

Global News Networks and Return Predictability ^{*}

Gustavo Freire[†], Ali Moin[‡], Alberto Quaini[§], Amar Soebhag[¶]

September 2025

Abstract

We employ a novel, high-dimensional global news corpus (GDELTA) to examine the effects of news-based sentiment on international equity market returns. Leveraging over 520 million articles from 14 developed countries, spanning 260 major themes and classified by source and target country, we construct granular sentiment linkages that capture how one country's media reports on another. Using machine learning, we construct out-of-sample market-timing strategies based on local and global sentiment measures and document significant economic gains. Further, global sentiment measures contain distinct, incremental information compared to local measures, indicating market integration via sensitivity to shared news information and sentiment.

Keywords: News Sentiment, Equity Markets, Return Predictability, Big Data, Machine Learning

JEL classification: G12, G14, G15, C55

^{*}We thank Guido Baltussen, Jefferson Duarte, Pasquale Della Corte, Oliver Hellum, Elise Gourier, Jens Jackwerth, Robin Lumsdaine, Liang Peng, Hashem Pesaran, Mirco Rubin, Olivier Scaillet, Paul Schneider, Simon Trimborn, Bart van Vliet, Tjeerd de Vries (discussant), Christina Dan Wang, Marvin N. Wright, and seminar participants at QFFE 2025 (Marseille), IAAE 2025 (Turin), EFA 2025 (Paris), Erasmus University Rotterdam and Robeco for helpful comments. The views expressed in this paper are solely those of the authors and are not necessarily shared by their institutional affiliations.

[†]freire@ese.eur.nl. Department of Econometrics, Erasmus University Rotterdam, Tinbergen Institute and Erasmus Research Institute of Management.

[‡]moin@ese.eur.nl. Department of Econometrics, Erasmus University Rotterdam.

[§]quaini@ese.eur.nl. Department of Econometrics, Erasmus University Rotterdam.

[¶]soebhag@ese.eur.nl (corresponding author). Robeco Quantitative Investing and Department of Finance, Erasmus University Rotterdam.

1 Introduction

Measuring sentiment in the economy is an important yet complex task. Both public and private institutions invest substantial resources in developing sentiment measures, which are closely monitored by market participants and serve as key inputs for financial decision-making.¹ In fact, investor sentiment has been shown to have strong implications for asset pricing. Theoretically, sentiment leads to temporal variations in asset prices resulting from the trading behavior of noise or liquidity traders (De Long, Shleifer, Summers, and Waldmann, 1990; Campbell, Grossman, and Wang, 1993). Empirically, numerous studies have shown that investor sentiment predicts both stock and market returns (for a recent literature review, see Zhou, 2018).

A recent trend in the literature is the use of textual analysis to extract news-based sentiment measures (Shapiro, Sudhof, and Wilson, 2022). However, most of these studies consider only English-based news outlets, focus either on aggregate media sentiment measures or a handful of news topics, and do not explicitly link the sentiment of an entity to a target. These limitations constrain prior research by overlooking a sizable portion of non-English news and masking the potentially heterogeneous impact of granular sentiment in predicting asset returns. They also prevent a systematic analysis of the sources and dynamics of media sentiment (e.g., sentiment about what and expressed by whom).

We offer a novel approach to measure sentiment that addresses these shortcomings and allows us to study the effects of sentiment networks in international financial markets. To illustrate, our approach allows to construct a sentiment measure for

¹Examples include the University of Michigan’s Survey of Consumers and Consumer Sentiment Index, the AAI Investor Sentiment Survey, as well as sentiment analysis tools provided by Ravenpack and Bloomberg.

the U.S. (the target) about the “stock market” (a topic) based on news from Japan (the source). To our knowledge, our work is the first to model such linkages between countries and construct a high-dimensional set of theme-specific sentiment measures across and within countries.

To do so, we use the most extensive international text corpora of news studied in the finance literature to date: *the Global Database of Events, Language, and Tone (GDELT)*, which collects global news from over 150,000 news outlets offering both English and translingual articles. GDELT has two unique features: (i) articles are classified into themes, and (ii) the source entity and target locations are provided. We collect information from 520 million articles across 14 developed countries from 2015 till 2023, and leverage upon these unique features of the GDELT data to construct this high-dimensional set of sentiment measures.

With this data, we examine the effect of local and global sentiment on asset prices using our high-dimensional set of sentiment series. Specifically, we test whether local news (covered by local media about the local market) carries predictive power for future local market returns. Subsequently, we focus on a high-dimensional global sentiment network. In this set-up, we allow both local and foreign sentiment, irrespective of the source and target entity, to predict market returns. This allows us to further examine and decompose whether specific countries are driving the return predictability, or whether there exist clusters of countries that influence each other.

To this end, we form out-of-sample return forecasts using a Random Forest model trained on a dynamically expanding window of historical sentiment data. We then evaluate the economic significance of the forecasting model by converting these forecasts into an out-of-sample market-timing strategy. By comparing performance met-

rics commonly used by the market-timing literature (Kelly, Malamud, and Zhou, 2024), we demonstrate that news sentiment linkages contain significant market timing ability in global equity markets.

We find that the out-of-sample market-timing strategy based on local sentiment measures significantly outperforms the buy-and-hold market strategy for our 14-country panel. Specifically, we find that strategies based on local sentiment achieve higher Sharpe ratios than buy-and-hold in 13 of 14 markets, and the average net-of-costs Sharpe ratio is 82.5% higher than the average buy-and-hold strategy. For example, in the U.S., the local-news strategy attains a Sharpe ratio of 1.34 (vs. 0.62 for the buy-and-hold strategy), a net Sharpe of 1.11, and a net CAPM alpha of 15.98% per annum. These returns are not driven by market exposure—betas are low (between -0.20 and 0.41)—and left-tail risk is reduced (maximum drawdowns fall by roughly one-third).

In addition, incorporating global sentiment linkages substantially enhances performance in terms of market timing. Strategies based on global sentiment linkages yield economically meaningful and statistically robust Sharpe ratios, significantly outperforming buy-and-hold benchmarks, even after transaction costs. Moreover, spanning regressions demonstrate that global sentiment measures provides significant incremental and distinct information beyond local sentiment measures alone. Specifically, we observe positive and economically sizable spanning alphas, ranging between approximately 10% and 30% per annum, in virtually every market, indicating enhanced risk-adjusted returns from integrating global signals with local strategies.

Further, we visualize these cross-country predictive relationships in a global news-based sentiment network (Figure 1). This network reveals a clear structure, promi-

nently featuring the United States as the central informational hub, exerting predictive influence globally, and China as a crucial secondary node. European countries such as the United Kingdom and Germany emerge as significant regional hubs, illustrating strong intra-European informational integration. Conversely, markets like Australia and Sweden exhibit relatively limited cross-country predictive influence, underscoring their peripheral role within the global sentiment network.

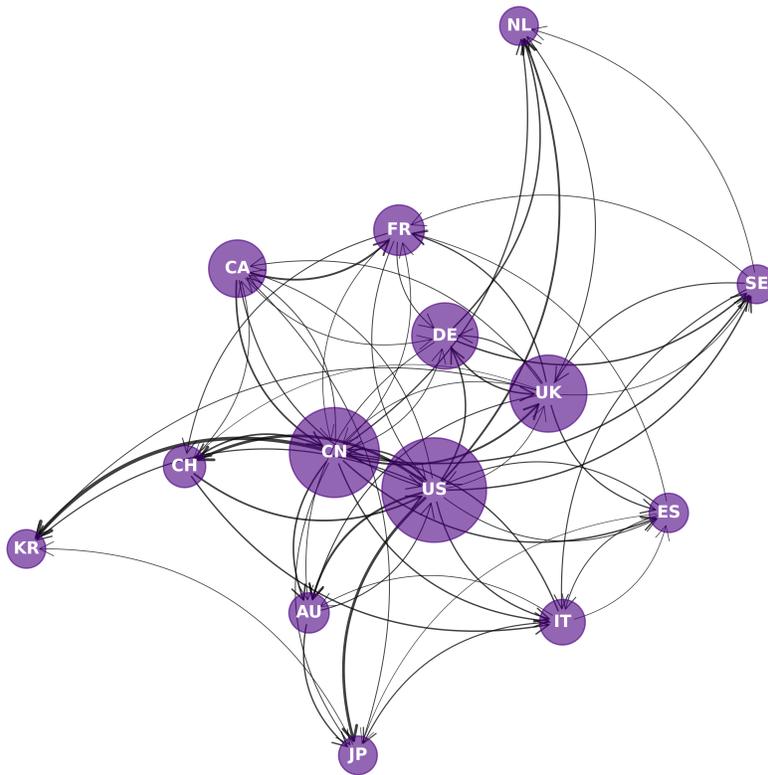
Overall, these results provide compelling evidence of informational market integration, demonstrating that international equity markets are sensitive to shared sentiment and news flows. The identified global sentiment network structure highlights that cross-border informational connections significantly drive equity return predictability worldwide.

Lastly, our results remain robust across a battery of robustness checks. Our findings hold when using various alternative sentiment lexicons, indicating that the choice of dictionary does not materially influence the predictive power of news sentiment. Moreover, controlling for past returns suggests that our findings are not simply driven by a short-term momentum effect, since adding lagged returns sometimes even weakens the predictive ability of local news sentiment. Finally, to address potential translation biases, we compare native-language Hedonometer ([Dodds, Harris, Kloumann, Bliss, and Danforth, 2011](#)) lexicons with their English-translated counterparts, finding that sentiment-based predictability often strengthens after translation.

Our study contributes to several strands of the literature. First, our findings relate to the literature that studies media sentiment and asset prices at the market level (see, e.g., [Tetlock, 2007](#); [Dougal, Engelberg, Garcia, and Parsons, 2012](#); [Garcia, 2013](#); [Manela and Moreira, 2017](#); [Heston and Sinha, 2017](#); [Calomiris and Mamaysky,](#)

Figure 1: Cross-country news-based predictability network

Directed network visualization of cross-country predictive influence. Nodes are countries; node size encodes each country's overall importance, computed by aggregating gain-based Random Forest variable importances (as both source and target in news themes) and nonlinearly transformed to enhance visual separation. Incoming arrows indicate the top five foreign contributions to forecasting a country's equity returns, with arrow thickness proportional to the strength of that influence. The layout is force-directed to spatially reflect connectivity structure. The out-of-sample forecasting period runs from February 2017 to December 2023. See Section 4.3 for details.



2019; Fraiberger, Lee, Puy, and Ranciere, 2021). For example, Fraiberger, Lee, Puy, and Ranciere (2021) show that international asset prices respond to foreign media coverage, indicating that sentiment crosses borders via global news channels. Our contribution to this literature is that we decompose media sentiment into granular sentiment measures that are topic and country-specific. More importantly, we model sentiment linkages, which is novel to the literature, and we are first to utilize GDELT within the field of Financial Economics.

Second, our study contributes to the recent literature on narrative-based asset pricing. This literature applies textual analysis to capture media exposure towards particular themes (see, e.g., Mai and Pukthuanthong, 2021; Bhargava, Lou, Ozik, Sadka, and Whitmore, 2022; Bybee, Kelly, Manela, and Xiu, 2023; Bybee, Kelly, and Su, 2023; Lee, Ozik, and Sadka, 2023; Adämmer and Schüssler, 2020). These studies demonstrate that narratives contain predictive information for future asset prices, but is typically limited to the United States. As in the previous studies, we classify news into hundreds of unique themes. As a contribution, we construct attention-weighted sentiment measures. As such, we do not only rely on the frequency of a topic, but also on the sentiment.² Our interpretable machine-learning method allows us to dynamically determine themes that drive stock market returns, whilst taking both sentiment and attention into account. Our findings indicate that theme importance fluctuates over time. For instance, health related themes become high in attention, negative in sentiment and key drivers for equity markets during the Covid-19 pandemic.

Third, we contribute to the literature on equity market integration. Numer-

²Cookson, Lu, Mullins, and Niessner (2024) highlights the importance of studying attention and sentiment jointly, and notes that prior studies focus on only one of the two.

ous studies have documented varying degrees of equity market integration (see, e.g., De Jong and De Roon, 2005; Karolyi and Stulz, 2003; Bekaert, Hodrick, and Zhang, 2009; Lewis, 2011; Carrieri, Chaieb, and Errunza, 2013; Hau, 2011). Our spanning tests reveal that global sentiment measures provide economically distinct and incremental market timing ability in virtually all countries in our sample, beyond information contained in local sentiment measures. In addition, the network representation underscores the global prominence of U.S. and existence of regional informational hubs, such as Europe. These findings strongly supports the existence of informational market integration through sensitivity to shared foreign news and sentiment, consistent with prior studies emphasizing the increasing prominence of common information in shaping international return dynamics (Pukthuanthong and Roll, 2009; Rapach, Strauss, and Zhou, 2013; Hellum, Pedersen, and Rønn-Nielsen, 2023; Choi, Jiang, and Zhang, 2025), and reinforcing the idea that equity markets operate as parts of an integrated informational network.

The paper is structured as follows. Section 2 describes the GDELTE database and the construction of our set of sentiment indices. Section 3 describes our machine learning methodology to construct sentiment-based return forecasts. Section 4 uses sentiment to predict the time series of market returns and decomposes this predictability. Section 5 considers a battery of robustness tests. Section 6 concludes.

2 Data

This section introduces our sample and provides descriptive statistics. Section 2.1 introduces the GDELTE database, and Section 2.2 describes the sample construction pro-

cedure. Section 2.3 details the methodology used to construct our high-dimensional set of sentiment measures. Finally, Section 2.4 briefly describes equity returns data.

2.1 The GDELT database

To develop our high-dimensional set of sentiment measures across themes and countries, we exploit an open-source, global news database: The Global Database of Events, Language, and Tone (GDELT), which was founded in 1994 and is supported by Google. GDELT provides two main datasets: the GDELT Event Database and the Global Knowledge Graph Database (GKG). We use the GKG to leverage on two important features of this database.

First, GKG collects global news from over 150,000 news outlets offering both English and translingual articles (covering over 100 different languages). This is important as it allows us to significantly broaden the coverage of global news in different languages, a feature that is often overlooked in prior literature. In Figure 2 we plot the proportion of English-based and translingual news for each month in our sample. We observe that translingual news outlets produce a significant portion (55% on average) of the total global news, which underlines the importance of including translingual news in sentiment measures.

Second, GKG labels each publicly available global news article to individuals, organizations, locations, emotions, and themes. For instance, GKG tracks the “Election” (“Theme”) in the U.S. (“Target”), which can be covered by U.S. media or other international outlets (“Source”). This feature is crucial as it allows us to construct sentiment measures between and within countries for a wide range of themes. Furthermore, it enables us to construct sentiment measures of one source country about

another target country for a specific theme, which we dub *sentiment linkages*. To our knowledge, modeling sentiment linkages is novel to the literature.

Although GDELT contains a wealth of global news information, it does not include the full text of news articles and instead provides historical links. Some of these links may be inaccessible for several reasons like broken URLs, paywalls, or changes in website structures. GDELT overcomes these challenges by offering a comprehensive set of sentiment metrics. This approach eliminates the need for individual researchers to deploy computationally intensive NLP techniques or large language models, making it easier to construct high-dimensional sentiment measures efficiently.

2.2 Sample construction

We collect GKG data from February 2015 through December 2023 for the following 14 countries: United States (US), China (CN), Germany (DE), Japan (JP), United Kingdom (UK), France (FR), Italy (IT), Canada (CA), Australia (AU), South Korea (KR), Spain (ES), Netherlands (NL), Switzerland (CH), and Sweden (SE). These countries collectively account for approximately 72% of the global GDP ([World Bank, 2023](#)) and provide us with a representative sample of major economies.

GKG classifies articles into themes based on the content of the article. Our study focuses on 260 major themes provided by GKG since its inception.³ These major themes exclude entertainment and lifestyle-related content (such as gossip, leisure, and sports), and focuses mostly on socioeconomic news events. Articles in GDELT may be tagged with multiple themes. Furthermore, GDELT updates its news database every 15 minutes and all news items are recorded in Coordinated Universal

³A brief description of each theme can be found at data.gdeltproject.org. This link has also been archived via the Wayback Machine on 2025-01-21 at web.archive.org.

Time (UTC). In total, our sample covers over 520 million web articles.

To provide insights in the composition of our sample, we show the article distribution by source country in Figure 3(A). We find that the U.S. plays a major role in the production of news articles, with a relative share of 49%, followed up by China (10.9%), U.K. (6.9%) and Germany (6.0%). We also show the article distribution by theme in Figure 3(B).⁴ There is no single dominating theme in our sample, stressing the diversity in themes that occur in global news. For instance, the “Leader” theme is tagged in 5.9% of all articles, followed by “General Health” (5.2%) and “Education” (5.1%). Articles with the “Stock Market” theme occur 1.2% of the time.

These illustrative numbers are over the full sample. However, themes are time-varying and dynamic. For example, some themes are short-lived, trending, and eventually fade away. In Figure 4, we show the top 20 themes per month in our sample, with the U.S. as the target country. We find that some topics structurally obtain high levels of attention, such as science (“Science”) and crime (“Soc_GeneralCrime”). Other topics fluctuate over time and are recurring, such as the general discussion of holidays (“Gen_Holiday”). In addition, some topics are short-lived and capture significant media attention, such as impeachment (“Impeachment”, heavily covered from September 2019 till February 2020), shortage (“Shortage”, January 2022), “Sanctions” (March 2022), and the COVID pandemic (“Health Pandemic”).

As mentioned in Section 2.1, GDEL T provides entity linkages within articles. This feature allows us to construct cross-country sentiment linkages. In Figure 5, we show the number of articles that each source country (s) publishes about each target country (k) as a fraction of the total number of news articles published by s .⁵

⁴For ease of exposition, we only show the top 30 themes over the full sample.

⁵Note that the rows do not sum to 1.00 since an article published by country s can have multiple

We observe three notable patterns. First, most of the produced news by a source country is about itself, i.e., national news. For example, 85% of the news articles published by the U.S. are targeted at the U.S. itself. Second, non-U.S. countries have a considerable share of news about the U.S., highlighting the importance of the U.S. in terms of attention. For example, 26% of the Japanese news in our sample mentions the United States. Third, countries that are geographically close tend to have stronger connections in terms of news targeting. For instance, 20% and 27% of the South Korean news is about China and Japan, respectively. Germany mentions France and the U.K. in 10% and 8% of all German articles, which is relatively higher compared to other countries for Germany. We explicitly take these features into consideration when we measure sentiment, as explained in the next section.

2.3 Sentiment measures

To measure sentiment for each article in our sample, we rely on measures provided by GDELT. Specifically, we use the SentiWordNet 3.0 (SWN) lexicon-based framework (Baccianella, Esuli, and Sebastiani, 2010). SWN is able to take the context of words in sentences and texts into account. Each “synset” (synonym set) of WordNet groups together words that are synonyms and share the same meaning in different contexts. Synsets are the fundamental building blocks of WordNet and serve as a way to organize the lexicon based on semantic relationships rather than just lexical ones. Each synset is associated with three sentiment scores (summing to one): positivity, negativity, and objectivity. We utilize both the positive and negative scores of an article. These scores are derived based on the context in which words appear, making

target countries.

SentiWordNet very effective for nuanced sentiment analysis across diverse themes and contexts.

A well-known alternative is the financial lexicon-based approach of [Loughran and McDonald \(2011\)](#) (LM). We choose not to use this approach because our dataset contains 260 themes, the majority of which contain non-financial texts, where the LM method would be less effective in accurately capturing sentiment. However, in [Section 5.1](#), we conduct a robustness test across different lexicons for the national news of the United States and obtain similar results.

We now describe the construction of our sentiment measures. Let t denote index days, $\lambda = 1, \dots, 260$ index themes, and p, q index source and target countries, respectively. For each such index, we define $S_{t,\lambda,p,q}$ as the average daily sentiment score, and $a_{t,\lambda,p,q}$ as the total number of articles regarding theme λ on day t from source country p about target country q .

To avoid cumbersome triple indices (p, q, λ) when describing sentiment linkages, we introduce a single index ℓ to denote each unique combination of $(p_\ell, q_\ell, \lambda_\ell)$. Formally, let $\mathcal{L} = \{(p, q, \lambda)\}$ be the set of all such triplets and let $\ell = (p_\ell, q_\ell, \lambda_\ell) \in \mathcal{L}$ refer to a specific link. For instance, rather than writing $S_{t,(p_\ell,q_\ell,\lambda_\ell)}$ for the average daily sentiment on day t from source country p about target country q , we write $S_{t,\ell}$.

We aggregate article-level sentiment scores into a daily measure, the attention-weighted sentiment (*AWS*) change, for each source-target-theme combination as follows.⁶ First, we compute the one-day change in sentiment for each linkage ℓ . Second, we compute the total daily news percentage for each theme λ within each source-target pair (p, q) . Lastly, we multiply the average sentiment change by their attention-based

⁶To avoid staleness and trends in levels, we consider changes in sentiment.

weights:

$$\text{AWS}_{t,\ell} = \frac{a_{t,\ell}}{\underbrace{\sum_{\lambda'=1}^{260} a_{t,\lambda',p\ell,q\ell}}_{\text{attention weight}}} \times \underbrace{\left(S_{t,\ell} - S_{t-1,\ell} \right)}_{\text{one-day change in sentiment}} .$$

Recent studies stress the role of narratives in asset pricing (e.g., [Bybee, Kelly, Manela, and Xiu, 2023](#)). Our approach takes both the direction and frequency of sentiment into account, thereby also accounting for time-varying narratives. Given that our sample has 14 countries and 260 themes, each with a positive and negative sentiment dimension, we obtain a high-dimensional set of 101,920 AWS sentiment linkages.

Figure 6 illustrates the net AWS measure (i.e., positive minus negative AWS sentiment) for the “Stock Market” (panel A) and “Health Pandemic” (panel B) themes for the United States. The upper plot shows that the sign of sentiment, i.e., the extent to which news are perceived as positive or negative about the stock market, varies over time. While the net AWS for the stock market was mostly positive until 2019, it turned significantly negative in 2020 due to the COVID-19 pandemic, eventually bouncing back to positive values towards the end of the sample in 2023. On the other hand, articles on pandemics received very little attention before 2020 due to the absence of major international outbreaks. In 2020, articles on the COVID-19 pandemic significantly increased and played a major negative narrative, reverting back over time as the pandemic subsided. This example illustrates how our AWS measure accounts for rising and trending themes in the news.

2.4 Market return data

Our daily market data consists of excess returns obtained via futures contract prices for 14 equity indexes from February 2015 through December 2023, obtained from Bloomberg. Table A1, in the appendix, details the sample universe, market closing times, timezones, index name, and the annualized Sharpe ratio of each instrument.

As previously noted, GDELT records news events in Coordinated Universal Time (UTC), while market return data are timestamped in local time. To prevent any forward-looking bias in our forecasting process, we convert each news event’s UTC timestamp to the local timezone of the corresponding market. If a news item occurs after the local market’s closing time, it is reassigned to the following trading day. This adjustment ensures that our analysis only incorporates information that would have been available to market participants at the relevant time.

3 Machine Learning Setup

Our decomposition of aggregate media sentiment by country, theme, and source-targets allows to study such time-variation across a wide-range of equity markets. Given the high-dimensionality of our set of sentiment predictors, we rely on machine learning techniques. This section describes and motivates the machine learning setup to predict international equity excess returns using our high-dimensional set of sentiment series.

3.1 The prediction problem and Random Forests

To frame our setting in a general prediction problem, assume that our set of sentiment series on day t can be summarized by a state vector x_t and that investors use this information on day t to predict equity returns. The general return prediction problem involves identifying a functional form $g(\cdot)$ of $E(r_{t+1})$ that maps the set x_t such that the maximum out-of-sample (OOS) predictive power relative to r_{t+1} is obtained. We describe this mapping between the excess market return and the set of sentiment measures as:

$$r_{t+1} = E(r_{t+1}) + \epsilon_{t+1} = g(x_t) + \epsilon_{t+1},$$

where r_{t+1} is the excess return of an equity market on day $t + 1$, and ϵ_{t+1} is the residual. We assume that the conditional expectation of r_{t+1} is determined by some function $g(\cdot)$ that takes a set of predictor variables x_t as input at time t .

There exist a variety of methods to estimate $g(\cdot)$ and produce out-of-sample predictions using this estimate. Traditionally, the simple linear regression model has been the mainstay of return predictability. Despite its popularity, linear regression is bound to fail in our setting, given the high-dimensionality of our set of predictors. Alternatively, in the presence of a large set of predictors, a popular strategy is to impose sparsity in the set of regressors via a penalty term.⁷ However, these methods do not inherently enable a model to capture complex nonlinear relationships in data.⁸

To this end, we estimate $g(\cdot)$ using a Random Forest (Breiman, 2001), which is

⁷For example, via Ridge regression, Lasso regression or Elastic Nets (Zou and Hastie, 2005).

⁸While it is possible to introduce non-linearity through feature engineering (such as polynomial transformations or interaction terms) regularization alone does not provide a mechanism for learning non-linear patterns.

an ensemble learning method that builds multiple decision trees and merges them to obtain more accurate and stable predictions. Random Forests are well-suited for our application due to their ability to model complex, non-linear interactions between predictors. The sentiment data we employ captures thematic insights from news, and the relationships between these themes are unlikely to be linear (or, at least, this cannot be ruled out a priori).

Alternatively, we could also train Neural Networks to model complex nonlinear relationships in our data. However, this demands meticulous feature engineering and tuning, making them overly sensitive to pre-processing and prone to overfitting.⁹

3.2 Model estimation and evaluation

To investigate the effects of news-based sentiment on asset prices, we predict daily market returns from a given country using information from a high-dimensional set of sentiment measures. We proceed in three stages. First, we fit Random Forests on the initial 18 months of data (from February 2015 through July 2016) to establish our base models. Second, we use the following 6 months (August 2016–January 2017) as a validation sample to select the optimal hyperparameters¹⁰. Finally, having locked in those hyperparameters, we retrain the model in an expanding-window fashion, adding 20 trading days at a time (roughly one calendar month), so that each new fit incorporates all data from February 2015 up to the end of the previous month.

⁹Grinsztajn, Oyallon, and Varoquaux (2022) argue that tree-based models often outperform Neural Networks on tabular datasets due to their inherent robustness and lower sensitivity to pre-processing.

¹⁰We implement a Random Forest with subagging (subsample aggregating, i.e., without-replacement row sampling) and choose hyperparameters by minimizing validation MSE. Tuned parameters include: (i) the feature-subset fraction per tree ($\in (0, 1]$), (ii) the subagging fraction of observations per tree ($\in (0, 1]$), (iii) the number of trees, and (iv) the maximum number of leaves per tree.

To evaluate our machine learning-based predictive performance, we rely on assessing economic significance via a market timing strategy. Although the out-of-sample R^2 (as in [Campbell and Thompson, 2008](#)) is commonly used in the literature, it is an incomplete measure of the economic significance. A market timer can still generate significant economic profits even when the out-of-sample R^2 is negative (see, e.g., [Kelly, Malamud, and Zhou, 2024](#); [Leitch and Tanner, 1991](#); [Cenesizoglu and Timmermann, 2012](#); [Rapach and Zhou, 2013](#)). To this end, we explore a real-time investment strategy as follows: at day t , we construct an out-of-sample return forecast for day $t + 1$. If the forecast is positive (negative) we take a position of +1 (-1) during the next day $t + 1$ and multiply this position with the realized return on day $t + 1$. The advantage of this approach is that strategy returns are not inflated by implicit leverage, since positions are restricted to -1 and +1.

We compute the annualized Sharpe ratios and net alphas with respect to the market of these real-time investment strategies. When reporting Sharpe ratios, we annualize them by multiplying by $\sqrt{252}$ and compute t-statistics for Sharpe ratios by multiplying them by \sqrt{T} . To obtain the alpha of our trading strategy, we regress the percentage returns of our trading strategy on the market percentage returns. When reporting alphas, we annualize them by multiplying by 252. Our target variables are excess returns obtained from futures instruments. These are known to be highly liquid, with low transaction costs. We account for trading costs in our analysis of net alphas by assuming a fixed transaction cost of 2 basis points (bps) per trade direction for equity futures.

4 Empirical results

Several studies show that news-based sentiment contains predictive information about future market returns (see, e.g., [Tetlock, 2007](#); [Garcia, 2013](#); [Manela and Moreira, 2017](#); [Mamaysky, 2020](#); [Fraiberger, Lee, Puy, and Ranciere, 2021](#)). This literature typically focuses on country-level measures of sentiment. Such aggregation masks the potentially heterogeneous impact of granular sentiment on the predictability of asset returns. In addition, aggregation does not allow for further understanding the sources and dynamics of the return predictability of media sentiment. Our classification of news into thousands of country- and theme-specific categories allows us to study which components are informative about market returns and when. Our data also allows us to pinpoint the source and target entities embedded in news items, enabling us to measure sentiment linkages between countries. By decomposing global news into these components, we can obtain a better understanding of the network structure of global news and its effect on international market returns.

In this section, we study market timing strategies based on predictions constructed from local and global sentiment measures. In addition, we shed light on which theme-specific sentiment measures drive the local market timing ability. Lastly, we show that U.S.-centered sentiment carries an outsized influence in domestic and global contexts.

4.1 Local news and market returns

We first examine the predictive ability of local sentiment measures for local equity market returns. We refer to local news as the coverage produced by a country’s media outlets that predominantly focuses on that same country. [Figure 5](#) shows that such

coverage constitutes the majority of a nation’s news output: between 65% and 85% of articles originating from a given country explicitly reference it, indicating that the source country is also among the target countries discussed. Whereas, previous studies have largely documented this local effect in the United States (see, e.g., [Baker and Wurgler, 2007](#); [Garcia, 2013](#)), we extend the analysis to test its predictive power in equity markets for a sample of 14 countries using 260 theme-specific sentiment measures per local market.

Given the high-dimensionality, we train these theme-specific sentiment measures on one-day ahead local equity market returns using a Random Forest model, as outlined in Section 3. Our setup results in a set of one-day-ahead out-of-sample return forecasts. We examine the economic significance of these forecasts by constructing a real-time trading strategy for each local equity market.

Table 1 presents the out-of-sample Sharpe ratios, *CAPM* alphas and other performance statistics for our local news sentiment strategy. We document strong predictive performance of local news sentiment on future equity returns across all countries. For example, in the U.S., a local news sentiment strategy achieves an out-of-sample Sharpe ratio of 1.34 (t-statistic = 3.59), almost double from the ‘buy and hold’ Sharpe ratio of 0.62. Furthermore, we find sizable *CAPM* alphas (versus the buy and hold benchmark) of 20.41% per annum.

Given that our strategy is rebalanced on a daily basis, transaction costs may substantially diminish the documented profitability. However, even after accounting for transaction costs, the net Sharpe ratio (1.11 for the U.S.) and net alpha (15.98%) remain significantly sizable. In 13 out of the 14 countries in our sample, we find that the net Sharpe ratio is larger than that of holding the underlying instrument.

Overall, we find that average net Sharpe ratio is 82.5% larger than the average benchmark Sharpe ratio. Furthermore, we find that local news sentiment strategies exhibit low betas relative to their underlying markets, ranging from -0.20 (France) to 0.41 (U.S.). In addition, trading on local news sentiment substantially reduces left-tail risk—measured by maximum drawdown—by approximately one third compared to the average buy-and-hold strategy, indicating improved downside protection.

Overall, our results suggest that local news about *the own country* contains significant predictive power for future equity returns. Furthermore, we show that the predictability of news sentiment on equity markets is pervasive and virtually present in all countries in our sample. Our results are consistent with prior empirical literature showing that news sentiment predicts asset prices. From a theoretical perspective, our results align with theories of investor sentiment.

4.2 Interpreting News Local Sentiment

So far, our results indicate that local news sentiment about a country’s own equity market contains significant predictive power for next-day returns. Specifically, our approach employs hundreds of theme-specific sentiment measures within a Random Forest framework. As is often the case with machine learning methods, it can be difficult to explain why they work (Nagel, 2021), particularly when using a large set of predictors. In this section, we shed light on which theme-specific sentiment measures drive the observed return predictability.

To measure the importance of theme-specific sentiment, we employ a gain-based approach within our Random Forest framework. This metric reflects the improvement in predictive accuracy attributable to each feature. Specifically, when a feature is used

to split a decision tree in the Random Forest, it helps reduce the prediction error. The gain metric quantifies the extent to which including that feature improves the model's accuracy at each split. We aggregate these gains across all trees and splits where the feature is used, resulting in an overall measure of its contribution to the model's predictions.

To facilitate comparisons over time, we normalize the raw importance values for each feature so that the sum of importance scores across all features at any given time equals 100%. This normalization allows for direct comparison of each feature's relative significance over time. Additionally, to enable meaningful comparisons across themes, we aggregate the variable importance of the positive and negative sentiment variants within each news-based theme. By summing the importance scores of both sentiment measures for a given theme, we obtain a unified importance score. This approach highlights which themes exert the most significant influence on equity returns over time.

Figure 7 presents the top five theme-specific sentiment ranked by the Random Forest model, that appear more than once in the sample, for the U.S. for a given quarter. This variable importance figure highlights key patterns in the contribution of theme-specific sentiment to forecast accuracy. The graph reveals that the predictive importance of different news themes fluctuates over time, reflecting changing market conditions and the relevance of specific topics at certain moments in time. For instance, health-related sentiment becomes critical during the Covid-19 pandemic, while propaganda appears as significant drivers in late 2022 around the start of the Russian invasion of Ukraine, possibly reflecting geopolitical or social developments affecting market returns. Themes like "Econ Debt" and "Econ Stockmarket" also

show intermittent importance spikes, indicating periods where economic or financial sentiment had stronger predictive power.

Overall, the time-varying importance suggests that the model’s predictive accuracy depends on dynamically adjusting to which news themes matter most at different times, rather than relying on a fixed set of features. This highlights the benefit of a flexible, data-driven approach to understanding how news sentiment influences equity markets.

4.3 Global news and market returns

The evidence reported in Section 4.1 indicates that local news sentiment about the domestic market effectively predicts one-day ahead local equity market returns. Beyond local news about the domestic market, global news coverage represents a potentially powerful source of incremental information, capturing cross-border sentiment linkages, i.e. news from country X reporting on developments in country Y . Specifically, we construct a high-dimensional set of theme-specific sentiment measures that accounts for both incoming and outgoing news from a given country, allowing us to explore the predictive role of global sentiment measures for individual equity markets.

Our local predictor set consists of 260 themes, each separately assessed with positive and negative sentiment scores, totaling 520 local measures per market. To incorporate global informational linkages, we expand our predictor space to include sentiment flows from each market to the other 13 countries and vice versa, yielding a high-dimensional dataset of 14560 sentiment metrics per country. We employ this extensive predictor set within our Random Forest model to generate one-day-ahead return forecasts for each national equity market. These out-of-sample forecasts form

the basis of our real-time trading strategies.

Table 2 summarizes the performance statistics of these global sentiment-based strategies. The empirical results clearly demonstrate that global sentiment measures provide economically substantial and statistically robust predictive power. The strategies consistently yield Sharpe ratios that significantly outperform the buy-and-hold benchmarks across all markets, even after accounting for realistic transaction costs. Importantly, strategy alphas remain economically meaningful post-costs, underscoring the persistent profitability derived from global sentiment information. Moreover, the estimated market betas are generally modest (predominantly below 0.2), indicating that strategy returns reflect distinct informational content embedded in global sentiment flows rather than simple market exposure. The strategies also feature lower maximum drawdowns relative to passive market positions, highlighting their resilience to tail events.

To assess whether the local news sentiment captures the incremental informational content embedded in global sentiment, we conduct spanning regressions. Specifically, for each market, we regress the returns from the global sentiment-based strategy on the returns from the local sentiment-based strategy. A statistically significant and economically large intercept (alpha) from these regressions indicates that global sentiment delivers orthogonal predictive content not encapsulated by local signals alone. Consequently, a significant spanning alpha suggests potential improvements in risk-adjusted portfolio performance when investors integrate global sentiment signals with existing local strategies.

Figure 8 presents the results of these spanning tests, depicting annualized gross spanning alphas and their associated t-statistics. Remarkably, we find significant and

economically large alphas across all markets, except Canada, with annualized magnitudes ranging from roughly 10% to 30%. These findings strongly suggest that international sentiment linkages convey incremental, economically valuable information beyond local sentiment, underscoring their potential for enhancing portfolio performance.

Collectively, these results provide compelling evidence of market integration from an informational perspective. The observed significance of global sentiment measures in timing market returns across nearly all countries illustrates that equity markets worldwide are sensitive to common information flows and sentiment. Such informational integration aligns with the theoretical predictions and empirical findings documented in the broader literature on international financial market linkages, further reinforcing the pivotal role of shared sentiment and information in driving global asset price dynamics.

4.4 Global news sentiment network

Thus far, our analysis does not explicitly identify which countries drive the return predictability observed across international markets. To shed light on these directional influences, we construct a news-based predictability network employing a force-directed visualization framework ([Fruchterman and Reingold, 1991](#)). Specifically, we estimate Random Forest models separately for each country's equity returns using the comprehensive global predictor set and extract gain-based variable importance metrics. These metrics quantify the predictive power of each sentiment linkage, defined by source and target countries, for forecasting returns. To isolate cross-country sentiment effects, we remove local sentiment measures (within-country predictions)

and aggregate the resulting importance scores by country pairs. Subsequently, we standardize these scores within each predictive model, ensuring comparability across countries.

Figure 1 and Figure A1 present our resulting global sentiment network visualization. Several noteworthy patterns emerge from this figure. First, the United States occupies a central and dominant position within the network, exerting substantial predictive influence on virtually all other countries, consistent with its documented role as a leading driver of global financial markets (Rapach, Strauss, and Zhou, 2013). China also emerges as a key country, ranking second in overall importance, suggesting significant bilateral informational flows between the U.S. and China influencing global market predictability.

Second, our visualization reveals pronounced regional clustering, especially within Europe, where the United Kingdom and Germany serve as primary hubs, closely connected with France, Italy, Spain, and Switzerland. This European cluster highlights strong regional informational integration. Conversely, certain markets, including Australia and Sweden, exert a limited cross-border predictive influence.

Third, we identify a distinct Asian cluster featuring China, South Korea, and Japan. While China significantly interacts with the U.S. and other key global players, Japan and South Korea exhibit relatively moderate predictive linkages primarily involving the U.S. Lastly, countries such as Australia and Canada display fewer and weaker connections within the global sentiment network, indicating their peripheral role relative to the U.S. and the major European economies.

Taken together, these results illustrate substantial heterogeneity in cross-country predictive influences, underscoring both the central role of major global players and

significant regional integration patterns in international news-driven equity return predictability.

5 Robustness

We examine the robustness of results from Section 4 by verifying whether our main results persist under various alternative specifications and potential sources of bias. Specifically, we examine the impact of using different sentiment lexicons, controlling for past returns, and accounting for possible translation biases in the measurement of non-English news sentiment.

5.1 Alternative sentiment measures

Our primary results rely on sentiment measures derived from the SentiWordNet 3.0 (SWN) lexicon-based framework. To assess the robustness of our findings to the choice of sentiment lexicon, we replicate the U.S. local effect Random Forest trading strategy using four alternative lexicons: (i) the financial lexicon of [Loughran and McDonald \(2011\)](#) (LM), (ii) the Subjectivity Lexicon ([Wilson, 2005](#)), (iii) ML-SENTICON ([Cruz, Troyano, Pontes, and Ortega, 2014](#)), and (iv) Lexicoder ([Young and Soroka, 2012](#)). The results, summarized in Table 3, indicate that all four lexicons yield broadly similar performance relative to our SWN-based strategy, consistently outperforming the buy-and-hold Sharpe ratio of S&P 500 futures (0.58).

Among the alternative lexicons, the LM lexicon produced the lowest Sharpe ratio (1.15). A possible explanation is that the LM lexicon is specifically designed for financial text and may not be well-suited for analyzing general news articles. Since a

significant portion of the news themes in our dataset extends beyond strictly financial topics, the LM lexicon may fail to capture sentiment nuances effectively.

Overall, these results suggest that our findings are not sensitive to the choice of sentiment lexicon. One reason for this robustness is that we estimate daily sentiment by averaging across a large number of news articles. This aggregation reduces the influence of individual lexicon-specific biases and minimizes potential measurement errors. Consequently, rather than relying on precise sentiment classifications for each individual article, our approach focuses on capturing broader daily sentiment fluctuations, making it less susceptible to discrepancies between lexicons.

5.2 Controlling for past returns

To further test the robustness of our findings, we extend the analysis from Section 4.1, which examined how a country’s local news predicts its market returns, by incorporating the previous day’s return as an additional predictor. Table 4 presents the updated results alongside those from Table 1, allowing for direct comparison.

The findings reveal that in the United States, China, Japan, Canada, and South Korea, adding lagged returns reduces predictive performance. In contrast, most other countries experience modest but consistent improvements in both statistical and economic significance when past returns are included. This pattern suggests that our primary results are not merely capturing short-term return dynamics. Instead, sentiment measures provide unique predictive value that is not fully explained by past market movements.

5.3 Translation bias

Measuring news sentiment often faces challenges when applied to multilingual contexts, as lexicons and sentiment measures may not transfer seamlessly across languages. One common issue is translation bias, where differences in sentiment measurement arise depending on whether texts are analyzed in their native language or translated into another (typically English) before analysis.

To address the translation bias we rely on sentiment scores based on Hedonometer lexicons provided by GDELT for the Chinese, German, French, Korean, and Spanish languages against their English-translated counterparts.¹¹ The Hedonometer is a sentiment analysis tool that assigns words a score along a single “happiness” dimension (Dodds, Harris, Kloumann, Bliss, and Danforth, 2011). We examine to what extent the out-of-sample market timing performance differs when we use sentiment scores (Hedonometer) based on the news in its native language versus translated to English.

In Table 5, we present out-of-sample market timing Sharpe ratios for news sentiment strategies using translated versus native news. Overall, we find that the market timing performance improves in four out of the five languages when the texts are first translated into English before applying the Hedonometer. For example, the Sharpe ratio when using native Spanish news to predict the Spanish equity market equals 0.92 and increases to 1.09 when using Spanish news translated to English. Our findings suggest that translation bias is unlikely to drive our main findings, as sentiment analysis does not uniformly degrade when passing through English translation—rather, it often improves. These results align with prior research that explores the impact of

¹¹GDELT does not provide SWN-based sentiment scores in the native language of the news, but only to the English-translated version of the news.

translation on sentiment analysis performance.

6 Conclusions

In this paper, we develop a novel, high-dimensional approach to measuring news-based sentiment, exploiting both thematic and cross-country perspectives in the Global Database of Events, Language, and Tone (GDELT). By systematically capturing how one country’s media reports on another across 260 distinct topics, our framework provides granular sentiment linkages that extend well beyond traditional one-dimensional or purely local sentiment measures. We demonstrate that these linkages can forecast equity market returns across 14 developed economies, with local sentiment especially influential in North America and Europe, and global (incoming and outgoing) sentiment linkages proving critical in markets such as Japan and China.

Our results underscore the importance of understanding how global information travels and how investors process this information—consistent with theories of market inefficiency and limits to arbitrage. Moreover, we show that these findings are robust to using alternative lexicons, controlling for lagged returns, and accounting for potential translation biases. Taken together, these conclusions highlight the global interplay of economic and geopolitical narratives in shaping market outcomes, emphasizing that additional insights arise when one examines both local sentiment and cross-border news flows. Future work might extend this analysis to firm-level or sector-specific linkages, or to more advanced natural language processing techniques, but our findings here suggest that high-dimensional sentiment measures offer a rich and underutilized source to understand global asset pricing.

References

- Adämmer, Philipp, and Rainer A Schüssler, 2020, Forecasting the equity premium: mind the news!, *Review of Finance* 24, 1313–1355.
- Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani, 2010, SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining, in *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, ed. by Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias Valletta, Malta. European Language Resources Association (ELRA).
- Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor sentiment in the stock market, *Journal of economic perspectives* 21, 129–151.
- Bekaert, Geert, Robert J Hodrick, and Xiaoyan Zhang, 2009, International stock return comovements, *The Journal of Finance* 64, 2591–2626.
- Bhargava, Rajeev, Xiaoxia Lou, Gideon Ozik, Ronnie Sadka, and Travis Whitmore, 2022, Quantifying narratives and their impact on financial markets, *Journal of Portfolio Management*, *Forthcoming*.
- Breiman, Leo, 2001, Random forests, *Machine learning* 45, 5–32.
- Bybee, Leland, Bryan Kelly, and Yinan Su, 2023, Narrative asset pricing: Interpretable systematic risk factors from news text, *The Review of Financial Studies* 36, 4759–4787.
- Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu, 2023, Business News and Business Cycles, *Journal of Finance* *Forthcoming*.
- Calomiris, Charles W, and Harry Mamaysky, 2019, How news and its context drive risk and returns around the world, *Journal of Financial Economics* 133, 299–336.
- Campbell, John Y, Sanford J Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *The Quarterly Journal of Economics* 108, 905–939.
- Campbell, John Y, and Samuel B Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *The Review of Financial Studies* 21, 1509–1531.
- Carrieri, Francesca, Ines Chaieb, and Vihang Errunza, 2013, Do implicit barriers matter for globalization?, *The Review of Financial Studies* 26, 1694–1739.

- Cenesizoglu, Tolga, and Allan Timmermann, 2012, Do return prediction models add economic value?, *Journal of Banking & Finance* 36, 2974–2987.
- Choi, Darwin, Wenxi Jiang, and Chao Zhang, 2025, Alpha go everywhere: Machine learning and international stock returns, *Review of Asset Pricing Studies* p. raaf005.
- Cookson, J Anthony, Runjing Lu, William Mullins, and Marina Niessner, 2024, The social signal, *Journal of Financial Economics* 158, 103870.
- Cruz, Fermín L, José A Troyano, Beatriz Pontes, and F Javier Ortega, 2014, Building layered, multilingual sentiment lexicons at synset and lemma levels, *Expert Systems with Applications* 41, 5984–5994.
- De Jong, Frank, and Frans A De Roon, 2005, Time-varying market integration and expected returns in emerging markets, *Journal of financial economics* 78, 583–613.
- De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann, 1990, Noise trader risk in financial markets, *Journal of political Economy* 98, 703–738.
- Dodds, Peter Sheridan, Kameron Decker Harris, Isabel M Kloumann, Catherine A Bliss, and Christopher M Danforth, 2011, Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter, *PloS one* 6, e26752.
- Dougal, Casey, Joseph Engelberg, Diego Garcia, and Christopher A Parsons, 2012, Journalists and the stock market, *The Review of Financial Studies* 25, 639–679.
- Fraiberger, Samuel P, Do Lee, Damien Puy, and Romain Ranciere, 2021, Media sentiment and international asset prices, *Journal of International Economics* 133, 103526.
- Fruchterman, Thomas MJ, and Edward M Reingold, 1991, Graph drawing by force-directed placement, *Software: Practice and experience* 21, 1129–1164.
- Garcia, Diego, 2013, Sentiment during recessions, *The Journal of Finance* 68, 1267–1300.
- Grinsztajn, Léo, Edouard Oyallon, and Gaël Varoquaux, 2022, Why do tree-based models still outperform deep learning on typical tabular data?, *Advances in neural information processing systems* 35, 507–520.
- Hau, Harald, 2011, Global versus local asset pricing: A new test of market integration, *The Review of Financial Studies* 24, 3891–3940.

- Hellum, Oliver, Lasse Heje Pedersen, and Anders Rønn-Nielsen, 2023, How Global is Predictability? The Power of Financial Transfer Learning, *The Power of Financial Transfer Learning* (November 1, 2023).
- Heston, Steven L, and Nitish Ranjan Sinha, 2017, News vs. sentiment: Predicting stock returns from news stories, *Financial Analysts Journal* 73, 67–83.
- Karolyi, G Andrew, and René M Stulz, 2003, Are financial assets priced locally or globally?, *Handbook of the Economics of Finance* 1, 975–1020.
- Kelly, Bryan, Semyon Malamud, and Kangying Zhou, 2024, The virtue of complexity in return prediction, *The Journal of Finance* 79, 459–503.
- Lee, Hojoon, Gideon Ozik, and Ronnie Sadka, 2023, Narrative Attention Pricing, .
- Leitch, Gordon, and J Ernest Tanner, 1991, Economic forecast evaluation: profits versus the conventional error measures, *The American Economic Review* pp. 580–590.
- Lewis, Karen K, 2011, Global asset pricing, *Annu. Rev. Financ. Econ.* 3, 435–466.
- Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, *The Journal of finance* 66, 35–65.
- Mai, Dat, and Kuntara Pukthuanthong, 2021, Economic narratives and market outcomes: a semi-supervised topic modeling approach, *Available at SSRN 3990324*.
- Mamaysky, Harry, 2020, News and Markets in the Time of COVID-19, *Journal of Financial and Quantitative Analysis* pp. 1–37.
- Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.
- Nagel, Stefan, 2021, *Machine learning in asset pricing*. (Princeton University Press).
- Pukthuanthong, Kuntara, and Richard Roll, 2009, Global market integration: An alternative measure and its application, *Journal of Financial Economics* 94, 214–232.
- Rapach, David, and Guofu Zhou, 2013, Forecasting stock returns, in *Handbook of economic forecasting* (Elsevier,).
- Rapach, David E, Jack K Strauss, and Guofu Zhou, 2013, International stock return predictability: What is the role of the United States?, *The Journal of Finance* 68, 1633–1662.

- Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson, 2022, Measuring news sentiment, *Journal of econometrics* 228, 221–243.
- Tetlock, Paul C, 2007, Giving content to investor sentiment: The role of media in the stock market, *The Journal of finance* 62, 1139–1168.
- Wilson, T, 2005, Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis, in *Proceedings of the conference on human language technology and empirical methods in natural language processing*.
- World Bank, 2023, World Development Indicators: GDP (current US\$), Accessed: 2024-08-06.
- Young, Lori, and Stuart Soroka, 2012, Affective news: The automated coding of sentiment in political texts, *Political communication* 29, 205–231.
- Zhou, Guofu, 2018, Measuring investor sentiment, *Annual Review of Financial Economics* 10, 239–259.
- Zou, Hui, and Trevor Hastie, 2005, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67, 301–320.

Figure 2: English versus Translingual articles over time

The figure shows the proportion of articles in GDELT for each month between February 2015 through December 2023 categorized by English-based articles versus translingual articles.

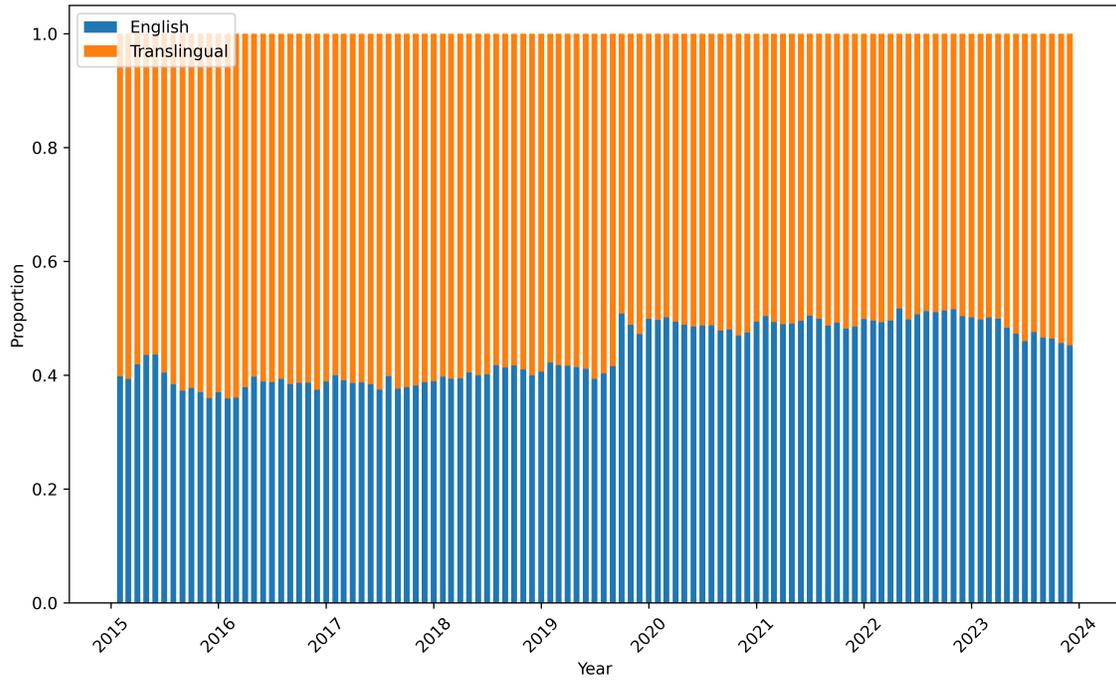


Figure 4: Time-varying coverage of themes

The figure shows the top 20 themes per month in our sample, where the U.S. is the target country. The sample period runs from February 2015 through December 2023.

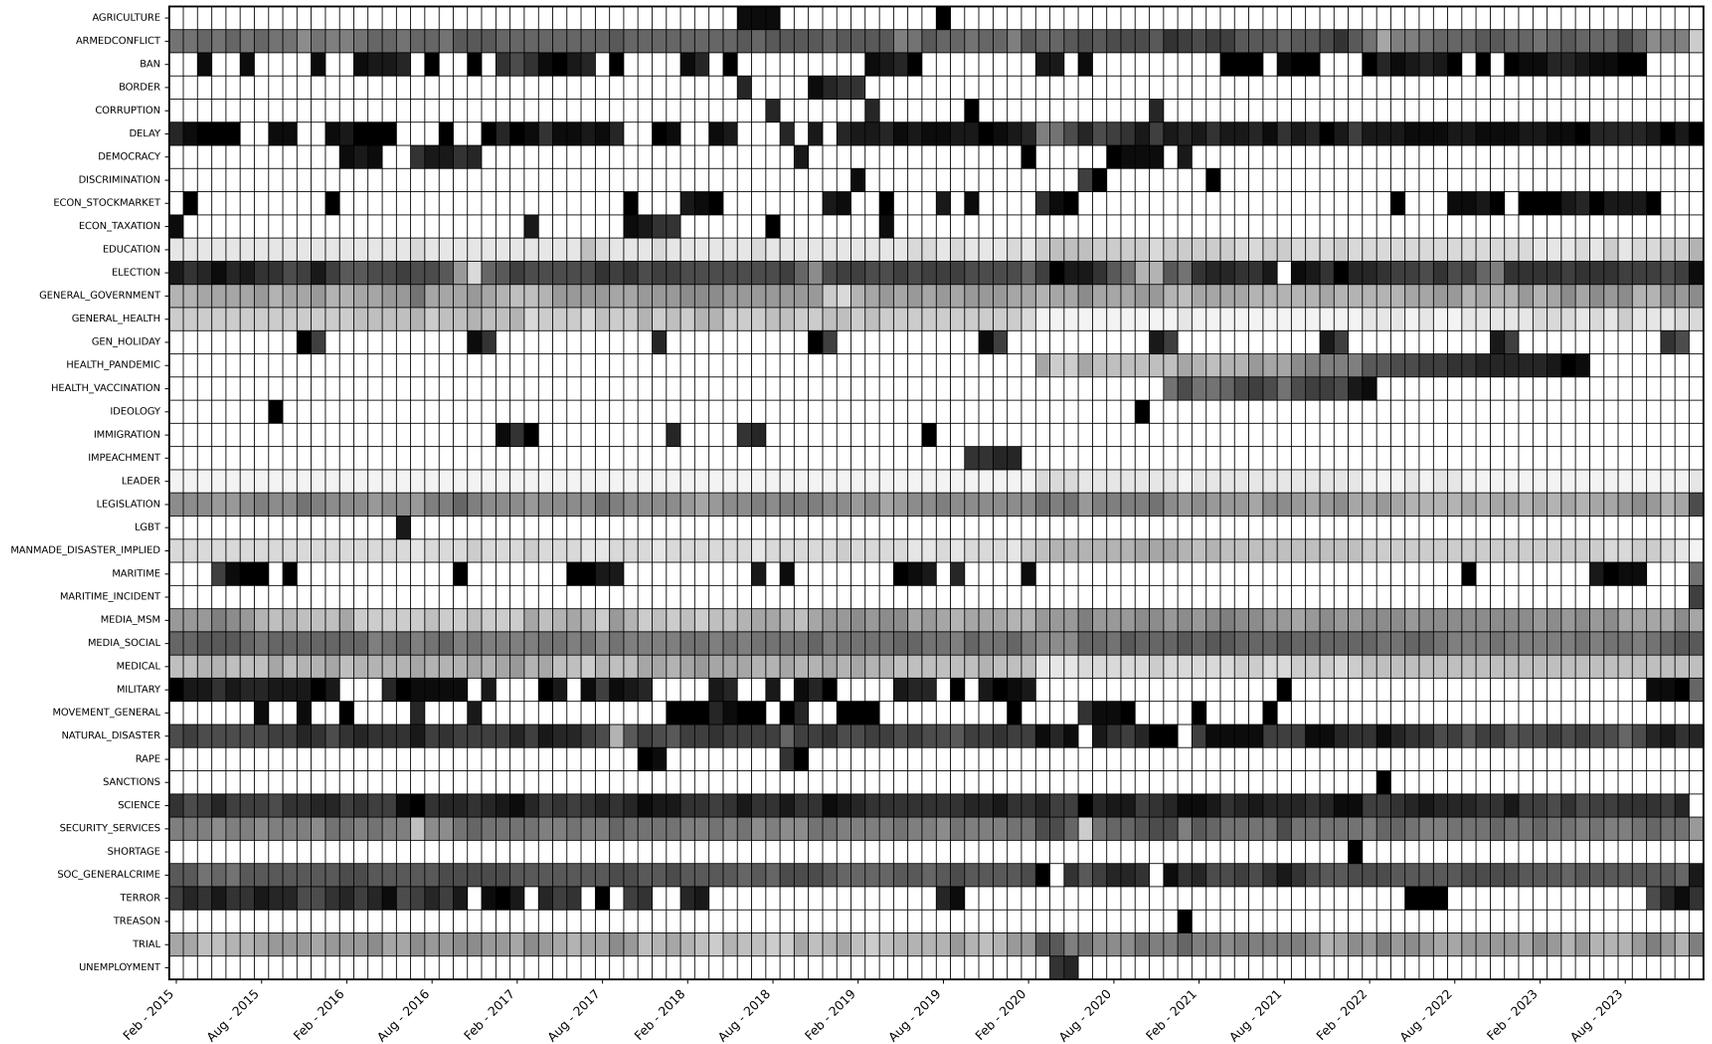


Figure 5: Cross-country news linkages

The figure shows the number of articles that each source country publishes about each target country as a fraction of the total number of news articles published by the source country. The sample period runs from February 2015 through December 2023.

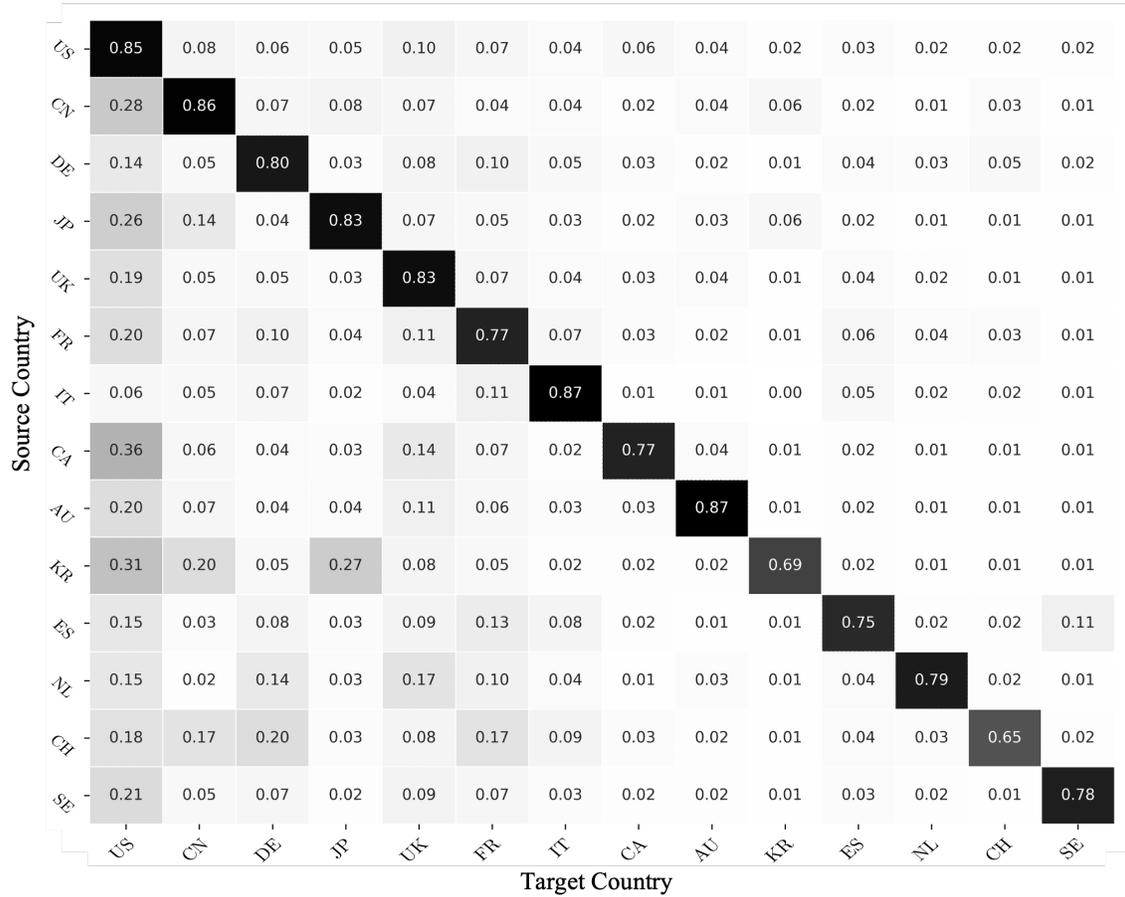


Figure 6: Net attention-weighted sentiment over time

The figure shows the net attention-weighted sentiment (AWS) over time for the "Stock Market" theme (panel A) and the "Health Pandemic" theme (panel B). The AWS series are standardized between -1 and 1 to enhance visual comparison across themes. The sample period runs from February 2015 till December 2023.

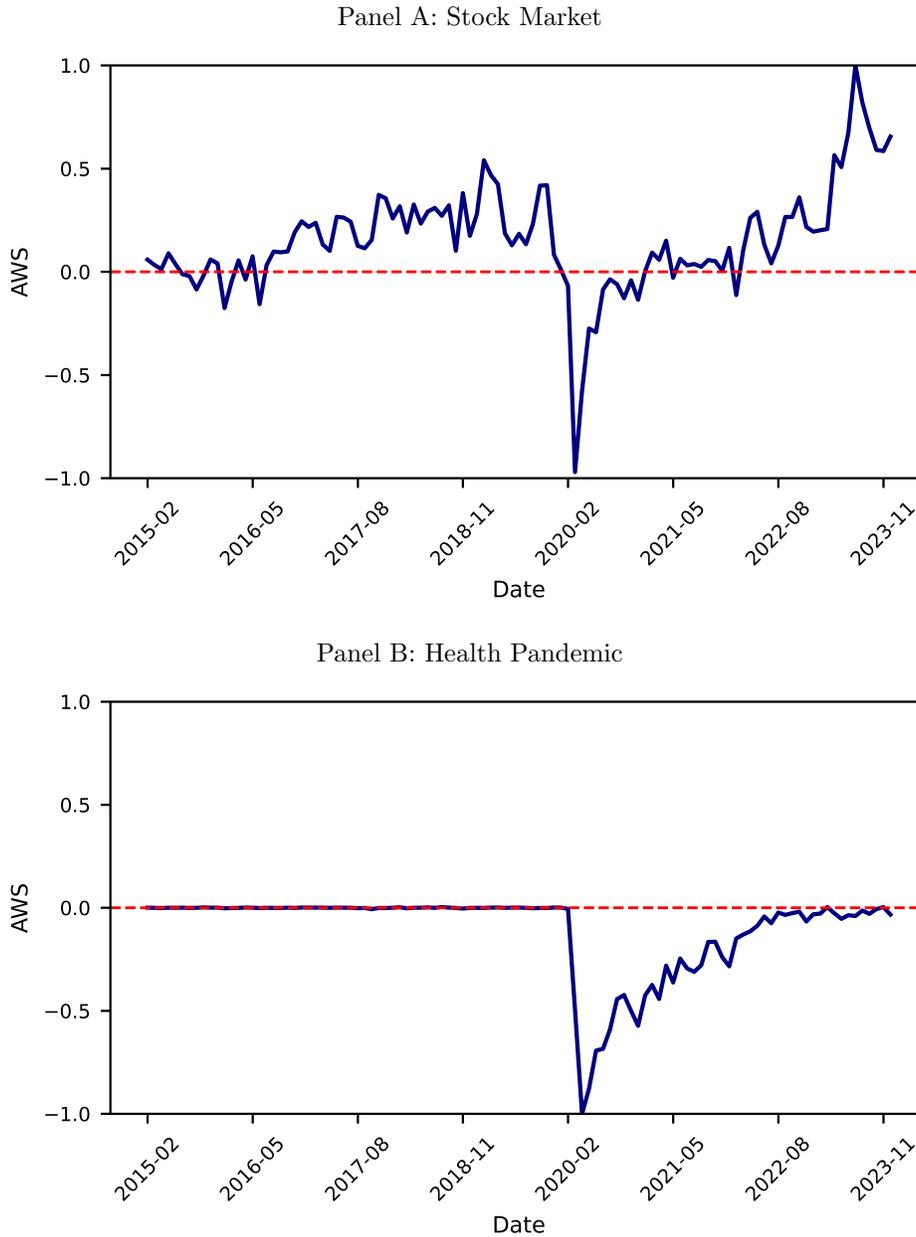


Figure 8: Global versus local: spanning regressions

The figure presents the annualized spanning alpha for each market. We regress the global sentiment strategy return on the local strategy return and report the intercept (annualized in % on the left y-axis) and the corresponding t-statistic (right y-axis). The sample period starts in February 2015. The forecasting period runs from February 2017 till December 2023.

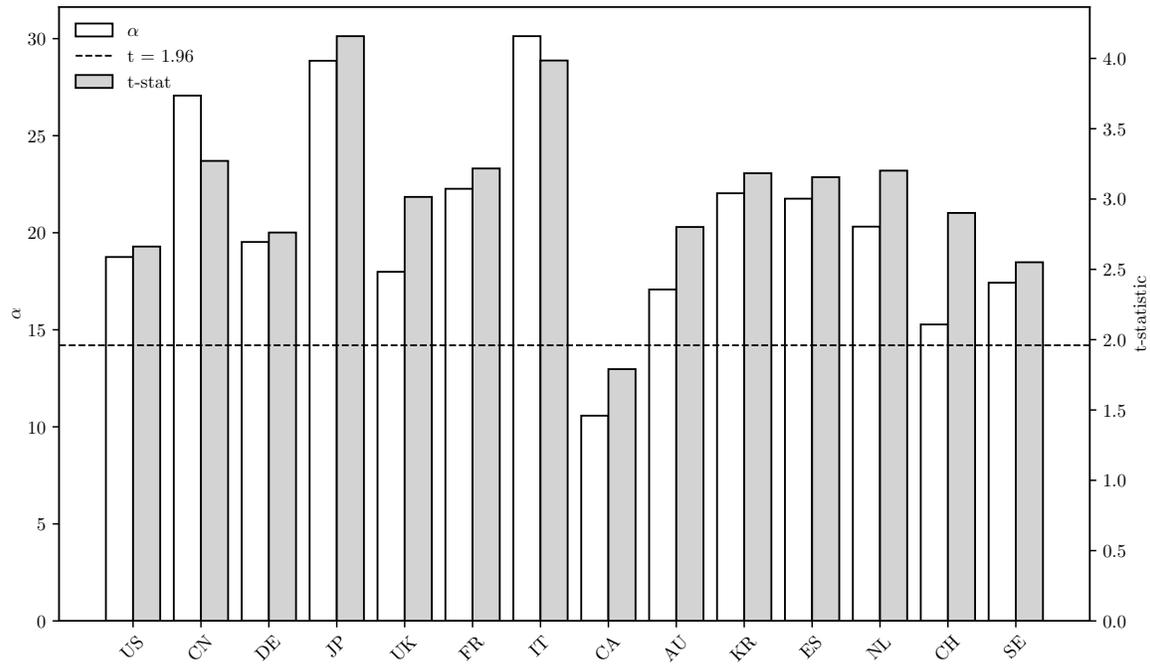


Table 1: Local news sentiment and equity market returns

This table presents performance statistics of out-of-sample local news sentiment strategies for equity markets. We report the Sharpe ratio (gross and net), the Sharpe ratio of the buy-and-hold strategy, CAPM alpha (gross and net), market beta, information ratio and maximum drawdown of the strategy and benchmark. The sample period begins in February 2015, and the forecasting period spans from February 2017 to December 2023.

	SR_{gross}	SR_{net}	$SR_{B\&H}$	α	α_{net}	β	IR	DD_{max}	$DD_{max,B\&H}$
US	1.34 (3.59)	1.11 (2.96)	0.62 (1.65)	20.41 (3.19)	15.98 (2.50)	0.41 (19.27)	0.67	-23.68	-34.15
CN	1.06 (2.83)	0.82 (2.20)	0.01 (0.02)	23.32 (2.84)	18.17 (2.21)	0.12 (5.16)	0.79	-35.05	-50.33
DE	0.99 (2.63)	0.73 (1.94)	0.34 (0.92)	18.43 (2.61)	13.54 (1.92)	0.02 (0.93)	0.46	-37.35	-38.53
JP	0.88 (2.36)	0.61 (1.64)	0.63 (1.68)	14.83 (2.15)	9.78 (1.42)	0.14 (5.88)	0.20	-33.96	-30.80
UK	1.12 (2.99)	0.81 (2.17)	0.29 (0.78)	17.99 (3.02)	13.10 (2.19)	-0.03 (-1.30)	0.57	-18.97	-35.18
FR	1.13 (3.03)	0.85 (2.27)	0.57 (1.52)	23.11 (3.41)	17.88 (2.64)	-0.20 (-8.79)	0.36	-27.77	-38.67
IT	1.03 (2.75)	0.80 (2.13)	0.59 (1.59)	20.46 (2.70)	15.76 (2.08)	0.03 (1.20)	0.31	-24.12	-41.06
CA	0.95 (2.55)	0.70 (1.87)	0.46 (1.22)	14.57 (2.41)	10.38 (1.72)	0.14 (5.82)	0.38	-18.58	-35.88
AU	1.35 (3.60)	1.05 (2.81)	0.48 (1.29)	20.94 (3.46)	16.13 (2.66)	0.13 (5.51)	0.65	-23.85	-36.61
KR	0.78 (2.10)	0.53 (1.40)	0.34 (0.90)	14.50 (2.09)	9.70 (1.39)	0.01 (0.42)	0.32	-20.92	-41.05
ES	1.11 (2.95)	0.83 (2.22)	0.31 (0.84)	20.93 (3.01)	15.82 (2.27)	-0.06 (-2.68)	0.54	-23.93	-39.62
NL	1.01 (2.70)	0.74 (1.97)	0.63 (1.67)	16.48 (2.60)	11.81 (1.86)	0.06 (2.76)	0.28	-19.45	-35.77
CH	1.48 (3.94)	1.13 (3.01)	0.56 (1.48)	21.08 (3.92)	16.05 (2.98)	0.01 (0.61)	0.65	-18.62	-27.11
SE	1.21 (3.24)	0.97 (2.60)	0.57 (1.52)	20.30 (3.02)	15.91 (2.37)	0.17 (7.43)	0.50	-17.94	-31.33

Table 2: Global news sentiment and equity market returns

This table presents performance statistics of out-of-sample local news sentiment strategies for equity markets. We report the Sharpe ratio (gross and net), the Sharpe ratio of the buy-and-hold strategy, CAPM alpha (gross and net), market beta, information ratio and maximum drawdown of the strategy and benchmark. The sample period begins in February 2015, and the forecasting period spans from February 2017 to December 2023.

	SR_{gross}	SR_{net}	$SR_{B\&H}$	α	α_{net}	β	IR	DD_{max}	$DD_{max,B\&H}$
US	1.07 (2.87)	0.80 (2.15)	0.62 (1.65)	18.01 (2.61)	12.96 (1.87)	0.19 (8.03)	0.36	-25.52	-34.15
CN	1.22 (3.27)	0.99 (2.64)	0.01 (0.02)	26.98 (3.27)	21.79 (2.64)	0.03 (1.40)	0.87	-27.83	-50.33
DE	1.02 (2.72)	0.76 (2.02)	0.34 (0.92)	18.26 (2.62)	13.33 (1.91)	0.15 (6.27)	0.52	-24.85	-38.53
JP	1.57 (4.18)	1.31 (3.50)	0.63 (1.68)	27.30 (3.98)	22.57 (3.29)	0.14 (6.20)	0.71	-19.07	-30.80
UK	1.21 (3.23)	0.87 (2.34)	0.29 (0.78)	18.53 (3.14)	13.19 (2.24)	0.16 (6.74)	0.70	-16.25	-35.18
FR	1.28 (3.43)	1.01 (2.70)	0.57 (1.52)	21.85 (3.21)	16.77 (2.46)	0.18 (7.74)	0.56	-16.34	-38.67
IT	1.55 (4.15)	1.31 (3.50)	0.59 (1.59)	31.33 (4.14)	26.43 (3.50)	0.00 (0.01)	0.67	-18.55	-41.06
CA	0.90 (2.40)	0.62 (1.65)	0.46 (1.22)	11.96 (2.10)	7.40 (1.30)	0.36 (16.43)	0.39	-14.72	-35.88
AU	1.18 (3.16)	0.89 (2.37)	0.48 (1.29)	20.21 (3.33)	15.39 (2.54)	-0.12 (-5.14)	0.47	-13.97	-36.61
KR	1.25 (3.35)	0.97 (2.58)	0.34 (0.90)	21.98 (3.23)	16.66 (2.45)	0.20 (8.50)	0.72	-31.19	-41.05
ES	1.34 (3.57)	1.04 (2.77)	0.31 (0.84)	25.37 (3.66)	19.85 (2.86)	-0.09 (-4.00)	0.69	-26.05	-39.62
NL	1.25 (3.35)	0.95 (2.55)	0.63 (1.67)	21.10 (3.32)	16.02 (2.52)	0.01 (0.51)	0.44	-24.15	-35.77
CH	1.39 (3.71)	1.06 (2.82)	0.56 (1.48)	18.56 (3.51)	13.77 (2.60)	0.18 (7.57)	0.65	-17.00	-27.11
SE	1.04 (2.77)	0.79 (2.11)	0.57 (1.52)	18.45 (2.70)	13.95 (2.04)	0.05 (1.97)	0.34	-19.34	-31.33

Table 3: Alternative sentiment measures

This table reports the annualized Sharpe ratio of a out-of-sample local news sentiment strategy applied to the United States using a range of sentiment measures. The out-of-sample period runs from February 2017 till December 2023.

	Sharpe Ratio	t-stat
SentiWordNet3	1.34	3.59
Loughran McDonald	1.15	3.07
Subjectivity Lexicon	1.31	3.50
ML-SENTICON	1.34	3.58
Lexicoder	1.33	3.56

Table 4: Local news and accounting for lagged returns

This table presents performance statistics of out-of-sample local news sentiment strategies for equity markets, accounting for lagged returns. We account for lagged returns by incorporating a one-day lagged return in the set of predictors. We report the Sharpe ratio (gross and net), the Sharpe ratio of the buy-and-hold strategy, CAPM alpha (gross and net), market beta, information ratio and maximum drawdown of the strategy and benchmark. The sample period begins in February 2015, and the forecasting period spans from February 2017 to December 2023.

	SR_{gross}	SR_{net}	$SR_{B\&H}$	α	α_{net}	β	IR	DD_{max}	$DD_{max,B\&H}$
US	1.50 (4.01)	1.26 (3.36)	0.62 (1.65)	25.96 (3.77)	21.38 (3.10)	0.19 (8.10)	0.69	-21.73	-34.15
CN	0.88 (2.36)	0.66 (1.76)	0.01 (0.02)	19.49 (2.36)	14.55 (1.76)	-0.07 (-2.91)	0.60	-43.38	-50.33
DE	0.57 (1.52)	0.31 (0.82)	0.34 (0.92)	10.48 (1.48)	5.58 (0.79)	0.03 (1.42)	0.16	-47.52	-38.53
JP	1.19 (3.18)	0.91 (2.44)	0.63 (1.68)	20.78 (3.01)	15.64 (2.26)	0.11 (4.69)	0.42	-21.66	-30.80
UK	0.75 (2.02)	0.43 (1.15)	0.29 (0.78)	11.37 (1.92)	6.18 (1.04)	0.14 (6.21)	0.35	-26.36	-35.18
FR	1.32 (3.54)	1.07 (2.86)	0.57 (1.52)	25.41 (3.69)	20.73 (3.01)	-0.09 (-3.84)	0.51	-22.08	-38.67
IT	1.12 (2.99)	0.89 (2.37)	0.59 (1.59)	22.62 (2.98)	17.96 (2.37)	-0.00 (-0.03)	0.37	-25.05	-41.06
CA	0.75 (2.00)	0.48 (1.29)	0.46 (1.22)	10.68 (1.78)	6.36 (1.06)	0.20 (8.82)	0.23	-32.69	-35.88
AU	1.09 (2.91)	0.80 (2.13)	0.48 (1.29)	15.98 (2.69)	11.22 (1.89)	0.23 (9.98)	0.49	-27.19	-36.61
KR	0.99 (2.66)	0.72 (1.92)	0.34 (0.90)	17.89 (2.59)	12.75 (1.84)	0.09 (3.81)	0.49	-18.65	-41.05
ES	1.27 (3.40)	1.02 (2.72)	0.31 (0.84)	22.97 (3.32)	18.30 (2.65)	0.11 (4.79)	0.72	-22.00	-39.62
NL	1.09 (2.91)	0.81 (2.17)	0.63 (1.67)	17.54 (2.77)	12.84 (2.03)	0.09 (3.63)	0.34	-25.87	-35.77
CH	1.35 (3.60)	1.01 (2.69)	0.56 (1.48)	17.74 (3.37)	12.88 (2.44)	0.20 (8.71)	0.62	-12.42	-27.11
SE	1.33 (3.55)	1.07 (2.87)	0.57 (1.52)	22.29 (3.32)	17.66 (2.63)	0.18 (7.88)	0.59	-17.76	-31.33

Table 5: Native versus translated local news

This table presents performance statistics of out-of-sample local news sentiment strategies for equity markets using either translated (to English) news or native news. We report the Sharpe ratio, CAPM alpha, market beta, and information ratio of the strategies. The columns with (ENG) indicate that the news was first translated to English and then the sentiment was calculated using the English version of Hedonometer. The columns without the (ENG) tag indicate that the sentiment was calculated directly using the native language version of Hedonometer. The out-of-sample period spans from February 2017 to December 2023.

	SR_{gross}^{ENG}	SR_{gross}	α^{ENG}	α	IR^{ENG}	IR
CN	1.09 (2.91)	1.12 (2.99)	24.03 (2.91)	24.73 (3.00)	0.77	0.77
DE	1.44 (3.84)	1.22 (3.26)	25.54 (3.72)	22.73 (3.23)	0.87	0.63
FR	0.86 (2.29)	0.88 (2.36)	14.55 (2.11)	13.80 (2.05)	0.22	0.25
KR	1.44 (3.85)	1.28 (3.42)	25.66 (3.75)	22.98 (3.33)	0.85	0.70
ES	0.95 (2.54)	0.92 (2.45)	17.91 (2.57)	17.91 (2.60)	0.44	0.40

Appendices

Table A1: Sample of equity futures contracts

Reported are the country name ("Country"), local close time ("Close"), timezone ("TZ"), name of the market index ("Index"), the Bloomberg ticker ("BB"), and the annualized Sharpe ratio ("Sharpe"). All data are obtained from Bloomberg from February 2015 till December 2023.

Country	Close	TZ	Index	BB	Sharpe
United States	17:00	America/Chicago	S&P 500	ES	0.58
China	16:30	Asia/Hong_Kong	HSI	HI	0.31
Germany	22:05	Europe/Berlin	DAX	GX	0.30
Japan	15:16	Asia/Tokyo	Nikkei 225	NK	0.51
United Kingdom	21:00	Europe/London	FTSE 100	Z	0.31
France	22:00	Europe/Paris	CAC 40	CF	0.49
Italy	22:00	Europe/Rome	FTSE MIB	ST	0.43
Canada	16:30	America/Montreal	TSX 60	PT	0.42
Australia	16:30	Australia/Sydney	ASX SPI 200	XP	0.38
South Korea	15:46	Asia/Seoul	KOSPI 200	KM	0.33
Spain	20:00	Europe/Madrid	IBEX 35	IB	0.22
Netherlands	22:00	Europe/Amsterdam	AEX	EO	0.55
Switzerland	22:05	Europe/Berlin	SMI	SM	0.44
Sweden	17:25	Europe/Stockholm	OMXS 30	QC	0.46

Figure A1: Global Importance over other countries

The figure shows the sum of standardized variable importance of a given country (on the x-axis) over other countries in the global model estimation. These sums are used to scale node sizes in Figure 1.

