-based measures: economic policy uncertainty (*EPU*) and sentiment (*SEN*). The results show that both *EPU* and *SEN* positively predict future market returns. Columns (5) and (6) include all text-based uncertainty measures *ENT*, *EPU*, and *SEN* in a single regression model, again without and with uncertainty controls. After controlling for these other measures, *ENT* still forecasts future market returns, and the standardized coefficients for *ENT* (from -2.060% to -2.203%) are nearly unchanged from columns (1) and (2). In all cases, introducing uncertainty controls (columns 2, 4, 6) does not meaningfully change the results without the uncertainty controls (columns 1, 3, 5).

We also replicate the analysis in Section 4.2 to test if any of the uncertainty measures are robust out-of-sample forecasters of aggregate market returns. The results for *BEX*, *JLN*, *PCI* and *VIX2* in Table 4 show that none of the uncertainty measures forecast market returns out-of-sample.

7.4 Extended Entropy

Our news-based entropy series begins in July 1997 because our news source begins in January 1996. We now describe our methodology for extending *ENT* back in time in order to expand the time frame of the in-sample tests in Tables 3 and 8.

Our extrapolation builds on work of van Binsbergen et al. (2024), who construct a 170-year-long measure of U.S. economic sentiment at both state and national levels. We posit that historical periods of high entropy should align with periods of high dispersion in regional sentiment because we expect both to be associated with unusual or difficult-to-process news. The regional sentiment in van Binsbergen et al. (2024) is observed annually. We find that in the post-July 1997 period, the correlation between the yearly average *ENT* and the standard deviation of regional sentiment is 0.34.

To extrapolate *ENT* backward in time, we first regress *ENT* on dispersion, regional

[29] Full regression results are in Tables A20 (scaled coefficients) and A21 (raw) of the Online Appendix.

sentiment deciles, and changes in national sentiment, using regional and national sentiment data from van Binsbergen et al. (2024). We then use the fitted linear model to predict *ENT* for periods prior to July 1997. Finally, we construct the time series $\widehat{ENT}$ by combining the predicted values for pre-July 1997 with the actual *ENT* data from July 1997 onward.

In testing the forecasting power of $\widehat{ENT}$, we also need extended time series of control variables. Because the cumulative factor returns over the previous 60 months, *RMW60* and *CMA60* from Fama and French (2015) begin in January 1970, this serves as the starting point for the robustness analysis in this section. *VIX* is available starting in January 1990, so prior to that we use the news implied volatility *NVIX* of Manela and Moreira (2017). *EPU* data starts in January 1985, and *SEN* starts in January 1980, but we are unable to identify a suitable method to extrapolate these variables. We therefore conducted additional robustness checks by running regressions starting in 1980 to include *SEN* as a control variable and from 1985 to incorporate both *EPU* and *SEN*. The results remain qualitatively unchanged from what we discuss below.

The results on robustness check with extended entropy are summarized in Table 10. Column (1) reports the coefficient estimate (and t-statistic) after scaling the coefficient estimate by the standard deviation of $\widehat{ENT}$, which replicates the full-sample part of Panel A of Table 3, except that we do not have *EPU* and *SEN* as control variables. A one standard deviation increase in extended entropy predicts a 3.754% decrease in the 12-month ahead cumulative market return. The remaining columns (2) to (8) report the scaled (by the standard deviation of $\widehat{ENT}$) impact on macro outcomes, which replicate the first row of corresponding columns (2) to (8) in Table 8, except that we do not have *EPU* as control variables. An increase in $\widehat{ENT}$ positively forecasts the unemployment rate and the VIX volatility index, and negatively forecasts industrial production. There is a weak negative forecasting relationship from extended entropy for S&P 500 earnings per share, and no relationship for *CPI_YOY*, *DGS10*, and *DGS10-2*. Thus, our extended entropy analysis preserves several, though not all, of the forecasting relations we observed for *ENT*.

8 Conclusion

This paper combines natural language processing with the tools of empirical asset pricing to investigate a novel aspect of how news affects prices at the aggregate level. We use a recurrent neural network to derive entropy, a measure of the novelty or unusualness of

aggregate news. This measure extends the literature (Glasserman and Mamaysky, 2019) by applying modern NLP tools to solve the sparsity and context problems with prior entropy measures.

We show that entropy negatively forecasts next twelve-month market returns, even after controlling for a multitude of known return forecasters. In particular, entropy does better than either economic policy uncertainty or news sentiment in our in-sample tests. In an out-of-sample market forecasting horse race, we find that, remarkably, entropy is the best forecaster of year-ahead market returns out of a large set of candidate variables.

Using a Fama-MacBeth GMM framework, we show that entropy has a negative risk price in the cross-section of portfolio returns. Together with the finding that entropy forecasts market returns negatively, it turns out entropy is the *only* factor in our study that is consistent with the Maio and Santa-Clara (2012) ICAPM sign property. Entropy thus proxies for a factor that negatively impacts investors' opportunity set and investors are willing to give up expected return to hedge against entropy innovations.

We show that entropy can be decomposed into one part that reflects the change in news flow and another part that reflects the change in the text model, and that it is the former, change-in-news, that accounts for entropy's out-of-sample forecasting properties and its cross-sectional risk pricing. We further show that the factor mimicking portfolio for entropy is among the least well-spanned out of 155 long-short factors we obtain from Jensen et al. (2022).

The forecasting power of entropy for future market returns likely stems from its ability to forecast future fundamental variables. Both entropy and its news innovation component positively forecast 12-month ahead unemployment and volatility, and negatively forecast industrial production, inflation, interest rates, and S&P500 corporate earnings. A theoretical model links entropy's ability to forecast a decline in economic conditions and the response of stock market prices: poor economic states are associated with lower dividends but they are also associated with less discounting of future dividends.

Our paper opens up several interesting directions for future work. First, the field of natural language processing is witnessing rapid advances which may produce more successful network architectures than recurrent neural networks for measuring news unusualness (e.g., Mikolov et al., 2010; Lee et al., 2017; Devlin et al., 2018; Xiong et al., 2019). Second, our analysis focuses on the impact of entropy at the aggregate level, while there may also be interesting applications at the level of individual stocks (Glasserman and Mamaysky, 2019). Finally, our entropy factor replicating portfolio should be a useful factor for asset pricing models.

A Theoretical Model

In this appendix, we develop a model in which an increase in entropy forecasts weaker economic conditions and a decline in the stock market, consistent with our empirical findings. The model summarizes the state of the economy through a scalar process. This economic state variable influences the news produced by a media sector, which responds to economic conditions. The state variable also influences the prices of risky assets in two ways: it affects the level of future dividends, and it enters into the stochastic discount factor used to price future dividends.

A.1 Model Formulation

Economic state

The state of the economy at time t is described by a scalar process θ_t, which evolves according to the equation

$$\theta_{t+1} = a\theta_t + \eta_{t+1}, \quad \eta_1, \eta_2, \ldots, \text{ i.i.d. } N(0, \sigma_\eta^2) \quad a \in (0, 1). \tag{18}$$

Larger values of θ_t correspond to better states of the economy. The stationary (or unconditional) distribution of θ_t is $N(0, v^2)$, with $v^2 = \sigma_\eta^2/(1 - a^2)$.

Media output

A media sector produces a sequence of "words" $X_{t,1}, X_{t,2}, \ldots$, in each period t. These "words" are encoded as positive real numbers. The encoding reflects the prevalence of the word in general usage, with numbers close to zero for very commonly occurring words and large numbers for more esoteric terms.

The media's selection of words in period t is influenced by the economic state θ_t. We posit the following mechanism for word selection:

$$X_{t,i} = \sigma(\theta_t)|\epsilon_{t,i}|, \quad \epsilon_{t,1}, \epsilon_{t,2} \ldots, \text{ i.i.d. } N(0, 1). \tag{19}$$

We choose a normal distribution for the $\epsilon_{t,i}$ for tractability, and we take their absolute values to make the $X_{t,i}$ positive. The economic state influences the selection of words through the factor $\sigma(\theta_t)$, for which we set

$$\sigma^2(\theta) = e^{-\lambda\theta}, \quad \lambda > 0. \tag{20}$$

Under this specification, weaker economic conditions (smaller values of θ) are associated with more variable media content. This feature is qualitatively consistent with the finding in Van Nieuwerburgh and Veldkamp (2006) and Patton and Timmermann (2010) that the dispersion of economic forecasts is greater in recessions, as greater disagreement among forecasters should lead to greater variation in the media's discussion of the economy. The specification in (20) also means that a weaker economy is associated with a higher frequency of words that are ordinarily less prevalent—those with larger numerical representations. A weak economy may result, in part, from an inability of policy-makers to navigate unfamiliar circumstances, and these circumstances may be associated with an increase in the frequency of less-familiar terms, like "subprime," "coronavirus," or "Evergrande."

Entropy

Under (19), the conditional density of each $X_{t,i}$, given the economic state $\theta_t = \theta$, is the half-normal density (the density of the absolute value of a mean-zero normal random variable)

$$f(x|\theta) = \frac{\sqrt{2}}{\sigma(\theta)\sqrt{\pi}} e^{-x^2/2\sigma^2(\theta)}, \quad x \geq 0. \tag{21}$$

In our empirical analysis, we use the empirical distribution of words in each month to calculate entropy. For our theoretical analysis, we will use the theoretical density (21) instead, viewing it as an approximation to the empirical distribution for a long sequence of words.

Define the cross-entropy measure

$$H(\theta_{t-1}, \theta_t) = -\int [\log f(x|\theta_{t-1})] f(x|\theta_t)\, dx. \tag{22}$$

This expression is the model counterpart of the empirical quantity in (3). In (3), the log-probabilities of words, as determined by the RNN trained from $t-6$ to $t-1$, are averaged with respect to their empirical frequencies at time t. In (22), we evaluate the log-densities of words according to the time-$(t-1)$ density $f(x|\theta_{t-1})$ and average them with respect to their time-t frequencies.[30]

[30]The entropy measure in (3) is calculated from probability mass functions, so it will always be positive because the log-probabilities are negative or zero. The log-density in (22) can be positive or negative, so $H(\theta_{t-1}, \theta_t)$ can be negative. This does not pose a problem because we are interested in relative levels of $H(\theta_{t-1}, \theta_t)$.

Noting that

$$\log f(x|\theta) = \log(\sqrt{2/\pi}) - \log\sigma(\theta) - \frac{x^2}{2\sigma^2(\theta)},$$

we can calculate (22) (shifting the time index for later convenience) as follows:

$$\begin{aligned} H(\theta_t, \theta_{t+1}) &= -\mathbb{E}[\log f(X_{t+1}|\theta_t)|\theta_t, \theta_{t+1}] \\ &= \log(\sqrt{\pi/2}) + \log\sigma(\theta_t) + \frac{\mathbb{E}[X_{t+1}^2|\theta_t, \theta_{t+1}]}{2\sigma^2(\theta_t)} \\ &= \log(\sqrt{\pi/2}) + \log\sigma(\theta_t) + \frac{\sigma^2(\theta_{t+1})}{2\sigma^2(\theta_t)} \qquad (23) \\ &= \log(\sqrt{\pi/2}) - \frac{\lambda}{2}\theta_t + \frac{1}{2}e^{-\lambda(\theta_{t+1}-\theta_t)}. \qquad (24) \end{aligned}$$

With $\lambda > 0$, changes in the distribution of media output are driven by changes in economic conditions, so entropy $H(\theta_t, \theta_{t+1})$ is a function of consecutive economic states. At $\lambda = 0$, media output does not respond to the state of the economy; entropy becomes constant and therefore uninformative about changes in economic conditions.

Stock price

A risky asset (a stock) pays a dividend at time $t+1$ given by

$$D_{t+1} = c_0 + c_1\theta_{t+1} + u_{t+1}, \quad u_1, u_2, \ldots \text{i.i.d. } N(0, \sigma_D^2), \qquad (25)$$

for some constants $c_0, c_1 > 0$. Thus, dividends are higher in better economic states. Dividends inherit some persistence from θ_{t+1}. The shocks u_t are independent of the η_t in (18).

We assume that claims on future dividends are priced through a stochastic discount factor (SDF). We posit an SDF proportional to $e^{-\gamma\theta_{t+1}}$, for some risk aversion parameter $\gamma > 0$. This captures the usual interpretation that marginal utilities of consumption are lower in good times. Under this SDF, dividends that covary positively with the future economic state θ_{t+1} are more heavily discounted at time t.[31] In particular, we assume that the time-t price of a claim paying D_{t+1} at time $t+1$ is given by

$$S(\theta_t) = \mathbb{E}[e^{-\gamma\theta_{t+1} - \gamma^2\sigma_\eta^2/2} D_{t+1}|\theta_t], \qquad (26)$$

for some $\gamma > 0$. The normalization factor of $e^{-\gamma^2\sigma_\eta^2/2}$ in the SDF is convenient but not

[31] Similarly, an asset whose dividends covary negatively with θ_{t+1} would sell for a higher price, consistent with the fact that it hedges a deterioration in economic fundamentals.

essential. With this normalization, the time-t price of a claim paying 1 at $t+1$ is given by

$$\mathbb{E}[e^{-\gamma\theta_{t+1}-\gamma^2\sigma_\eta^2/2}|\theta_t] = e^{-\gamma a\theta_t-\gamma^2\sigma_\eta^2/2}\mathbb{E}[e^{-\gamma\eta_{t+1}}] = e^{-\gamma a\theta_t},$$

implying an interest rate of $\gamma a\theta_t$. In particular, poorer economic conditions are associated with lower interest rates. Applying a similar calculation to (26) yields

$$\begin{aligned} S(\theta_t) &= e^{-\gamma a\theta_t-\gamma^2\sigma_\eta^2/2}\mathbb{E}[e^{-\gamma\eta_{t+1}}D_{t+1}|\theta_t] \\ &= e^{-\gamma a\theta_t-\gamma^2\sigma_\eta^2/2}\mathbb{E}[e^{-\gamma\eta_{t+1}}(c_0+c_1\theta_{t+1}+u_{t+1})|\theta_t] \\ &= e^{-\gamma a\theta_t-\gamma^2\sigma_\eta^2/2}(c_0+c_1a\theta_t)\mathbb{E}[e^{-\gamma\eta_{t+1}}] + e^{-\gamma a\theta_t-\gamma^2\sigma_\eta^2/2}c_1\mathbb{E}[e^{-\gamma\eta_{t+1}}\eta_{t+1}] \\ &= e^{-\gamma a\theta_t}(c_0+c_1a\theta_t-c_1\gamma\sigma_\eta^2). \end{aligned} \tag{27}$$

Define the "return"

$$R_{t,t+1} = D_{t+1} - S(\theta_t) = c_0 + c_1\theta_{t+1} - (c_0 - c_1\gamma\sigma_\eta^2)e^{-\gamma a\theta_t} - c_1a\theta_t e^{-\gamma a\theta_t} + u_{t+1}; \tag{28}$$

this is the profit from buying the single-period stock at time t at a price of $S(\theta_t)$ and receiving the dividend D_{t+1} at time $t+1$.

A.2 Model Implications

Our empirical analysis documents that higher entropy forecasts (i) a decline in economic conditions, and (ii) negative stock market returns. We now turn to support for these findings in our model.

Define

$$C_E(\lambda) = \mathsf{cov}[H(\theta_t,\theta_{t+1}),\theta_{t+2}] \tag{29}$$

$$C_M(\lambda) = \mathsf{cov}[H(\theta_t,\theta_{t+1}),R_{t+1,t+2}], \tag{30}$$

where the dependence on λ enters through $H(\theta_t,\theta_{t+1})$, as in (24). If $C_E(\lambda) < 0$, then higher entropy forecasts weaker economic conditions. If $C_M(\lambda) < 0$, then higher entropy forecasts a lower market return. Thus, settings in which both covariance terms are negative are settings consistent with our empirical findings.

We have defined the covariances in (29) and (30) using the level of entropy. If we used

the change in entropy instead, we would get

$$\begin{aligned}\mathsf{cov}[H(\theta_t,\theta_{t+1})-H(\theta_{t-1},\theta_t),\theta_{t+2}] &= C_E(\lambda)-\mathsf{cov}[H(\theta_{t-1},\theta_t),a\theta_{t+1}+\eta_{t+2}]\\ &= C_E(\lambda)-a\mathsf{cov}[H(\theta_{t-1},\theta_t),\theta_{t+1}]\\ &= (1-a)C_E(\lambda),\end{aligned}$$

which has the same sign as $C_E(\lambda)$. A similar argument applies to $C_M(\lambda)$. Since our interest is in the signs of the covariances and not their magnitudes, it suffices to consider covariances with the level of entropy, as in (29) and (30).

The first of these is always negative—higher entropy forecasts weaker economic conditions:

Proposition 1 *For all $a \in (0,1)$, all $\sigma_\eta > 0$, and all $\lambda > 0$, we have $C_E(\lambda) < 0$.*

The covariance between entropy and market return is more complicated, primarily because the stock price depends on the SDF. (Indeed, with $\gamma = 0$ the covariance is always zero.) We can decompose $C_M(\lambda)$ into the difference between $\mathsf{cov}[H(\theta_t,\theta_{t+1}),D_{t+2}]$ and $\mathsf{cov}[H(\theta_t,\theta_{t+1}),S(\theta_{t+1})]$ The firm term is negative because $C_E(\lambda)$ is negative and D_{t+2} loads positively on θ_{t+2}: if entropy forecasts a weaker economy, it also forecasts lower dividends. But the covariance between entropy and the stock price is ambiguous because the stock price depends non-monotonically on the economic state θ_{t+1}. A better economic state leads to higher dividends, but these dividends are more heavily discounted through the SDF. For some parameter values, $C_M(\lambda)$ can be positive.

To find a parameter range that guarantees a negative covariance, we first note that $C_M(0) = 0$ because $H(\theta_t,\theta_{t+1})$ is a constant if $\lambda = 0$. If we can show that (over a range of parameters) $\partial C_M(0)/\partial\lambda < 0$, it will follow that $C_M(\lambda)$ is strictly negative (for those parameters) in an interval $(0,\lambda_*)$, for some $\lambda_* > 0$. The following result provides an explicit sufficient condition, which simplifies a more general expression in the next section.

Proposition 2 *Suppose $a \in (0,1)$ and $c_0, c_1, \gamma, \sigma_\eta^2 > 0$. If*

$$c_1/c_0 < \frac{\gamma(1-a^2)}{1-a^2+\sigma_\eta^2\gamma^2}, \tag{31}$$

then $\partial C_M(0)/\partial\lambda < 0$, and therefore $C_M(\lambda) < 0$ for all $\lambda \in (0,\lambda_*)$, for some $\lambda_* > 0$.

To help shed light on condition (31), we discuss reasonable values for the model's parameters. Because common measures of economic conditions are highly persistent, we take $a = 0.9$ in (18). To help pin down other parameters, note from (25) and (27) that the discount — recall that $\gamma a \theta_t$ is the interest rate — in the stock price is

$$e^{-\gamma a \theta_t} \mathbb{E}[D_{t+1}|\theta_t] - S(\theta_t) = e^{-\gamma a \theta_t} c_1 \gamma \sigma_\eta^2.$$

Normalizing by $e^{-\gamma a \theta_t} \mathbb{E}[D_{t+1}] = e^{-\gamma a \theta_t} c_0$ (since $\mathbb{E}[D_{t+1}|\theta_t]$ and $S(\theta_t)$ may be negative), yields $(c_1/c_0)\gamma\sigma_\eta^2$, which we interpret as the model's equity risk premium. The ratio c_1/c_0 is the exposure of the dividend D_{t+1} in (25) to θ_{t+1}, normalized to an expected dividend of 1.

Similarly, to get the systematic component of dividend volatility (ignoring the idiosyncratic shocks u_t in (25)), we calculate the standard deviation of $c_1(\theta_{t+1} - \theta_t)$ and then normalize the standard deviation by the average dividend $\mathbb{E}[D_{t+1}] = c_0$ to get a volatility of $(c_1/c_0)\sigma_\eta\sqrt{2/(1+a)}$.

Setting $c_0 = 1$, $c_1 = 0.25$, $\sigma_\eta^2 = 0.4$, and $\gamma = 0.5$ yields an equity risk premium of 5% and a (systematic) dividend volatility of 16%. These are broadly plausible magnitudes. With these parameters, the left side of (31) is 0.25, the right side is 0.33, and the inequality is satisfied.

The condition in (31) can also be written as

$$\frac{c_1 \sigma_\eta^2}{1 - a^2} \gamma^2 - c_0 \gamma + c_1 < 0. \tag{32}$$

The left side is a quadratic function of γ. If $c_1/c_0 < \sqrt{1-a^2}/(2\sigma_\eta)$, then the quadratic has two positive roots, and the inequality is satisfied for all γ between these roots.

Figure 9 plots $C_M(\lambda)$ for serval combinations of c_1 and γ. The plots are nearly linear in λ, suggesting that the derivative $\partial C_M(0)/\partial\lambda$ is a good indicator of the sign of $C_M(\lambda)$ for a wide range of λ values.

The only combination producing a positive covariance is $c_1 = 0.4$, $\gamma = 1$, which has both parameters at the largest values compared. This case violates our sufficient condition. More generally, (32) becomes more difficult to satisfy at larger values of c_1 and γ.

In fact, (32) is violated for all three combinations with $c_1 = 0.4$, even though two of these combinations yield negative covariances. Thus, our sufficient condition is not a necessary condition, particularly at larger values of c_1. An exact expression for $\partial C_M(0)/\partial\lambda$,

and thus a necessary and sufficient condition for this derivative to be negative, is derived in Section A.3.2.

A.3 Supporting Analysis

A.3.1 Proposition 1

As a preliminary calculation, we have, for any $\mu_0, \mu_1 \in \mathbb{R}$,

$$\begin{aligned}
\mathbb{E}[\theta_{t+1} e^{\mu_1 \theta_{t+1}} e^{\mu_0 \theta_t}] &= \partial_{\mu_1} \mathbb{E}[e^{\mu_1 \theta_{t+1}} e^{\mu_0 \theta_t}] \\
&= \partial_{\mu_1} \mathbb{E}[e^{(a\mu_1+\mu_0)\theta_t} e^{\mu_1 \eta_{t+1}}] \\
&= \partial_{\mu_1} e^{(a\mu_1+\mu_0)^2 v^2/2} e^{\mu_1^2 \sigma_\eta^2/2} \\
&= [(a\mu_1+\mu_0)av^2 + \mu_1 \sigma_\eta^2] e^{(a\mu_1+\mu_0)^2 v^2/2} e^{\mu_1^2 \sigma_\eta^2/2}.
\end{aligned}$$

With $\mu_0 = \lambda$ and $\mu_1 = -\lambda$, we get

$$\begin{aligned}
\mathbb{E}[\theta_{t+1} e^{-\lambda(\theta_{t+1}-\theta_t)}] &= \lambda[(1-a)av^2 - \sigma_\eta^2] e^{\lambda^2(1-a)^2 v^2/2 + \lambda^2 \sigma_\eta^2/2} \\
&= -\frac{\lambda \sigma_\eta^2}{1+a} e^{\lambda^2 \sigma_\eta^2/(1+a)} < 0,
\end{aligned}$$

where the second step uses $v^2 = \sigma_\eta^2/(1-a^2)$. Thus,

$$\begin{aligned}
\mathsf{cov}[H(\theta_t, \theta_{t+1}), \theta_{t+2}] &= -\frac{\lambda}{2}\mathbb{E}[\theta_t \theta_{t+2}] + \frac{1}{2}\mathbb{E}[e^{-\lambda(\theta_{t+1}-\theta_t)}\theta_{t+2}] \\
&= -\frac{\lambda}{2}a^2 v^2 + \frac{a}{2}\mathbb{E}[e^{-\lambda(\theta_{t+1}-\theta_t)}\theta_{t+1}] < 0.
\end{aligned}$$

A.3.2 Proposition 2

Define

$$\begin{aligned}
R(\theta_{t+1}) &= \mathbb{E}[R_{t+1,t+2}|\theta_{t+1}] \\
&= c_1 a \theta_{t+1} - S(\theta_{t+1}) \\
&= c_1 a \theta_{t+1} + (c_1 \gamma \sigma_\eta^2 - c_0) e^{-a\gamma\theta_{t+1}} - c_1 a \theta_{t+1} e^{-a\gamma\theta_{t+1}}. \qquad (33)
\end{aligned}$$

Then

$$\mathsf{cov}[H(\theta_t, \theta_{t+1}), R_{t+1,t+2}] = \mathsf{cov}[H(\theta_t, \theta_{t+1}), R(\theta_{t+1})] = \mathbb{E}[H(\theta_t, \theta_{t+1})(R(\theta_{t+1}) - \mathbb{E}[R(\theta_{t+1})]].$$

So

$$\begin{aligned}\partial C_M/\partial\lambda &= \partial_\lambda \mathbb{E}[H(\theta_t,\theta_{t+1})(R(\theta_{t+1})-\mathbb{E}[R(\theta_{t+1})]] \\ &= \mathbb{E}[\partial_\lambda H(\theta_t,\theta_{t+1})(R(\theta_{t+1})-\mathbb{E}[R(\theta_{t+1})]].\end{aligned}$$

At $\lambda = 0$, (24) yields

$$\partial_\lambda H(\theta_t,\theta_{t+1}) = -\frac{\theta_t}{2} - \frac{1}{2}(\theta_{t+1}-\theta_t) = -\frac{\theta_{t+1}}{2},$$

so

$$\begin{aligned}\partial C_M(0)/\partial\lambda &= -\frac{1}{2}\mathbb{E}[\theta_{t+1}R(\theta_{t+1})] \\ &= -\frac{1}{2}\left(c_1 a\mathbb{E}[\theta_{t+1}^2] + (c_1\gamma\sigma_\eta^2 - c_0)\mathbb{E}[\theta_{t+1}e^{-a\gamma\theta_{t+1}}] - c_1 a\mathbb{E}[\theta_{t+1}^2 e^{-a\gamma\theta_{t+1}}]\right) \\ &= -\frac{1}{2}\left(c_1 a v^2 - (c_1\gamma\sigma_\eta^2 - c_0)\gamma a v^2 e^{a^2\gamma^2 v^2/2} - c_1 a v^2(1+a^2\gamma^2 v^2)e^{a^2\gamma^2 v^2/2}\right) \\ &= -\frac{av^2}{2}\left(c_1 - [(c_1\gamma\sigma_\eta^2 - c_0)\gamma + c_1(1+a^2\gamma^2 v^2)]e^{a^2\gamma^2 v^2/2}\right).\end{aligned}$$

A sufficient condition for $\partial C_M(0)/\partial\lambda$ to be negative is that the term in square brackets be negative; i.e.,

$$c_1(\sigma_\eta^2 + a^2 v^2)\gamma^2 - c_0\gamma + c_1 < 0.$$

Making the substitution $v^2 = \sigma_\eta^2/(1-a^2)$ yields the equivalent condition

$$\frac{c_1\sigma_\eta^2}{1-a^2}\gamma^2 - c_0\gamma + c_1 < 0,$$

which is equivalent to (31).

References

Aït-Sahalia, Yacine, Felix Matthys, Emilio Osambela, and Ronnie Sircar (2021) "When uncertainty and volatility are disconnected: implications for asset pricing and portfolio performance,"Technical report, National Bureau of Economic Research.

Ang, Andrew, Robert J Hodrick, Yuhang Xing, and Xiaoyan Zhang (2006) "The cross-section of volatility and expected returns," *Journal of Finance*, 61 (1), 259–299.

Azzimonti, Marina (2018) "Partisan conflict and private investment," *Journal of Monetary Economics*, 93, 114–131.

Baker, Scott R, Nicholas Bloom, and Steven J Davis (2016) "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, 131 (4), 1593–1636.

Ball, Ray, Joseph Gerakos, Juhani T. Linnainmaa, and Valeri Nikolaev (2016) "Accruals, cash flows, and operating profitability in the cross section of stock returns," *Journal of Financial Economics*, 121 (1), 28–45.

Bekaert, Geert, Eric C Engstrom, and Nancy R Xu (2022) "The time variation in risk appetite and uncertainty," *Management Science*, 68 (6), 3975–4004.

van Binsbergen, Jules H, Svetlana Bryzgalova, Mayukh Mukhopadhyay, and Varun Sharma (2024) "(Almost) 200 Years of News-Based Economic Sentiment,"Technical report, National Bureau of Economic Research.

Boons, Martijn (2016) "State variables, macroeconomic activity, and the cross section of individual stocks," *Journal of Financial Economics*, 119 (3), 489–511.

Breeden, Douglas T, Michael R Gibbons, and Robert H Litzenberger (1989) "Empirical tests of the consumption-oriented CAPM," *Journal of Finance*, 44 (2), 231–262.

Brenner, Menachem and Yehuda Izhakian (2018) "Asset pricing and ambiguity: Empirical evidence," *Journal of Financial Economics*, 130 (3), 503–531.

Brogaard, Jonathan and Andrew Detzel (2015) "The asset-pricing implications of government economic policy uncertainty," *Management Science*, 61 (1), 3–18.

Brooks, Jordan, Noah Feilbogen, Yao Hua Ooi, and Adam Akant (2023) "Economic Trend,"Technical report, AQR.

Calomiris, Charles W and Harry Mamaysky (2019) "How news and its context drive risk and returns around the world," *Journal of Financial Economics*, 133 (2), 299–336.

Campbell, John Y and Robert J Shiller (1988) "Stock prices, earnings, and expected dividends," *Journal of Finance*, 43 (3), 661–676.

Campbell, John Y and Samuel B Thompson (2008) "Predicting excess stock returns out of sample: Can anything beat the historical average?" *Review of Financial Studies*, 21 (4), 1509–1531.

Charniak, E (2019) *An Introduction to Deep Learning*: MIT Press.

Chen, Andrew and Tom Zimmerman (2021) "Open source cross-sectional asset pricing,"Technical report, SSRN.

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

Cohen, Lauren, Christopher Malloy, and Quoc Nguyen (2020) "Lazy prices," *Journal of Finance*, 75 (3), 1371–1415.

David, Alexander and Pietro Veronesi (2022) "A Survey of Alternative Measures of Macroeconomic Uncertainty: Which Measures Forecast Real Variables and Explain Fluctuations in Asset Volatilities Better?" *Annual Review of Financial Economics*, 14, 439–463.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018) "Bert: Pre-training of deep bidirectional transformers for language understanding,"Technical report, arXiv preprint arXiv:1810.04805.

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

Fama, Eugene F. and James D. MacBeth (1973) "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81 (3), 607–636.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu (2020) "Taming the factor zoo: A test of new factors," *Journal of Finance*, 75 (3), 1327–1370.

Gabaix, Xavier and Ralph S.J. Koijen (2021) "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis,"Technical report, working paper.

Glasserman, Paul, Fulin Li, and Harry Mamaysky (2025) "Time Variation in the News–Returns Relationship," *Journal of Financial and Quantitative Analysis*, 60 (1), 258–294.

Glasserman, Paul and Caden Lin (2024) "Assessing Look-Ahead Bias in Stock Return Predictions Generated By GPT Sentiment Analysis," *Journal of Financial Data Science*, 6 (1), 25–42.

Glasserman, Paul and Harry Mamaysky (2019) "Does unusual news forecast market stress?" *Journal of Financial and Quantitative Analysis*, 54 (5), 1937–1974.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016) *Deep learning*: MIT press.

Gromb, Denis and Dimitri Vayanos (2010) "Limits of Arbitrage," *Annual Review of Financial Economics*, 2 (1), 251–275.

Hansen, Lars Peter and Thomas J Sargent (2008) "Risk sensitivity, model uncertainty, and asset pricing," in *Robustness*, 307–326: Princeton University Press.

Harvey, Campbell R, Yan Liu, and Heqing Zhu (2016) "... and the cross-section of expected returns," *Review of Financial Studies*, 29 (1), 5–68.

Hoberg, Gerard and Gordon Phillips (2016) "Text-based network industries and endogenous product differentiation," *Journal of Political Economy*, 124 (5), 1423–1465.

Hoberg, Gerard and Gordon M Phillips (2018) "Text-based industry momentum," *Journal of Financial and Quantitative Analysis*, 53 (6), 2355–2388.

Hochreiter, Sepp and Jürgen Schmidhuber (1997) "Long short-term memory," *Neural Computation*, 9 (8), 1735–1780.

Jensen, Theis Ingerslev, Bryan T Kelly, and Lasse Heje Pedersen (2021) "Is there a replication crisis in finance?"Technical report, National Bureau of Economic Research.

——— (2022) "Is There A Replication Crisis In Finance?" *Journal of Finance*, 78 (5), 2465–2518.

Jiang, Fuwei, Joshua Lee, Xiumin Martin, and Guofu Zhou (2019) "Manager sentiment and stock returns," *Journal of Financial Economics*, 132 (1), 126–149.

Jurado, Kyle, Sydney C Ludvigson, and Serena Ng (2015) "Measuring uncertainty," *American Economic Review*, 105 (3), 1177–1216.

Jurafsky, D and J Martin (2023) "Speech and Language Processing, 3rd Edition."

Keskar, Nitish Shirish, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang (2016) "On large-batch training for deep learning: Generalization gap and sharp minima,"Technical report, arXiv preprint arXiv:1609.04836.

Kingma, Diederik P and Jimmy Ba (2014) "Adam: A method for stochastic optimization,"Technical report, arXiv preprint arXiv:1412.6980.

Lamont, Owen, Christopher Polk, and Jesús Saaá-Requejo (2001) "Financial constraints and stock returns," *Review of Financial Studies*, 14 (2), 529–554.

Lee, Kenton, Luheng He, Mike Lewis, and Luke Zettlemoyer (2017) "End-to-end neural coreference resolution,"Technical report, arXiv preprint arXiv:1707.07045.

Lettau, Martin and Sydney Ludvigson (2001) "Consumption, aggregate wealth, and expected stock returns," *Journal of Finance*, 56 (3), 815–849.

Levy, Bradford (2025) "Caution Ahead: Numerical Reasoning and Look-ahead Bias in AI Models,"Technical report, Chicago Booth School of Business.

Maio, Paulo and Pedro Santa-Clara (2012) "Multifactor models and their consistency with the ICAPM," *Journal of Financial Economics*, 106 (3), 586–613.

Manela, Asaf and Alan Moreira (2017) "News implied volatility and disaster concerns," *Journal of Financial Economics*, 123 (1), 137–162.

Martin, Ian (2017) "What is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132 (1), 367–433.

Merton, Robert C (1973) "An intertemporal capital asset pricing model," *Econometrica: Journal of the Econometric Society*, 867–887.

Mikolov, Tomas, Martin Karafiát, Lukas Burget, Jan Cernockỳ, and Sanjeev Khudanpur (2010) "Recurrent neural network based language model," in *Interspeech*, 2 (3), 1045–1048, Makuhari.

Newey, Whitney K. and Kenneth D. West (1987) "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55 (3), 703–708.

Patton, Andrew J and Allan Timmermann (2010) "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics*, 57 (7), 803–820.

Pennington, Jeffrey, Richard Socher, and Christopher D Manning (2014) "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Politis, Dimitris N and Joseph P Romano (1994) "The stationary bootstrap," *Journal of the American Statistical association*, 89 (428), 1303–1313.

Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever et al. (2019) "Language models are unsupervised multitask learners," *OpenAI blog*, 1 (8), 9.

Sarkar, Suproteem K and Keyon Vafa (2024) "Lookahead bias in pretrained language models,"Technical report, working paper.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson (2022) "Measuring news sentiment," *Journal of econometrics*, 228 (2), 221–243.

Sims, Christopher A. (2003) "Implications of rational inattention," *Journal of Monetary Economics*, 50 (3), 665–690.

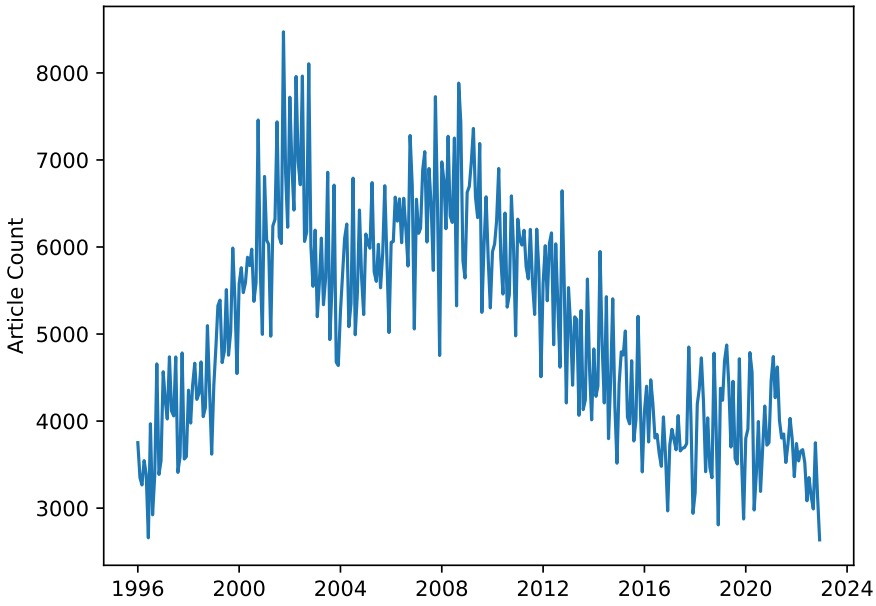
Tetlock, Paul C (2011) "All the news that's fit to reprint: Do investors react to stale information?" *Review of Financial Studies*, 24 (5), 1481–1512.

Van Nieuwerburgh, Stijn and Laura Veldkamp (2006) "Learning asymmetries in real business cycles," *Journal of monetary Economics*, 53 (4), 753–772.

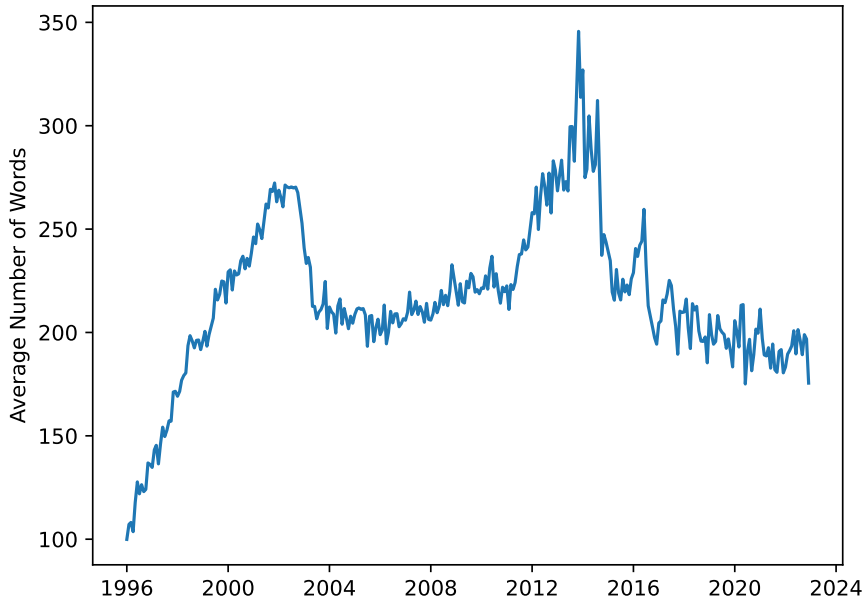
Welch, Ivo and Amit Goyal (2008) "A comprehensive look at the empirical performance of equity premium prediction," *Review of Financial Studies*, 21 (4), 1455–1508.

Werbos, Paul J (1990) "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, 78 (10), 1550–1560.

Xiong, Wenhan, Jingfei Du, William Yang Wang, and Veselin Stoyanov (2019) "Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model,"Technical report, arXiv preprint arXiv:1912.09637.



(a) Time series of number of articles in each month.



(b) Time series of average article length in each month.

Figure 1: Thomson Reuters News Feed Direct archive articles characteristics from January 1996 to December 2022.

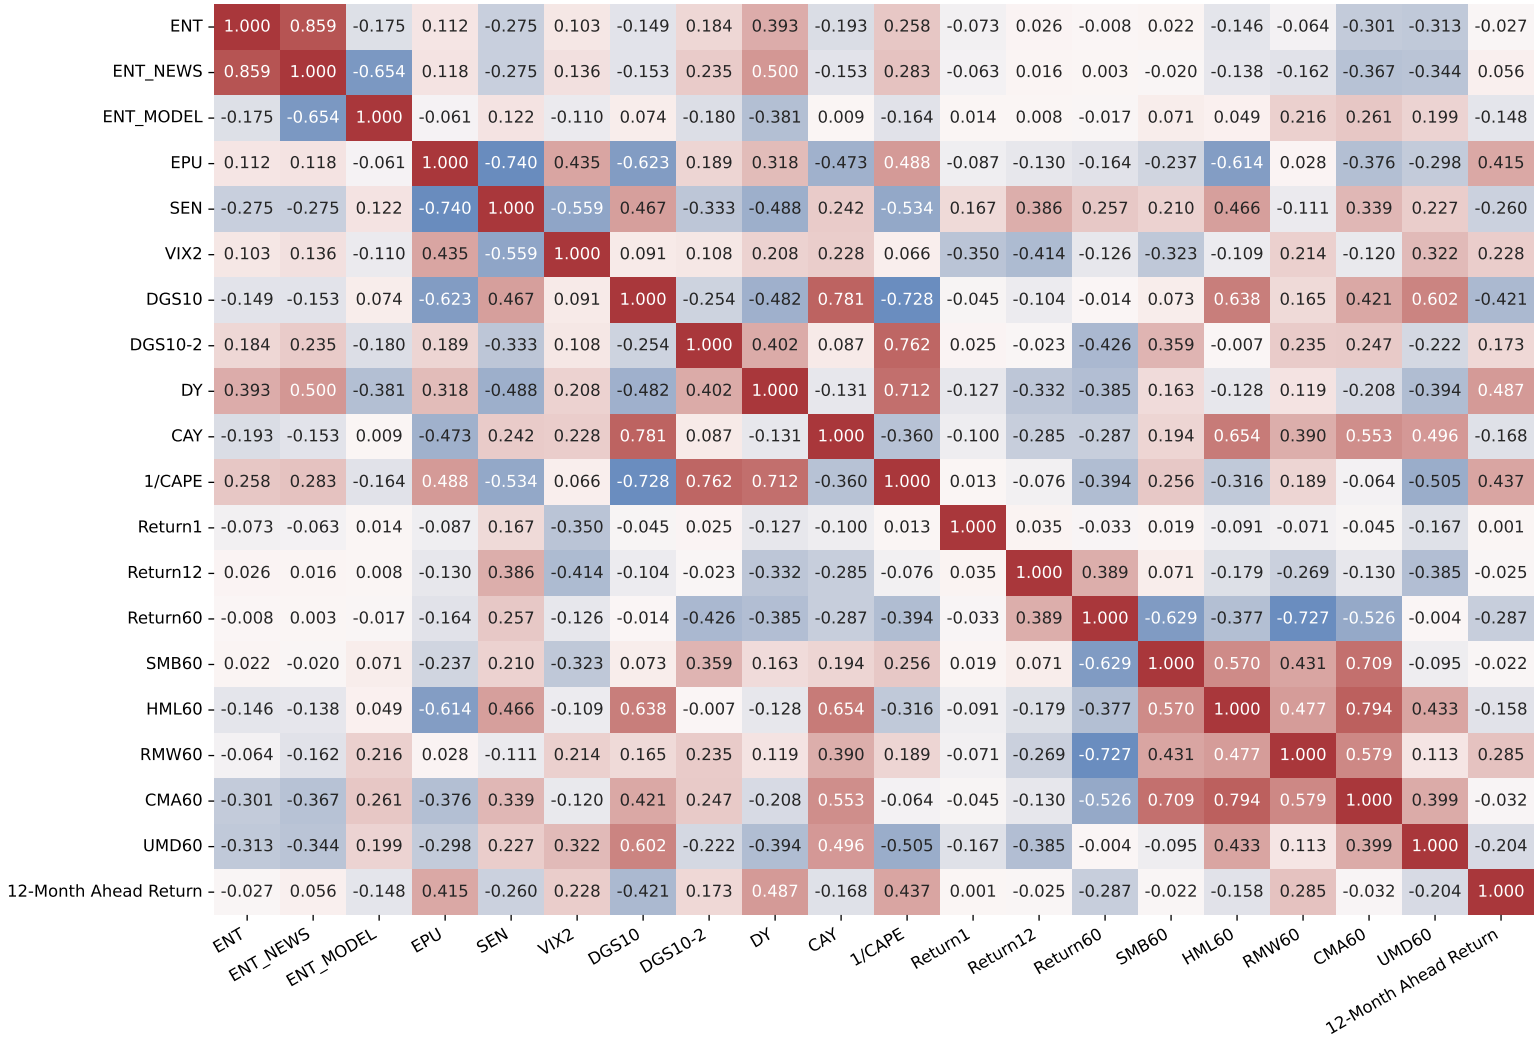


Figure 2: Correlations at monthly observation using data from July 1997 to December 2022, inclusive.

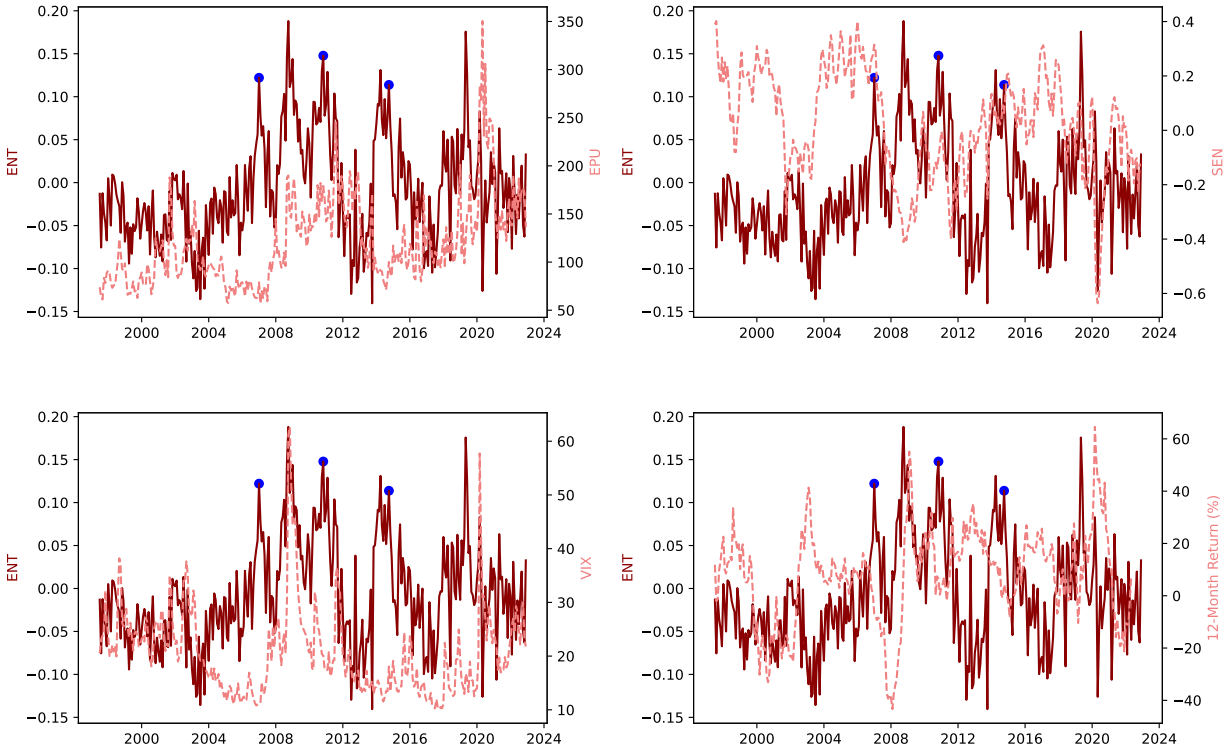


Figure 3: The top-left panel plots *ENT* and *EPU*. The top-right panel plots *ENT* and *SEN*. The bottom-left panel plots *ENT* and the *VIX* index. The bottom-right panel plots *ENT* and 12-month ahead cumulative returns on the CRSP value-weighted index (ends on December 2021). In all cases, *ENT* is plotted as the solid line, with the other series plotted as the dashed line. Data are shown on a monthly frequency. The blue dots on each series correspond to the months January 2007, November 2010, and October 2014, which are discussed in Section 3.4.

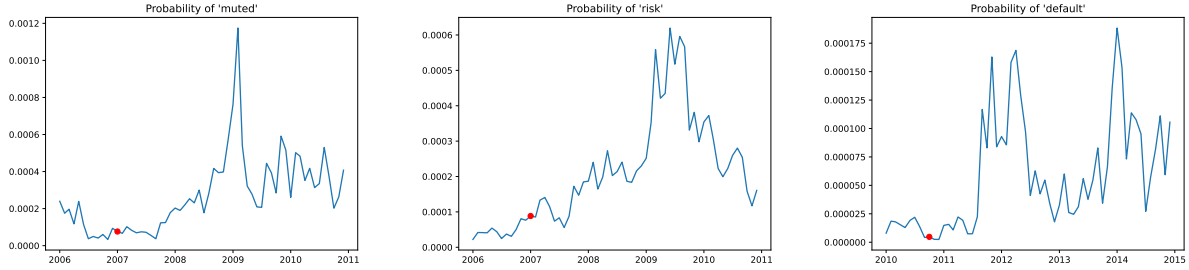


Figure 4: The left panel plots the probability of "muted" following "Fannie Mae and Freddie Mac growth will be." The middle panel plots the probability of "risk" following "lack of progress in reining in mortgage lenders Fannie Mae and Freddie Mac leaves the economy at." The right panel plots the probability of "default" following "growing concerns Ireland will be forced to." The red dot in each panel marks the date the corresponding sentence appears in the database.



Figure 5: The most probable 100 words following the string of words "Fannie Mae and Freddie Mac growth will be" under the December 2006 model (left panel) and under the December 2008 model (right panel). The size of each word is proportional to its probability.

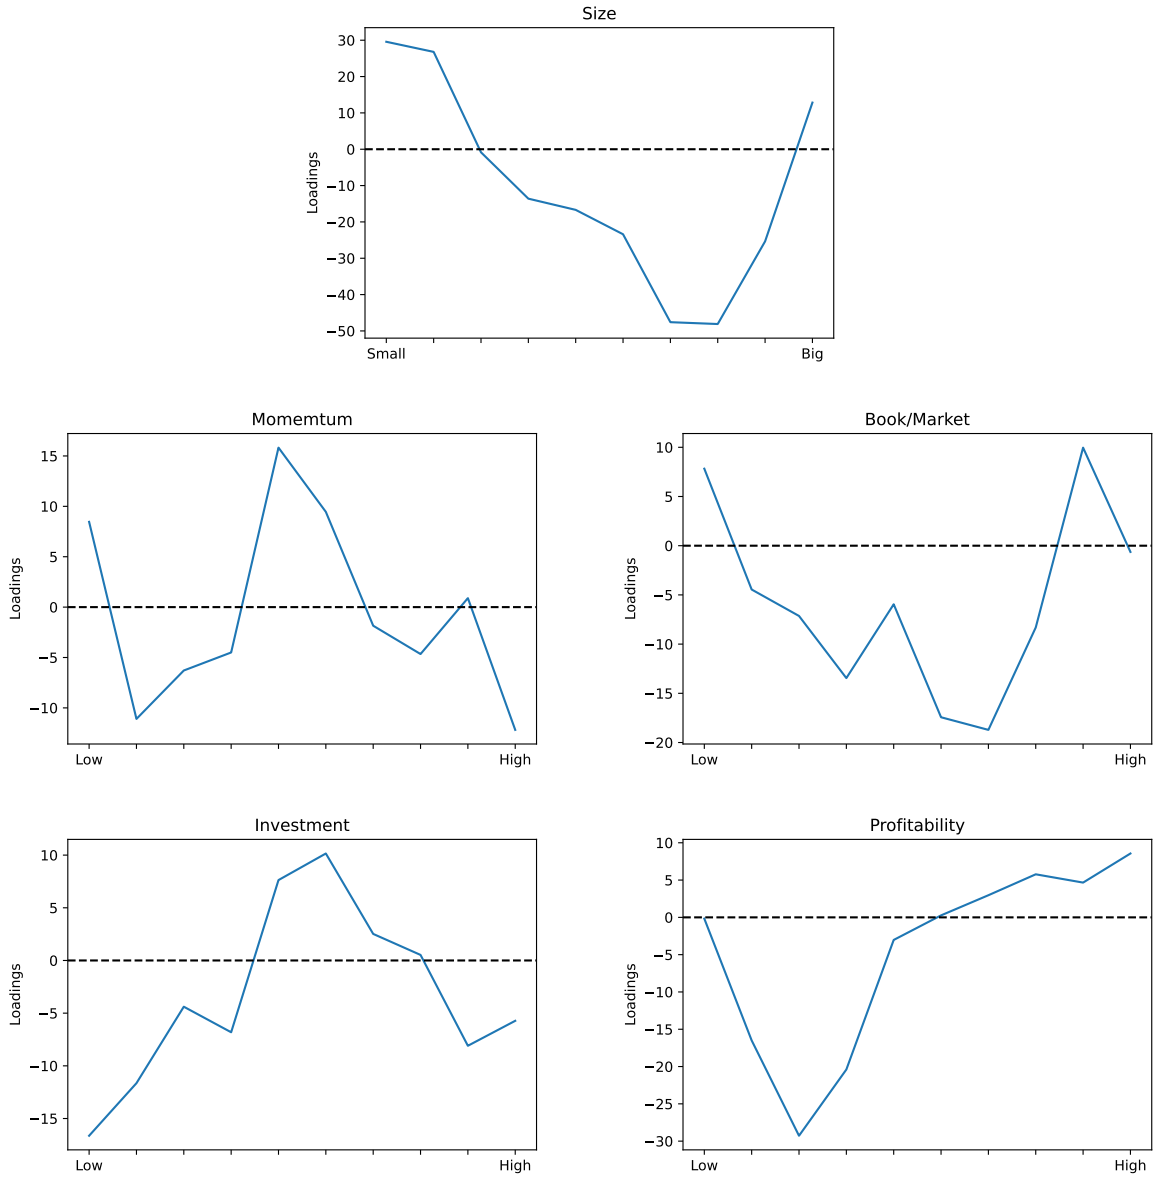


Figure 6: The first panel shows the F_{ENT} loadings of 10 portfolios sorted on size. The other panels show F_{ENT} loadings of portfolios sorted on momentum, book/market, investment, and operating profitability, respectively. For each panel, we construct F_{ENT} using 100 base assets: 25 portfolios formed on size-momentum, size-book/market, size-investment, and size-profitability. The methodology is described in Section 5.3.

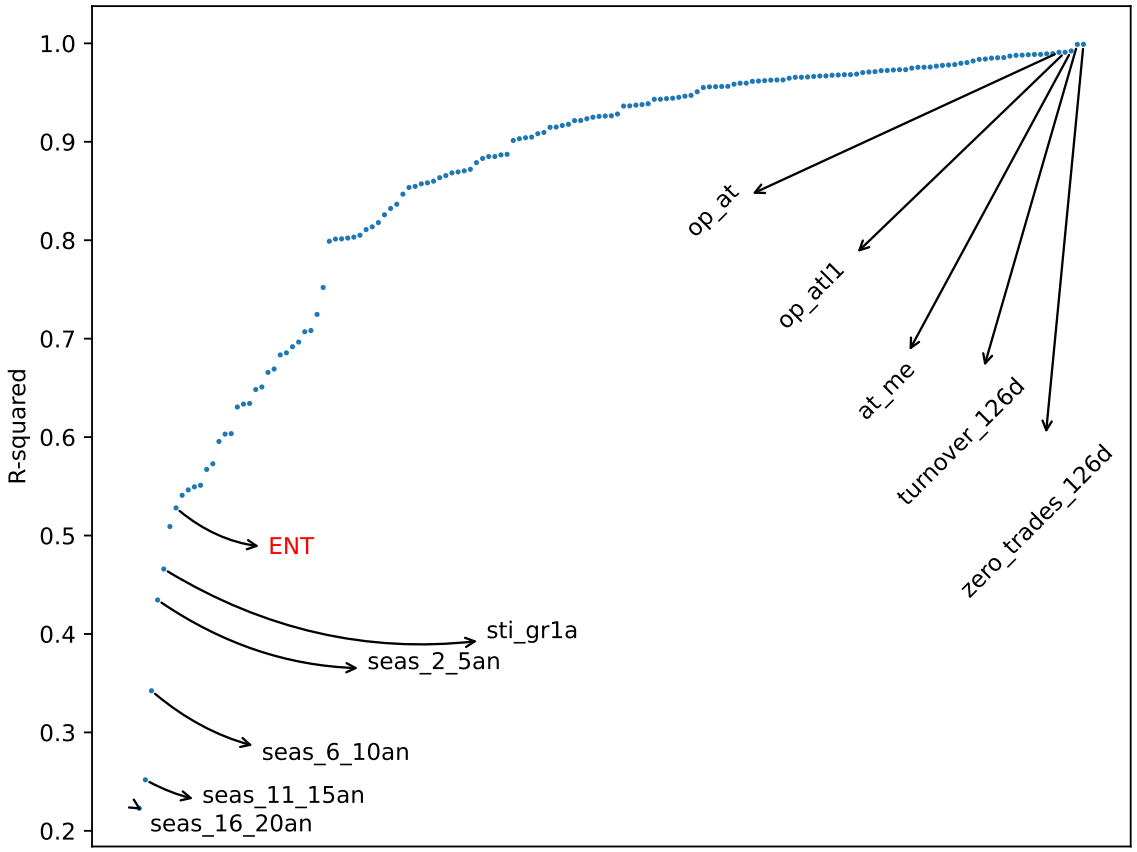


Figure 7: In-sample R^2, ordered from lowest to highest, of regressing each 153 value-weighted factors in Jensen et al. (2022) plus F_{ENT} and F_{EPU} on all other factors using the OLS model. Top and bottom five factors based on this spanning R-squared measure are annotated. Five factors with the highest R-squared are zero_trades_126d (number of zero trades with turnover as tiebreaker, 6 months), turnover_126d (share turnover), at_me (assets to market), op_atl1 (Ball et al. 2016 operating profit scaled by lagged assets), and op_at (Ball et al. 2016 operating profit to assets). Five factors with the lowest R-squared are seas_2_5an (year 2-5 lagged returns, annual), seas_6_10an (year 6-10 lagged returns, annual), sti_gr1a (change in short-term investments), seas_11_15an (year 11-15 lagged returns, annual), seas_16_20an (year 16-20 lagged returns, annual). The entropy mimicking portfolio, F_{ENT}, ranks as the seventh lowest R-squared among all 155 factors.

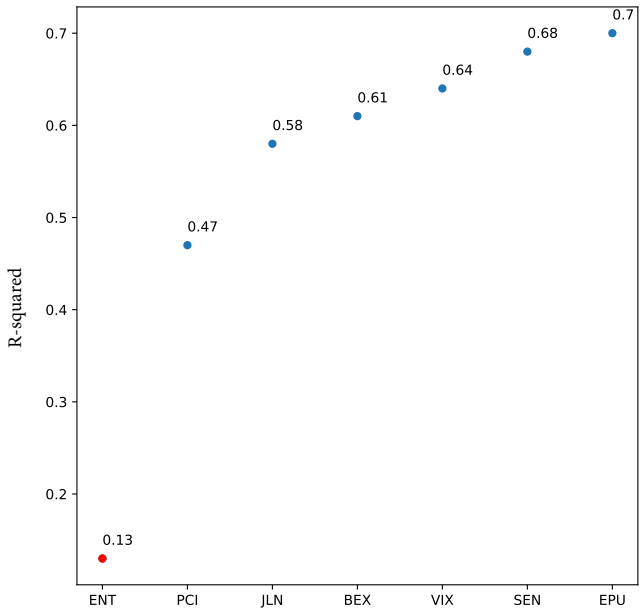


Figure 8: In-sample R^2, ordered from lowest to highest, of regressing each variable on all other variables using ordinary least squares. Included measures of macroeconomic uncertainty, ranked from lowest to highest, are entropy (*ENT*, marked in red), the Azzimonti (2018) Partisan Conflict Index (*PCI*), the Jurado et al. (2015) (*JLN*) uncertainty measure, the Bekaert et al. (2022) (*BEX*) uncertainty measure, the Chicago Board Options Exchange's CBOE Volatility Index (*VIX*), the Shapiro et al. (2022) news sentiment index (*SEN*), and the Baker et al. (2016) Economic Policy Uncertainty (*EPU*).

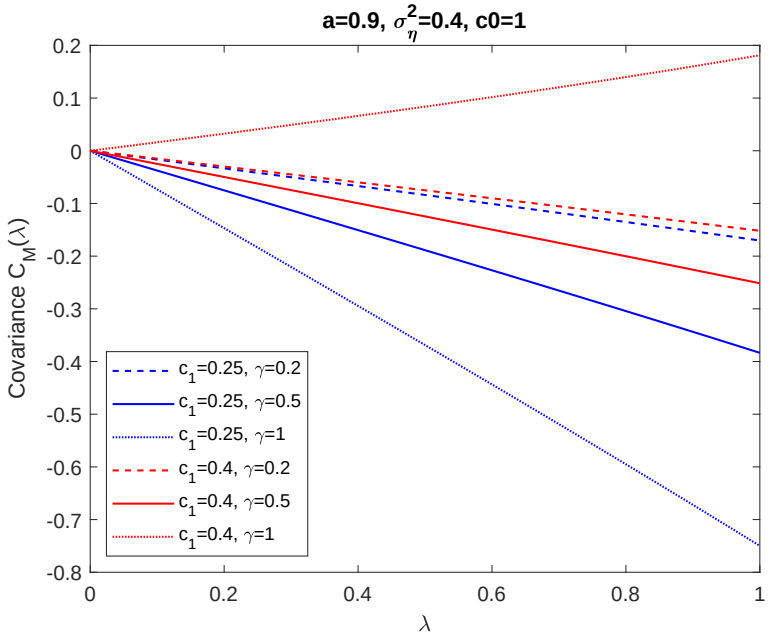


Figure 9: Plot of $C_M(\lambda)$ versus λ for several choices of γ and c_1, with $a = 0.9$, $\sigma_\eta^2 = 0.4$, and $c_0 = 1$.

Table 1: This table presents summary statistics of variables at the monthly frequency used in this study, including entropy measure (*ENT*), the first part of the decomposition of entropy (*ENT_NEWS*), the second part of the decomposition of entropy (*ENT_MODEL*), economic policy uncertainty (*EPU*), sentiment (*SEN*), squared implied volatility (*VIX2*), market yield on U.S. treasury securities at 10-year constant maturity (*DGS10*), market yield on U.S. treasury securities at 10-year constant maturity minus 2-year constant maturity (*DGS10-2*), the dividend yield (*DY*), the difference between actual consumption and the consumption level predicted by wealth and income (*CAY*), the inverse of the cyclically adjusted price-to-earnings ratio (*1/CAPE*), the market return of the previous month (*Return1*), the cumulative market return of the previous 12 months excluding the most recent month (*Return12*), the cumulative returns of the Fama-French five factors and momentum factor over the previous 60 months (*Return60*, *SMB60*, *HML60*, *RMW60*, *CMA60*, and *UMD60*).

	mean	std	min	25%	50%	75%	max
ENT	−0.0073	0.0613	−0.1404	−0.0499	−0.0135	0.0322	0.1881
ENT_NEWS	0.1542	0.0795	−0.0310	0.0965	0.1384	0.2023	0.4708
ENT_MODEL	−0.1615	0.0413	−0.3006	−0.1830	−0.1563	−0.1343	−0.0733
EPU	120.1456	45.0389	57.2026	88.2728	109.7164	144.1681	350.4598
SEN	0.0225	0.1964	−0.6361	−0.1022	0.0377	0.1775	0.4019
BEX	0.4059	0.2150	0.0878	0.2820	0.3624	0.4875	1.6233
JLN	0.6704	0.1245	0.5299	0.5924	0.6310	0.6920	1.2166
PCI	116.1411	38.9802	34.7400	84.6175	107.6350	143.0900	271.2900
VIX2	492.9679	474.4772	102.5248	209.2536	375.2209	600.2726	3927.3970
DGS10	3.4231	1.4588	0.5500	2.2100	3.3450	4.5675	6.6800
DGS10-2	1.1074	0.9207	−0.7000	0.2325	1.0850	1.9075	2.8400
DY	1.8094	0.3790	1.1100	1.6100	1.8200	2.0000	3.6000
CAY	−0.0074	0.0204	−0.0510	−0.0194	−0.0042	0.0079	0.0227
1/CAPE	0.0195	0.0252	−0.0418	0.0016	0.0266	0.0387	0.0727
Return1	0.6438	4.6863	−17.2300	−2.0200	1.2300	3.4800	13.6500
Return12	8.0261	17.0219	−42.8302	0.0748	9.8582	18.0511	59.6603
Return60	51.4848	52.2836	−36.2647	1.3241	55.0996	95.4152	178.0090
SMB60	9.4205	25.0681	−50.7400	−7.8025	7.7350	20.6425	78.2400
HML60	6.6328	27.9585	−53.4300	−12.5600	2.2200	26.3650	91.0700
RMW60	19.8381	15.8996	−25.4300	7.1425	20.2550	30.6275	80.4300
CMA60	12.8265	22.5871	−19.2300	−2.5200	7.1000	21.6575	76.0000
UMD60	24.4751	34.9969	−49.7600	2.1925	18.9050	48.9200	108.9900

Table 2: Illustrative examples of high entropy articles. The table shows articles from three distinct months (first column) when *ENT* were high: January 2007, November 2010, and October 2014. Within each month, articles are ranked from the highest to lowest based on their raw entropy scores (last column). The displayed headlines (second column) and first sentences (third column) are from selected articles, whose entropy scores are within the top 1% for their respective month.

Month	Headline	Body (first sentence)	Score
2007-01	Freddie sells more than half $7 bln notes overseas	Freddie Mac says it sold 58% of its $4 billion 2-year notes to overseas investors, with central bankers taking 46% of the offering.	6.200
2007-01	US House sends business issues to Senate for action	Raising the minimum wage, changing Medicare drug purchasing and other initiatives important to the business community are landing on the U.S. Senate's doorstep, having won rapid passage since Jan. 9 in the House of Representatives.	6.016
2007-01	UBS sees narrow GSE debt spreads through 2007	Agency and mortgage securities yield spreads will stay narrow through most of this year, UBS strategists Laurie Goodman and Ivan Hrazdira said in an investor conference call.	5.936
2007-01	Fed's Poole says lag in GSE reform leaves crisis risk	Lack of progress in reining in mortgage lenders Fannie Mae and Freddie Mac leaves the economy at risk of possible financial crisis, St. Louis Federal Reserve Bank President William Poole said on Wednesday.	5.754
2010-11	U.S. as Currency Manipulator? It's a Bit Rich	Breakingviews U.S. Editor Rob Cox says the Fed's QE policy fits its mandate, despite China's criticism over what it called an "indirect currency manipulation."	6.384
2010-11	Ireland Makes Concessions to Restore Confidence	Following S&P's downgrade of Ireland's credit rating to single A, the Ireland government said it will reduce current spending, cut minimum wage and maintain its debated corporate tax at 12.5 percent.	6.266
2010-11	QE2 Critics Put Easing on Agenda at G20 Meeting in Seoul	Emerging market economies critical of the Fed's asset-purchase program have said quantitative easing threatens to flood their economies with excess liquidity.	6.191
2010-11	Euro Slumps on Irish Debt Fears	The euro falls to a 5-week low on growing concerns Ireland will be forced to default on its debts.	5.912
2014-10	Wall St declines after Fed ends bond-buying program	U.S. stocks fell on Wednesday, adding to their earlier declines after the Federal Reserve ended its monthly bond purchase program, as had been expected.	4.524
2014-10	Futures lower, investors look ahead to GDP data	U.S. stock index futures were slightly lower on Thursday as investors looked ahead to a report on economic growth and continued to digest recent comments from the Federal Reserve.	4.187
2014-10	Wall St flat after GDP, but Visa lifts Dow	U.S. stocks were mostly flat on Thursday, as a strong read on third-quarter economic growth raised new questions about monetary policy, though strong results at Visa single-handedly put the Dow in positive territory.	4.116

Table 3: In-sample prediction of h-month ahead cumulative market returns for $h \in \{12, 9, 6, 3\}$ months using equation (6) run using monthly data. Returns are measured in percent. *Return60*, *SMB60*, ..., *UMD60* convert factor returns to state variables as explained in Section 2. The pre-COVID period ends in 2019, i.e., the year-ahead returns on the left-hand side of the regression in equation (6) do not extend past 2019. The coefficient estimates have been normalized by the standard deviation of the right-hand side variables. Robust t-statistics are in parentheses and are based on Newey and West (1987) standard errors with four lags.

	Full sample			Pre-COVID		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 12-Month Ahead Cumulative Return						
ENT	−2.387***		−2.276**	−2.703***		−2.699**
	(−2.579)		(−2.408)	(−2.583)		(−2.575)
EPU		2.973***	2.730**		0.328	0.113
		(2.679)	(2.468)		(0.296)	(0.108)
SEN	2.115	3.464**	2.961*	4.411***	5.043***	4.429**
	(1.312)	(2.090)	(1.784)	(2.629)	(2.942)	(2.572)
Panel B: 9-Month Ahead Cumulative Return						
ENT	−2.284***		−2.242**	−2.155**		−2.178**
	(−2.624)		(−2.526)	(−2.258)		(−2.296)
EPU		1.295	1.057		−0.402	−0.574
		(1.218)	(0.997)		(−0.382)	(−0.549)
SEN	1.398**	2.821**	2.112**	2.906**	3.645**	2.926**
	(1.927)	(2.271)	(2.024)	(2.200)	(2.329)	(2.075)
Panel C: 6-Month Ahead Cumulative Return						
ENT	−1.450*		−1.399*	−1.735*		−1.720*
	(−1.768)		(−1.663)	(−1.784)		(−1.749)
EPU		1.426	1.277		0.514	0.377
		(1.258)	(1.106)		(0.457)	(0.323)
SEN	4.461***	5.166***	4.857***	4.446***	4.898***	4.506***
	(2.828)	(3.205)	(3.020)	(2.682)	(2.932)	(2.670)
Panel D: 3-Month Ahead Cumulative Return						
ENT	−1.221**		−1.197**	−1.193*		−1.170
	(−2.031)		(−1.947)	(−1.659)		(−1.601)
EPU		0.733	0.607		0.666	0.571
		(0.882)	(0.727)		(0.820)	(0.677)
SEN	3.292***	3.745***	3.480***	3.469***	3.826***	3.560***
	(3.138)	(3.559)	(3.220)	(3.132)	(3.581)	(3.184)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: This table shows out-of-sample R-squareds, calculated using (8), of 12-month ahead market return forecasts using monthly rolling estimates of (7) in h-month windows. The different hs correspond to the columns of the table. The data series are explained in Section 2. Note that the raw consumption-wealth ratio (CAY) (Lettau and Ludvigson, 2001) ended in September 2019. The methodology, including how the standard errors of the out-of-sample R-squareds are calculated, is explained further in Section 4.2. Stars indicate p-value thresholds; p-values measure the fraction of simulated out-of-sample R-squareds which are larger than the observed ones.

Out-of-Sample R-Squareds

	Prediction Window in Months				
	12	15	18	21	24
ENT	-0.019^{*}	0.024^{**}	0.047^{**}	0.051^{**}	0.046^{**}
ENT_NEWS	0.069^{***}	0.087^{***}	0.096^{***}	0.091^{***}	0.060^{**}
ENT_MODEL	-0.074^{**}	-0.061^{*}	-0.059	-0.057	-0.084
EPU	-0.272	-0.322	-0.414	-0.417	-0.478
SEN	-0.595	-0.642	-0.678	-0.668	-0.607
VIX2	-1.320	-1.626	-1.813	-2.130	-2.422
DGS10	-2.340	-1.527	-0.823	-0.555	-0.450
DGS10-2	-3.144	-2.162	-2.213	-2.482	-2.022
DY	-2.001	-1.982	-1.722	-1.580	-1.767
CAY	-0.395	-0.377	-0.344	-0.276	-0.211
1/CAPE	-1.597^{**}	-1.964	-2.227	-2.269	-2.170
Return1	-0.059	-0.081	-0.075	-0.050	-0.060
Return12	-1.805	-1.691	-1.157	-0.855	-0.749
Return60	0.038^{***}	0.001^{***}	0.075^{***}	0.087^{***}	0.027^{***}
SMB60	-0.427^{**}	-0.532	-0.568	-0.573	-0.584
HML60	-0.799^{*}	-0.797	-0.826	-0.935	-1.082
RMW60	-0.170^{***}	-0.124^{***}	0.020^{***}	0.054^{***}	0.049^{***}
CMA60	-0.913	-0.977	-1.023	-1.143	-1.232
UMD60	-0.732	-0.705	-0.546	-0.481	-0.460
BEX	-1.206	-1.326	-1.715	-2.034	-2.237
JLN	-2.037	-1.999	-1.887	-1.754	-1.493
PCI	-0.133	-0.125	-0.158	-0.163	-0.216

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: This table presents summary statistics of variables at the daily frequency used in this study, including Fama-French five factors market return minus risk-free rate (MKT), small minus big (SMB), high minus low (HML), robust minus weak (RMW), conservative minus aggressive (CMA), and momentum factor (UMD). F_{EPU}, F_{ENT} , F_{ENT_NEWS} , and F_{ENT_MODEL} represent factor-mimicking portfolio for economic policy uncertainty, entropy, news updates, and model updates. We further subscript these labels to indicate the base assets used to construct the portfolios. For example, "Momentum" uses the 25 Fama-French portfolios sorted on size and momentum as base assets, "BM" uses the 25 portfolios sorted on size and book-to-market, and so on.

	mean	std	min	25%	50%	75%	max
MKT	0.0294	1.2682	−12.0000	−0.5100	0.0700	0.6300	11.3500
SMB	0.0090	0.6425	−4.5500	−0.3500	0.0200	0.3800	5.7100
HML	0.0062	0.7828	−5.0000	−0.3400	−0.0100	0.3200	6.7400
RMW	0.0173	0.5487	−3.0100	−0.2500	0.0100	0.2800	4.5200
CMA	0.0136	0.4606	−5.8700	−0.2100	0.0000	0.2200	2.5300
UMD	0.0171	1.0637	−14.3700	−0.4200	0.0700	0.5200	7.1200
$F_{EPU_Momentum}$	−0.0954	2.5804	−22.3755	−1.4273	−0.1415	1.2136	20.5997
$F_{ENT_Momentum}$	−0.0001	0.0035	−0.0246	−0.0018	−0.0000	0.0016	0.0367
$F_{ENT_NEWS_Momentum}$	−0.0001	0.0029	−0.0218	−0.0015	−0.0001	0.0014	0.0325
$F_{ENT_MODEL_Momentum}$	−0.0000	0.0018	−0.0129	−0.0009	0.0000	0.0010	0.0102
F_{EPU_BM}	−0.0547	2.0073	−13.2736	−1.0804	−0.0717	0.9715	16.7301
F_{ENT_BM}	−0.0001	0.0027	−0.0183	−0.0016	−0.0001	0.0014	0.0218
$F_{ENT_NEWS_BM}$	−0.0001	0.0028	−0.0164	−0.0016	−0.0001	0.0015	0.0187
$F_{ENT_MODEL_BM}$	−0.0000	0.0017	−0.0120	−0.0010	0.0000	0.0009	0.0100
F_{EPU_INV}	−0.0207	2.2522	−18.3677	−1.2390	−0.0525	1.1925	18.1130
F_{ENT_INV}	−0.0001	0.0029	−0.0287	−0.0017	−0.0002	0.0014	0.0199
$F_{ENT_NEWS_INV}$	−0.0001	0.0034	−0.0290	−0.0020	−0.0002	0.0017	0.0276
$F_{ENT_MODEL_INV}$	0.0000	0.0019	−0.0114	−0.0011	−0.0000	0.0011	0.0093
F_{EPU_OP}	−0.0039	1.9306	−15.5703	−0.9966	−0.0204	0.9854	14.9994
F_{ENT_OP}	−0.0001	0.0032	−0.0236	−0.0018	−0.0001	0.0016	0.0181
$F_{ENT_NEWS_OP}$	−0.0001	0.0032	−0.0232	−0.0018	−0.0001	0.0016	0.0207
$F_{ENT_MODEL_OP}$	−0.0000	0.0015	−0.0100	−0.0009	−0.0000	0.0008	0.0087

Table 6: Fama-MacBeth regressions use *MKT* (market minus risk-free rate), *SMB* (small minus big), *HML* (high minus low), *RMW* (robust minus weak), *CMA* (conservative minus aggressive), *UMD* (momentum factor), all of which are obtained from French's website, along with factor mimicking portfolios with base assets corresponding to the test assets. The data for the Fama-MacBeth analysis are daily. The first four columns represent different base assets: size–momentum, size–book/market, size–investment, and size–profitability. The last column shows the sign of the corresponding coefficient in column (3) of Table A1 for forecasting 12-month ahead market returns if that coefficient is significant at 5% level. Risk premia coefficients are in percent per month. The risk premia coefficients are scaled by the standard deviations of betas in the first stage of Fama-MacBeth. Generalized method of moments t-statistics are in parentheses.

Fama-MacBeth Factor Risk Premia (scaled coefficients)

	Base Assets the Same as Test Assets: Size -				
	Momentum	Book/Market	Investment	Profitability	Mkt-RF
	(1)	(2)	(3)	(4)	(5)
MKT	0.048*	0.083**	0.064**	0.056**	
	(1.681)	(2.189)	(2.023)	(1.997)	
SMB	0.100	0.067	0.008	0.079	−
	(1.371)	(0.907)	(0.108)	(1.051)	
HML	0.078*	0.048	0.136***	0.041	−
	(1.754)	(0.601)	(2.879)	(0.905)	
RMW	−0.066	0.062	−0.018	0.119*	
	(−1.050)	(1.283)	(−0.337)	(1.953)	
CMA	0.078**	−0.007	0.065	0.018	+
	(2.142)	(−0.166)	(1.498)	(0.438)	
UMD	0.149	0.083***	0.064**	0.029	
	(1.342)	(2.896)	(2.346)	(0.928)	
F_{EPU}	−0.063*	−0.083**	−0.008	−0.013	+
	(−1.746)	(−2.342)	(−0.326)	(−0.432)	
F_{ENT}	−0.060*	−0.066**	−0.091***	−0.082**	−
	(−1.757)	(−2.398)	(−2.669)	(−2.175)	
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 7: Fama-MacBeth regressions use *MKT* (market minus risk-free rate), *SMB* (small minus big), *HML* (high minus low), *RMW* (robust minus weak), *CMA* (conservative minus aggressive), *UMD* (momentum factor), all of which are obtained from French's website, along with factor mimicking portfolios with base assets corresponding to the test assets. The data for the Fama-MacBeth analysis are daily. The first four columns represent different base assets: size–momentum, size–book/market, size–investment, and size–profitability. The last column shows the sign of the corresponding coefficient in Table A5 for forecasting 12-month ahead market returns if that coefficient is significant at 5% level. Risk premia coefficients are in percent per month. The risk premia coefficients are scaled by the standard deviations of betas in the first stage of Fama-MacBeth. Generalized method of moments t-statistics are in parentheses.

Fama-MacBeth Factor Risk Premia (scaled coefficients)

	Base Assets the Same as Test Assets: Size -				
	Momentum	Book/Market	Investment	Profitability	Mkt-RF
	(1)	(2)	(3)	(4)	(5)
MKT	0.050*	0.077**	0.053**	0.051**	
	(1.720)	(2.157)	(2.022)	(2.000)	
SMB	0.093	0.071	0.010	0.082	–
	(1.287)	(0.966)	(0.131)	(1.095)	
HML	0.081*	0.054	0.133***	0.035	
	(1.780)	(0.684)	(2.766)	(0.758)	
RMW	−0.065	0.067	−0.014	0.122**	
	(−1.038)	(1.316)	(−0.265)	(1.978)	
CMA	0.071**	−0.013	0.063	0.009	+
	(2.055)	(−0.326)	(1.457)	(0.216)	
UMD	0.148	0.086***	0.066**	0.030	
	(1.335)	(3.032)	(2.332)	(0.937)	
F_{EPU}	−0.063*	−0.083**	−0.013	−0.009	+
	(−1.764)	(−2.297)	(−0.484)	(−0.311)	
F_{ENT_NEWS}	−0.044	−0.060**	−0.055**	−0.059*	
	(−1.426)	(−2.237)	(−2.000)	(−1.771)	

Note: *p<0.1; **p<0.05; ***p<0.01

Table 8: In-sample predictions of 12-month ahead economic policy uncertainty (*EPU*) and fundamental variables including unemployment rate (*UNRATE*), year-over-year change of industrial production (*INDPRO_YOY*), year-over-year change of the consumer price index (*CPI_YOY*), interest rates (*DGS10* and *DGS10-2*), the Chicago Board Options Exchange's CBOE Volatility Index (*VIX*), and the S&P500 real earnings per share (*EPS*). The columns correspond to different dependent variables in (16). Each specification includes the lagged value of the dependent variable in question and the lagged value of all other macro variables as controls. The lagged value of one entropy measure (*ENT*, *ENT_NEWS*, or *ENT_MODEL*) is also used as a control, corresponding to each row. The coefficient estimates have been normalized by the standard deviation of entropy, *ENT_NEWS*, or *ENT_MODEL* respectively. This table summarizes results in Online Appendix Tables A12, A13, and A14. Robust t-statistics are in parentheses and are based on Newey and West (1987) standard errors with four lags.

Macro Forecasting Using Entropy and Other Macro Controls

	12-Month Ahead Macro Variables							
	EPU	UNRATE	INDPRO_YOY	CPI_YOY	DGS10	DGS10-2	VIX	EPS
ENT	12.249	0.700***	−1.503***	−0.361**	−0.274***	0.084	1.579***	−4.971**
	(1.539)	(3.720)	(−3.534)	(−2.087)	(−3.933)	(0.092)	(3.238)	(−2.169)
ENT_NEWS	7.777	0.627***	−1.344***	−0.358*	−0.182**	0.128*	2.105***	−5.373**
	(1.236)	(4.188)	(−3.033)	(−1.875)	(−2.529)	(1.710)	(3.719)	(−2.039)
ENT_MODEL	4.519	−0.069	0.142	0.096	−0.087	−0.103	−1.388**	2.128
	(0.612)	(−0.444)	(0.360)	(0.651)	(−1.505)	(−1.602)	(−1.985)	(0.889)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 9: In-sample predictions of 12-month ahead market returns, measured in percent, using the specification in (17). Uncertainty measures include Bekaert et al. (2022) (*BEX*), Jurado et al. (2015) (*JLN*), the Azzimonti (2018) Partisan Conflict Index (*PCI*), and squared implied volatility (*VIX2*). The other control variables are explained in Section 2. Full results are shown in Online Appendix Table A20. The coefficient estimates have been normalized by the standard deviation of the right-hand side variables. Unscaled full results are shown in Online Appendix Table A21. Robust t-statistics are in parentheses and are based on Newey and West (1987) standard errors with four lags.

In-Sample Return Forecasting with Uncertainty Controls

	12-Month Ahead Cumulative Return					
	(1)	(2)	(3)	(4)	(5)	(6)
ENT	−2.355**	−2.218**			−2.203**	−2.060**
	(−2.522)	(−2.339)			(−2.262)	(−2.101)
EPU			4.632***	3.013**	4.435***	2.964**
			(3.959)	(2.188)	(3.875)	(2.197)
SEN			2.391	3.249**	1.879	2.917*
			(1.371)	(2.021)	(1.080)	(1.819)
Other Uncertainty	No	Yes	No	Yes	No	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.658	0.678	0.664	0.678	0.672	0.685

Note: *p<0.1; **p<0.05; ***p<0.01

Table 10: In-sample predictions of 12-month ahead cumulative market returns and fundamental variables—including unemployment rate (*UNRATE*), year-over-year change of industrial production (*INDPRO_YOY*), year-over-year change of the consumer price index (*CPI_YOY*), interest rates (*DGS10* and *DGS10-2*, the 2s-10s curve), the Chicago Board Options Exchange's CBOE Volatility Index (*VIX*) combined with *NVIX* prior to 1990, and the S&P500 real earnings per share (*EPS*)—using extended entropy $\widehat{ENT}$. The first column replicates Table 3 Panel A Column (3) without *EPU* and *SEN* as controls. The remaining columns correspond to different dependent variables in (16). Each specification includes the lagged value of the dependent variable in question and the lagged value of all other macro variables as controls, again without *EPU* and *SEN*. The coefficient estimates have been normalized by the standard deviation of $\widehat{ENT}$. Robust t-statistics are in parentheses and are based on Newey and West (1987) standard errors with four lags. The analysis starts in January 1970, as discussed in Section 7.4.

In-Sample Return and Macro Forecasting Using Extended Entropy and Other Controls

	12-Month Ahead Macro Variables							
	Return	UNRATE	INDPRO_YOY	CPI_YOY	DGS10	DGS10-2	VIX	EPS
$\widehat{ENT}$	-3.754^{***}	0.271^{***}	-1.268^{***}	0.108	-0.002	0.055	0.875^{**}	-1.738
	(-2.842)	(2.819)	(-3.299)	(0.560)	(-0.019)	(0.674)	(2.043)	(-0.799)

Note: *p<0.1; **p<0.05; ***p<0.01